Arcserve® Unified Data Protection Solutions Guide

Version 8.0

arcserve®

Legal Notice

This Documentation, which includes embedded help systems and electronically distributed materials, (hereinafter referred to as the “Documentation”) is for your informational purposes only and is subject to change or withdrawal by Arcserve at any time. This Documentation is proprietary information of Arcserve and may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of Arcserve.

If you are a licensed user of the software product(s) addressed in the Documentation, you may print or otherwise make available a reasonable number of copies of the Documentation for internal use by you and your employees in connection with that software, provided that all Arcserve copyright notices and legends are affixed to each reproduced copy.

The right to print or otherwise make available copies of the Documentation is limited to the period during which the applicable license for such software remains in full force and effect. Should the license terminate for any reason, it is your responsibility to certify in writing to Arcserve that all copies and partial copies of the Documentation have been returned to Arcserve or destroyed.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, ARCSERVE PROVIDES THIS DOCUMENTATION “AS IS” WITHOUT WARRANTY OF ANY KIND, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT WILL ARCSERVE BE LIABLE TO YOU OR ANY THIRD PARTY FOR ANY LOSS OR DAMAGE, DIRECT OR INDIRECT, FROM THE USE OF THIS DOCUMENTATION, INCLUDING WITHOUT LIMITATION, LOST PROFITS, LOST INVESTMENT, BUSINESS INTERRUPTION, GOODWILL, OR LOST DATA, EVEN IF ARCSERVE IS EXPRESSLY ADVISED IN ADVANCE OF THE POSSIBILITY OF SUCH LOSS OR DAMAGE.

The use of any software product referenced in the Documentation is governed by the applicable license agreement and such license agreement is not modified in any way by the terms of this notice.

The manufacturer of this Documentation is Arcserve.

Provided with “Restricted Rights.” Use, duplication or disclosure by the United States Government is subject to the restrictions set forth in FAR Sections 12.212, 52.227-14, and 52.227-19(c)(1) - (2) and DFARS Section 252.227-7014(b)(3), as applicable, or their successors.

© 2014-2021 Arcserve, including its affiliates and subsidiaries. All rights reserved. Any third party trademarks or copyrights are the property of their respective owners.

Arcserve Product References

This document references the following Arcserve products:

- Arcserve® Unified Data Protection
- Arcserve® Unified Data Protection Agent for Windows
- Arcserve® Unified Data Protection Agent for Linux
- Arcserve® Backup
- Arcserve® Continuous Availability

Contact Arcserve Support

The Arcserve Support team offers a rich set of resources for resolving your technical issues and provides easy access to important product information.

Contact Support

With Arcserve Support:

- You can get in direct touch with the same library of information that is shared internally by our Arcserve Support experts. This site provides you with access to our knowledge-base (KB) documents. From here you easily search for and find the product-related KB articles which contain field-tested solutions for many top issues and common problems.
- You can use our Live Chat link to instantly launch a real-time conversation between you and the Arcserve Support team. With Live Chat, you can get immediate answers to your concerns and questions, while still maintaining access to the product.
- You can participate in the Arcserve Global User Community to ask and answer questions, share tips and tricks, discuss best practices and participate in conversations with your peers.
- You can open a support ticket. By opening a support ticket online, you can expect a callback from one of our experts in the product area you are inquiring about.
- You can access other helpful resources appropriate for your Arcserve product.

Contents

Chapter 1: Features and Enhancements 27
New Features 28
Feature Enhancements 30
Linux Agent Enhancements 31
Database and Platform Support 32
Security and Third-party Enhancements 33
Arcserve Appliance Enhancements 34
Dropped Support 35
Chapter 2: Understanding Arcserve UDP 37
Introduction 38
How Arcserve UDP Works 40
Instant Virtual Machine (IVM) versus Virtual Standby (VSB) Machine 42
User Security 45
Roles for Arcserve UDP Services 46
User Privileges for Arcserve UDP Functions 47
Chapter 3: Installing Arcserve UDP 55
How to Install Arcserve UDP 56
Review the Prerequisites and Considerations 57
Decide the Installation Type 58
Install Arcserve Unified Data Protection Using the Setup Wizard 59
Install Arcserve Unified Data Protection Using the Command Line 63
Install Arcserve UDP Using the Unified Installer 67
Verify the Installation 69
Communication Ports Used by Arcserve UDP 70
How the Installation Process Affects Operating Systems 86
How to Install Arcserve UDP Updates 92
Review the Considerations for Installing Updates 93
Specify Updates Preference 94
Check and Install the Updates 100
Update Agents on Remote Nodes Using Gateway 101
Update RPS on Remote Server Using Gateway 102
(Optional) Install Arcserve UDP Updates Silently 103
Verify that the Updates are Successfully Installed 104
How to Uninstall Arcserve UDP 105

Standard Uninstall 106
Silent Uninstall 107
Remove Components Left Behind by the Uninstaller 109
Chapter 4: Upgrading to Arcserve UDP Version 8.0 111
Supported Versions for Upgrade 112
Upgrade Sequence on the Arcserve Appliance 113
Upgrade Arcserve Appliance that performs as Arcserve UDP Console and RPS 114
Upgrade Arcserve Appliance that performs as Arcserve UDP RPS only 115
Upgrade Steps When Two or More Arcserve Appliances Are Used in the Environment 116
Upgrade Arcserve UDP Linux Agent on Arcserve Appliance 117
Upgrade Arcserve Backup on Arcserve Appliance 118
Backward Compatibility Support Policy 119
Backward Compatibility Support Policy for Arcserve UDP Version 7.x 120
Backward Compatibility Support Policy for Arcserve UDP Version 6.5 Update 4 123
Backward Compatibility Support for Linux Backup Server 126
Upgrade Sequence for UDP Console, RPS, and Agent 127
How to Upgrade to Arcserve UDP 8.0 from a Previous Release 128
How to Upgrade to Arcserve UDP 8.0 Using the Single Installer 131
How to Upgrade to Arcserve UDP 8.0 or Arcserve Backup 18.0 from Previous Version of Arcserve UDP or Arcserve Backup to Enable Copy to Tape 135
How to Migrate Arcserve UDP Console from One Server to Another 137
How to Upgrade Gateway to the Same Version as its Registered Console 138
How to Perform Arcserve UDP Console Migration 138
Best Practices to protect Arcserve UDP server 141
Chapter 5: Exploring and Configuring Arcserve UDP 143
Arcserve Unified Data Protection User Interface 144
Navigating Arcserve UDP 145
Tabs 146
Job Monitor Dialog 161
How to Configure Arcserve UDP 162
Configure Server Communication Protocol 163
Configure Database 165
Configure Arcserve Backup Data Synchronization 169
Configure SRM 170
Node Discovery Configuration 172
Configure Email and Alert 174
Configure Proxy Settings 176

Update Configuration 178
Configure Administrator Account 180
Remote Deployment Settings 181
Map the Plan to the User Account 182
User Management 184
Configure the Console Timeout Duration 186
How to Migrate Arcserve r16.5 Recovery Points to Arcserve UDP 187
Create a Data Store to Replicate Data from an Arcserve r16.5 Recovery Point 188
Replicate Arcserve r16.5 Data to the UDP Data Store 189
Set up Configuration Wizard 190
Chapter 6: Working With Key Features of Arcserve UDP 193
Understanding Nutanix Feature in Arcserve UDP 194
Understanding OneDrive Feature in Arcserve UDP 195
Understanding Exchange Online Feature in Arcserve UDP 196
Understanding Hardware Snapshot Feature in UDP 197
Understanding Cloud Feature in UDP 198
Understanding UNC/NFS Path Feature in UDP 200
Understanding SharePoint Online Feature in UDP 201
Prerequisites 201
Chapter 7: Using Arcserve UDP Role-based Administration 203
Access the User Management Console 204
Add User, Delete User, and Change Password 206
Configure User Management 207
Pre-defined Roles 208
Assign a Pre-defined Role 212
View and Cancel a Role 213
Add a New Role 214
Search Users and Roles 216
Integrate Arcserve UDP with Active Directory 217
How to Integrate Arcserve UDP 8.0 with Active Directory Using Windows Groups 218
How to Integrate Arcserve UDP 8.0 with Active Directory Using Active Directory Groups 221
Access Arcserve UDP Using Integrated Windows authentication 228
Troubleshooting for Integrated windows authentication (IWA) 229
Arcserve UDP Console does not support IWA if Console Database Connection uses Windows Authentication Mode 231
Troubleshooting 233
Chapter 8: Adding and Managing Source Nodes 235

How to Add Nodes to the Console 236
Review the Prerequisites 237
Add Nodes 238
Discover Nodes 243
Import Nodes 245
How to Manage Nodes 251
Review the Prerequisites 252
Update Hypervisor Information 253
Specify the Hypervisor 254
Update VM Information 256
Update Nodes 257
Export Node 261
Pause Node 262
Resume Node 263
Synchronize Data 264
Delete Nodes from the Console 265
Deploy Agent to Nodes 266
Perform Preflight Checks for Your Backup Jobs 267
How to Add and Manage UNC/NFS Path 281
Add UNC/NFS Path 282
Update UNC/NFS Path 285
Delete UNC/NFS Path 286
How to Add and Manage Node Groups 287
Review the Prerequisites 289
Add Node Groups 290
Modify Node Groups 291
Delete Node Groups 292
How to Manage Nodes for Cloud 293
Download Recovery Point from Cloud 294
Upload Recovery Point to Cloud 296
Copy Recovery Point to Local Disk or Network Share 299
Add a Storage Array 301
Add a Storage Array for NetApp 302
Add Details of HPE RMC Managing HPE 3PAR Storeserv storage array 307
Add a Nimble Storage Array 309
How to Add and Manage a Site 311

Specify the Site Name ... 313
Share the Registration Instruction ... 314
Verify and Add the Site ... 315
Modify the Console URL ... 316
Modify the Site ... 317
Delete a Site ... 318
Set up a Proxy Server for the Gateway ... 319
Managing Exchange Online Nodes ... 322
Add an Exchange Online Node ... 323
Update an Exchange Online Node ... 325
Delete an Exchange Online Node ... 326
Public Folder Mailbox Support for Exchange Online Protection ... 327
Manage SharePoint Online Nodes ... 328
Add a SharePoint Online Node ... 329
Update a SharePoint Online Node ... 330
Delete a SharePoint Online Node ... 331
Manage OneDrive Node ... 331
Add a OneDrive Node ... 332
Delete a OneDrive Online Node ... 332
Chapter 9: Adding and Managing Destinations ... 333
How to Add a Destination ... 334
How to Add a Destination ... 335
Add Arcserve Backup Servers ... 347
Add a Remote Console ... 349
Add a Cloud Account ... 351
Add an Arcserve Cloud Account ... 383
How to Manage a Data Store ... 386
Review the Prerequisites ... 387
Modify a Data Store ... 388
Start a Data Store ... 397
Stop a Data Store ... 398
Delete a Data Store from the Console ... 399
Delete Node Data from a Data Store ... 400
Monitor Data Store Space Capacity ... 401
Browse Recovery Points in a Data Store ... 408
FQDN Support for a Data Store ... 409

Run a Manual or On-demand Merge Job 410
Modify Concurrent Active Nodes Limit for Manual Backup 411
Troubleshooting: How to Use a Data Store When One or Multiple folders are Full 412
How to Manage a Recovery Point Server 414
Review the Prerequisites 415
Update a Recovery Point Server 416
Delete a Recovery Point Server from the Console 417
Import a Data Store 418
Install/Upgrade Recovery Point Server 420
How to Manage Arcserve Backup Servers 421
Update an Arcserve Backup Server 422
Delete an Arcserve Backup Server 423
How to Manage Shared Folder 424
Modify a Shared Folder 425
Delete a Shared Folder 426
How to Manage a Remote Console 427
Modify a Remote Console 428
Delete a Remote Console 429
Test Remote Console Connection 430
Chapter 10: Creating Plans to Protect Data 433
Understanding Plan and Tasks 434
How to Create a Windows Backup Plan 438
Review the Prerequisites and Considerations 439
Create a Backup Plan with a Backup Task 449
(Optional) Perform a Manual Backup 469
Verify the Backup 471
How to Create a Nutanix Backup Plan 472
How to Create a Virtual Standby to Nutanix AHV 473
How to Create an Instant Virtual Machine on Nutanix AHV 473
How to Create an Assured Recovery Plan to Protect Nutanix AHV for Linux Node 473
How to Perform Backup of Cluster Shared Volume 473
Modify Agent Machine Settings to Backup CSV Volume 474
How to Perform Backup of Volumes without Driver Letter 476
Review Prerequisites 476
Mount the Volume without Drive Letter in NTFS Folder 476
Specify Volume without Drive Letter in Registry 477

How to Perform Backup of Selected Volume 478
How to Create a Linux Backup Plan 479
Review the Prerequisites and Considerations 480
Create a Backup Plan 481
(Optional) Perform a Manual Backup 501
Verify the Backup 503
Troubleshooting 504
How to Create a Host-Based Virtual Machine Backup Plan 506
Review the Prerequisites and Considerations 508
Create a Host-Based Backup Plan 513
(Optional) Perform a Manual Backup 548
Verify the Plan 550
How to Create a Virtual Standby to AWS EC2 Plan 551
Review the Prerequisites and Considerations 552
Create a Plan with a Backup Task 560
Add a Virtual Standby to EC2 Task to the Plan 575
(Optional) Run the Virtual Standby to EC2 Task Manually 581
Pause and Resume Virtual Standby Job 582
Verify the Plan 583
Terminate EC2 Resources 583
How to Create a Virtual Standby to Microsoft Azure Plan 585
Review the Prerequisites and Considerations 586
Create a Plan with a Backup Task 587
Add a Virtual Standby Task to the Plan 588
(Optional) Run the Virtual Standby Job Manually 595
Pause and Resume Heartbeat 596
Pause and Resume Virtual Standby Job 597
Verify the Plan 598
How to Create a Virtual Standby Plan 599
Review the Prerequisites and Considerations 600
Create a Plan with a Backup Task 602
Add a Virtual Standby Task to the Plan 617
(Optional) Run the Virtual Standby Job Manually 631
Pause and Resume Heartbeat 632
Pause and Resume Virtual Standby Job 633
Verify the Plan 634

Applying Best Practices 635
How to Create an Assured Recovery Plan 637
Review the Prerequisites and Considerations 638
Add an Assured Recovery Task to the Plan 639
(Optional) Run the Assured Recovery Test Job Manually 653
How to Create a UNC/NFS Path Backup Plan 655
Review the Prerequisites and Considerations 656
Create a File Share 656
Create a Backup Plan with a UNC/NFS Path Task 661
(Optional) Perform a Manual Backup 681
Other tasks on UNC/NFS Plan 683
Verify the Backup 684
How to Protect Virtual Standby Machines 685
Review the Prerequisites and Considerations 686
Power On Virtual Standby Machines 687
Protect Virtual Standby Machines After it is Powered On 693
Verify if the Virtual Standby Machine is Protected 694
How to Protect Instant Virtual Machines 695
Review the Prerequisites and Considerations 696
Power On Instant Virtual Machines from Recovery Point 697
Protect an Instant Virtual Machine when Powered On 698
Verify if the Instant Virtual Machine is Protected 699
How to Replicate Data Between Data Stores Managed from a UDP Console 700
Review the Prerequisites and Considerations 701
Create a Backup Task 702
Add a Replicate Task to the Plan 703
(Optional) Perform a Manual Replication 706
Verify the Plan 707
How to Replicate Data Between Data Stores Managed From Different UDP Consoles 708
Review the Prerequisites 711
Create a User Account for the Source Console 712
Create a Plan to Define the Destination Data Store 713
Map the Plan to the User Account 716
Send the Plan and User Account Details to the Source Administrator 718
Receive the Plan and User Account Details from the Destination Administrator 719
Create a Replication Plan to Send Data to the Destination Console 720

Verify the Data is Replicated 724
Applying Best Practices 725
How to Perform an Offline Data Replication Using RPS Jumpstart 728
Review the Prerequisites 731
Create a Temporary Data Store on an External Device 732
Replicate Source Data to the Temporary Data Store 733
Delete the Temporary Data Store from the Source Console 734
Send the External Device to the Destination Location 735
Receive the External Device 736
Import the Temporary Data Store from the External Device 737
Create a Destination Data Store 738
Replicate Data from the Temporary Data Store to the Destination Data Store 739
Verify that the Data is Replicated 740
(Optional) Set the Concurrent Node Count for RPS Jumpstart 741
How to Perform Reverse Replication 741
Review the Prerequisites 741
Perform Ad-hoc Replication 742
Perform Reverse Replication from Remote Console to On-Premises RPS Server 749
How to Create a Copy Recovery Points Plan 756
Review the Prerequisites and Considerations 757
Create a Plan with a Backup Task 758
Add a Copy Recovery Points Task to the Plan 773
Verify the Plan 780
How to Create a File Copy Plan 781
Review the Prerequisites and Considerations 782
Create a Plan with a Backup Task 783
Add a File Copy Task to the Plan 797
(Optional) Perform a Manual File Copy 807
Verify the Plan 808
How to Create a File Archive Plan 809
Review the Prerequisites and Considerations 810
Create a Plan with a Backup Task 811
Add a File Archive Task to the Plan 825
(Optional) Perform a Manual File Archive 836
Verify the Plan 837
How to Create a Copy to Tape Plan 838

Review the Prerequisites and Considerations 839
Create a Plan with a Backup Task 840
Add a Copy to Tape Task to the Plan 841
How to Create and Manage an Instant Virtual Machine on Microsoft Azure 849
Best Practices for an Instant Virtual Machine on Microsoft Azure 849
Review the Prerequisites for an Instant Virtual Machine on Microsoft Azure 851
Create an Instant Virtual Machine on Microsoft Azure 852
Manage an Instant Virtual Machine 861
How to Create and Manage an Instant Virtual Machine on Hyper-V and VMware ESX Servers 864
Review the Prerequisites for an Instant Virtual Machine 866
Create an Instant Virtual Machine 869
Manage an Instant Virtual Machine 885
How to Create and Manage an Instant Virtual Machine on Amazon EC2 891
Review the Prerequisites for an Instant Virtual Machine on Amazon EC2 892
Create an Instant Virtual Machine on Amazon EC2 893
Manage an Instant Virtual Machine 902
How to Create a Plan for Replication Across Sites 908
Create a Plan for Replication Across Sites 910
How to Protect Microsoft Office 365 Data 912
Modern Authentication 912
How to Create an Exchange Online Backup Plan 916
How to Create a SharePoint Online Backup Plan 961
How to Create a Microsoft Office 365 OneDrive Backup Plan 989
Managing Microsoft Teams Backup and Restore 1007
Review the Prerequisites and Considerations 1010
Add the Required Role and Group to the Exchange Online Backup Account to Perform Backup and Restore 1012
Create an Exchange Online Backup Plan 1014
Specify the Source 1016
Specify the Destination 1024
Specify the Schedule 1025
Understanding Advanced Scheduling and Retention 1028
Specify the Advanced Settings 1030
Review the Prerequisites 1050
Create a SharePoint Online Backup Plan 1051
Specify the Source 1053

Specify the Destination 1057
Specify the Schedule 1058
Specify the Advanced Settings 1061
How to Create and Configure Concurrent Instant Virtual Machine(s) on Multiple Hypervisors 1076
Select a source recovery point server(s) 1078
Select VM(s) Location 1079
Select Recovery server(s) 1084
Configure Virtual Machine(s) Settings 1085
Verify Summary 1090
How to Configure Concurrent Virtual Standby Machines 1092
Chapter 12: Using Hardware Snapshot for Backup 1095
How to use Hardware Snapshot for Backup 1096
Use Hardware Snapshot for VMware Agentless Backup 1097
Considerations for NetApp iSCSI/FC Support for VMware 1099
Considerations for NFS Support for VMware 1103
Considerations for Nimble Storage When CHAP Authentication is enabled 1106
Use Hardware Snapshot for Hyper-V Agentless Backup 1108
Use Hardware Snapshot for Agent-based Backup 1110
Verify that the Backup has Used Hardware Snapshot 1111
Chapter 13: Restoring Protected Data 1113
How to Restore From a Recovery Point 1114
Review the Restore Prerequisites and Considerations 1116
Specify the Recovery Point Information to Restore 1123
Restore the Recovery Point Content 1135
Verify that Content is Restored 1136
How to Restore From a File Copy 1137
Review the Restore Prerequisites and Considerations 1139
Specify the File Copy Information to Restore 1141
Restore the File Copy Content 1151
Verify that Content was Restored 1153
How to Restore From a File Archive 1154
Review the Restore Prerequisites and Considerations 1155
Specify the File Copy Information to Restore 1156
Restore the Recovery Point Content 1166
Verify that Content is Restored 1168
How to Restore Files/Folders 1169

Review the Restore Prerequisites and Considerations 1171
Specify the File/Folder Information to Restore 1179
Restore the File/Folder 1195
Verify that the File/Folder was Restored 1197
How to Perform a Bare Metal Recovery Using a Virtual Standby VM or Instant VM 1198
Review the BMR Prerequisites and Considerations 1199
Define BMR Options 1201
Verify that the BMR was Successful 1223
BMR Reference Information 1224
Troubleshooting BMR Issues 1232
How to Restore a Virtual Machine 1239
Review the Restore Prerequisites and Considerations 1241
Specify the Virtual Machine Information to Restore 1242
Restore the Virtual Machine 1257
Verify that the Virtual Machine was Restored 1258
How to Use Exchange Granular Restore (GRT) Utility 1259
Introduction 1260
Review the Prerequisites and Considerations 1261
How to Restore Microsoft Exchange Data Using Exchange Granular Restore (GRT) Utility 1262
How to Restore a Microsoft Exchange Application 1268
Review the Restore Prerequisites and Considerations 1270
Specify the Microsoft Exchange Information to Restore 1272
Restore the Microsoft Exchange Application 1278
Verify that the Microsoft Exchange Application was Restored 1280
How to Restore Exchange Data on a VMware Virtual Machine 1281
How to Download File/Folders without Restore 1282
How to Download File/Folders without Restore for Linux Nodes 1283
How to Restore a Microsoft SQL Server Application 1285
Review the Restore Prerequisites and Considerations 1287
Specify the Microsoft SQL Server Information to Restore 1290
Restore the Microsoft SQL Server Application 1296
Verify that the Microsoft SQL Server Application was Restored 1298
How to Restore From a UNC/NFS Path 1299
Review the Restore Prerequisites and Considerations 1300
Specify the UNC/NFS Path Information to Restore 1301

Restore the Files/Folders and Content Located on UNC/NFS Path 1311
Verify that Content is Restored 1313
How to Restore an Oracle Database 1314
Review the Prerequisites and Considerations 1316
Restore the Server Parameter File 1317
Restore the Parameter File 1318
Restore the Archived Redo Logs 1319
Restore the Tablespaces or Data Files 1320
Restore System, or Undo Tablespaces or Data Files 1322
Restore All Tablespaces and Data Files 1324
Restore Control Files 1326
Restore the Entire Database (Tablespaces and Control Files) 1328
Recover the Oracle Database Using Bare Metal Recovery 1330
How to Perform a File-Level Recovery on Linux Nodes 1331
Review the Prerequisites 1333
(Optional) Recover Data from the iSCSI Volume to the Target Machine 1334
Specify the Recovery Point for Agent-Based Backup 1336
Specify the Target Machine Details 1342
Specify the Advanced Settings 1345
Create and Run the Restore Job 1349
Verify that Files are Restored 1351
How to Perform a File-Level Recovery from Host-Based Agentless Backup for Linux Nodes 1352
Review the Prerequisites 1353
Specify the Recovery Point 1354
Specify the Target Machine Details 1358
Specify the Advanced Settings 1361
Create and Run the Restore Job 1365
Verify that Files are Restored 1366
How to Perform IVM migration (from Linux recovery point) from Cloud to Local 1367
Review the Prerequisites and Consideration for IVM Migration 1368
Perform IVM Migration from Cloud to Local 1369
How to Perform a Bare Metal Recovery (BMR) for Linux Machines 1370
Review the BMR Prerequisites 1373
Get the IP Address of the Target Machine Using the Live CD 1374
(Optional) Recover Data to the iSCSI Volume of the Target Machine 1376
(Optional) Recover Data from the iSCSI Volume to the Target Machine 1378

Review the Backup Server ... 1380
Specify the Recovery Points ... 1381
Specify the Target Machine Details ... 1383
Specify the Advanced Settings ... 1385
Create and Run the Restore Job ... 1390
Verify that the Target Node is Restored ... 1398
How to Perform a Migration BMR for Linux Machines ... 1399
Review the Prerequisites for Migration BMR ... 1400
Perform a BMR to the Temporary Machine ... 1401
Perform a Migration BMR ... 1403
Verify that the Target Node is Restored ... 1404
How to Perform a BMR Using a Backup ... 1405
Review the BMR Prerequisites and Considerations ... 1407
Define BMR Options ... 1409
Verify that the BMR was Successful ... 1428
BMR Reference Information ... 1429
Troubleshooting BMR Issues ... 1437
How to Restore Microsoft Clustered Nodes and Shared Disks ... 1445
Review the Prerequisites ... 1446
Restore Files of a Cluster Shared Disk ... 1447
Restore a Specific Node in a Cluster ... 1448
Restore a Corrupted Cluster Shared Disk ... 1449
Restore the Entire Clustered Nodes and Shared Disk ... 1450
How to Restore an Active Directory ... 1451
Review the Restore Prerequisites and Considerations ... 1453
Restore an Active Directory ... 1454
Verify that the Active Directory was Restored ... 1458
How to Restore Active Directory Data Using Arcserve UDP Active Directory Object Level Restore Utility ... 1458
How to Restore Microsoft Office 365 Data ... 1466
How to Restore Exchange Online Mailbox Data ... 1467
How to Restore SharePoint Online Site Collection Data ... 1481
How to Restore OneDrive Data ... 1491
How to Perform Point-In-Time Restore ... 1501
Prerequisites ... 1501
Considerations ... 1501
Perform Point-In-Time Restore ... 1502

How to Restore Cluster Shared Volume 1506
Review Prerequisites and Considerations 1506
Specify CSV Information to Restore 1513
Restore CSV Content 1522
Chapter 14: Managing Tape Backup and Restore 1525
How to Back Up a Deduplication Data Store to a Tape 1526
How to Restore a Deduplication Data Store From a Tape 1527
Review the Prerequisites 1528
Restore From a Tape Media to an Alternate Location 1529
Import Restored Data Store to the RPS 1530
Chapter 15: Using the PowerShell Interface 1533
How to use the PowerShell Interface 1534
Review the Prerequisite 1535
Using the PowerShell Interface for Arcserve UDP 1536
PowerShell Syntax and Parameters 1537
PowerShell Examples 1550
Configuration file and properties of PowerCLI 1554
Chapter 16: Protecting Oracle System using Recovery Manager (RMAN) 1557
Review the Prerequisites and Considerations 1558
How to Create a Backup Plan for Oracle Database (Windows and Linux) 1559
Create a Backup Plan with a Backup Task 1560
Add an Assured Recovery Task to the Plan 1582
(Optional) Perform a Manual Backup 1591
Verify the Backup Job 1593
How to Restore an Oracle Database using RMAN 1594
Select the Oracle Database Items to Restore 1595
Define the Restore Options 1602
Limitations 1621
Chapter 17: Protecting Microsoft SharePoint Environment 1623
Installation Consideration for Microsoft SharePoint Environment 1624
How to Back Up Microsoft SharePoint Servers 1625
Review the Considerations 1626
Perform a Backup of the SharePoint Server 1627
How to Restore a SharePoint Environment 1629
Restore using Instant VM 1630
Restore using Arcserve UDP Agent UI 1640

Restore using Mount Database from Recovery Point 1645
Create an Isolated Network for SharePoint Recovery 1669
How to Create an Isolated Network for SharePoint Recovery for VMware VM 1670
How to Create an Isolated Network for SharePoint Recovery for Hyper-V VM 1674
Chapter 18: Generating Arcserve UDP Reports 1677
Understanding Arcserve UDP Reports 1678
Arcserve UDP Reports 1679
Using Filters and Actions 1692
Working with Arcserve UDP Reports 1694
Create an SLA Profile 1695
Schedule Emails 1698
Send Report by Email 1702
Generate a Report 1704
Customize Retention Days for a Job Status Report 1705
Raw Data Size in Managed Capacity Report for Host-based Agentless VM Backup 1706
View Actual Values at Source and Destination for Job Types 1710
Chapter 19: Managing High Availability 1713
How High Availability Works 1714
Manage High Availability Control Services 1715
Manage High Availability Licenses 1716
Manage Scenarios 1717
Remote Installation 1731
High Availability Reports 1734
Chapter 20: Using the VSphereConfig Utility 1735
Enable or Disable the VSPhereBackupNFC tag Using Command Prompt 1735
Chapter 21: Using the Diagnostic Utility 1743
Collect the Diagnostic Information 1744
Collect Diagnostic Information from a Standalone Agent 1746
Upload Diagnostic Information to Arcserve Website Using FTP 1747
Unzip Agent Logs 1749
Unzip the Console Logs 1750
Skip Log History Folder Contents Using Registry 1751
Collect Log from Gateway Machine Using Command Line 1751
Collect Logs from RPS using Command Line 1752
Collect Stub logs From Hyper-V Having Agent Installed 1752
Collect Hyper-V Event Viewer Messages 1753

Collect Ca_lic Folder and OLF Files ... 1754
Collect Directory Listing of the Backup Destination/Datastore Directories ... 1754
Chapter 22: Manually Installing Sophos Intercept X Advanced for Server on Arcserve UDP ... 1757
Chapter 23: Troubleshooting ... 1761
Arcserve UDP Communication Failure Related ... 1762
Arcserve UDP Cannot Communicate with Windows Nodes ... 1763
Unable to Receive E-mail Alerts from Gmail Account ... 1764
Arcserve UDP Cannot Communicate with the Arcserve UDP Linux Backup Server on Remote Nodes ... 1765
Arcserve UDP Cannot Communicate with the Arcserve UDP Recovery Point Server on Remote Nodes ... 1766
Arcserve UDP Cannot Communicate with the Arcserve Backup Server on Remote Nodes ... 1767
Arcserve UDP Cannot Communicate with the Remote Site ... 1768
Plan, Job, and Settings Related ... 1769
Backup Job Failure after Changing the Console Hostname/IP Address ... 1770
How to Add Encryption Password for an Existing Encrypted Destination ... 1771
Unable to Apply Backup Settings to Node ... 1772
Plan Deployment Fails after Changing Password of Agentless Backup Proxy Machine ... 1773
Settings Disabled when Opening Agent UI ... 1774
Pause and Resume Fails when Agent Is Not Connected to the Network ... 1775
The Arcserve UDP Agent Service Runs Slowly ... 1776
Configure the Registry to Rerun a Copy to Tape Job ... 1778
Configure the Registry to Copy Multiple Recovery Points of Same Type to Tape in the Same Job ... 1779
File/Folder Missed from Backup of NFS Shared Folder or File/Folder Name is Converted to Junk ... 1780
NFS Shared Folder Backup Fails ... 1780
Instant VM Related ... 1781
Failed to Create an Instant VM in VMware Due to Duplicate NFS Data store Name ... 1782
Failed to Create an Instant VM with Windows 2008 R2 as the Recovery Server for VMware or Windows 2008 R2 Hyper-V Server ... 1783
Failed to Boot an Instant VM if the source machine is a Windows 2008/2012/2016 AD Server ... 1784
Failed to Power on VM after Restoring Hyper-V ... 1786
Instant VM Job Fails Due to Windows NFS Service Error ... 1787
Instant VM File Folder cannot be Accessed or Deleted even with Administrator Privilege. ... 1788
Instant Virtual Machine Fails to Boot in Hyper-V After Restarting the Recovery Server ... 1789

Failed to Create VMware NFS Data store and Displayed Unable to Resolve Hostname Error ... 1790
Failed to Deploy Integration service to Guest VM in Hyper-V ... 1792
Linux Instant VM failed on Non-English Hyper-V Server ... 1793
Extra license consumption for UNC Paths residing on a VM under a licensed Hyper-V ... 1794
License failure error appears when you change the license edition/type and create an Instant VM ... 1794
Linux Agent Related ... 1795
Backup destination settings disabled when opening the Linux agent UI ... 1796
Job Status, Job History, and Activity Log are Not Visible ... 1797
System Volume of SUSE15 with XFS file system Failed to Restore ... 1798
Restore Related ... 1799
Unable to Restore Files ... 1800
Add the Restored Content Database to Original Web Application ... 1801
Database Unable to Mount while Restoring the Microsoft Exchange Database ... 1804
The From Field Does Not Display Correctly for Emails Sent by Users Having "On Behalf Of" Permissions for a Shared Mailbox ... 1806
Restore Jobs Fail After Light Integration Backups ... 1807
Restore to NFS Shared Folder Changes the File Name to Junk ... 1807
Gateway, RPS, Data Store, Console, and Database Related ... 1809
Data Store Name is Already in Use ... 1810
Unable to Connect to data store due to DNS Issue ... 1811
Data Store Switches to the Restore Only Mode ... 1812
Error when RPS Version is Lower than Console Version ... 1814
Adding the Same Resource in Different Sites is not Supported ... 1815
How to enable Log Truncations when SQL Database is in Full Recovery Mode ... 1816
Browse Recovery Points Does Not Show Available Recovery Points When RPS Configured with FQDN ... 1817
Access Denied when Adding RPS ... 1818
Unable to Change to UDP View for Recovery Point ... 1820
Access Denied when Adding or Updating Nodes ... 1821
UDP Console Does not Open when SQL Administrator Password is Changed ... 1823
Failed to Mount the Recovery Points Due to Timeout ... 1824
How to Update the Gateway Server Credentials ... 1825
How to Update the Gateway when the Gateway Proxy Credentials are Changed ... 1827
Console Displays Identity Service is Starting Message ... 1829
Virtual Machine Backup and Restore Related ... 1830
Add Permissions for VDDK at vCenter Server Level ... 1832

Backup job for VM template is always converted to full backup and backup data size is the provision size of virtual disk ... 1839
Independent disks are skipped by backup job of VM template ... 1841
Backup job for VM on SMB 3.0 share fails with error message ... 1843
VM recovery job fails when restoring VM to a Windows default file share ... 1844
Volume Information Unavailable for Recovery Point ... 1845
Convert Incremental Backup to a Verify Backup Because the Virtual Machine Snapshots Either Changed from the Last Backup Job or Needs Consolidation ... 1846
Agentless backup for VMware VM fails if CD/DVD device of VM is connected to an ISO image which resides on a disconnected NFS datastore ... 1847
Agentless Host-based Backup for Hyper-V VM Fails after Upgrading Arcserve UDP ... 1848
Host-based Agentless Backup Fails in VMware ESXi 6.0 ... 1849
Agentless Host-based Backup crashes when using Windows 2003 R2 64-bit as Backup Proxy ... 1851
Host-based Agentless Backup Does Not Use HotAdd Transport Mode ... 1852
Host-based Agentless Backup or Restore Job Uses NBD or NBDSSL Transport Mode even when SAN Mode is Possible ... 1853
Failed to Create a Snapshot for Hyper-V Virtual Machines When Multiple Jobs are Running ... 1855
Hyper-V failed to Create VSS Snapshot ... 1857
Failed to Perform Backup of a Virtual Disk due to System error=[The device is not ready(21)] ... 1858
Backup Job Fails ... 1859
Failing to import VMware VMs from vCenter ... 1861
Backups Fail Due to ESXi License ... 1862
HotAdd Transport Mode does not Work when Attempting to Back Up a VMware VM ... 1863
More Granular Control for VSS over Windows Guest OS Quiescings ... 1865
Unable to configure the Virtual Machine to Enable the "Disk.EnableUUID" Parameter ... 1867
Recovery Operations Fail When Recovering Data Using the HOTADD or SAN Transport Mode ... 1869
Recover VM Operation fails when a non-default port is specified ... 1871
Scheduled Incremental or Full Backup Job Fails for Hyper-V VM ... 1872
Hyper-V VSS NTDS Writer Fails While Taking the VSS snapshot in the VM ... 1873
MAC Address Changes are Not Retained After VM Recovery ... 1874
VM Recovery Fails with Error - Unable to Open VMDK Files ... 1875
Problems Caused by Duplicate VM UUID ... 1876
File System Catalog Job Fails or Recovery Point Check Fails for Host-Based Agentless Backup ... 1878
Incremental Backup Converts to Verify Backup or the Backup Size Increases in Hyper-V 1880
Host-based Backup Fails for Hyper-V VM That Has a Special Differencing Disk Configuration ... 1881

Backup Job Fails for a VMware Virtual Machine 1882
Disable Re-scan of HBA Adapters During Incremental Backup 1884
Disable the Consecutive Snapshot Creation in VMware VM for a Backup 1885
When Restored from a Higher Version of ESXi host to a Lower version ESXi host, VM gets stuck at booting stage 1886
RAM Utilization reaches 99% When Backup Jobs are Submitted to VM 1887
Hyper-V Restore Job Fails, Cannot Connect to Utility on Host 1888
Automatic Protection Fails to Detect and Protect VM 1890
Set Read block size when backing up VMDK file 1892
Lun Space Reservation Inherited while Performing Lun Clone 1892
Virtual Standby Related 1893
Operating System Not Found 1894
Virtual Standby Jobs Fail Due to Internal Errors 1895
Virtual Standby Jobs Fail Using the HotAdd Transport Mode 1898
Virtual Standby Jobs to Hyper-V Systems Fail 1900
Issue Related to Duplicate Agent UUID 1901
Modifications to Virtual Private Cloud in Virtual Standby Task Do Not Reflect in Amazon EC2 1901
Terminal EC2 Resources option is not shown 1902
Copy Recovery Points Related 1903
Configure the Registry for Copy Recovery Point Job 1904
Bandwidth Congestion with Copy Recovery Point to Cloud Jobs 1905
Merge Job Skipped 1907
Arcserve UDP Reports Related 1909
How to Display Charts to View Arcserve UDP Reports 1910
Microsoft Office 365 Exchange Online Related 1912
Unable to View the List of Mailboxes on the Select the destination Dialog 1912
Access or Permission Denied for Exchange Online Group 1913
Chapter 24: Appendix 1915
Command Line Tool for Deduplication Data Store 1916
How to Show Recovery Point Check Option 1920
How to apply different version of VDDK other than the built-in Version in Arcserve UDP 1923
Modifying the VDDK Manually for Virtual Standby Backup (VSB) 1923
Modifying the VDDK Manually for Host Based Backup (HBBU) 1924
How to Migrate Arcserve D2D r16.5 Backup Data of Two Servers Having Same Host Name to RPS Data Store 1926
How to Deploy Arcserve UDP in Microsoft Azure 1927

Understanding Process Flow ... 1927
Best Practices ... 1928
Considerations ... 1928
Planning the Arcserve UDP Deployment in Microsoft Azure ... 1929
Deploy Arcserve UDP in Microsoft Azure ... 1937
Arcserve UDP Terms and Definitions ... 1948
Agent-Based Backup ... 1948
Compression ... 1948
configuration ... 1949
dashboard ... 1949
Data Store ... 1949
Destination ... 1949
Discovered Nodes ... 1949
Encryption ... 1949
Global Deduplication ... 1952
Host-Based Agentless Backup ... 1952
HOTADD Transport Mode ... 1952
Job ... 1952
jobs ... 1952
NBD Transport Mode ... 1952
NBDSSL Transport Mode ... 1953
Nodes ... 1953
Plan ... 1953
Protected Nodes ... 1953
Recent Event ... 1953
Recovery Point ... 1953
Recovery Point Server ... 1953
Replicate ... 1954
Resources ... 1954
SAN Transport Mode ... 1954
Systems ... 1954
Tasks ... 1954
Unprotected nodes ... 1954
Data Deduplication ... 1955
Types of Data Deduplication ... 1956
How Data Deduplication Works ... 1957

When Should You Use Deduplication 1959
Configuring Deduplication Data Stores in Arcserve UDP 1960
Deduplication, Encryption, and Compression 1962
Deduplication Limitations 1963

Chapter 1: Features and Enhancements

This section explains the features or enhancements provided in each release of Arcserve UDP. The Arcserve UDP solution provides an all-inclusive solution for next-generation storage problems of organizations that are trying to protect their data in a rapidly changing virtual, cloud, and services world. The solution does this by providing a single user interface to a wide range of functionality, addressing multi-site business continuity and disaster preparedness problems.

This section provides information about new features, enhancements, and support.

This section contains the following topics:

New Features

Here is a list of the new features available in Arcserve UDP 8.0.

- **Backup to and Restore from Hidden Volume:** Allowed to backup and restore data from hidden volumes. Hidden volume is a volume that is mounted without a drive letter or a folder path in the OS. Hidden volumes can be accessed using corresponding volume GUID.
- **Export and Import Settings:** Allowed export and import settings using the JSON file in Windows Agent.
- **OneDrive Support:** Allowed restore options to the original or alternate locations.
- **SharePoint Online Support:** The following features are available for SharePoint Online Support:
  - Allowed to backup lists and libraries from personal sites.
  - Allowed to add multiple site collections all at once.
  - Allowed to add the newly added site collections automatically inside the plan.
- **Oracle RMAN Support:** Arcserve UDP now provides a way of protecting Oracle system using Recovery Manager (RMAN). RMAN is an Oracle utility that performs backup, restore and recovery tasks on databases and automates the administration of backup strategies. For more information, see Protecting Oracle System using Recovery Manager (RMAN).
- **Office 365 Group Support:** Enabled group support for Office 365 users.
- **Reverse Replication:** Enabled replication of data (recovery point) from Cloud, MSP, or Remotely Managed RPS to local UDP or on-premises RPS server. For more information, see How to Perform Reverse Replication.
- **Nutanix Files:** Arcserve UDP now provides a way of protecting Nutanix Files using the Arcserve UDP UNC/NFS option. Nutanix Files is a software-defined, scale-out file storage solution that provides a repository for unstructured data, such as home directories, user profiles, departmental shares, application logs, backups, and archives. For more information, see Protecting UNC/NFS Path.
- **Teams Support:** Data backup and restore support is enabled for MS Teams, which is a business communication chat based platform. For more information, see Managing Microsoft Teams Backup and Restore.

- **Immutable Storage support for Copy Recovery Point:** Allowed to create immutable storage on Amazon S3 using the Object Lock feature.
- **Adobe Flash Replacement:** Adobe Flash is no longer required and installing .NET Framework allows you to view charts.

Feature Enhancements

- **PowerCLI changes:** The following options are enabled for the log file(s) UDPPowerCLI.log:
  - Allowed to change the location of the UDP PowerCLI log file(s).
  - Allowed to control the maximum size of the log file(s) using circulation of logs.
  - The time stamp in the log file(s) is changed to 24-hour format.
  - Allowed Retention Policy based on Retention Period for UDP PowerCLI log files(s).
  - Allowed Retention Policy based on Retain No of Files for UDP PowerCLI log files(s).
- **Snapshot Enhancement:** Enabled the Arcserve UDP Agent backup plan (Windows/Linux) and the Host-based Agentless Backup plan (Nutanix/ESXI/Hyper-V) to run the post snapshot script even when the snapshot fails.
- **Modern Authentication:** Enabled modern authentication for Office 365 backup and restore.
- **Specify Network for Restore and Replication:** Enabled restore and replication using a network to establish communication between Windows Agent and Recovery Point Server.
- **Specify Network for Linux:** Enabled backup and restore using a network to establish communication between Linux Agent and Recovery Point Server. However, *Specify Network for Linux* supports only RPS as a local data store destination.
- **Concurrent IVM:** Allowed to create and configure concurrent Instant Virtual Machine(s) on multiple hypervisors using the Arcserve UDP console.
- **Concurrent VSB:** Allowed to configure concurrent Virtual Standby Machine(s) using the Arcserve UDP console.

Linux Agent Enhancements

- New supported platforms:
  - Red Hat Enterprise Linux 8.3
  - CentOS 7.9, 8.3
  - Oracle Linux 8.3
  - Debian Linux 9.12, 9.13, 10.6, 10.7
- **Specify Network for Linux:** Enabled backup and restore using a network to establish communication between Linux Agent and Recovery Point Server. However, *Specify Network for Linux* supports only RPS as a local data store destination.
- **Concurrent IVM:** Allowed to create and configure concurrent Instant Virtual Machine(s) on multiple hypervisors using the Arcserve UDP console.
- **Snapshot Enhancement:** Enabled the Arcserve UDP Agent Linux backup plan to run the post snapshot script even when the snapshot fails.
- **Adobe Flash Replacement:** Adobe Flash is no longer required to view charts on the Linux Agent.

Database and Platform Support

- Nutanix AOS 5.19
- Red Hat Enterprise Linux 8.3
- CentOS 7.9, 8.3
- Oracle Linux 8.3
- Debian Linux 9.12, 9.13, 10.6, 10.7
- VMware vSphere 7.0 Update 1
- Oracle Database 19c
- Nutanix Files 3.8.z, 3.7.z, 3.6.z

Security and Third-party Enhancements

- **Open JDK support:** Open JDK version is upgraded to 1.8.0.265.b01
- **VMware Virtual Disk Development Kit support:** VDDK version is upgraded to 7.0. 1.40833
- **VMware vSphere Web Services SDK support:** SDK version is upgraded to 6.7.1
- **WSO2 support:** WSO2 version is upgraded to v5.7.0
- **Apache support:** Apache version is upgraded to 2.4.46
- **OpenSSL support:** OpenSSL version is upgraded to 1.1.1i
- **Tomcat support:** Tomcat version is upgraded to 9.0.37
- **SQLite support:** SQLite version is upgraded to 3.26
- **AWS SDK for Java:** AWS SDK for Java version is upgraded to 1.11.852
- **AWS SDK for C++:** AWS SDK for C ++ version is upgraded to 1.6.53

Arcserve Appliance Enhancements

- Existing Appliance customers can upgrade to Arcserve UDP 8. For more information, see Installation.

Dropped Support

- vSphere 5.5 support. We strongly recommend upgrading the environment (vCenter/ESXi) to the latest VMware supported version. For more information about the end of general support for vSphere 5.5, see the KB article.
- Nutanix AHV 5.5.3.1.

Chapter 2: Understanding Arcserve UDP

This section contains the following topics:

Introduction

Arcserve Unified Data Protection is a comprehensive solution to protect complex IT environments. The solution protects your data residing in various types of nodes such as Windows, Linux, and virtual machines on VMware ESX servers, Microsoft Hyper-V servers or Nutanix AHV servers. You can back up data to either a local machine or a recovery point server. A recovery point server is a central server where backups from multiple sources are stored.

Arcserve UDP provides the following capabilities:

- Protects various type of source nodes, including Agent Based, Agentless Based, CIFS, Exchange Online, SharePoint Online, OneDrive and so on.
- Backs up data to recovery point servers
- Replicates backup data to local and remote recovery point servers
- Monitors Arcserve High Availability
- Archives data
- Copies:
  - Selected source files to a secondary backup location
  - Recovery points to local and Cloud locations such as Share Folder and AWS EC2
  - Recovery points to tape
- Creates
  - Virtual standby machines from backup data to local hypervisor (Hyper-V, ESX, Nutanix), AWS EC2 or Microsoft Azure
  - Instant virtual machine to local hypervisor (Hyper-V, ESX, Nutanix) for Windows
  - Instant virtual machine to local hypervisor (Hyper-V, ESX, Nutanix), AWS EC2 or Microsoft Azure for Linux
- Restores
  - Backup data and performs Bare Metal Recovery (BMR)
  - Microsoft Exchange email and non-email objects using the Arcserve UDP Exchange Granular Restore utility.

    Note: For more details on the supported specifications, functions, and other features, see the Exchange Granular Restore user guide (esr.pdf).

- Supports
  - Role-based administration
  - Hardware snapshot
  - Assured Recovery Test for recovery points
  - SLA report for RPO and RTO

Arcserve UDP replicates backup data that is saved as recovery points from one server to another recovery point server. You can also create virtual machines from the backup data that can act as standby machines when the source node fails. The standby virtual machine is created by converting recovery points to VMware ESX, Microsoft Hyper-V virtual machine or Nutanix AHV format.

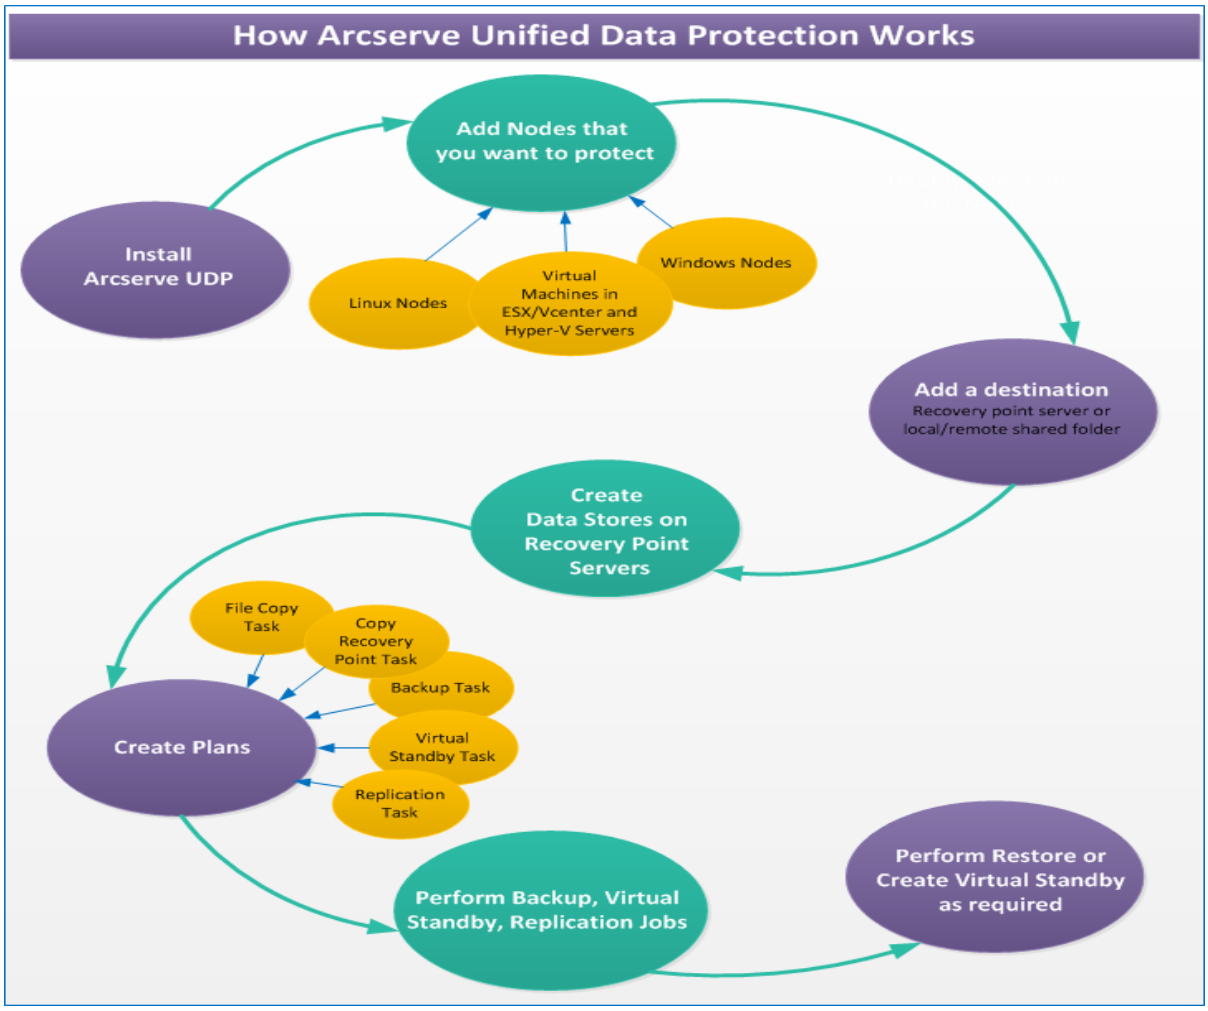
The Arcserve UDP solution provides integration with Arcserve High Availability. After you create scenarios in Arcserve High Availability, you can manage and monitor your scenarios and perform operations like adding or deleting destination machines.

How Arcserve UDP Works

Arcserve UDP is a unified data protection solution that lets you protect your computer systems. Use the following high-level steps to protect your systems using Arcserve UDP.

1. Install Arcserve Unified Data Protection.
2. Add Nodes that you want to protect. You can add Windows or Linux nodes and virtual machines in ESX/vCenter, Hyper-V servers, and Nutanix AHV servers.
3. Add a destination. A destination could be a recovery point server, local folder, or remote shared folder.
4. Create data stores on the recovery point server. A data store is a physical area on a disk. You can create deduplication and non-deduplication data stores.
5. Create a plan. A plan is a group of tasks to manage backup, replication, copy recovery point, copy to tape, creation of virtual standby machines, or assured recovery test. You can also add UNC Path, Office 365 Exchange online, SharePoint Online node or OneDrive node and create related tasks.
6. Perform jobs such as backup, create virtual standby, and replicate.
7. Perform a simple restore or a bare metal recovery.

The following diagram illustrates the high-level steps that you need to perform to protect data:


How Arcserve Unified Data Protection Works
Add Nodes that you want to protect
Install Arcserve UDP
Linux Nodes
Virtual Machines in ESX/Vcenter and Hyper-V Servers
Windows Nodes
Add a destination
Recovery point server or local/remote shared folder
Create Data Stores on Recovery Point Servers
File Copy Task
Copy Recovery Point Task
Backup Task
Virtual Standby Task
Replication Task
Create Plans
Perform Backup, Virtual Standby, Replication Jobs
Perform Restore or Create Virtual Standby as required

Instant Virtual Machine (IVM) versus Virtual Standby (VSB) Machine

While restoring data after a disaster or during disaster recovery training, you may need to start virtualized instances of servers that were previously protected by Arcserve UDP.

Arcserve UDP provides the following two features that allow you to start the virtual machine from recovery points:

- I**nstant Virtual Machine (IVM)**: Creates a virtual machine instantly from a recovery point. Using an Instant virtual machine helps you get an immediate access to data and applications present in the Arcserve UDP backup sessions. An Instant virtual machine eliminates the downtime associated with a traditional restore or conversion of the backup session to a physical or virtual machine.

  For more details on Instant Virtual Machine, see How to Create and Manage an Instant Virtual Machine.

- **Virtual Standby Machine (VSB)**: Converts the recovery points to virtual machine formats and prepares a snapshot to easily recover your data when needed. This feature provides the high availability capability also and ensures that the virtual machine can take over immediately when the source machine fails. The standby virtual machine is created by converting the recovery points to any VMware, a Hyper-V virtual machine or Nutanix AHV virtual machine format.

  For more details on Virtual Standby, see How to Create a Virtual Standby Plan.

To determine which feature works best, you need to consider your RTO (Recovery Time Objective) and scenario. The following table provides the comparison between the IVM and VSB features:

FEATURES	IVM	VSB
Power standby VMs from the latest recovery point	Yes (no conversion needed)	Yes, ONLY if a VSB task was added to the backup plan. For example, Advance planning required)
Requires backup time processing	Not required	Required, it is necessary to add a VSB task to the plan that is used to back up the source machine.
VM Boot time	A slower process (up to 30%) due to I/O redirection.	Same time as any other VM on the same hypervisor.
Disk space requirements	Minimal storage space to host	Yes, storage space is con-

	child disk or store changes when running a VM.	sumed on the destination hypervisor where VSB standby VM is maintained. Requires storage space equal to or greater than the size of the source machine.
High Availability (HA) option	N/A	AVAILABLE Monitors source machine and can start VSB VM if the source machine becomes unavailable.
VM performance	May run slower compared to regular VMs (up to 30%) due to I/O redirection, however the performance depends on the nature of the application workload.	Performance is the same as the regular VMs.
Management/Configuration	Managed from the UDP console, can start or stop the IVM on demand when access is needed by the user.	Added as a task to a plan so that all the backed up data is converted to a VM format automatically. VSB task is applied to all nodes that are protected by the plan.
Persisting data and migrating VM to production	The virtual disk of the IVM refers to the data blocks in the recovery point from which the VM was started. Therefore, when the IVM accesses data blocks within its virtual disk the data is actually requested from RPS (this process is transparent to the user). Such I/O redirection introduces some additional performance footprint. If you plan to use IVM in production then we recommend to make IVM persistent and hydrate the virtual machine virtual disk with the actual data. Hydration of the IVM can be achieved by copying/replicating the VM. Depending on the type of Hypervisor used in your production environment, to persist the IVM	The virtual disk or disks of the VSB VM already contain most of the recent data from the corresponding recovery point. Since I/O redirection does not occur (which is the same as the IVM), the performance of the VSB VM is the same as regular VMs where there is no dependency on the RPS or recovery point (compared to the IVM scenario).

	data you can use VMware Storage vMotion or Hyper-V VM Storage Migration/Replication to copy the IVM where the data becomes permanent.	

User Security

This section contains the following topics:

- Roles for Arcserve UDP Services
- User Privileges for Arcserve UDP Functions

Roles for Arcserve UDP Services

The following table describes the roles for Arcserve UDP Services:

Service	Description	Running on Role
Arcserve Event Log Watch	Provides Licensing SDK service to accept license keys.	Console/RPS/Agent
Arcserve UDP Agent Explorer Extension Service	Provide backend service for UDP View in Windows Explorer.	RPS/Agent
Arcserve UDP Agent Service	Provides Web UI and Web Service for Arcserve UDP Agent, including backup / restore job submission.	RPS/Agent
Arcserve UDP Identity Service	Provide authentication and authorization service for Arcserve Unified Data Protection Console.	Console
Arcserve UDP Management Port Sharing Service	Provide port sharing service which allows console, gateway and identity server to share one port.	Console
Arcserve UDP Management Service	Provides Web Service for Arcserve Unified Data Protection central management Console.	Console
Arcserve Remote Management Gateway Service	Provide a capability to handle request bi-directionally between Gateway and Console across intranet or internet.	Gateway
Arcserve UDP RPS Data Store Service	Provides Web Service for UDP Data Store management including creation, modification, delete, start, and stop.	RPS
Arcserve UDP RPS Port Sharing Service	Expose only one port for RPS server to handle communication to Agent web UI / service, RPS web service, and replication job.	RPS
Arcserve UDP Update Service	Detect Arcserve UDP update and download if available.	Console/RPS/Agent

User Privileges for Arcserve UDP Functions

The following table describes the user privileges for Arcserve UDP functions:

Functions	User	Privilege	Comments
Installation	**Local administrators group**	**Local administrators group**	
Console and Gateway	**Local administrators group**	**Local administrators group**	
Recovery Point Server	**Local administrators group**	**Local administrators group**	
Windows Client Backup	**Local administrators group**	**Local administrators group** Security Policies: • Act as part of operating system • Log on locally • Log on as a service • Log on as Batch Job	**Many backup-related operations like VSS snapshot requires admin privilege.**
Network Share for non-dedupe data store	If RPS UAC is enabled, domain account or built-in administrator.		
SQL log truncation	Local administrators group	Local administrator with SQL sysadmin, or db_ owner fixed database role.	Log truncation requires query backup database, back log, and query shrink (shrink DB).
Exchange log truncation	Domain administrators group	Domain administrators group	Need to have access to exchange DB
Active Directory protection	Domain administrators group	Domain administrators group	
Windows Client Restore			
Network Share for non-dedupe data store	If RPS UAC is enabled, domain account or built-in administrator.		
SQL	Local administrators group	Local admin with SQL	

		sysadmin, or db_owner fixed database role.	
Exchange	Domain administrators group	Domain administrators group	Need to have access to Exchange DB
Active Directory	Domain administrators group	Domain administrators group	
Exchange Granular Restore Utility	Restoring to mailbox: the account used to restore should have impersonate privilege on the target mailbox. For other restore options, the account does not need special requirement.	Restoring to mailbox: the account which is used to restore should have impersonate privilege on the target mailbox. For other restore options, there is no special requirement on the account.	
Host-based Agentless Backup			
Add VM node from vCenter/ESXi	• vCenter: built-in administrator • ESXi: root		For vCenter, if non-built-in administrator is used, refer to link.
Add VM node from Nutanix AHV	Cluster Admin or User Admin	Cluster Admin or User Admin	
Add VM node from Hyper-V	• Standalone Hyper-V: built-in local administrator, built-in domain administrator, or domain account which is member of the local Administrators group • Hyper-V cluster: built-in domain administrator or domain account which is member of the local Administrators group.		If other administrative account is used, UAC remote access needs to be disabled. Refer to link.
Switch VMware Snapshot Quiescing Method in plan	Built-in local administrator or built-in domain administrator **Note:** Required credentials here are set by Update Node		If other administrative account is used, UAC needs to be disabled. Refer to

			link.
Application DB level restore for Hyper-V VM/Nutanix VM	Built-in local administrator, built-in domain administrator, or domain account which is member of the local Administrators group **Notes:** • Required credentials here are set by Update Node • If VM guest OS is client version Windows (such as Windows 10), need to manually configure firewall to allow Windows Management Instrumentation(WMI)		If other administrative account is used, UAC remote access needs to be disabled. Refer to link.
PFC	• VMware VM: built-in local administrator or built-in domain administrator • Hyper-V VM: built-in local administrator, built-in domain administrator, or domain account which is member of the local Administrators group • Nutanix VM: built-in local administrator, built-in domain administrator, or domain account which is member of the local Administrators group **Notes:** • Required credentials here are set by Update Node		If other administrative account is used: • For VMware VM, UAC needs to be disabled. See link. • For Hyper-V VM, UAC remote access needs to be disabled. See link.

	• For Hyper-V VM/Nutanix VM, if VM guest OS is client version Windows (like Windows 10), need to manually configure firewall to allow Windows Management Instrumentation (WMI)		
Pre / Post Command	• VMware VM: built-in local administrator or built-in domain administrator • Hyper-V VM: built-in local administrator, built-in domain administrator, or domain account which is member of the local Administrators group • Nutanix VM: built-in local administrator, built-in domain administrator, or domain account which is member of the local Administrators group **Notes:** • Required credentials here are set by Update Node and on the Advanced tab of a Plan. • For Hyper-V VM, if VM guest OS is client version Windows (like Windows 10), need to manually configure firewall to allow Windows Management Instrumentation		For the usage of the credentials that are set by Update Node and on the Advanced tab of a Plan, refer to link.

	(WMI) • For Hyper-V VM/Nutanix VM, if VM guest OS is client version Windows (Such as, Windows 10), then manually configure firewall to allow Windows Management Instrumentation (WMI)		
SQL log truncation	Same as Pre / Post Command		Same as Pre / Post Command
Exchange log truncation	Same as Pre / Post Command		Same as Pre / Post Command
File-level restore to original location	Built-in local administrator, built-in domain administrator, or domain account which is member of the local Administrators group **Notes:** • Required credentials here are set by Update Node • For Hyper-V VM/Nutanix VM, if VM guest OS is client version Windows (like Windows 10), need to manually configure firewall to allow Windows Management Instrumentation (WMI)		If other administrative account is used, UAC remote access needs to be disabled as per link.
Virtual StandBy			
For Hyper-V	• Built-in local administrator • Built-in domain administrator • Domain account which is member of	Local administrators group	If local administrative account is used, UAC remote access needs to be disabled. See link.

	the local Administrators group • Local account which is member of the local Administrator group		
For Nutanix	Cluster Admin	Cluster Admin	
For VMware	• vCenter: built-in administrator • ESXi: root		For vCenter, if non-built-in administrator is used, refer to link.
Instant Virtual Machine/Assured Recovery			
For Hyper-V	• Built-in local administrator • Built-in domain administrator • Domain account which is member of the local Administrators group • Local account which is member of the local Administrator group		If local administrative account is used, UAC remote access needs to be disabled. See link.
For VMware	• vCenter: built-in administrator • ESXi: root	Local administrators group	For vCenter, if non-built-in administrator is used, refer to link.
File Copy & Archive	**Local administrators group**	**Local administrators group**	
Copy Recovery Point to Cloud	**Local administrators group**	**Local administrators group**	
UNC/NFS Path protection	**Any user could login and be impersonated**	**Read permission to the UNC/NFS Path**	
Virtual StandBy to AWS EC2	**The Amazon IAM users who have the required permissions to interaction with AWS API**		For AWS EC2, refer to this link.
Virtual StandBy to Microsoft	**Application**	**Contributor role of selected subscription**	

Azure			
Linux			
Install	root	Read, Write, Execution	
Console registration	console admin		
Agent-based Backup			
-Network Share	storage administrator	Read, Write	
--Node Connection	root/non-root/sudo	Read, Write, Execution	
File Level Restore			
-Network Share	storage administrator	Read, Write	
--Node Connection	root/non-root/sudo	Read, Write, Execution	root user can restore to anywhere; other users can restore only to their owned directories
BMR		Access information to hardware	
Migration BMR			
Instant VM for Hyper-V			
Instant VM for VMware			
Instant VM for Nutanix AHV	cluster admin	cluster admin	
Instant VM to Amazon EC2	**IAM User**	**Full Access of EC2**	
Instant VM to Microsoft Azure	**Application**	**Contributor role of selected subscription**	
Exchange Online protection	**Any Exchange Online account**	**Has Application Impersonation privilege on the protected accounts**	
SharePoint Online protection	**SharePoint Online Site Collection Administrator**	**SharePoint Online Site Collection Administrator**	
OneDrive	**Azure Active Directory Administrators**	**Azure Active Directory Administrators**	

Chapter 3: Installing Arcserve UDP

This section contains the following topics:

How to Install Arcserve UDP

Arcserve UDP - Full: After the installation, you log in to the Arcserve UDP Console (Console) and perform data management functions. The Console lets you manage and monitor nodes, recovery point servers, backups, restore, and replication.

Arcserve UDP - Agent: Installs only the Arcserve UDP Agent. Install the agent to the nodes that you want to protect. Perform this step only when you want to install the agent manually to a node. Typically, the agent is deployed automatically to nodes from the Console when you create a plan.

What To Do Next?

1. Review the Prerequisites and Considerations
2. Decide the Installation Type
3. Install Arcserve Unified Data Protection Using the Setup Wizard
4. Install Arcserve Unified Data Protection Using the Command Line
5. Install Arcserve Unified Data Protection Using the Unified Installer
6. Verify the Installation
7. (Optional) Communication Ports Used
8. (Optional) How the Installation Process Affects Operating Systems

Review the Prerequisites and Considerations

Review the following installation prerequisites and considerations before installing Arcserve UDP:

Prerequisites

- Review the Arcserve UDP Release Notes 8.0. The Release Notes contains a description of system requirements, supported operating systems, and a list of issues known to exist with this release.
- Verify that your systems meet the software and hardware requirements necessary to install Arcserve Unified Data Protection components.
- Verify that your Windows account has administrator privileges or any other equal privileges to install software on the systems where you plan to install Arcserve Unified Data Protection components.
- Verify that you have the user names and passwords of the systems where you are installing Arcserve Unified Data Protection components.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Considerations

Before the installation, you should determine how to set up your Arcserve Unified Data Protection installation:

- The system where you want to install the Console.
- The nodes that you want to protect.
- The number of recovery point servers that would serve as backup destinations.
- The number of replica servers that would replicate recovery point servers.

Decide the Installation Type

You can install Arcserve UDP using one of the following methods:

- **Standard Installation Using the Setup Wizard:** This method lets you install Arcserve Unified Data Protection using the Setup Wizard. In this method, you are prompted at each step to choose the desired options.
- **Silent Installation Using the Command Line:** This method lets you perform an unattended installation using the Windows command line.
- **Installation using the Unified Installer:** This method lets you install Arcserve UDP, Arcserve Backup, and Arcserve RHA from a single installer. You can choose to install all the three products at once or install each product separately. You can download the installer from the Arcserve website.

Install Arcserve Unified Data Protection Using the Setup Wizard

Using Arcserve Unified Data Protection, you can centrally manage and monitor nodes, recovery point servers, virtual machines in vCenter or ESX servers or Microsoft Hyper-V servers, replica servers, and Arcserve Unified Data Protection reports.

Install Arcserve Unified Data Protection on a server from where you can manage protected nodes and other Arcserve Unified Data Protection components.

Follow these steps:

1. Access the Arcserve UDP installation package from either the Arcserve web site or the product CD.

   Note: If one of the supported non-English operating systems is detected, you need to select the language for the product installation.

2. Double-click the installation package.

   The **License Agreement** dialog opens.

3. Read and accept the terms of the licensing agreement and click **Next**.

   The **Installation Type** dialog opens.

4. Select one of the installation types.

   Standard Installation

   Lets you install either the agent or all the Arcserve Unified Data Protection components.

   Arcserve Unified Data Protection- Agent

   Installs Arcserve UDP Agent only.

   For more information, see *Install Arcserve Unified Data Protection Agent for Windows Using the Installation Wizard* in the *Arcserve UDP Agent for Windows User Guide*.

   Arcserve Unified Data Protection- Full

   Installs Arcserve Unified Data Protection Console, Recovery Point Server, and Agent.

   Advanced Installation

   Lets you install one or more of the following Arcserve Unified Data Protection components:

- Arcserve UDP Agent
- Arcserve UDP Recovery Point Server
- Arcserve UDP Console

5. Specify if you want to install the Arcserve Unified Data Protection Agent for Windows change tracking driver.

   By default, this option is selected.

   - Without this driver installed, Arcserve Unified Data Protection Agent for Windows always performs full backup.
   - With this driver installed, you would still need to have a valid Arcserve Unified Data Protection Agent for Windows license to perform a local backup.
   - This driver is not required if this agent is used as Virtual Standby monitor or host-based VM backup proxy server.

     Note: You can install this driver at any time after the installation is complete by running the InstallDriver.bat utility from the following location:

     <Arcserve UDP install folder>\Engine\BIN\DRIVER

6. Click **Next**.

   The **Destination Folder** dialog opens.

7. Specify the folder where you want to install Arcserve Unified Data Protection and click **Next**.

   The **Configuration** dialog opens.

8. On the **Configuration** dialog, specify the following information:

   a. Select the protocol.

      Note: For a secure communication, select the HTTPS protocol.

   b. Enter the port number for the agent, if applicable. Typically, the port number is 8014.

   c. Enter the port number for the Console, if applicable. Typically, the port number is 8015.

   d. Enter the Windows Administrator name and password.

   e. Specify if you want to display the Arcserve UDP agent monitor for all users or only the current user.

9. Click **Next**.

   The **Database Settings** dialog opens.

10. On the **Database Settings** dialog, click the **Database** drop-down list to choose a database type. You can specify one of the following:

    - Microsoft SQL Server 2014 Express (included)
    - Microsoft SQL Server

      Important! When you have more than 500 nodes to manage from the Console, make sure that you select Microsoft SQLServer and not SQLExpress.

    After you specify a database, the required options for the specified database are displayed on the **Database Settings** dialog. Provide database settings details for one of the databases that you select.

    Microsoft SQL Server 2014 Express (included):

    On the **Database Settings** dialog, complete the following:

    a. Specify the location where you want to install Microsoft SQL Server 2014 Express. You can accept the default path or specify an alternative path.

    b. Specify the location where you want to install the data file for the Arcserve Unified Data Protection default database. You can accept the default path or specify an alternative path.

       Note: Microsoft SQL Server 2014 Express does not support remote communication. Therefore, install the default database and the data file on the computer where you are installing the application.

    Microsoft SQL Server Databases

    On the **Database Settings** dialog, complete the following:

    a. **SQL Server Type**: Specify the type of communication that the application should use to communicate with the SQL Server database.

       Local: Specify Local when the application and SQL Server are installed on the same computer.

       Remote: Specify Remote when the application and SQL Server are installed on different computers.

    b. **SQL Server Name**: If the SQL Server Type is Remote, specify the remote SQL Server name. If the SQL Server is local, select the server from the drop-down list.

    c. **Security**: Specify the type of credentials that you want to use to authenticate SQL Server.

       Use Windows Security: Authenticates using your Windows credentials. You can login using Arcserve UDP Console credentials.

Use SQL Server Security: Authenticates using SQL Server credentials. Enter the Login ID and Password to access the SQL Server account.

11. Click **Next**. The **Firewall Exceptions** dialog opens.

    The **Firewall Exceptions** dialog lists the services and programs to be registered to Windows Firewall as exceptions for Arcserve UDP.

    Note: Firewall exceptions are required if you want to configure and manage Arcserve UDP from remote machines.

12. Click **Next**. The **Message** dialog opens.

13. Click **Next**. The **Summary** dialog opens.

14. Click **Install** to launch the installation process.

    The **Installation Progress** dialog is displayed indicating the status of the installation. When the installation is complete, the **Installation Report** dialog is displayed.

    (Optional) If you want to check for any latest product updates, follow these steps:

    a. Select **Check for an update immediately** and click **Finish**.

       The **Check for Updates** dialog opens.

    b. Select the server from where you want to download updates and click **Download and Install Updates**.

       The **Update Process** dialog is displayed indicating the download status.

       When the update is complete, an alert message is displayed.

    (Optional) To install the Arcserve UDP Agent for Linux, follow the instruction in the **Install arcserve Unified Data Protection Agent for Linux** section.

15. Click **Finish**.

    Arcserve UDP is installed on your computer.

Install Arcserve Unified Data Protection Using the Command Line

You can install Arcserve Unified Data Protection silently. A silent installation eliminates the need for user interaction. The following steps describe how to install the application silently using the Windows Command Line:

Follow these steps:

1. Open the Windows Command Line on the computer where you want to start the silent installation process.
2. Download the self-extracting installation package to your computer.

   Start the silent installation process using the following Command Line syntax:

   Arcserve_Unified_Data_Protection.exe -s -a -q -Products:<ProductList> -Path:<INSTALLDIR> -User:<UserName> -Password:<Password> -Https:<HTTPS> -ConsolePort:<Port Number> -AgentPort:<Port Number> -Driver:<DRIVER> -MonitorFlag:<MONITORFLAG> -StopUA:<STOPUA> -SummaryPath:<SUMMARYPATH> -AutoReboot:<AUTOREBOOT>

   Example:

   Arcserve_Unified_Data_Protection.exe -s -a -q -Products:Agent -User:administrator -Password:test
3. Configure the silent installation using the following syntax and arguments:

   Important: If the parameters include any of the following special characters, enclose the parameters in quotes:

   - <space>
   - &()[]{}^=;!'+,`~

   For example: If the password is abc^*123, the input should be -Password:"abc^*123".

   -s

   Runs the executable file package in the silent mode.

   -a

   Specifies additional command line options.

   -q

   Installs the application in the silent mode.

   -Products:<ProductList>

(Optional) Specifies the components to install silently. If you do not specify a value for this argument, the silent installation process installs all components. You can specify the following components:

Agent: Installs the Data Protection Agent component.

RPS: Installs the Recovery Point Server component.

Console: Installs the Console component.

All: Installs all the components of Arcserve Unified Data Protection.

Example:

For Install Data Protection Agent:

-Products:Agent

For Install Recovery Point Server:

-Products:Agent,RPS

For Install Data Protection Agent, Recovery Point Server and Data Protection Console:

-Products:Agent,RPS,Console

For Install all the components in the build:

-Products:All

-User:<UserName>

Specifies the user name that you want to use to install and run the application.

Note: The user name is of the administrator or an account with administrative privileges.

-Password:<Password>

Specifies the password of the user name.

-Https:<HTTPS>

(Optional) Specifies the communication protocol. The options are 0 and 1. Use 0 for http and 1 for https.
Default: 0
Example:

-https:1

-Path:<INSTALLDIR>

(Optional) Specifies the target installation path of Data Protection Agent.
Example:

-Path:C:\Program Files\Arcserve\Unified Data Protection

Note: If the value for INSTALLDIR contains a space, enclose the path with quotation marks. Additionally, the path cannot end with a backslash character.

-ConsolePort:<Port Number>

(Optional) Specifies the communication port number for the Console.

Default: 8015

Example:

-ConsolePort:8015

Note: Use this option when you want to install the Console.

-AgentPort:<Port Number>

(Optional) Specifies the communication port number to access Arcserve Unified Data Protection Agent.

Default: 8014
Example:

-AgentPort:8014

Note: Use this option when you want to install the Arcserve Unified Data Protection Agent.

-Driver:<DRIVER>

(Optional) Specifies whether to install Arcserve Unified Data Protection Agent change tracking driver. The options are 0 and 1.

0: Does not install the driver
1: Installs the driver

Default:1
Example:

-driver:1

-MonitorFlag:<MONITORFLAG>

(Optional) Specifies the Arcserve Unified Data Protection Agent monitor display to users. The options are 0 and 1.

0: Displays the agent monitor to all users.
1: Displays the agent monitor only to the current user.

Default: 0.

Example:

-MonitorFlag:0

-StopUA:< STOPUA >

(Optional) Specifies to stop the Arcserve Universal Agent service.

0: Does not stop the Arcserve Universal Agent service if it is running during the installation process.
1: Stops the Arcserve Universal Agent service if it is running during the installation process.
Default: 0

Example:

-StopUA:1

Note: Use this option while upgrading to a new version. Verify that you set the value to 1 or stop the service before starting the upgrade process. This helps ensure that the installation does not fail.

-SummaryPath:<SUMMARYPATH>

(Optional) Specifies the target path to generate the summary file of the installation.

Example:

-SummaryPath:C:\Result

Note: If the value for SUMMARYPATH contains a space, enclose the path with quotation marks. Additionally, the path cannot end with a backslash character.

-AutoReboot:<AUTOREBOOT>

(Optional) Let Setup reboot the machine after installation if the installation requires a reboot. The options are 0 and 1.

0: Does not reboot the machine.

1: Reboots the machine if the installation requires a reboot.

Default: 0

Example:

-AutoReboot:1

Note: If the installation does not require a reboot, Setup will not reboot the machine even if this parameter is set to 1.

You have successfully completed the silent installation.

Install Arcserve UDP Using the Unified Installer

Arcserve UDP lets you install all its components using a single, unified installer. Based on your requirement, the installer suggests the best license to meet your requirements and then downloads and installs the components.

Follow these steps:

1. Download the ASDownloader file from the Arcserve web site.

   Note: If one of the supported non-English operating systems is detected, you need to select the language for the product installation.

2. Double-click the installation package.

   The **License Agreement** dialog opens.

3. Read and accept the terms of the licensing agreement and click **Next**.

   The **Getting Started** dialog opens.

4. Click **Next**.

   The **Choose Components to Download** dialog opens.

5. Select one or more of the following options depending on your requirement and click **Next**.

   Arcserve Unified Data Protection

   Installs Arcserve UDP. Arcserve UDP lets you protect Windows and Linux physical and virtual nodes. You can manage all your data protection needs from a single Console. You can use the global source-side deduplication, replication, remote replication and other features to manage your data.

   Arcserve Backup

   Installs Arcserve Backup. When coupled with Arcserve UDP, you can manage the tape backup from the Arcserve UDP Console and leverage all the benefits of Arcserve UDP.

   The **Product Download** dialog opens.

6. Click **Download**.

   The product begins to download in a zip format. You can check the download status on the progress bar. You can also pause and resume the download. Depending on the bandwidth and the number of components for download, this may take a while.

   Until the download is complete, the Next button is inactive.

7. When the download completes, click **Next**.

   The **Installation Method** dialog opens.

8. Select one of the installation types.

   Express Installation

   Installs the components using the default configuration. Click **View Default Configuration and Components** to see the components that are installed.

   Advanced Installation

   Lets you install each component separately.

9. Click **Next**.

   The **Account Configuration** dialog opens.

10. Specify the user name and password and click **Install**.

    The installation begins. You can see the progress on the dialog. On completion of the installation, close the wizard.

    Arcserve UDP is installed on your computer.

Verify the Installation

Follow these steps:

1. Verify that the Arcserve UDP icon appears in the system tray.
2. Verify that the agent and server services are up and running from the Windows Services Manager.

   You have successfully installed Arcserve Unified Data Protection and you are ready to back up your Windows machine.

Communication Ports Used by Arcserve UDP

This section provides information about ports used for the following components:

- Components installed on Microsoft Windows
- Components installed on Linux
- Components installed on Hypervisor

Ports listed are required for backup and other jobs when you have a LAN environment.

*Port sharing is supported for replication jobs. All data on different ports can be forwarded to port 8014 (default port for the UDP Server, can be modified during installation). When a replication job runs between two recovery point servers across WAN, only port 8014 is required to be open.

Similarly, for remote replications, the Remote administrator needs to open or forward port 8014 (for data replication) and port 8015 (default port for the UDP console, can be modified during installation) for local recovery point servers to get the assigned replication plan.

Components installed on Microsoft Windows

This section provides information about ports used in UDP Console and UDP Recovery Point Server (RPS):

- UDP Console
- UDP Recovery Point Server (RPS)
- UDP Windows Agent

UDP Console

The following table lists the ports used by Arcserve UDP Console:

Port Number	Port Type	Initiated by	Listening Process	Internal / External Port	Description
1433	TCP	Remote Java	sqlsvr.exe	External	Default communication port between the UDP console and Microsoft SQL Server databases when they reside on different computers. **Note:** You can modify the default communication port when installing SQL Server.
6052	TCP	Arcserve Backup Global Dashboard	Arcserve.CommunicationFoundation.WindowsService.exe	External	Communication that lets the UDP Console and the Arcserve Backup Global Dashboard Primary server synchronize data. **Note:** This port is only

					needed when you want to synchronize Arcserve Backup Global Dashboard data to UDP Console.
6054	TCP	Arcserve Backup Primary server	Arcserve.CommunicationFoundation.WindowsService.exe	External	Communication that lets the Console and the Arcserve Backup Primary server synchronize data. **Note:** This port is only needed when you want to synchronize Arcserve Backup Global Dashboard data to UDP Console.
8012	TCP	UDP Console	java.exe	Internal	Default port internally used by UDP Console Identity Service. **Note:** The port could not be customized and can be ignored for the firewall setting. If 8012 is occu-

					pied by other programs, UDP setup program will dynamically assign another available port.
8015	TCP	UDP Console UDP Gateway	httpd.exe	External	Default HTTP/HTTPS communication port to visit UDP Console and UDP Gateway. **Note:** You can modify the default communication port when you install the UDP components.
8029	TCP	UDP Console	tomcat9.exe	Internal	Default port internally used by UDP Console Management Service. **Note:** The port could not be customized and can be ignored for the firewall setting. If 8029 is occupied by other programs, UDP setup pro-

					gram will dynamically assign another available port.
8030	TCP	UDP Console	tomcat9.exe	Internal	Default port internally used by UDP Console Management Service. **Note:** The port could not be customized and can be ignored for the firewall setting. If 8030 is occupied by other programs, UDP setup program will dynamically assign another available port.
18007	TCP	TOMCAT	tomcat9.exe	Internal	Internally Used by Tomcat Management Service. **Note:** This port can only be changed by modifying TOMCAT configuration files. This port can be ignored for the firewall setting.

UDP Recovery Point Server (RPS)

The following table lists the ports used by Arcserve UDP Recovery Point Server (RPS):

Port Number	Port Type	Initiated by	Listening Process	Internal / External Port	Description
8014	TCP	UDP	httpd.exe	External	Default HTTP/HTTPS communication port to visit UDP RPS and UDP Agent **Notes:** ■ This port is the default shared port and the only port you must open when you use the UDP RPS as the replication destination. Do not open port 5000-5060 used by data stores with global deduplication enabled. ■ You can modify the default communication port when you install the UDP components.
8016	TCP	UDP	tomcat9.exe	Internal	Internally used by UDP RPS Web Services to communicate with the UDP RPS Port Sharing Service on the same server. **Note:** The port could not be customized and can be ignored for the firewall setting.
5000-5060	TCP	UDP	GDDServer.exe	Internal	This port range is reserved for UDP RPS Data Store Service. One UDP RPS Deduplication data store will use 4 free ports and one non-deduplication data store will use 1 free port, and they both start from 5000. It is needed when data store is processed for backup or restore. If you use RPS as the replication target only, you don't need to open them in the firewall configuration. **Note:** Change the following to customize the port range in Registry: ■ HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\DataStore

					▪ Key Name: PortRangeForGDD ▪ Type: Reg_SZ ▪ Default Value: 5000-5060 Only data store created after registry changes will use the newly changed port range.
18005	TCP	TOMCAT	tomcat9.exe	Internal	To shut down Tomcat used by the UDP RPS or Agent. **Note:** This port can only be changed by modifying TOMCAT configuration files. This port can be ignored for the firewall setting
7788	TCP	UDP	Sync_util_ d.exe	Internal	Default HTTP/HTTPS communication port to accept replication in requests.
445	TCP			External	Used by SMB service of Windows OS. The port is used when the RPS host the data store on a local disk. The data store exposes the shared folder as the backup destination for the UDP Agent to back up data.

UDP Windows Agent

The following table lists the ports used by Arcserve UDP Windows Agent:

Port Number	Port Type	Initiated by	Listening Process	Internal / External Port	Description
8014	TCP	UDP Windows Agent	tomcat9.exe	External	Default HTTP/HTTPS communication port to visit UDP RPS and UDP Agent. **Note:** You can modify the default communication port when you install the UDP components.
18005	TCP	TOMCAT	tomcat9.exe	Internal	To shut down Tomcat used by the UDP RPS or Agent. **Note:** This port can only be changed by modifying TOMCAT configuration files. This port can be ignored for the firewall setting.
4090	TCP	UDP Windows Agent	HATransServer.exe	External	To transfer data for Virtual Standby task in the proxy mode. **Note:** This port is only needed when you specify this UDP Windows Agent as Virtual Standby Monitor.
135	TCP			External	Communication port for RPC service on Windows OS. **Note:** This port is only needed whenever UDP Console remotely deploys UDP Windows Agent to this Agent machine. If the UDP Windows Agent is installed by running setup locally, this port is not required.
445	TCP			External	Communication port for SMB service to enable Shared Folder on Windows OS. **Note:** This port is only needed

					whenever UDP Console remotely deploys UDP Windows Agent to this Agent machine. If the UDP Windows Agent is installed by running setup locally, this port is not required.

Components installed on Linux

This section provides information about ports used in Linux Backup Server and Linux Nodes protected by Linux Backup Server remotely.

Linux Backup Server

The following table lists the ports used by Linux Backup Server:

Port Number	Port Type	Initiated by	Listening Process	Internal / External Port	Description
67	UDP	UDP Linux	bootpd	External	Incoming, used for PXE boot server. Only required if user wants to use the PXE boot feature. **Note:** This port number cannot be customized.
69	UDP	UDP Linux	tftpd	External	Incoming, used for the PXE boot server. Only required if user wants to use the PXE boot feature. **Note:** This port number cannot be customized.
8014	TCP	UDP Linux	java	External	Both incoming and outgoing. Default HTTP/HTTPS communication port to visit UDP Agent for Linux. **Note:** You can modify the default communication port when you install the UDP components.
8021	TCP	UDP Linux	cresvc	External	Incoming, used for backup service.
18005	TCP	UDP Linux	java	Internal	Used by Tomcat. Ignore this port for the firewall setting. **Note:** This port can only be changed by modifying TOMCAT configuration files.
50000 or 50000+	TCP	UDP Linux	ssh	External	UDP Linux third-party dependency. Required only when running Linux Migration BMR from Cloud to local. One available port from port 50000 is selected and used. By default, UDP Linux opens the selected port in the system.
22	TCP	SSH service	sshd	External	UDP Linux third-party dependency. Default for SSH service, however, you can change this port. This port is required for both incoming and out-

					going communications.
8016	TCP	UDP Linux	d2ddss	External	Incoming, used for Instant VM or Instant BMR data service. Only required if user wants to use the Instant VM or Instant BMR feature.
111	TCP	Port mapper	rpcbind	External	UDP Linux third-party dependency. Default for port mapper. Only required when running Linux Instant VM from agentless backup recovery point to the vSphere ESX server.
2049	TCP	NFS server	nfsd	External	UDP Linux third-party dependency. Default for NFS server. Only required when running Linux Instant VM from agentless backup recovery point to the vSphere ESX server.
Dynamic port	TCP	NFS mount service	rpc.mountd	External	UDP Linux third-party dependency. See rpc.mountd's man page for how to make it listen to fixed port. Only required when run Linux Instant VM from agentless backup recovery point to the vSphere ESX server.

Linux Nodes Protected by Linux Backup Server Remotely

The following table lists the ports used by Linux Nodes Protected by Linux Backup Server Remotely:

Port Number	Port Type	Initiated by	Listening Process	Internal / External Port	Description
22	TCP	SSH service		External	UDP Linux third-party dependency. Default for SSH service, however, you can change this port. This port is required for both incoming and outgoing communications.

Components installed on Hypervisor

This section provides information about the ports used for Hyper-V host.

Hyper-V host

The following table lists the ports used by Hyper-V host:

Port Number	Port Type	Initiated by	Listening Process	Internal/External	Description
135	TCP			External	Used by WMI service of Windows OS. UDP uses WMI to interact with Hyper-V host in some situations.
445	TCP			External	Used by SMB service of Windows OS. UDP uses SMB to interact with Hyper-V host in some situations.
27000	TCP	UDP CBT Service	cbt_rep.exe	External	Used by UDP Host-based Backup CBT service. You do not need to register this port to the firewall exception because UDP automatically registers this port during backup. You need to verify that no other application is configured with the same port.
5895 or 5986	TCP	WinRM service		External	5895 for WinRM in HTTP protocol, 5896 for WinRM in HTTPS protocol. Only one of them is required. Only required when running Linux Instant VM to Hyper-V.
1024 ~ 65535	TCP			External	▪ Used while importing VMs from the Hyper-V host and/or the cluster for backup plan. ▪ By default, VM restore job randomly chooses an available port in the range of 1024 and 65535. You can manually specify the range by registry values set in Hyper-V host. For more information, see link.

How the Installation Process Affects Operating Systems

The following installation processes update various Windows Operating Systems:

- Installation of Unsigned Binary Files
- Installation of Binary Files with Incorrect File Version
- Installation of Binary Files Without OS in Manifest

Installation of Unsigned Binary Files

Binary Name	Source	Binary Name	Source
AxShockwaveFlashObjects.dll	Adobe	freetype.dll	Oracle
ShockwaveFlashObjects.dll	Adobe	hprof.dll	Oracle
httpd.exe	Apache	instrument.dll	Oracle
libapr-1.dll	Apache	j2pcsc.dll	Oracle
libapriconv-1.dll	Apache	j2pkcs11.dll	Oracle
libaprutil-1.dll	Apache	jaas_nt.dll	Oracle
libeay32.dll	Apache	jabswitch.exe	Oracle
libexpat.dll	Apache	java-rmi.exe	Oracle
libhttpd.dll	Apache	java.dll	Oracle
openssl.exe	Apache	java.exe	Oracle
pcre.dll	Apache	JavaAccessBridge-64.dll	Oracle
rotatelogs.exe	Apache	javaw.exe	Oracle
ssleay32.dll	Apache	java_crw_demo.dll	Oracle
tcnative-1.dll	Apache	jawt.dll	Oracle
CAPatch.dll	APM	JAWTAccessBridge-64.dll	Oracle
AsyncClient.net.dll	Axcient	jdwp.dll	Oracle
doclib.dll	Axcient	jjs.exe	Oracle
docxlib.dll	Axcient	jli.dll	Oracle
DspchConnector.dll	Axcient	jpeg.dll	Oracle
esr.exe	Axcient	jsdt.dll	Oracle
esrdf.dll	Axcient	jsound.dll	Oracle
esrsdll.dll	Axcient	jsoundds.dll	Oracle
eswrapper.dll	Axcient	keytool.exe	Oracle
html2text.dll	Axcient	kinit.exe	Oracle
licensemanager.dll	Axcient	klist.exe	Oracle
mhdll.dll	Axcient	ktab.exe	Oracle
pdflib.dll	Axcient	lcms.dll	Oracle
pptlib.dll	Axcient	management.dll	Oracle
pptxlib.dll	Axcient	mlib_image.dll	Oracle
protection.dll	Axcient	net.dll	Oracle
pstgen.dll	Axcient	nio.dll	Oracle
resources.dll	Axcient	npt.dll	Oracle
rtf2html.dll	Axcient	ojdkbuild_giflib.dll	Oracle
rtflib.dll	Axcient	ojdkbuild_libjpeg-turbo.dll	Oracle

SourceLibrary.dll	Axcient	ojdkbuild_libpng.dll	Oracle
uicommon.dll	Axcient	ojdkbuild_nss.dll	Oracle
xlslib.dll	Axcient	orbd.exe	Oracle
xlsxlib.dll	Axcient	pack200.exe	Oracle
libbind9.dll	Bind	policytool.exe	Oracle
libdns.dll	Bind	rmid.exe	Oracle
libisc.dll	Bind	rmiregistry.exe	Oracle
libisccfg.dll	Bind	sawindbg.dll	Oracle
liblwres.dll	Bind	servertool.exe	Oracle
libxml2.dll	Bind	splashscreen.dll	Oracle
msvcm80.dll	Bind	sunec.dll	Oracle
win_nsupdate.exe	Bind	sunmscapi.dll	Oracle
msvcm90.dll	Microsoft	tnameserv.exe	Oracle
RDXCleanerX64.EXE	Microsoft	unpack.dll	Oracle
RDXForceCleanX64.EXE	Microsoft	unpack200.exe	Oracle
MSCHRT20.OCX	Microsoft	verify.dll	Oracle
Microsoft.Exchange.WebServices.dll	Office365	w2k_lsa_auth.dll	Oracle
SQLite.CodeFirst.dll	Office365	WindowsAccessBridge-64.dll	Oracle
System.Data.SQLite.dll	Office365	zip.dll	Oracle
System.Data.SQLite.EF6.dll	Office365	jvm.dll	Oracle
System.Data.SQLite.Linq.dll	Office365	JavaAccessBridge-32.dll	Oracle
System.Management.Automation.dll	Office365	JavaAccessBridge.dll	Oracle
SQLite.Interop.dll	Office365	JAWTAccessBridge-32.dll	Oracle
libxml.dll	OpenSSL	JAWTAccessBridge.dll	Oracle
attach.dll	Oracle	WindowsAccessBridge-32.dll	Oracle
awt.dll	Oracle	WindowsAccessBridge.dll	Oracle
dt_shmem.dll	Oracle	plink.exe	Putty
dt_socket.dll	Oracle	sqlite3.exe	SQLite
fontmanager.dll	Oracle	MinHook.x64.dll	Tsuda Kageyu
javacpl.cpl	Oracle	zlib10.dll	Zlib

Installation of Binary Files with Incorrect File Version

Binary Name	Source	Binary Name	Source
AxShockwaveFlashObjects.dll	Adobe	libisccfg.dll	bind
ShockwaveFlashObjects.dll	Adobe	liblwres.dll	bind
openssl.exe	Apache	libxml2.dll	bind
UpdateData.exe	Arcserve licensing	win_nsupdate.exe	bind
AsyncClient.net.dll	Axcient	libxml.dll	NetApp
doclib.dll	Axcient	decora-sse.dll	Oracle
docxlib.dll	Axcient	fxplugins.dll	Oracle
DspchConnector.dll	Axcient	glass.dll	Oracle
esr.exe	Axcient	glib-lite.dll	Oracle
esrdf.dll	Axcient	gstreamer-lite.dll	Oracle
esrsdll.dll	Axcient	javafx-font.dll	Oracle
eswrapper.dll	Axcient	javafx-iio.dll	Oracle
html2text.dll	Axcient	jfxmedia.dll	Oracle
licensemanager.dll	Axcient	jfxwebkit.dll	Oracle
mhdll.dll	Axcient	libxml2.dll	Oracle
pdflib.dll	Axcient	libxslt.dll	Oracle
pptlib.dll	Axcient	prism-d3d.dll	Oracle
pptxlib.dll	Axcient	sqlite3.exe	sqlite
protection.dll	Axcient	libcurl.dll	VMware
pstgen.dll	Axcient	libexpat.dll	VMware
resources.dll	Axcient	liblber.dll	VMware
rtf2html.dll	Axcient	libldap.dll	VMware
rtflib.dll	Axcient	libldap_r.dll	VMware
uicommon.dll	Axcient	libxml2.dll	VMware
xlslib.dll	Axcient	ssoclient.dll	VMware
xlsxlib.dll	Axcient	vddkReporter.exe	VMware
libbind9.dll	bind	zlib1.dll	zlib
libdns.dll	bind	zlib10.dll	zlib
libisc.dll	bind		

Installation of Binary Files Without OS in Manifest

Binary Name	Source
openssl.exe	Apache
win_nsupdate.exe	bind
plink.exe	putty
sqlite3.exe	sqlite
vddkReporter.exe	VMware

Antivirus Configuration

Antivirus software can interfere with the smooth running of Arcserve UDP Agent (Windows) by either temporarily blocking access to files or by quarantining or deleting files that are incorrectly classified as suspicious or dangerous. You can configure an antivirus software to exclude files, or folders and avoid spending time on scanning data that does not need to be protected.

Configuring your antivirus software properly is required to exclude backup destination. Proper configuration lets you prevent interference with backup and restore operations, or any other processes like merge and catalog generation.

Paths to be excluded for antivirus scanning:

Antivirus scanning for	Paths to exclude
RPS	Data Store Folder Data Destination Path Index Destination Path Hash Destination Path
Windows Agent not managed by Console	Backup Destination

How to Install Arcserve UDP Updates

The process of getting and installing Arcserve UDP updates is a two-step process that involves checking and downloading the update, and then installing the update.

Note: All updates that are released for Arcserve UDP are cumulative. As a result, each update also includes all previously released updates to ensure that your computer is always up-to-date. The **Help About** dialog displays the update level that is installed on a computer. If necessary, you can use this information to build another server with the same configuration / patch level.

Perform the following tasks to install Arcserve UDP updates:

1. Review the Considerations for Installing Updates
2. Specify Update Preferences
3. Check and Install the Updates
4. (Optional) Install Arcserve UDP Updates Silently
5. Verify that the Updates are Successfully Installed

Review the Considerations for Installing Updates

Review the following considerations before installing Arcserve UDP updates:

- When installing an Arcserve UDP update or an Arcserve Unified Data Protection Agent for Windows update, it is important to maintain optimal performance between the Console, the Recovery Point Server (RPS), and Agents. As a result, when the update is installed in an environment that contains both a Console and an Agent, you must always install the update on the Console first, and then on the RPS, and finally on the Agent. (For the Agent that is installed on the Console or the RPS, the update will be automatically installed on that Agent at the same time).
- If necessary, download available updates from Arcserve either directly to a client machine or to a staging server first and then to a client machine.
- If necessary, use your workstation node as a staging server for downloading Arcserve UDP updates.
- Verify that the Update preference settings are properly configured.
  - Updates can be installed either through the user interface or silently using the command line.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

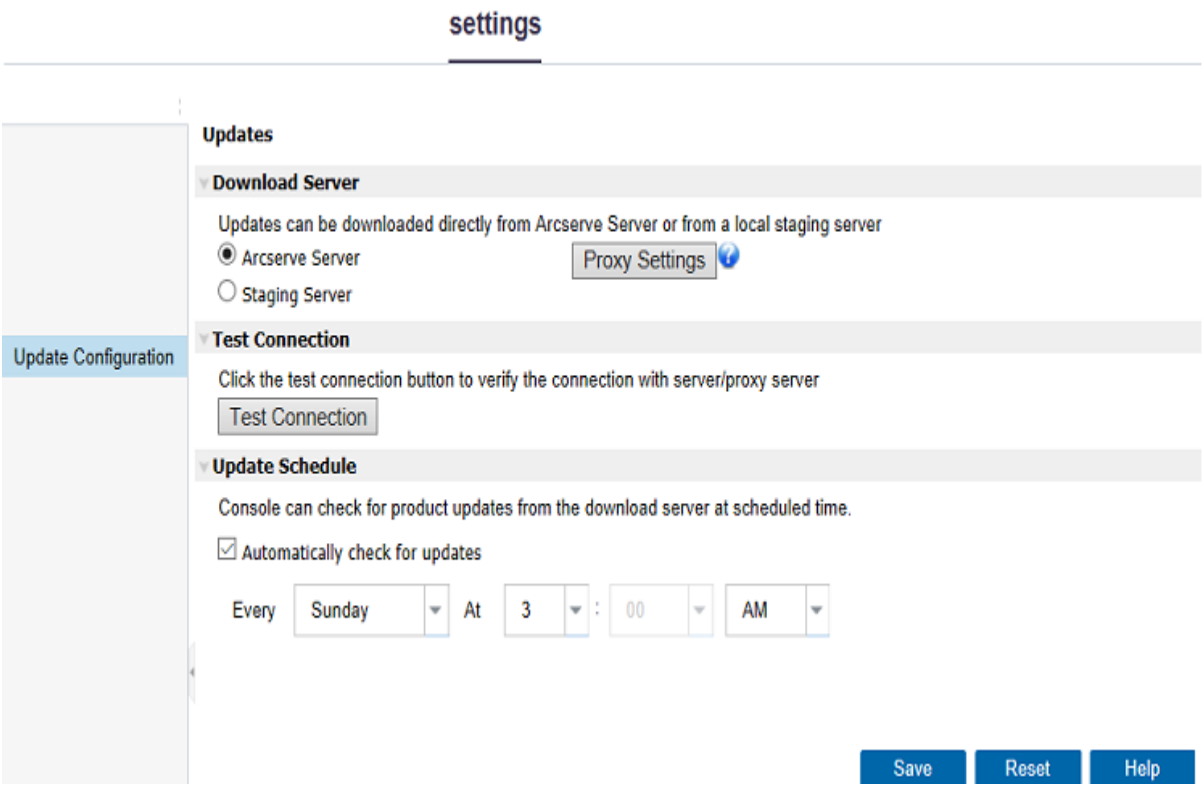
Specify Updates Preference

Using Arcserve UDP you can specify your updates preference.

Follow these steps:

1. From the Arcserve UDP Console, click the **settings** tab.
2. From the left pane, click **Update Configuration**.

   The **Updates** page is displayed on the right pane.



3. Specify your **Updates** preference settings.

   Download Server

   Specifies the source server from where your Arcserve UDP server will connect to and download available updates.

   - **Arcserve Server**

     Specifies that updates are downloaded from the Arcserve server directly to your local server.

     This is the default setting.

   - **Staging Server**

     Specifies that updates are downloaded from the staging server.

Note: If required, you can create a staging server. For more information, see How to Create a Staging Server.

If you specify more than one staging server, the first listed server is designated as the primary staging server. Arcserve UDP initially attempts to connect to the primary staging server. If the first listed server is not available, then the next listed server becomes the primary staging server. The same sequence is continued until the last listed server becomes the primary staging server. (The Staging Server list is limited to the maximum of 5 servers).

- You can use the **Move Up** and **Move Down** buttons to change the staging server sequence.
- You can use the **Delete** button to remove a server from this listing.
- You can use the **Add Server** button to add a new server to this listing. When you click the **Add Server** button, the **Staging Server** dialog opens, allowing you to specify the name of the added staging server.

When you select the staging server as your download server, then:

- If the specified staging server has any update, then the Arcserve UDP Console can get update from this staging server.
- If the specified staging server has no update, then the Arcserve UDP Console is unable to download update from this staging server. The log displays the following message:

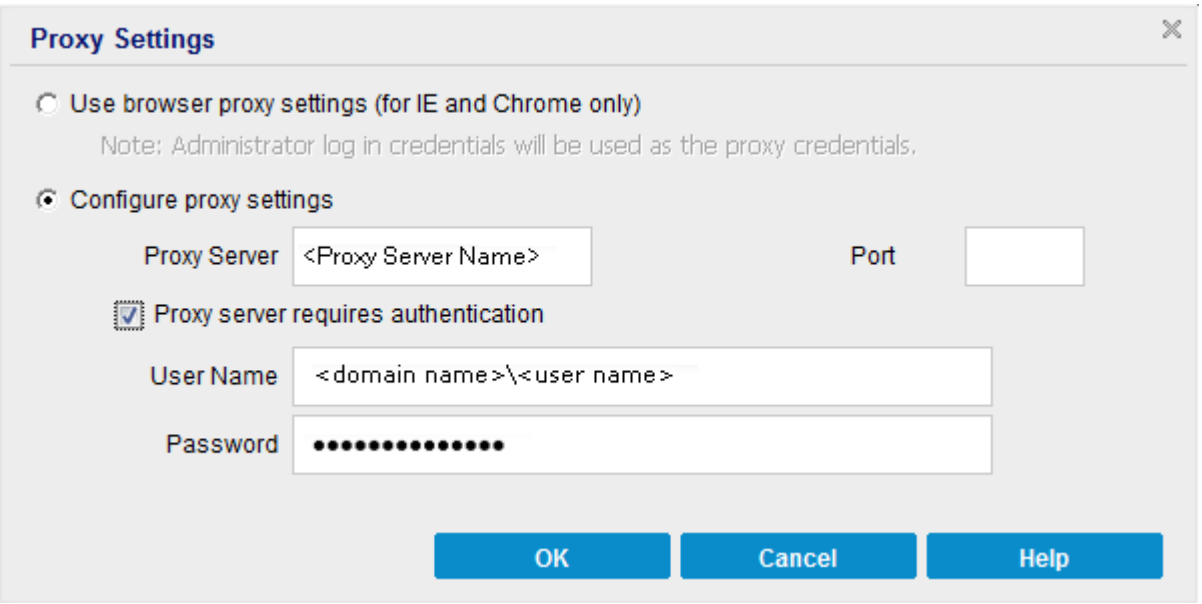
  No new update available.

- **Proxy Settings**

  Note: This **Proxy Server** option is only available when you select the Arcserve Server as the download server.

  Select Proxy Settings

  When you select this option the **Proxy Settings** dialog opens.



Use browser proxy settings

This selection is only applicable to Windows Internet Explorer (IE) and Google Chrome.

When selected, directs Arcserve UDP to automatically detect and use the same proxy settings that are applied to the browser to connect to the Arcserve server for Arcserve UDP update information.

Configure proxy settings

When selected, enables the specified proxy server to connect to the Arcserve server for Arcserve UDP update information. If you select this option, you must also include the IP address (or machine name) of the proxy server and the corresponding port number that is used by the proxy server for internet connections.

In addition, you can also specify if your proxy server will require authentication. When selected, specifies that authentication information (User ID and Password) are required to use the proxy server.

Note: The format for user name should be a fully qualified domain user name in the form of "<domain name>\<user name>".

Test Connection

Lets you test the following connections and displays a status message when completed:

- If you selected "Arcserve server" as the download server, tests the connection between the machine and the Arcserve server through the specified proxy server.

- If you selected "Staging Server" as the download server, tests the connection between the machine and the specified staging server. The test connection button is used to test the availability of each listed staging server, and a corresponding status is displayed in the **Connection Status** field. If none of the configured staging servers are available, the following message is displayed at the top of the Arcserve UDP Console: Update server unavailable.

  Note: The test connection is automatically performed when you open the **Update Configuration** page from the **settings** tab in the Arcserve UDP Console. When this auto test is performed, it will check the latest connection status of the previously configured download server (either Arcserve server or Staging Server(s), whichever is selected). If you previously configured more than one staging server, then this auto test is performed on all staging servers to get the latest connection status.

Update Schedule

Specifies when to check for (and download) new Arcserve UDP updates.

4. Click **Save**.

   Your Updates preference settings are saved.

How to create a Staging Server

Staging server is a node on which the Arcserve UDP Agent or Console is installed. The node needs to finish downloading updates from Arcserve download server to work as a staging server to provide updates for others.

Adding a Staging Server:

You can add Staging server on any node that meets the following two requirements:

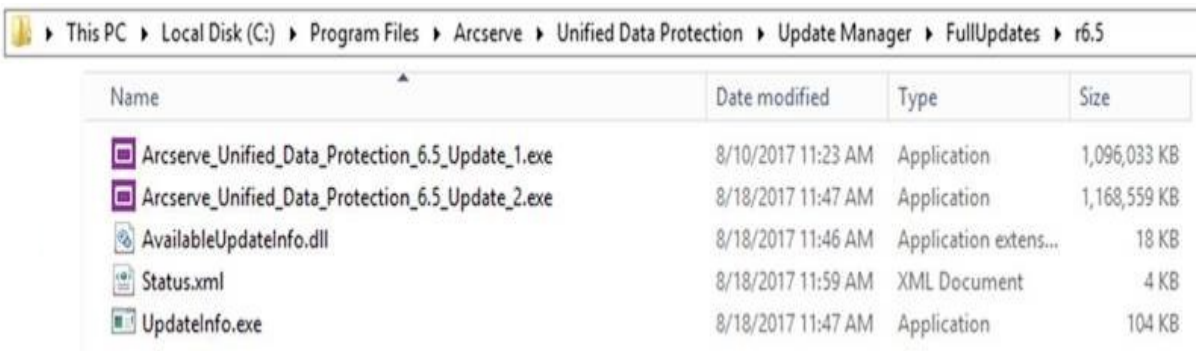
- The node has either Arcserve UDP agent or Arcserve UDP Console installed.
  - Arcserve UDP console can only download updates from console staging server.
  - Arcserve UDP agent can download updates from the console or agent staging server.
- The node has successfully downloaded updates from Arcserve download server at least once.

Note: Other configuration is not required if the target staging server meets the prerequisites.

Configure Update from Staging Server

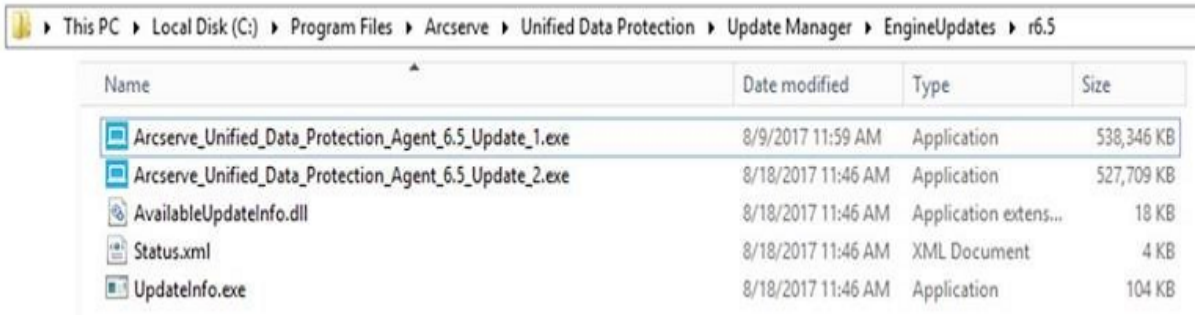
- When Staging server is selected per Arcserve UDP Console installed on that node, consider the following notes:
  - Default Port: 8015
  - Arcserve UDP Console can get updates from this Staging server
  - Arcserve UDP Agent can get updates from this Staging server

  Note: From the directory < UDP Installed path\Update Manager\FullUpdates> on the target staging server verify the available latest Update version. View example below.



- When Staging server is selected per Arcserve UDP Agent installed on that node, consider the following notes:
  - Default Port: 8014
  - Arcserve UDP Console cannot get updates from this Staging server
  - Arcserve UDP Agent can get updates from this Staging server

  Note: From the directory <UDP Installed path\Update Manager\EngineUpdates> on the target staging server verify the available latest Update version. View example below.



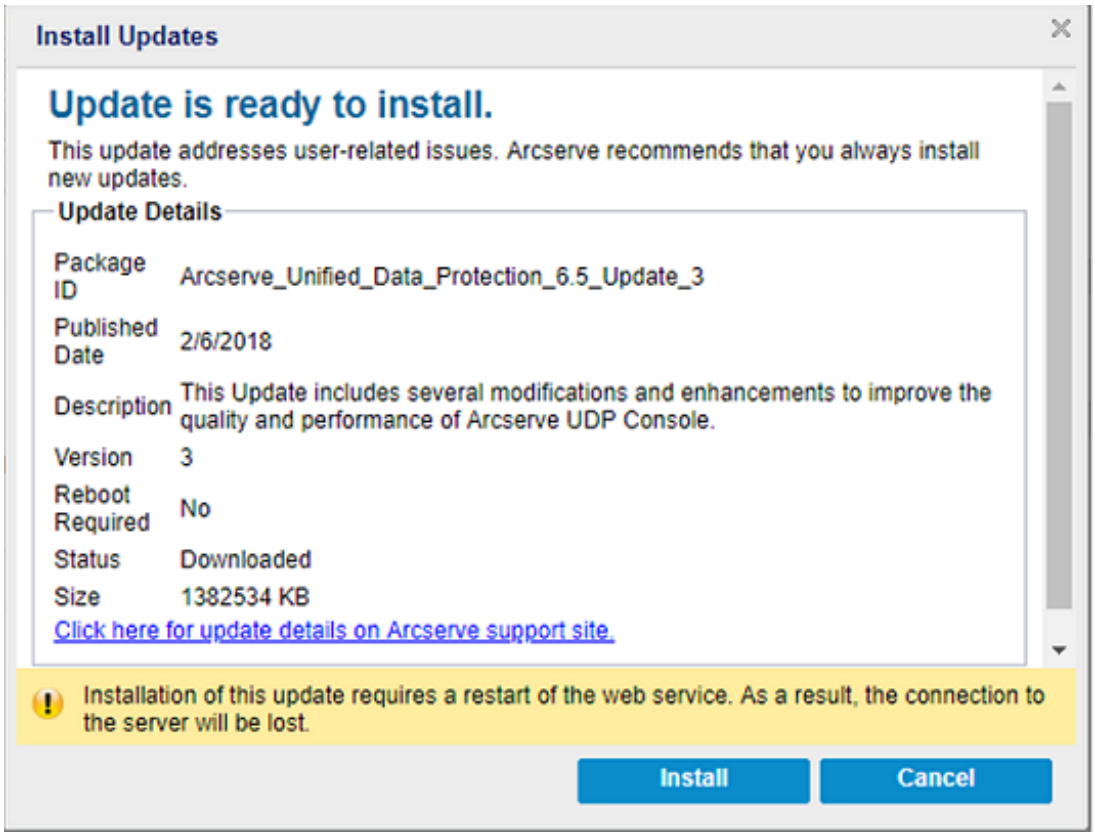
Check and Install the Updates

From the UDP Console, you can determine if any new updates are available.

Follow these steps:

1. Click **Check for Updates** from the **Help** drop-down menu. When a new update is available, a message is displayed at the top bar. Also, the **Update Installation** dialog is displayed.
2. If you enable the update schedule, and when a new update is available, it is automatically downloaded to the UDP server. A **New Update Available** link is displayed on the top bar to provide a visual indication that a new update is ready to install.
3. Click the **New Update Available link** on the top bar.

   The **Install Updates** dialog opens to display information that is related to the available update. The dialog includes information such as description, download status, size, reboot requirement, and a link to the Arcserve server for additional update details.



4. Click **Install**.

   Installation of Arcserve UDP updates starts.

Update Agents on Remote Nodes Using Gateway

The remote nodes and server in a Site interact with the Console using a gateway. Using Arcserve UDP, you can discover and deploy the latest version of the Agent to nodes. To upgrade or install Arcserve UDP Agent on a node in a site, use Install/Upgrade Agent.

Follow these steps:

1. Click the **resources** tab on the Console.
2. From the left **Navigation** pane, select a site from the drop-down list.
3. The **Nodes: All Nodes** page is displayed.
4. Select one or more nodes.
5. From the center pane, click the **Actions** drop-down list, and then click **Install/Upgrade Agent**.

   The details of Install or upgrade is displayed on the center pane.
6. Verify the details and click **OK**.

   The node is installed or upgraded with the latest version of Arcserve UDP Agent.

Update RPS on Remote Server Using Gateway

The remote nodes and server in a Site interact with the Console using a gateway. Using Arcserve UDP, you can discover and deploy the latest version of the RPS component to recovery point servers. After you deploy the RPS component, the server is ready to store the backup sessions and serves as a recovery point server. To upgrade or install the Arcserve UDP RPS component on a recovery point servers in a site, use Install/Upgrade Recovery Point Server.

Follow these steps:

1. Click the **resources** tab on the Console.
2. From the left **Navigation** pane, select a site from the drop-down list.
3. Click **Recovery Point Servers**.

   The **Destination: Recovery Point Server** page is displayed.
4. Perform one of the following actions:
   - Right-click a recovery point server.
   - Select a recovery point server, and from the center pane click the **Actions** drop-down list.

   A list of options is displayed.
5. Click **Install/Upgrade Recovery Point Server**.

   The **Installation and Upgrade** page is displayed.
6. Modify the deployment settings, and click **OK** to deploy the recovery point server on the selected node.

   The recovery point server deployment starts. You can view the deployment progress on the right pane.

(Optional) Install Arcserve UDP Updates Silently

Silent update installation allows you to perform an unattended update installation and does not prompt you for any input.

Follow these steps:

1. Launch the Arcserve UDP Update silent installation.

   "<UpdateExeFile>" /s /v"<Additional Arguments>"

2. Configure the silent installation using the following syntax and arguments:

   UpdateExeFile

   Specifies to run the self-extracting executable file.

   s

   Specifies to run the self-extracting executable file using the silent mode.

   v

   Specifies any additional arguments for update installation.

   Additional Arguments

   /s

   Specifies to run the update installation using the silent mode.

   /AutoReboot

   Specifies to perform an automatic reboot after the update is installed. If a reboot is required to complete the update, the machine will reboot automatically without any notification.

Examples

- To install an update using the silent mode and reboot automatically after completion, use the following command:

  "<UpdateExeFile>" /s /v"/s /AutoReboot"

- To install an update using the silent mode and not reboot automatically after completion, use the following command:

  "<UpdateExeFile>" /s /v"/s"

Verify that the Updates are Successfully Installed

From the Arcserve UDP Console, select **Help**, click **About**, and then verify that the about Arcserve UDP dialog displays the latest version updated.

How to Uninstall Arcserve UDP

You can uninstall Arcserve Unified Data Protection using the following methods:

- **Standard uninstallation**: Use this method to uninstall using the Windows Control Panel.
- **Silent uninstallation**: Use this method to perform an unattended uninstallation using the Windows Command Line.

Standard Uninstall

You can uninstall the following components.

- Arcserve Unified Data Protection Console
- Arcserve Unified Data Protection Recovery Point Server
- Arcserve Unified Data Protection Agent

Follow these steps:

1. Open the Windows Control Panel.
2. Click Uninstall a program.

   The Uninstall or change a program dialog opens.
3. Select Arcserve Unified Data Protection and click Uninstall.

   The Arcserve Unified Data Protection Uninstall Application dialog opens.
4. Select the components to uninstall and click Next.

   The **Messages** dialog opens.
5. Click Next.

   The **Remove Components** dialog opens.
6. Click **Remove**.

   The selected components are uninstalled from the computer.

Silent Uninstall

A silent uninstallation eliminates the need for user interaction while performing uninstall.

Follow these steps:

1. Log into the computer to uninstall Arcserve UDP components.

   Note: Log into the computer using an administrative account.

2. Open the Windows command line and run the following command that corresponds with the specified operating system:

   - **x86 operating system:**

     To uninstall all components

     %ProgramFiles%\Arcserve\SharedComponents\Arcserve Unified Data Protection\Setup\uninstall /q /ALL

     To uninstall selected components

     %ProgramFiles%\Arcserve\SharedComponents\Arcserve Unified Data Protection\Setup\uninstall /q /p <Product Code>

   - **x64 operating system:**

     To uninstall all components

     %ProgramFiles(x86)%\Arcserve\SharedComponents\Arcserve Unified Data Protection\Setup\uninstall /q /ALL

     To uninstall selected components

     %ProgramFiles(x86)%\Arcserve\SharedComponents\Arcserve Unified Data Protection\Setup\uninstall /q /p <Product Code>

   The following values explain the return codes:

   0 = Uninstall was successful.

   3010 = Uninstall was successful, but a reboot is required.

   Other = Uninstall failed.

   Usage:

   The table listed below defines the product code that you must specify for the Arcserve Unified Data Protection component that you want to uninstall.

   Example:

   The following syntax lets you uninstall Arcserve Unified Data Protection Recovery Point Server silently.

"%ProgramFiles(x86)%\Arcserve\SharedComponents\Arcserve Unified Data Protection\Setup\uninstall.exe" /q /p {CAAD8172-1858-4DC7-AE81-C887FA6AFB19}

Component	<Product Code>
Arcserve Unified Data Protection Agent (x86 platforms)	{CAAD8AEA-A455-4A9F-9B48-C3838976646A}
Arcserve Unified Data Protection Agent (x64 platforms)	{CAAD1E08-FC33-462F-B5F8-DE9B765F2C1E}
Arcserve Unified Data Protection Recovery Point Server	{CAAD8172-1858-4DC7-AE81-C887FA6AFB19}
Arcserve Unified Data Protection Console	{CAAD3E40-C804-4FF0-B1C0-26D534D438C0}
Arcserve Unified Data Protection Gateway	{FB95E75D-494F-4146-9B35-F867434B264A}

After the command is executed, Arcserve UDP components are uninstalled.

Remove Components Left Behind by the Uninstaller

Important!

1. The Arcserve licensing is shared by all Arcserve products. Ensure that you do not have any other Arcserve product installed on your machine or else you may lose the licensing for all Arcserve products installed on that machine.
2. If the components are removed, any programs that are installed after Arcserve Unified Data Protection Agent for Windows and depend on these components may not function properly.

   If you want to manually remove these components, perform the following steps:

   Remove Arcserve Licensing Component manually

1. Go to *C:\Program Files (x86)\Arcserve\SharedComponents\CA_LIC* directory.
2. Find the zip file named *lic98_uninstaller.zip* and unzip that file to some other location (for example: C:\temp).
3. Go to the location where the files were extracted and locate two script files that are named *rmlic.exe* and *rmlicense.bat*.
4. Click on *rmlicense.bat* to execute the script which uninstalls the components.
5. Manually delete the following folders:
   - C:\Program Files (x86)\Arcserve
   - C:\Program Files\Arcserve
   - Folder where you extracted the zip file.
6. Remove the registry key for the Arcserve Licensing component.
   - For x64 platform: HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Arcserve\License
   - For x86 platform: HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\License

   Remove Microsoft Visual C++ and Microsoft SQL Server Express manually

1. Access the standard Add or Remove Programs application located in the Windows Control Panel (Control Panel, Programs and Features, Remove Programs).
2. Select *Microsoft Visual C++ 2013 Redistributable (x86) - 12.0.30501* and then click Uninstall.
3. Select *Microsoft Visual C++ 2013 Redistributable (x64) - 12.0.30501* and then click Uninstall.
4. Select Microsoft SQL Server 2014 (64-bit) and then click Uninstall.

5. To remove only the Arcserve UDP database, select "ARCSERVE_APP" and click Uninstall.

Chapter 4: Upgrading to Arcserve UDP Version 8.0

This section contains the following topics:

Supported Versions for Upgrade

Arcserve UDP version 8.0 provides upgrade and backward compatibility according to the following rules:

- Arcserve UDP 8.0 supports upgrade from Arcserve UDP version 7.0 Update 2, 7.0 Update 1, 7.0 final version and Arcserve UDP version 6.5 Update 4.

  Notes:

  - For other previous versions of Arcserve UDP, such as Arcserve UDP 6.5 with Update 3, Update 2, Update 1, or final version, first upgrade to UDP 6.5 Update 4, and then you can upgrade to Arcserve UDP version 8.0.
  - For other previous versions of Arcserve UDP, such as Arcserve UDP version 6.0 Update 3 for Windows, first upgrade to Arcserve UDP version 7.0. Then, you can upgrade to Arcserve UDP version 8.0.
  - For other previous versions of Arcserve UDP, such as Arcserve UDP version 6.0 Update 2, Update 1 or final version for Windows, first upgrade to Arcserve UDP version 6.5 Update 4. Then, you can upgrade to Arcserve UDP version 8.0.
  - For other previous versions of Arcserve UDP, such as Arcserve UDP version 6.0 Update 1 or final version for Linux, first upgrade to Arcserve UDP version 6.5 Update 4. Then, you can upgrade to Arcserve UDP version 8.0.
  - For other previous versions of Arcserve UDP, such as Arcserve UDP 5.0 Update 4, first upgrade to UDP 6.5 Update 4, and then you can upgrade to Arcserve UDP version 8.0.
  - For other previous versions of Arcserve UDP, such as Arcserve UDP 5.0 with Update 3, Update 2, Update 1, or final version, first upgrade to UDP 6.0 Update 3, then upgrade to UDP 7.0, and finally you can upgrade to Arcserve UDP version 8.0.

- Arcserve UDP version 8.0 does not support upgrade from Arcserve D2D r16.5, Arcserve Central Protection Management r16.5, Arcserve Data Protection Console r16.5, Arcserve Central Reporting r16.5, and Arcserve Central Virtual Standby r16.5.

  For more information, refer to How to Upgrade License of Arcserve UDP.

Upgrade Sequence on the Arcserve Appliance

The Arcserve UDP Version 8.0 supports upgrade on the Appliance. The upgrade could involve one of the following sequences:

- Upgrade Arcserve UDP
  - Upgrade the Arcserve Appliance that performs as Arcserve UDP Console and RPS
  - Upgrade the Arcserve Appliance that performs as Arcserve UDP RPS only
  - Upgrade Steps When Two or More Arcserve Appliances Are Used in the Environment
- Upgrade the Arcserve UDP Linux Agent on the Arcserve Appliance
- Upgrade the Arcserve Backup on the Arcserve Appliance

Upgrade Arcserve Appliance that performs as Arcserve UDP Console and RPS

Upgrade Arcserve Appliance and then follow the upgrade sequence described to upgrade the environment.

Upgrade Arcserve Appliance that performs as Arcserve UDP RPS only

Upgrade the complete productive environment. For more information, see upgrade sequence.

Upgrade Steps When Two or More Arcserve Appliances Are Used in the Environment

Upgrade the whole productive environment. For more information, see upgrade sequence.

Upgrade Arcserve UDP Linux Agent on Arcserve Appliance

1. Upgrade the Arcserve UDP Console that manages the Linux Backup Server environment.
2. Upgrade the Linux Backup Server on the Arcserve Appliance. For details, refer to the *Arcserve Unified Data Protection Agent for Linux Online Help*.

Upgrade Arcserve Backup on Arcserve Appliance

Refer to *Arcserve Backup Implementation Guide* to complete upgrade on Arcserve Appliance.

Backward Compatibility Support Policy

We recommend upgrading all components to Arcserve UDP version 8.0 in the complete environment so that new features in Arcserve UDP version 8.0 can be explored. The existing backup plans are not impacted if the upgrade does not complete for all the components at the same time.

Backward Compatibility supports Arcserve UDP version 7.0, 7.0 Update 1, 7.0 Update 2, and Arcserve UDP 6.5 Update 4.

This section contains the following topics:

Backward Compatibility Support Policy for Arcserve UDP Version 7.x

Backward Compatibility Support between different UDP Console through the Replicate to a remotely-managed RPS task

If you plan to perform the *Replicate to a remotely-managed RPS* task, we recommend to upgrade the Console and RPS at destination before upgrading the Console / RPS / Agent at source.

- Replication is supported if the source RPS has version 8.0, while the destination RPS still has version UDP 7.0, 7.0 Update 1, and 7.0 Update 2.
- Replication is supported if the source RPS has version UDP 7.0, 7.0 Update 1, and 7.0 Update 2, and the destination RPS has version 8.0.
- Replication behaviors:
  - If you upgrade to UDP 8.0 on both sites, you can replicate from Cloud Hybrid to on-premises without a public IP (new UDP behavior).
  - If you upgrade UDP 8.0 only on one of the sites and the other site uses a supported version of UDP, you can replicate from Cloud Hybrid to on-premises using the public IP.

Backward Compatibility Support for Windows RPS / Agent (or agentless backup proxy)

- Arcserve UDP 8.0 Console supports recovery point servers (RPS) and Agents (or agentless backup proxies) of version UDP 7.0, 7.0 Update 1, and 7.0 Update 2 for backup if the plan does not change after an upgrade.
- Arcserve UDP 8.0 Console and RPS support Agents (or agentless backup proxies) of version UDP 7.0, 7.0 Update 1, and 7.0 Update 2 for backup if the plan does not change after an upgrade.
- Arcserve UDP 8.0 Console supports Deploying existing plan without any changes when RPS and Agents (or agentless backup proxies) are still running on version UDP 7.0, 7.0 Update 1, and 7.0 Update 2. Agent is not automatically upgraded.
- Arcserve UDP 8.0 Console support actions such as Pause or Resume existing plan without any modification when RPS and Agents (or agentless backup proxies) are still running on the version UDP 7.0, 7.0 Update 1, and 7.0 Update 2. Agent is not automatically upgraded.

- Arcserve UDP 8.0 Console supports Creating new plan or modifying existing plan to configure legacy features that was supported already in previous release when RPS and Agents (or agentless backup proxies) are still running on version UDP 7.0, 7.0 Update 1, and 7.0 Update 2. Agent is not automatically upgraded.
- Arcserve UDP 8.0 Console supports Creating new plan or Modifying existing plan to include 8.0 new features when associated RPS, Proxy or Agent are already upgraded to expected version.

  For example:

  - Specify Network for backup is supported when Console, RPS, Proxy and Agent are upgraded to UDP 7.0, 7.0 Update 1, 7.0 Update 2 or 8.0.
- Arcserve UDP 8.0 Console supports Update Node when RPS and Agents (or agentless backup proxies) are still running on version UDP 7.0, 7.0 Update 1, and 7.0 Update 2. Plan will be deployed. Agent is not automatically upgraded.
- Arcserve UDP 8.0 Console supports adding a version UDP 7.0, 7.0 Update 1, and 7.0 Update 2 of RPS server. But, the RPS is not automatically upgraded to Arcserve UDP version 8.0 when added.
- Arcserve UDP 8.0 Console supports adding a previous version of Agent (or agentless backup proxy) without automatically upgrading the Agent.
- Arcserve UDP 8.0 Console supports adding node to existing plan.
  - Adding clean node to the existing plan is supported when associated RPS upgrades to 8.0. Otherwise, plan is saved but plan deployment fails.
  - Adding previous version node to the existing plan is supported if the version is similar to the running RPS version. Agent is not automatically upgraded.
  - Adding Arcserve UDP 8.0 node to the existing plan is supported when associated RPS upgrades to 8.0.
- Gateway is automatically upgraded to match the version of Arcserve UDP Console.
- Replication backward compatibility policy:
  - Replication is supported when Console is upgraded to Arcserve UDP 8.0 but all the associated Recovery Point servers (RPSs) are still running on UDP 7.0, 7.0 Update 1, and 7.0 Update 2.
    - Replication is supported from Arcserve UDP 8.0 source RPS to UDP 7.0, 7.0 Update 1, and 7.0 Update 2 destination RPS.

- Replication is supported from Arcserve UDP UDP 7.0, 7.0 Update 1, and 7.0 Update 2 source RPS to 8.0 destination RPS.

- Virtual Standby backward compatibility policy:
  - VSB Monitor/Proxy version should be similar to the Agent version if the backup destination is a shared folder.
  - VSB Monitor/Proxy version should be similar to the RPS version if the backup destination is RPS.
- Instant Virtual Machine backward compatibility policy:

  Instant VM recovery server version should be similar to the Arcserve UDP console version.

Backward Compatibility Support Policy for Arcserve UDP Version 6.5 Update 4

Backward Compatibility Support between different UDP Console through the Replicate to a remotely-managed RPS task

If you plan to perform the *Replicate to a remotely-managed RPS* task, we recommend to upgrade the Console and RPS at destination before upgrading the Console / RPS / Agent at source.

- Replication is supported if the source RPS has version 8.0, while the destination RPS still has version 6.5 Update 4.
- Replication is supported if the source RPS is version 6.5 Update 4 and the destination RPS is version 8.0, and patch P00001738 is applied. If the source RPS is version 6.5 Update 1,2,3, upgrade to Update 4 first and then apply patch P00001738.

Backward Compatibility Support for Windows RPS / Agent (or agentless backup proxy)

- Arcserve UDP 8.0 Console supports recovery point servers (RPS) and Agents (or agentless backup proxies) of version 6.5 Update 4 for backup if the plan does not change after an upgrade.
- Arcserve UDP 8.0 Console and RPS support Agents (or agentless backup proxies) of version 6.5 Update 4 for backup if the plan does not change after an upgrade.
- Arcserve UDP 8.0 Console supports Deploying existing plan without any changes when RPS and Agents (or agentless backup proxies) are still running on version 6.5 Update 4. Agent is not automatically upgraded.
- Arcserve UDP 8.0 Console support actions such as Pause or Resume existing plan without any modification when RPS and Agents (or agentless backup proxies) are still running on the version 6.5 Update 4. Agent is not automatically upgraded.
- Arcserve UDP 8.0 Console supports Creating new plan or modifying existing plan to configure legacy features that was supported already in previous release when RPS and Agents (or agentless backup proxies) are still running on version 6.5 Update 4. Agent is not automatically upgraded.
- Arcserve UDP 8.0 Console supports Creating new plan or Modifying existing plan to include 8.0 New features when associated RPS, Proxy or Agent are already upgraded to expected version.

For example:

- Specify Network for backup is supported when Console, RPS, Proxy and Agent are upgraded to 7.0 Update 1, Update 2 or 8.0.

- Arcserve UDP 8.0 Console supports Update Node when RPS and Agents (or agentless backup proxies) are still running on version 6.5 Update 4. Plan will be deployed. Agent is not automatically upgraded.
- Arcserve UDP 8.0 Console supports adding a version 6.5 Update 4 of RPS server. But, the RPS is not automatically upgraded to Arcserve UDP 8.0 when added.
- Arcserve UDP 8.0 Console supports adding a previous version of Agent (or agentless backup proxy) without automatically upgrading the Agent.
- Arcserve UDP 8.0 Console supports adding node to existing plan.
  - Adding clean node to the existing plan is supported when associated RPS upgrades to 8.0. Otherwise, plan is saved but plan deployment fails.
  - Adding previous version node to the existing plan is supported if the version is similar to the running RPS version. Agent is not automatically upgraded.
  - Adding Arcserve UDP 8.0 node to the existing plan is supported when associated RPS upgrades to 8.0.
- Gateway is automatically upgraded to match the version of Arcserve UDP Console.
- Replication backward compatibility policy:
  - Replication is supported when Console is upgraded to Arcserve UDP 8.0, but all associated RPS are still running on 6.5 Update x.
    - Replication is supported from Arcserve UDP 6.5 Update x Source RPS to 6.5 Update 4 destination RPS.
    - Replication is supported from Arcserve UDP 6.5 Update 4 Source RPS to 6.5 Update x destination RPS
  - Replication is supported from Arcserve UDP 6.5 Update x Source RPS to 8.0, destination RPS.
  - Replication is supported from Arcserve UDP 8.0, Source RPS to 6.5 Update 4 destination RPS.
  - Replication is not supported from Arcserve UDP 8.0, Source RPS to 6.5 Update 3, Update 2, Update 1 or final version destination RPS.
- Virtual Standby backward compatibility policy:

- VSB Monitor/Proxy version should be similar to the Agent version if the backup destination is a shared folder.
- VSB Monitor/Proxy version should be similar to the RPS version if the backup destination is RPS.

- Instant Virtual Machine backward compatibility policy:

  Instant VM recovery server version should be similar to the Arcserve UDP console version.

Backward Compatibility Support for Linux Backup Server

For Linux Backup Server 7.x and 6.5.x:

- Arcserve UDP 8.0 Console supports Linux Backup Server of version 6.5.x and 7.x when the recovery point servers (RPS) has Arcserve UDP 8.0. All jobs run smoothly, except for the new feature or enhancement in 8.0.
- Arcserve UDP 8.0 Console supports recovery point servers (RPS) and Linux Backup server of version 6.5.x and 7.x. All jobs run smoothly, except for the new feature or enhancement in 8.0.
- After upgrading the Linux Backup server to Arcserve UDP version 8.0, all new features for 8.0 are supported.

Upgrade Sequence for UDP Console, RPS, and Agent

Based on the Backward Compatibility Support Policy, planning your upgrade in the following sequence helps a smooth upgrade:

1. Upgrade Arcserve UDP Console.
2. Upgrade Arcserve UDP RPS (DR site).
3. Upgrade Arcserve UDP RPS (Data Center).
4. Upgrade Arcserve UDP Agentless Proxy and some or all Agents in Data Center.
5. Upgrade Arcserve UDP RPS (Remote site).
6. Upgrade Arcserve UDP Agentless Proxy and some or all Agents at the remote site.

   Note: Repeat Step 5 and 6 for each remote location.

   Notes:

   - First, upgrade MSP on Console or RPS if the MSP replication task is configured.
   - Gateway is automatically upgraded to match version of Console.
   - According to the replication backward support policy, always upgrade the target RPS before the source RPS.

How to Upgrade to Arcserve UDP 8.0 from a Previous Release

Upgrade an installation means to reinstall features or components to a higher release. The upgrade process lets you retain most of your current settings and migrate the information stored in the previous Arcserve UDP database to the new Arcserve UDP database.

Upgrade Consideration

- Verify if hardware requirement of Arcserve UDP Version 8.0 is met. For more information, see the *System Information* in *Release Notes v8.0*.
- Arcserve UDP version 8.0 supported Platform, Hypervisor, OS, or Application Version is required. For detailed information, see the Compatibility Matrix.
- Before upgrading from 6.5.3 or previous supported versions, avoid pausing an existing plan. If the plan is paused before an upgrade, you cannot resume the plan until all related RPS are upgraded.
- Purchase the product key for Arcserve UDP version 8.0 and keep it ready.
- Remove the previous version Arcserve Exchange Granular Restore (AEGR) standalone utility for Arcserve UDP version 5.0. If detected, the installation wizard prompts to remove.
- Ideally, the older plans should work properly at every upgrade step described below.

Follow these steps to upgrade:

1. For MSP, upgrade the previous Arcserve UDP Console to Arcserve UDP 8.0.

   This step is required only if **Task: Replicate from remotely-managed RPS** is configured in the previous release.

   Notes:

   - If the node has only the Console installed, upgrade the Console. Stopping any running jobs is not required to upgrade the Console.
   - If the node has Console and RPS installed, plan to upgrade the Console when no jobs are running before the upgrade. If necessary, stop the data store of RPS to cancel the running jobs before upgrading the Console.
   - For more information on how to upgrade the Console, see Install Arcserve Unified Data Protection Using the Setup Wizard or Install Arcserve UDP Using the Unified Installer.

2. For MSP, upgrade the previous Arcserve UDP RPS to Arcserve UDP version 8.0.

This step is required only if **Task: Replicate from remotely-managed RPS** is configured in the previous release.

Notes:

- Plan to upgrade the RPS when no jobs are running before the upgrade. If necessary, stop the data store of RPS to cancel the running jobs before upgrading the RPS.
- For more information about how to upgrade the RPS, see Install/Upgrade Recovery Point Server.

3. For Customer, upgrade the previous Arcserve UDP Console to Arcserve UDP version 8.0.

   Notes:

   - If the node has only the Console installed, upgrade the Console. Stopping any running jobs is not required to upgrade the Console.
   - If the node has Console and RPS installed, and jobs are running on the RPS, stop the data store of the RPS to cancel first the running jobs. Then, upgrade the Console.
   - For more information about how to upgrade the Console, see Install Arcserve Unified Data Protection Using the Setup Wizard or Install Arcserve UDP Using the Unified Installer.

4. For Customer, upgrade the previous Arcserve UDP Replication target RPS to Arcserve UDP 8.0.

   Notes:

   - Arcserve UDP 8.0 supports replication from Arcserve UDP Version 6.5.x source RPS to Arcserve UDP 8.0 target RPS. When the plan includes multiple replication tasks, first upgrade the last replication target RPS.
   - Replication from Arcserve UDP 8.0 source RPS to a lower version of the target RPS is supported only if the target RPS is of Arcserve UDP 6.5 Update 4 version.
   - For more information about how to upgrade the RPS, see Install/Upgrade Recovery Point Server.

5. For Customer, upgrade the previous Replication Source RPS to Arcserve UDP 8.0.

   Note: For more information about how to upgrade the RPS, see Install/Upgrade Recovery Point Server.

6. For Customer, upgrade the previous Proxy and Agents to Arcserve UDP 8.0.

   Notes:

- For more information about how to upgrade the Windows agentless proxy and agent, see Deploy Agent to Nodes.
- For more information about how to upgrade the Linux Backup Server, see How to Upgrade Arcserve UDP Agent (Linux).

7. For Customer, upgrade the previous VSB Monitor or Instant VM Recovery Server to Arcserve UDP 8.0.

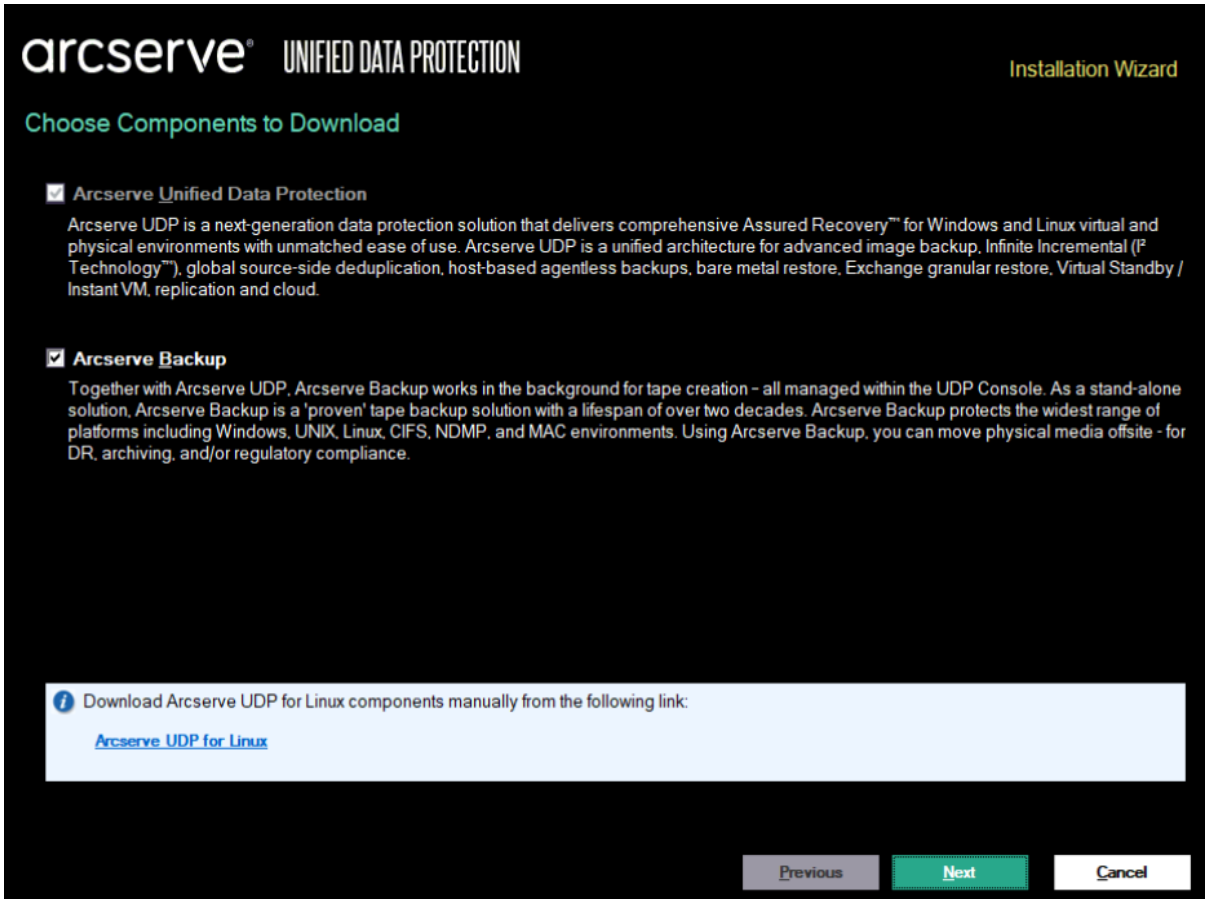
   Note: For more information about how to upgrade the VSB Monitor, see Install Arcserve Unified Data Protection Using the Setup Wizard or Install Arcserve UDP Using the Unified Installer.

How to Upgrade to Arcserve UDP 8.0 Using the Single Installer

You can upgrade to Arcserve UDP version 8.0 using the single installer. Review the supported versions before you upgrade.

Follow these steps:

1. Download the ASDownloader.exe from the Arcserve download link.
2. Run the ASDownloader.exe, accept the License Agreement, and click **Next**.
3. Click **Next** from the **Get Started** screen.
4. From **Choose Components to Download** screen, select the components that you want to download.



Considerations for the component selection:

- To upgrade the previous version of Arcserve UDP, select the **Arcserve Unified Data Protection** option.

- To upgrade the previous version of Arcserve Backup, select the **Arcserve Backup** option.

Note: Some components are automatically selected as per dependency when you select the **Arcserve Backup** option. If you do not want to install the components on the same machine, see How to Upgrade to Arcserve UDP 8.0 from a Previous Release and use the Setup wizard to upgrade only the existing components. From here, you can also download the latest Linux setup for Arcserve UDP.

5. After the components are downloaded, select the installation method to upgrade the selected components.

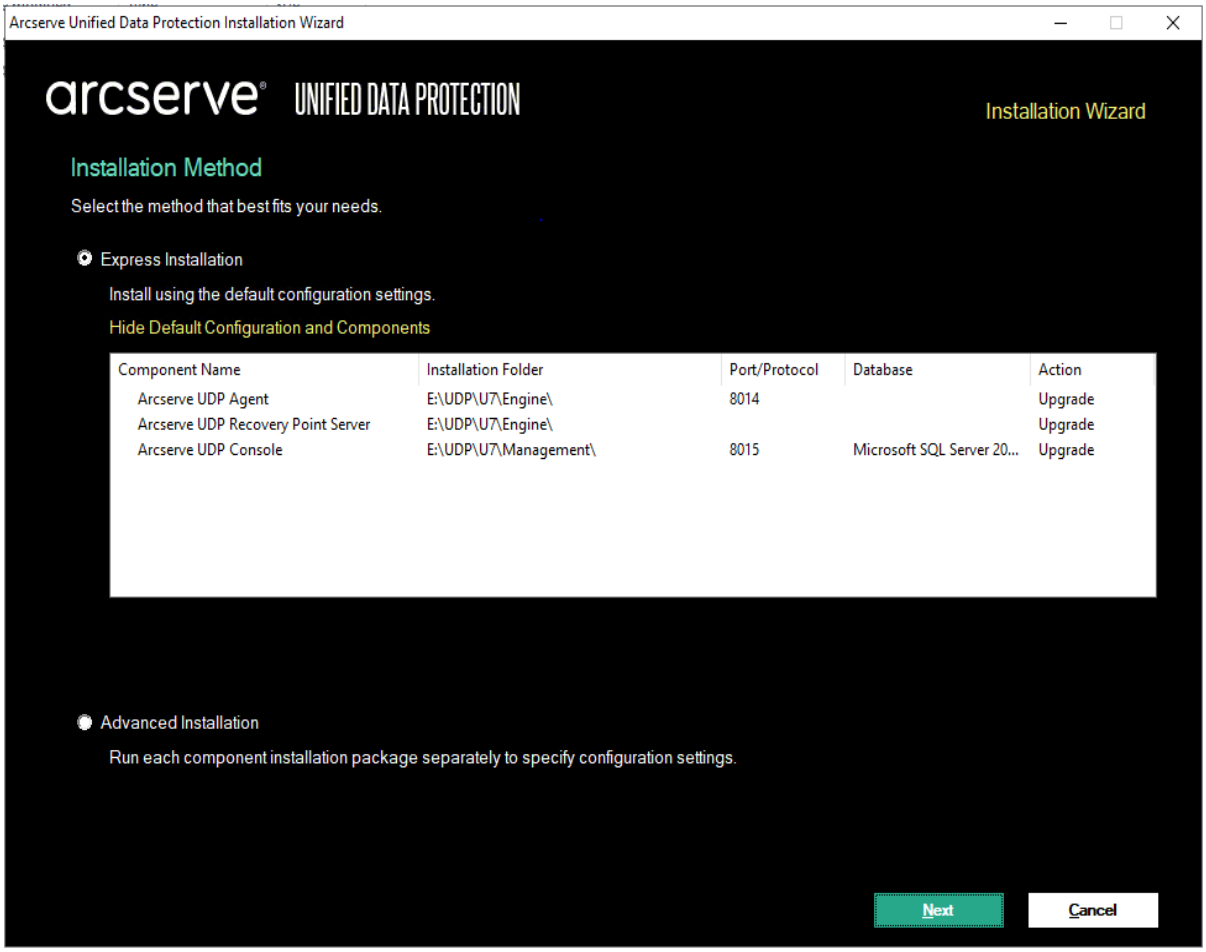
Express Installation

We recommend Express Installation when you plan to install all the components with their default configuration on the same machine.

Notes:

- Verify all the prerequisites for the upgrade. For example, .Net 4.5.1 is required for Arcserve Backup. Follow the instructions displayed on the Installer until all the prerequisites are met.
- Verify the Compatibility Matrix requirement.
- The Express Installation method installs the components as per the default configurations, such as path, protocol, port, and Database. When you upgrade using the Express Installation method, the configuration of the pre-

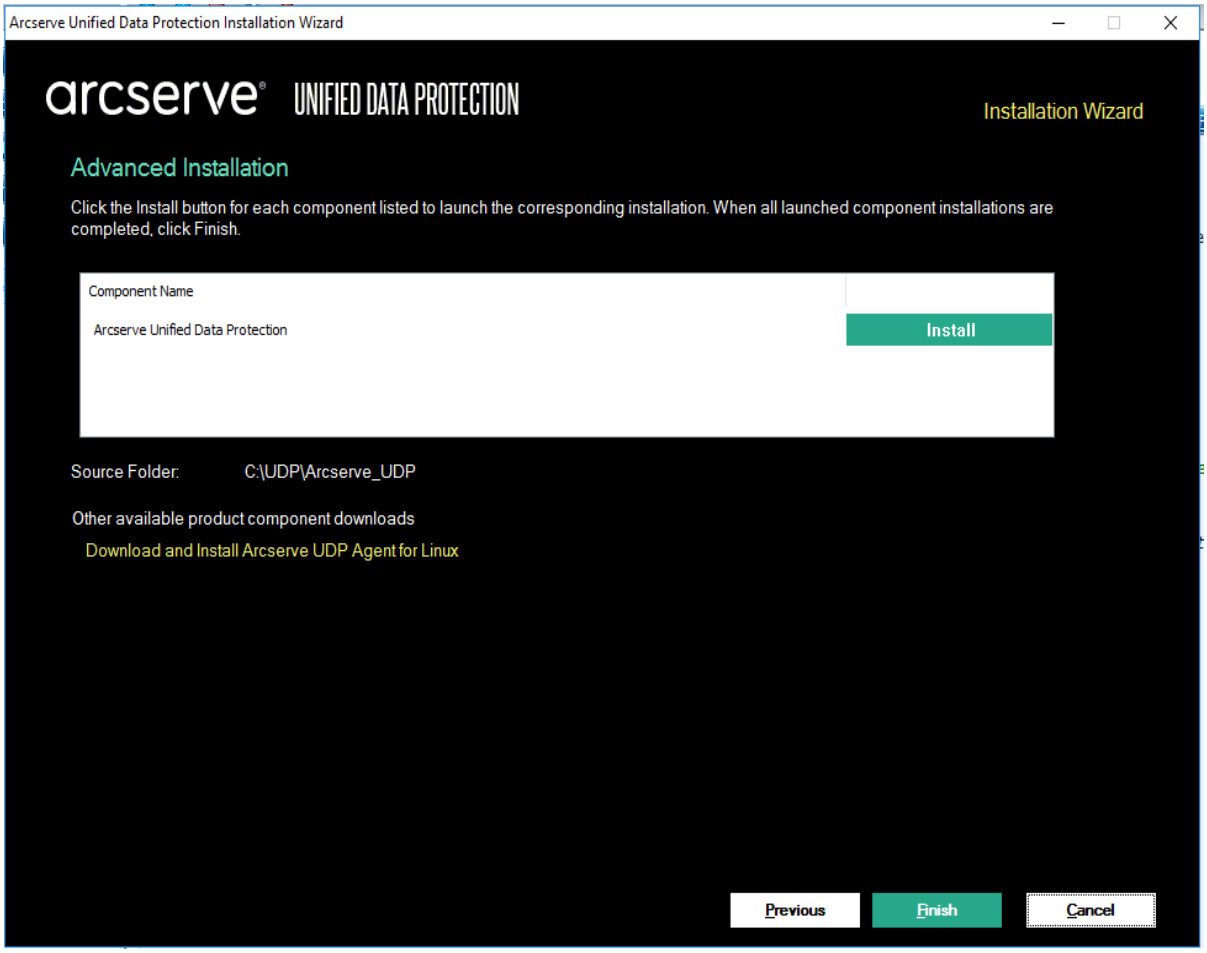
vious release is used.



Advance Installation

We recommend Advance Installation for installation or upgrade of each component separately on the same machine. Click **Install** to upgrade or install each component separately.

Note: To upgrade Arcserve Unified Data Protection Agent for Linux using the Unified Installer, download the images by referring the links shared at the end of the Installer. After copying the images to the Linux Backup Server, follow the steps mentioned in *How to Upgrade Arcserve UDP Agent (Linux)* in Agent for Linux Online Help.



Arcserve UDP is successfully upgraded using the Unified Installer.

Note: For more information, see How to Upgrade to Arcserve UDP 8.0 from a Previous Release.

How to Upgrade to Arcserve UDP 8.0 or Arcserve Backup 18.0 from Previous Version of Arcserve UDP or Arcserve Backup to Enable Copy to Tape

Arcserve UDP version 8.0 integrates with Arcserve Backup 18.0 and provides copy backup data to a tape media destination feature. You can modify an existing plan or create a plan with the **Copy to Tape** task in Arcserve UDP Console.

Upgrade Consideration

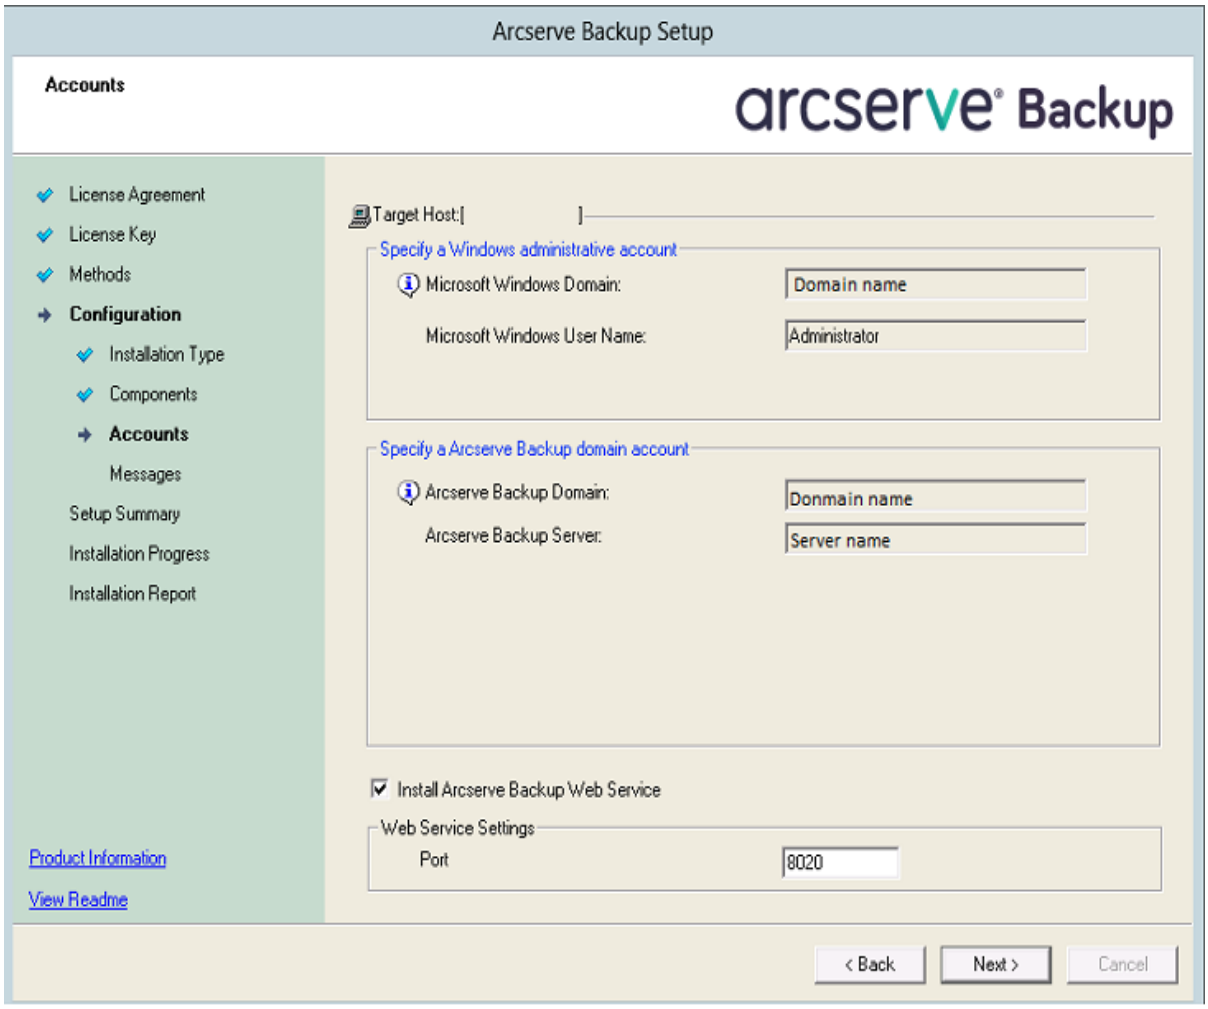
- To make sure that existing plan with the Copy to Tape task runs smoothly, upgrade Arcserve Backup to 18.0 before upgrading Arcserve UDP to 8.0.
- Verify Arcserve UDP version 8.0 hardware requirement. For more information about the detailed hardware requirement, see the *System Information* in *Release Notes 8.0*.
- Arcserve UDP version 8.0 supported Platform, Hypervisor, OS, or Application Version is required. For detailed information, see the Compatibility Matrix.
- Before upgrading UDP 6.5.3 or previous supported versions, avoid pausing an existing plan. If the plan is paused before an upgrade, you cannot resume the plan until the related RPS is upgraded.
- Purchase the product key for Arcserve UDP Version 8.0 and keep it ready.
- Remove the previous version of Arcserve Exchange Granular Restore (AEGR) standalone utility for Arcserve UDP version 5.0. If detected, the installation wizard prompts to remove.
- Ideally, the older plans should work properly at every upgrade step described below.

Copy to Tape Consideration

- Prepare the Arcserve Backup 18.0 server by installing or upgrading from any previously supported upgrade path. Arcserve Backup is required before configuring the **Copy to Tape** task in an Arcserve UDP plan.
- Verify if you meet the Arcserve Backup 18.0 requirement. For more information, see the Compatibility Matrix.
- Purchase the product key for Arcserve Backup 18.0 and keep it ready.

Follow these steps to upgrade and enable the Copy to Tape feature:

1. Upgrade to Arcserve Backup 18.0 from a previous release or install Arcserve Backup 18.0 with the Install **Arcserve Backup Web Service** option enabled.



Note: For more information about installing and upgrading Arcserve Backup, see the *Arcserve Backup 18.0 Implementation Guide*.

After completion of upgrade, perform the following tasks:

- Verify if the job engine has started in the Arcserve Backup manager.
- In Windows Services Console, verify that **Arcserve Backup Web Service** is in the running status.

2. Upgrade to Arcserve UDP version 8.0 from a previous release.

   Note: For more information on how to upgrade from a previous release, see How to Upgrade to Arcserve UDP 8.0 from a Previous Release.

3. Configure the **Copy to Tape** task in the Arcserve UDP Console for an existing or a new plan.

   Note: For more information about configuring the Copy to Tape task, see How to Create a Copy to Tape Plan.

   The plan should work for each step.

How to Migrate Arcserve UDP Console from One Server to Another

For more information, click link.

How to Upgrade Gateway to the Same Version as its Registered Console

Arcserve UDP Gateway must match the version of its registered Console. After Arcserve UDP Console is upgraded, the related gateway version is verified when the service starts. If the Gateway version does not match the Console version, Arcserve UDP automatically triggers an auto update for the gateway servers.

If the gateway is unavailable for auto update, you can manually upgrade the gateway later.

Follow these steps to manually upgrade the gateway:

1. Log into the Arcserve UDP Console.
2. Click the **resources** tab, navigate to **Infrastructure**, and click **Sites**.
3. Select the site that you want to upgrade.
4. From the **Actions** drop-down list, click **Upgrade Gateway**.

   A confirmation dialog opens.
5. Click **Yes**.

   The gateway is upgraded with the latest version of Arcserve UDP Gateway.

How to Perform Arcserve UDP Console Migration

You can migrate one Arcserve UDP Console to another Arcserve UDP Console using *ConsoleMigration.exe*. From Arcserve UDP v6.5 Update 2 onwards, you can migrate the Arcserve UDP Console between any two Arcserve UDP consoles.

Use *ConsoleMigration.exe* for BackupDB and RecoverDB. The following screen shot displays the usage of *ConsoleMigration.exe*:



To complete the migration process, follow these steps:

1. On old Arcserve UDP Console, perform backup for the Arcserve UDP database.

```
C:\Program Files\Arcserve\Unified Data Protection\Management\BIN\Appliance>ConsoleMigration.exe
-BackupDB

Backed up DB and version files completed.

DB and version files were created at "C:\Program Files\Arcserve\Unified Data
Protection\Management\BIN\Appliance\DB_Migration".
```

The *DB_Migration* folder is created successfully.

2. On the new Arcserve UDP Console, copy the *DB_Migration* folder to the following path:

   <UDP_Home> *Management**BIN**Appliance*\

3. On the new Arcserve UDP Console, perform the steps mentioned in the screen below to recover the Arcserve UDP Console database.

```
C:\Program Files\Arcserve\Unified Data Protection\Management\BIN\Appliance>ConsoleMigration.exe
-RecoverDB

Are you sure you want to recover the backup DB file? (y/n): y

Stopping Arcserve UDP management service, please wait...

Recovering backup DB file...

Updating nodes, please wait...

Please update nodes manually from UDP console, if you still encounter disconnected nodes.

The disconnected nodes(if exist) will be saved at "C:\Program Files\Arcserve\Unified Data Protec
tion\Management\BIN\Appliance\DB_Migration\logs".

Console migration completed. Console use DB "localhost\ARCSERVE_APP".
```

Note: In Arcserve UDP Console, if any site other than Local Site exists, follow the steps mentioned in *NewRegistrationText.txt* file to register the site again.

You have completed migration of Arcserve UDP Console to the new Arcserve UDP Console successfully.

You can use this tool to perform console migration for Arcserve UDP Console connected with remote SQL database. After the migration is complete, the migrated Arcserve UDP Console is configured to connect with the same remote SQL database.

Note: In UDP v6.5 Update 2 and Update 3, when you use Console Migration Tool to perform console migration between two Arcserve UDP Consoles then both the Arcserve UDP Console versions and SQL database versions should be consistent on the two Arcserve UDP Console systems. Otherwise, the console migration fails and an

error message is displayed in the command line and log files under the following path:

<UDP_Home> \Management\BIN\Appliance\logs

In Arcserve UDP v6.5 Update 4, ***-force*** option is introduced in ***ConsoleMigration.exe*** command to force the recovery backup database file migration to the target console under the following conditions:

1. When you want to perform console migration between two consoles where the source console uses SQL Server Enterprise edition and the target console uses SQL Server Express edition. In this case, the minimum required Database size of the source UDP console is 4000 MB.
2. When you want to perform console migration from a console that uses an advanced version of SQL Server database to a console that uses an older version of SQL Server database. For example, migrating from a console using SQL Server 2016 to a console using SQL Server 2014.

Best Practices to protect Arcserve UDP server

You can migrate the Console from the old server to a new server without losing any data.

- Do not have backup solution in the same Active Directory domain.
- Use backup accounts.
- Do not use Domain administrator accounts to access user related tasks. For example, administrator that uses administrator privileges for reading emails, browsing online, and so on.
- Separate the UDP Console from RPS and create VSB for the UDP Console.
- If possible, use dedicated backup network / vlan.
- If possible, disable SMBv1.

  Note: Cannot disable when using Red Hat / Centos v6.6 and older as these use SMB v1.

Chapter 5: Exploring and Configuring Arcserve UDP

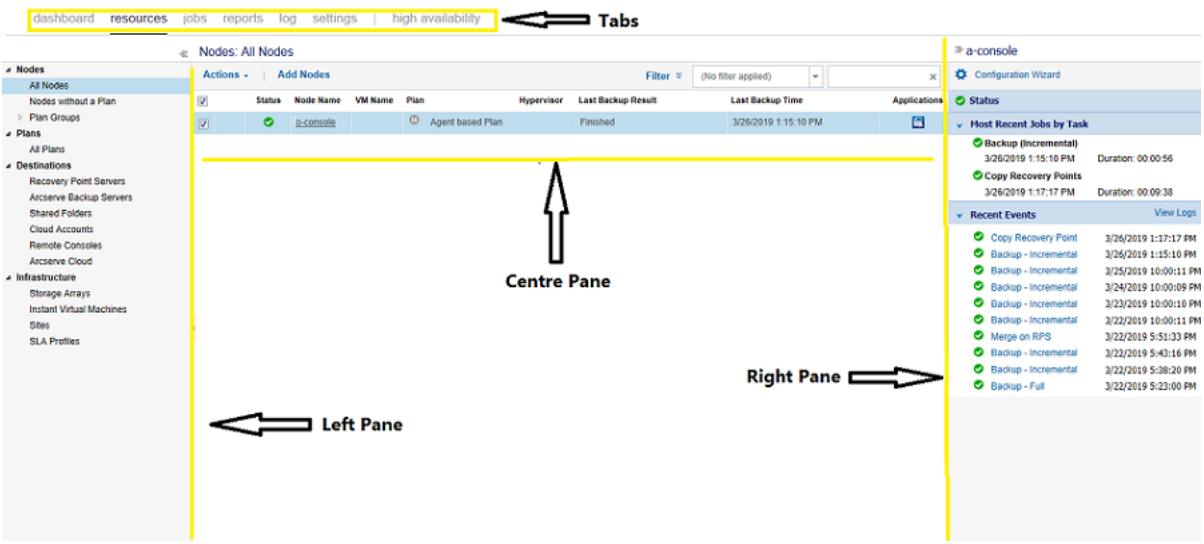
This section contains the following topics:

Arcserve Unified Data Protection User Interface

Before you use Arcserve Unified Data Protection, become familiar with the user interface. The Arcserve Unified Data Protection interface lets you perform the following tasks:

- Manage and monitor jobs
- Add and manage source nodes
- Add and manage destination recovery point servers
- Manage plans to create backup schedules
- Obtain data protection statistics
- View error and warning logs
- Manage and monitor Arcserve High Availability.
- Configure data protection settings
- Restore Backup Data

Navigating Arcserve UDP



Tabs

Lets you navigate to the various functions of Arcserve Unified Data Protection.

Panes

When you navigate to each tab, the displayed screen is divided into the following panes. Each pane is used to perform related actions.

Left Pane

Lets you navigate to various functions and operations. The result of each click is displayed in the center pane.

Center Pane

Lets you perform most of the actions in this pane such as adding, deleting, and modifying. This pane also displays the result and status of each activity such as jobs, plans, and reports. Most of your actions are performed on this pane. The information displayed on this page is mostly the result of the options that you selected in the left pane.

Right Pane

Displays a summary of the items you selected on the center pane. For example, on the Jobs tab, if you selected a job from the center pane, then a brief summary of the job such as job monitor (if there is a running job) and job details like source node name, task, destination Recovery Point Server, and destination data store is displayed in the right pane.

Tabs

The Arcserve Unified Data Protection solution provides the following tabs to perform data protection functions:

- dashboard
- resources
- jobs
- reports
- log
- settings
- high availability

dashboard

The **dashboard** tab lets you view graphical representation of the latest task status and data storage of the last seven days. Using the **dashboard** tab, you can perform the following action:

- View the time of the last update. Clicking the *Refresh* icon displays the latest data on the dashboard.
- Clicking the RTO bar chart opens directly the RTO Reports Page.
- View the last task status of nodes or plans according to the filters you select in the **Last Task Status** chart.



You can view the graphs for the following options on the dashboard tab:

Last Task Status

Last Task Status refers to the latest task status and provides you multiple filters to view the status. Based on your selection from the filter option, you can view the last task status. For example, select **All Nodes** to see the last task status of all nodes or select any plan to see the last task status of the nodes protected by this plan. Then, filter further with a specific task type. When you select **All Nodes**, you can see the following status:

- **Successful** indicates that the nodes are successfully backed up.
- **Failed** indicates that the last backup is not successful.
- **Canceled** indicates that the last backup was stopped.
- **Missed** indicates that the last backup was not performed as scheduled.

- **Not Connected** indicates that the Arcserve UDP Console failed to connect to the node.
- **Incomplete** indicates that the restore job was not complete.

When you click each slice (the status) from the pie chart, the **resources** page opens and the associated nodes are displayed. For example, if you click **Successful** from the pie chart, the **resources** page opens. The **resources** page displays the nodes that do not have any plan. Also, the **Successful** filter is preselected on the **resources** page.

Actual, Restorable and Raw Data Size: 7 days

The graph refers to the Raw Data vs Restorable Data vs Actual Data Storage of the last seven days.

Restorable Data Size

Refers to the size of data that you can restore from the data stored on all target storage that day. Regardless of the backup type, full backup or incremental backup, every recovery point is used to restore a full source. This size is the total size of all available size of restored source.

For example: If Raw data size is 30 GB and after compression and Deduplication, the size is 22 GB and Restorable data is 30 GB for the first full backup. For a change of 4 GB incremental change, the Restorable data size is 30+34 GB= 64 GB.

Actual Data Storage Size

Refers to the size of the final data. After compressing the raw data and removing duplicated data blocks from the raw data and adding some metadata, the raw data becomes the final data and is saved to the target storage.

Raw data Size

Refers to the data read from the source and transferred to the target.

SLA Report for RTO and RPO

Arcserve UDP introduces Service Level Agreement (SLA) report to help organizations generate compliance reports related to Recovery Point Objective (RPO) and Recovery Time Objective (RTO).

RTO Report

Arcserve UDP RTO report is a compliance report that displays the comparison of Recovery Time Actual and Recovery Time Objective values for all the executed recovery type of jobs such as File system restore, VM recovery, BMR, Instant VM, and Assured Recovery. The Bar chart uses

multiple colors to display status of jobs. You can further drill-down to see node level status filtered by RTO met, not met, not tested, and not defined status.

RPO Report

Arcserve UDP RPO report displays the total number of nodes with available recovery points during the specified time period in the bar view categorized by Age of latest recovery points (15 minutes, last hour, 12 hours, last day, and so on), Age of oldest recovery points (30 days and older), and monthly distribution (Jan - Dec). You can further drill-down to see node level status for the selected category.

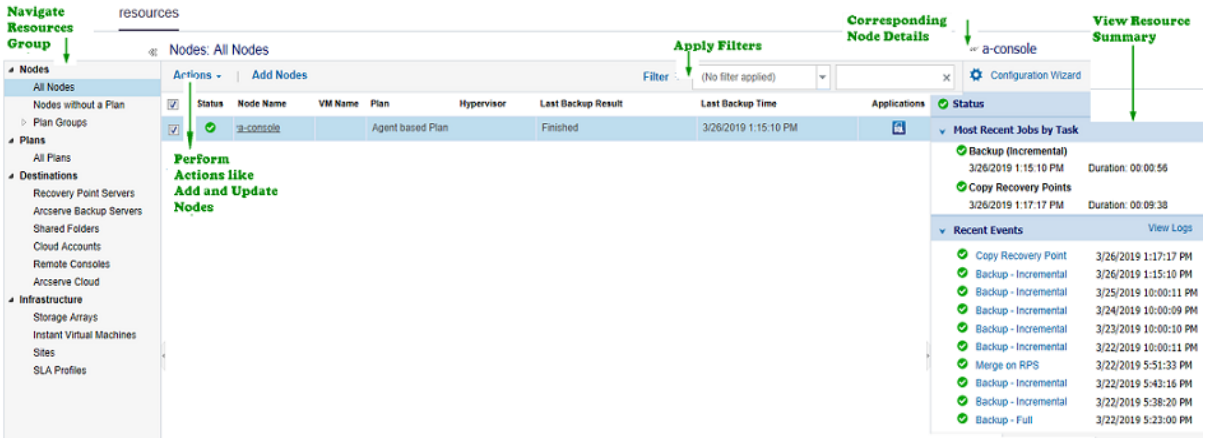
Note: From dashboard, you can directly view the RPO report of a specific month. The report gets populated directly from Arcserve UDP Dashboard. Clicking inside the RPO bar graph on a month at Dashboard displays RPO report screen for that specific month on the RPO screen.

resources

The **resources** tab lets you manage the Arcserve Unified Data Protection resources: Nodes, Destinations, Virtual Standby, and Plans. Use this tab to add resources to Arcserve Unified Data Protection such as nodes that you want to protect or recovery point servers for backup. You also use this tab to create plans and tasks for backup, virtual standby, and replication. Using the resources tab, you can perform the following options:

- Node Management
- Destination Management
- Plan Management
- Infrastructure Management

Note: Only the **resources** tab on the Console has site-awareness. The other tabs on the Console display consolidated data for all the sites. For more information on Sites, see How to Add and Manage a Site.



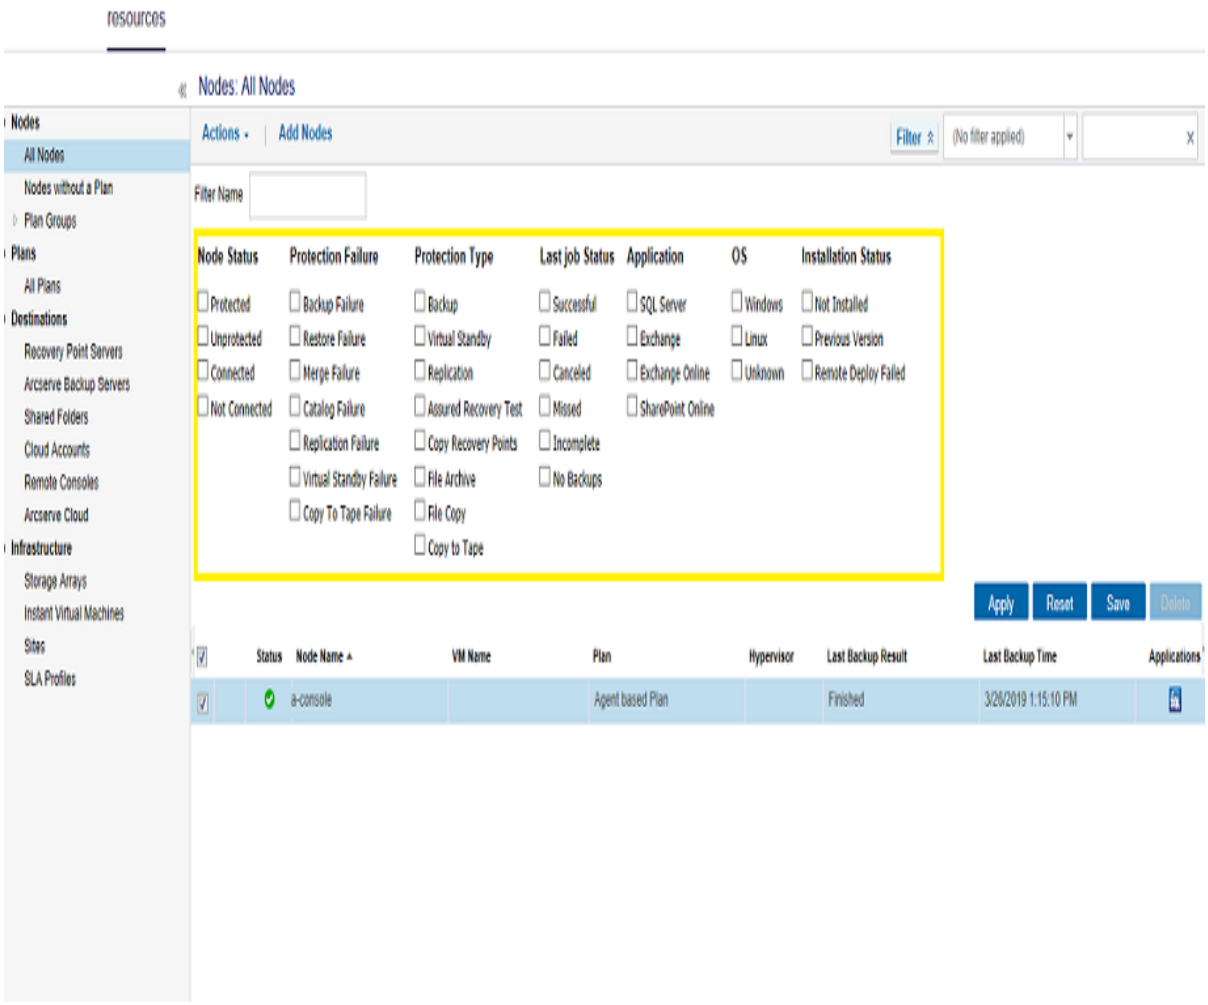
Node Management

The node management view lets you manage all the nodes and apply filters to refine the node search. When you select specific node in center pane, you can see the status and recent events about the node in right pane. You can apply various filters from the center pane. You can create node groups on the left pane to group specific nodes.

When you select a node from the center pane, the node status and recent events are displayed in the right pane.

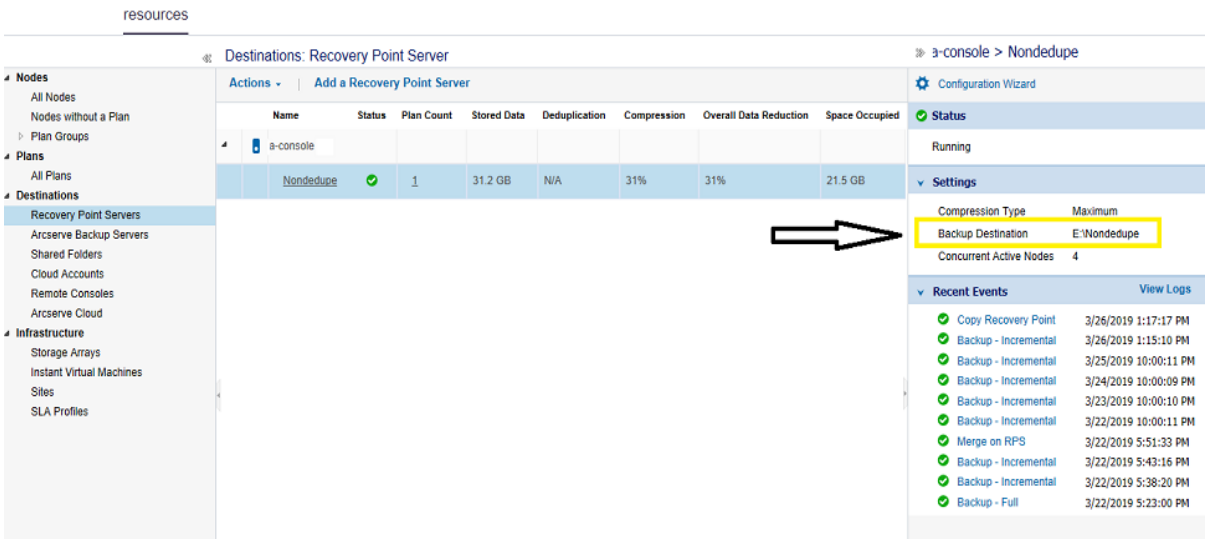
You can perform operations on nodes by clicking the Actions drop-down menu from the center pane. Such operations that you can perform through Actions in center pane is applied to all source nodes. Such operations that you can perform through Actions in the right pane is only applied to the node that you select in the center pane.

From Resources Tab, click **All Nodes** under Node in the Left Pane. Then, from centre pane, select a node and click on **Filter** to view Node Management options. You can select check box of desired options and click **Apply**.


resources
Nodes: All Nodes
Nodes
All Nodes
Nodes without a Plan
Plan Groups
Plans
All Plans
Destinations
Recovery Point Servers
Arcserve Backup Servers
Shared Folders
Cloud Accounts
Remote Consoles
Arcserve Cloud
Infrastructure
Storage Arrays
Instant Virtual Machines
Sites
SLA Profiles
Actions
Add Nodes
Filter
(No filter applied)
Filter Name
Node Status
Protected
Unprotected
Connected
Not Connected
Protection Failure
Backup Failure
Restore Failure
Merge Failure
Catalog Failure
Replication Failure
Virtual Standby Failure
Copy To Tape Failure
Protection Type
Backup
Virtual Standby
Replication
Assured Recovery Test
Copy Recovery Points
File Archive
File Copy
Copy to Tape
Last job Status
Successful
Failed
Canceled
Missed
Incomplete
No Backups
Application
SQL Server
Exchange
Exchange Online
SharePoint Online
OS
Windows
Linux
Unknown
Installation Status
Not Installed
Previous Version
Remote Deploy Failed
Apply
Reset
Save
Delete
Status
Node Name
VM Name
Plan
Hypervisor
Last Backup Result
Last Backup Time
Applications
a-console
Agent based Plan
Finished
3/26/2019 1:15:10 PM

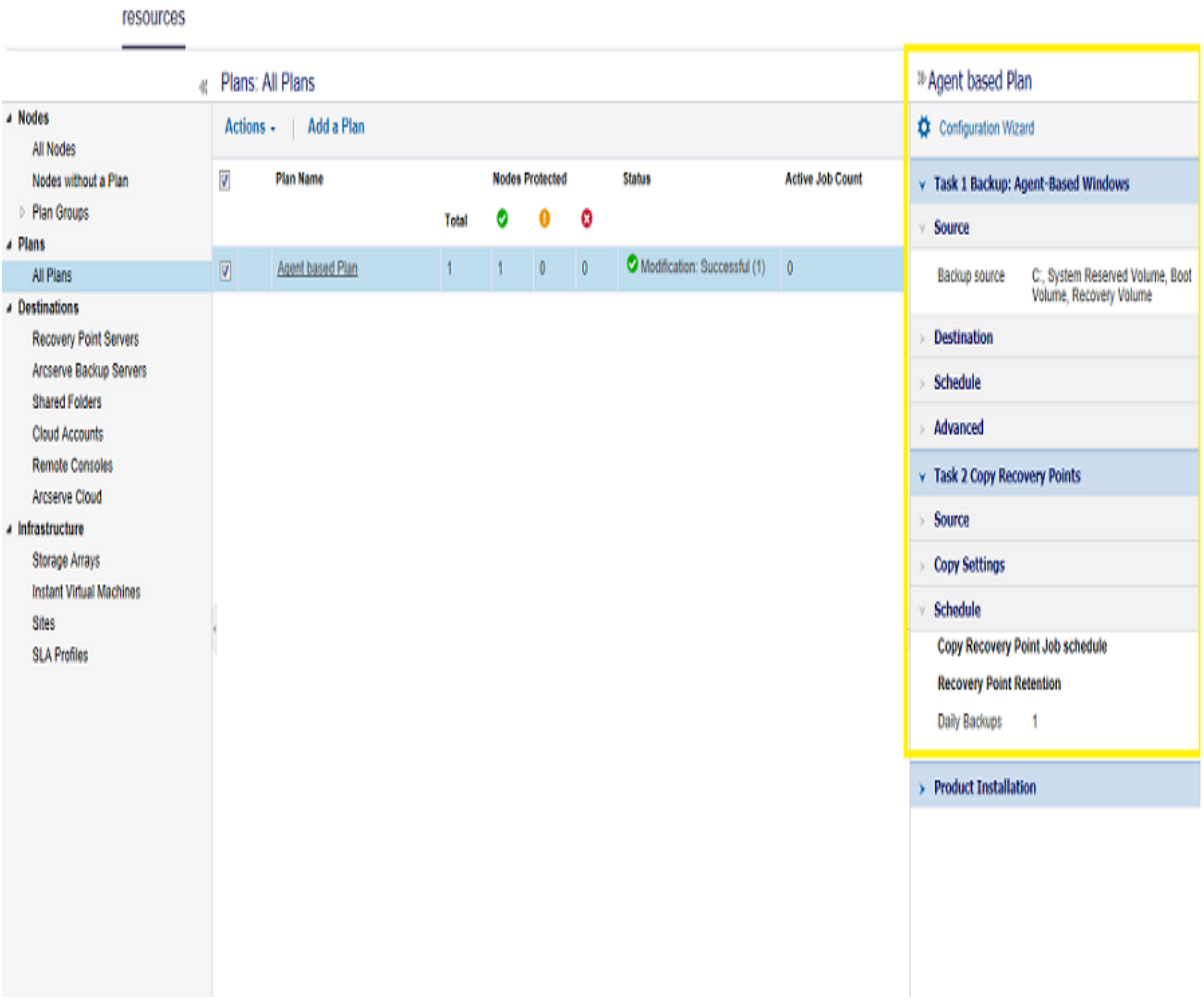
Destination Management

The destination management view lets you manage the destination recovery point servers. When you select a server from the center pane, its recent events are displayed in the right pane. When you select a data store, its status and settings are displayed in the right pane.



Plan Management

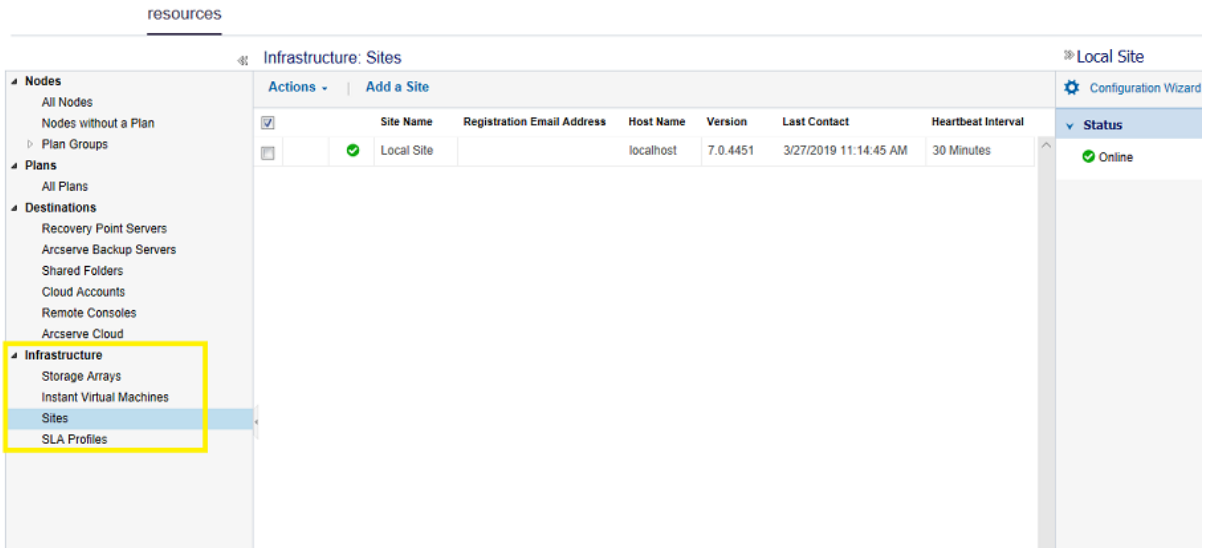
The plan management view lets you manage all your plans. You can create, modify, delete, deploy, pause, and resume plans from this view. The plan details appear in the right pane. You can click on required fields to view details and manage.



Infrastructure Management

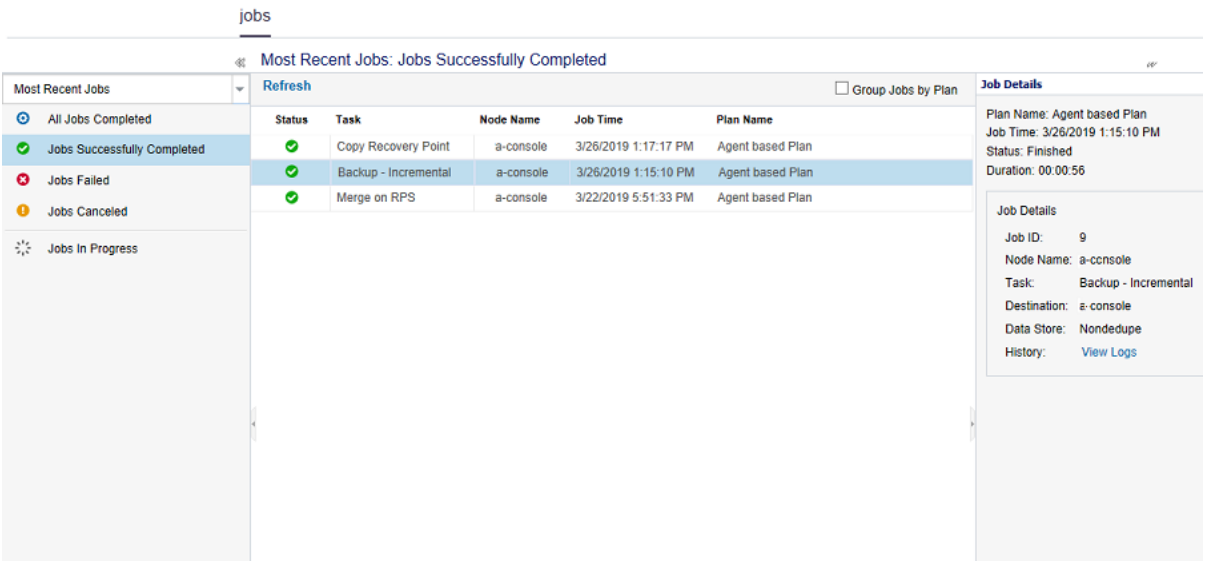
The infrastructure management view lets you manage storage arrays, instant virtual machines, and remote sites.

Following screenshot highlights Infrastructure management pane:



jobs

The **jobs** tab displays the status of the jobs for a specific period. Apply filters to categorize the results or group the jobs by plan.

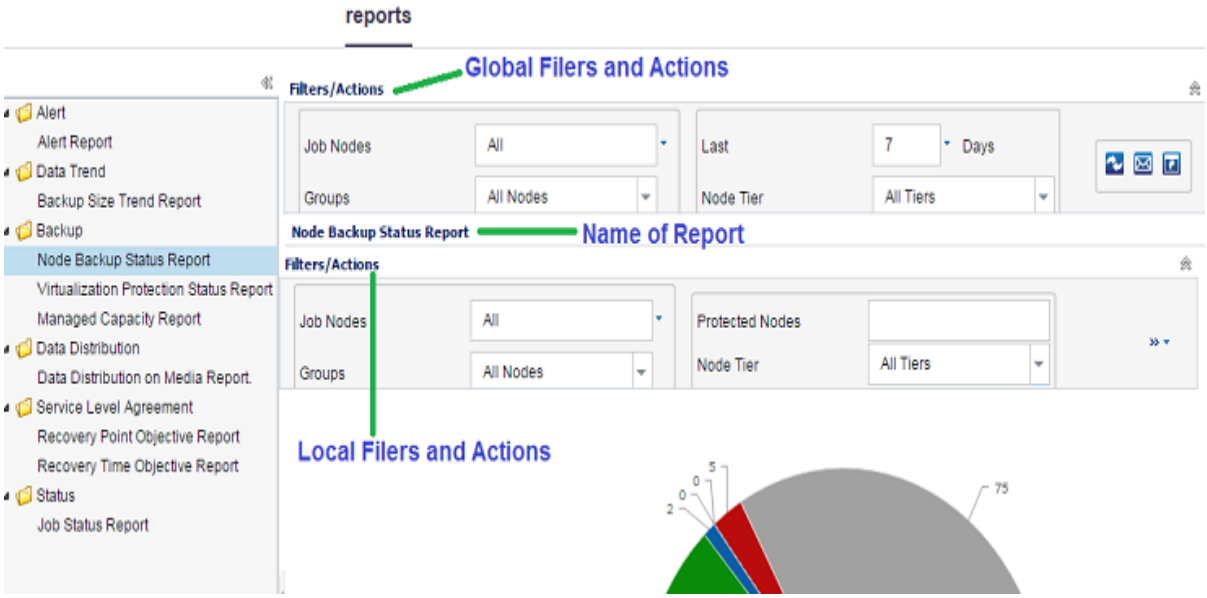


When a job is in progress, the right pane displays the job monitor that contains the progress of the job. Click **Job Details** on the right pane to open the job monitor. You can see the job monitors only if the job is in progress.

To cancel a job, open the job monitor and click **Cancel**.

reports

The **report** tab displays a list of reports that you can generate. You can apply filters to your reports to get specific reports. The reports are generated in CSV, PDF, or HTML formats. For more information about these reports, see How to Generate Arcserve UDP Reports.



log

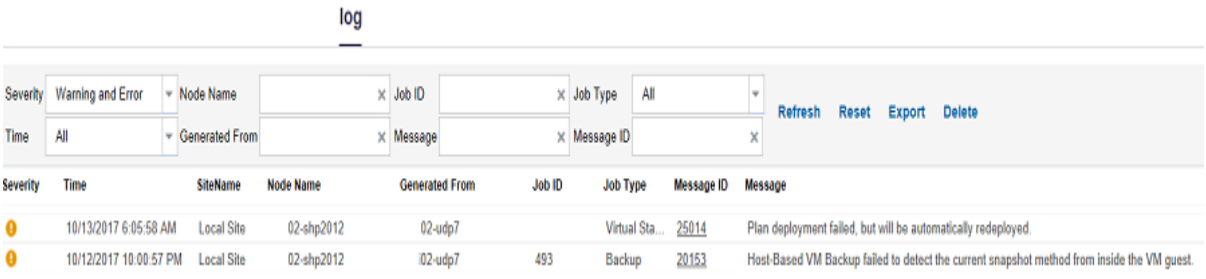
The **log** tab displays all activity logs for the protected nodes, destination servers, data stores, and plans. You can view logs and apply various filters such as severity, specific node, logs generated from the machine, job IDs, and log content.

Message ID provides a hyperlink to access detailed documentation. Click the hyper-link in the MessageID column to view the description and solution for that mes-sage.

Note: The Activity Logs generated by Console and Linux Backup Server / Agent and the Copy to Tape job do not have Message ID.

You can search the activity logs using a combination of the available filters or one of the following options:

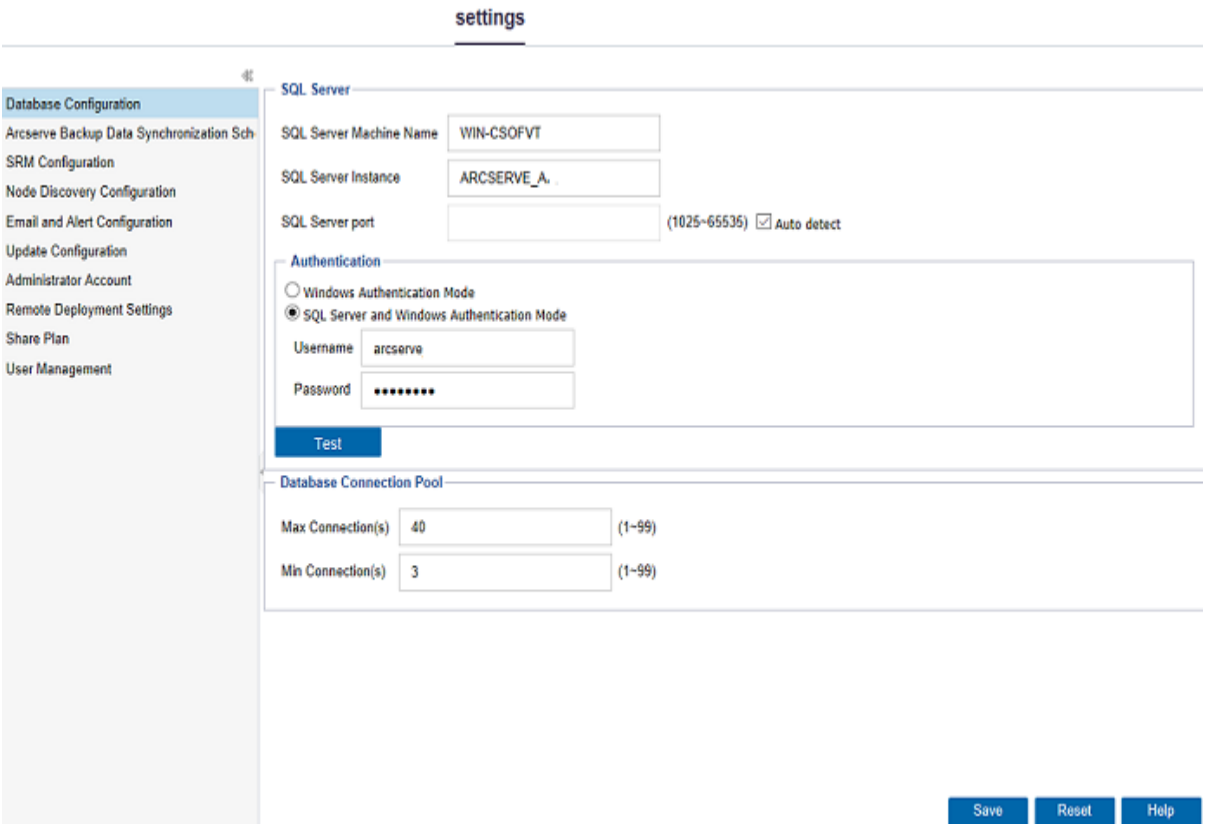
- Select Severity types to view all the logs related to the selected type.
- Enter other details (such as Node Name, Job ID), and click **Search**.
- **Refresh:** Display the latest logs available according to defined filters.
- **Reset:** Uses only the default filters and only display Warning and Error for all types of job.
- **Export:** Downloads current job log as *.zip (that is activitylog_export_2017_10_12_15_02_27_586.zip) to Windows system Downloads directory.
- **Delete:** Removes all log records or all log records older than the specific date.



settings

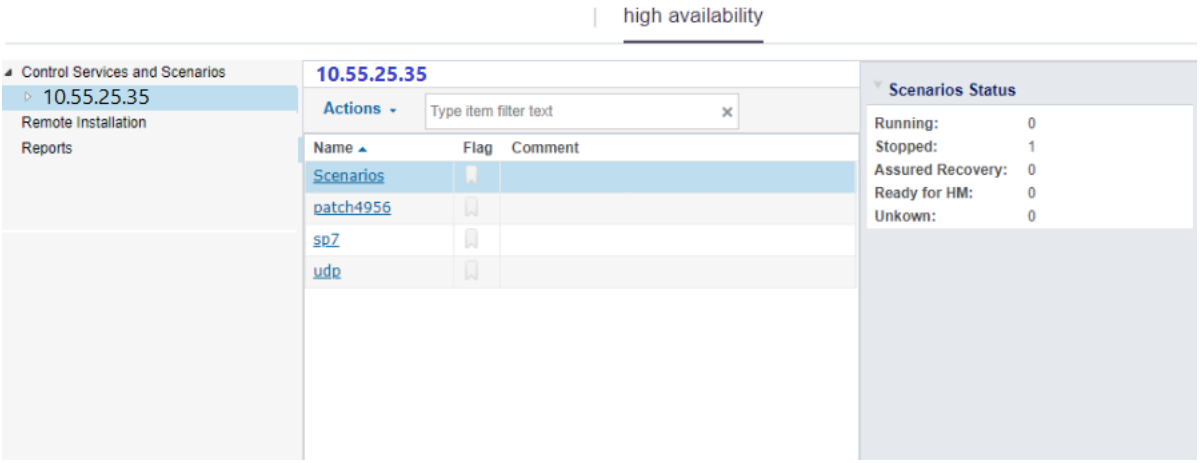
The **settings** tab lets you configure certain preferences such as which email server to use, set up administrator user ID and password, and define the default node deployment path.

For more information about the **settings** tab, see How to Configure Arcserve UDP.



high availability

The **high availability** tab lets you manage and control Arcserve High Availability functions. If required enter credentials in the **Add Control Service** dialog and view details.

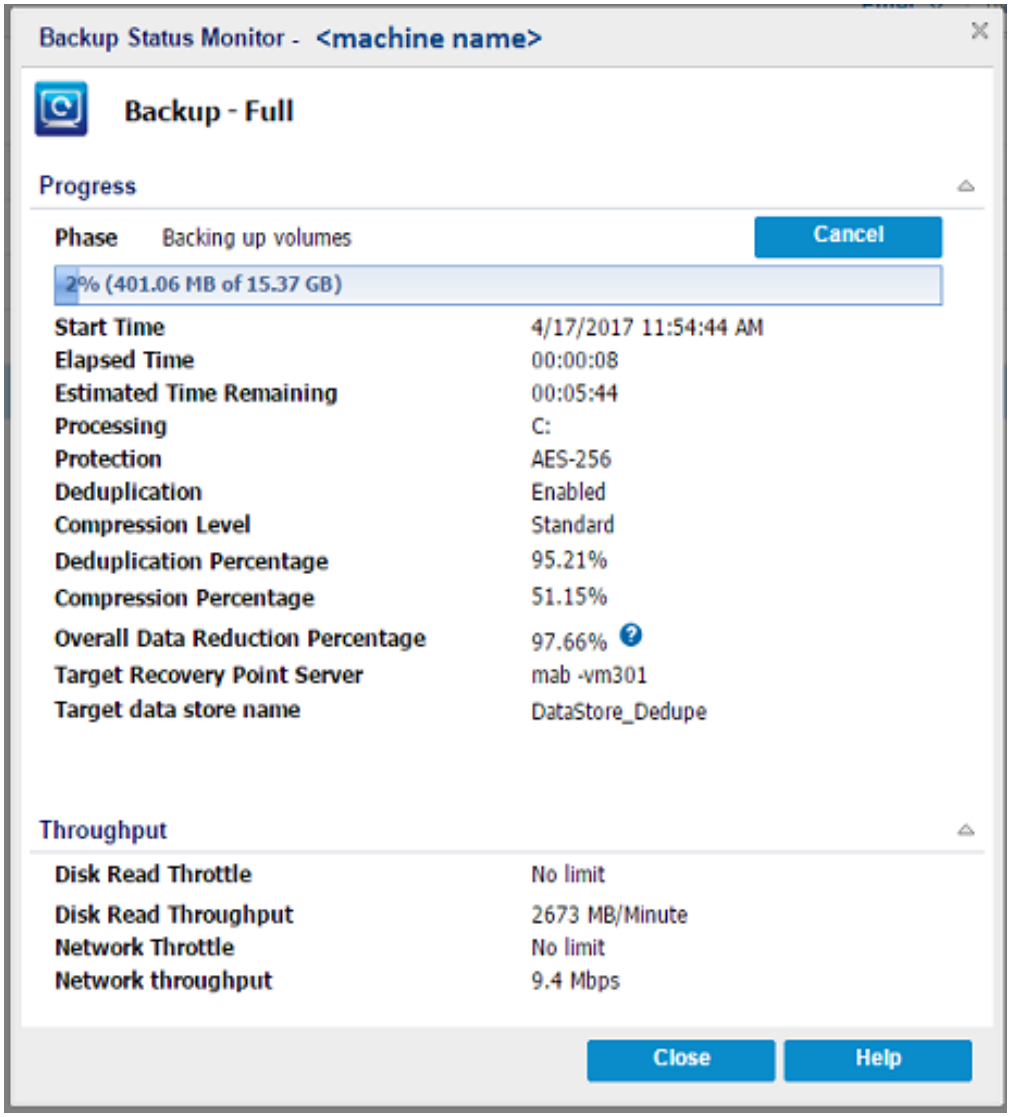


Job Monitor Dialog

The Job Monitor dialog lets you view the status of a job. When a job is running, this panel expands to display information about the ongoing event such as the estimated time remaining to complete the job, the percentage and size of the job already completed, and the total size of the job when completed.

When a job is running, from the right pane, expand **Recent Events** and click the **Detail** button to open the status monitors and display more detailed information about the current running job.

You can click the **Cancel** button to stop the current job.



How to Configure Arcserve UDP

Using Arcserve Unified Data Protection, you can specify the following Arcserve Unified Data Protection configuration settings.

- Server Communication Protocol
- Database Settings
- Arcserve Backup Data Synchronization
- SRM Configuration
- Node Discovery Configuration
- Email and Alert Configuration
- Configure Proxy Settings
- Update Configuration
- Administrator Account
- Remote Deployment Settings
- Share Plan
- User Management
- Configure the Console Timeout Duration

Configure Server Communication Protocol

The Arcserve Unified Data Protection solution uses the Hypertext Transfer Protocol (HTTP) for communication among all of its components. If you are concerned about the security of passwords that are communicated between these components, you can change the HTTP protocol to Hypertext Transfer Protocol Secure (HTTPS). When you do not need this extra level of security, you can change the protocol being used to HTTP.

Note: When you change the protocol to HTTPS, a warning displays in the web browser. The warning appears because of a self-signed security certificate that prompts you to ignore the warning and proceed or add that certificate to the browser to prevent the warning from reappearing.

Follow these steps:

1. Log into the computer where the Arcserve Unified Data Protection Console is installed using an administrative account or an account with administrative privileges.

   Note: If you do not log in using an administrative account or an account with administrative privileges, configure the Command Line to run using the Run as Administrator privilege.

2. Open the Windows Command Line.
3. Perform one of the following tasks:

   - To change the protocol from HTTP to HTTPS:

     Launch the "changeToHttps.bat" utility tool from the following default location.

     Note: The location of the BIN folder can vary depending upon where you installed the Arcserve Unified Data Protection Console.

     C:\Program Files\Arcserve\Unified Data Protection\Management\BIN

     When the protocol is successfully changed, the following message displays:

     The communication protocol was changed to HTTPS.

   - To change the protocol from HTTPS to HTTP:

     Launch the "changeToHttp.bat" utility tool from the following default location

     Note: The location of the BIN folder can vary depending upon where you installed the Arcserve Unified Data Protection Console.

     C:\Program Files\Arcserve\Unified Data Protection\Management\BIN

When the protocol is successfully changed, the following message displays:

The communication protocol was changed to HTTP.

4. Restart the browser and reconnect to Arcserve Unified Data Protection Console.

   Note: To update the communication protocol used by the Arcserve Unified Data Protection Recovery Point Server and the Arcserve Unified Data Protection Agent to communicate with the Arcserve Unified Data Protection Console, update the node directly from the Console.

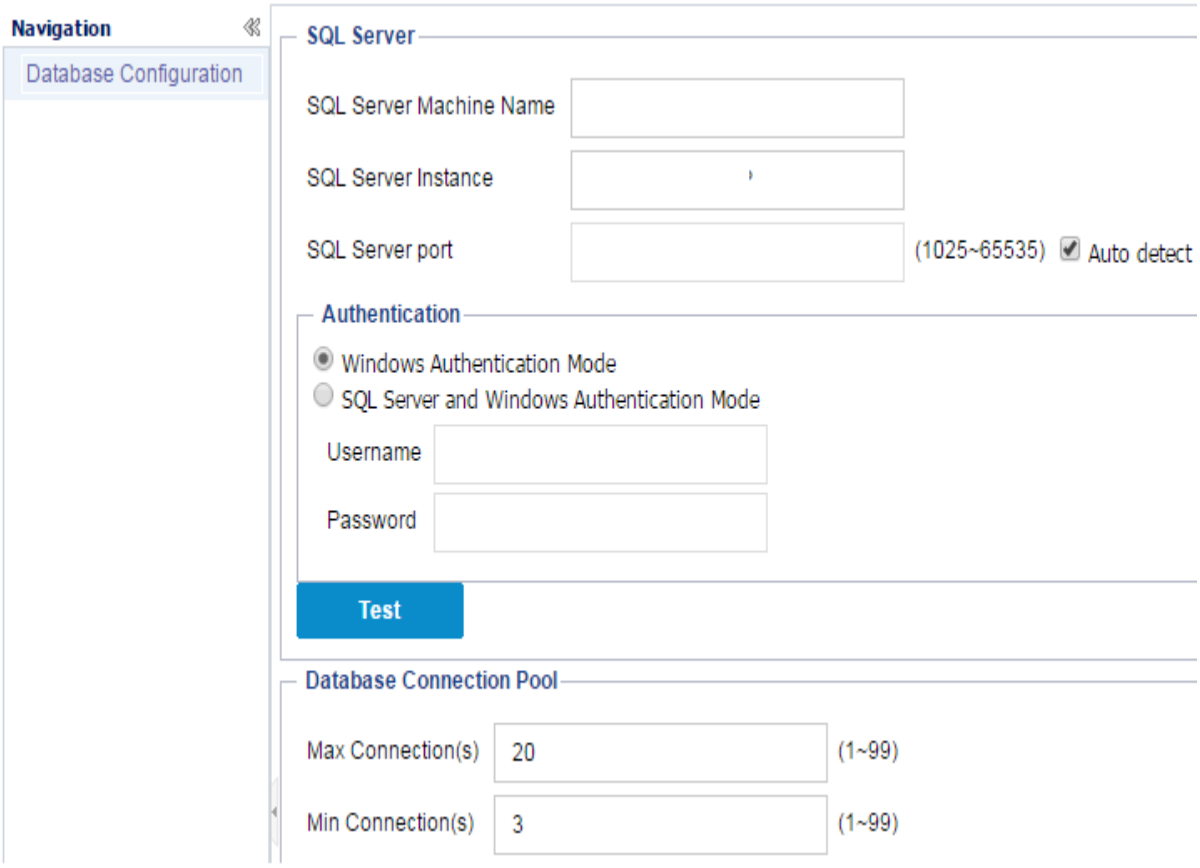
Configure Database

The **Database Configuration** page lets you enter details about the database. The database configuration requires details about SQL Server, number of connections, and authentication mode.

Note: You can re-create the database before configuring. Delete the Arcserve Unified Data Protection database using the procedure described in Re-create the Arcserve Unified Data Protection Database, and then configure the database.

Follow these steps:

1. From the Console, click the **settings** tab.
2. From the left pane, click **Database Configuration**.



To configure, complete the following fields on the configuration pane, and click **Save**.

SQL Server Machine Name

Specify the name of the server that hosts the SQL Server instance.

SQL Server Instance

Specify the name of the SQL Server instance.

SQL Server port

Specify the port number for this instance or enable the **Auto detect** option. 1025 to 65535 is the range of options for the port number.

Auto detect

Selecting the check box lets the application find the port number.

Authentication

Select one of the Authentication Modes from the following options:

Windows Authentication Mode: Default mode:

(Optional) **Test**: Click Test to verify that the application can communicate with the Microsoft SQL Server instance.

SQL Server and Windows Authentication Mode:

Select the option and enter details in the User Name and Password fields.

Database Connection Pool values

For Maximum and Minimum Connections, enter a value from 1 to 99.

The Database Server configuration is set.

Use **Reset** to clear all of the specified values and load the original data.

Re-create the Arcserve Unified Data Protection Database

For various reasons, you may want to re-create the Arcserve Unified Data Protection database. For example, your current database consumes more than 10 GB of data. To re-create the database, first you need to delete the existing Arcserve Unified Data Protection database and then configure a new database to replace the deleted database. The procedure applies to Microsoft SQL Server and Microsoft SQL Server Express Edition databases.

Important! When you delete the Arcserve Unified Data Protection database, all current data is lost.

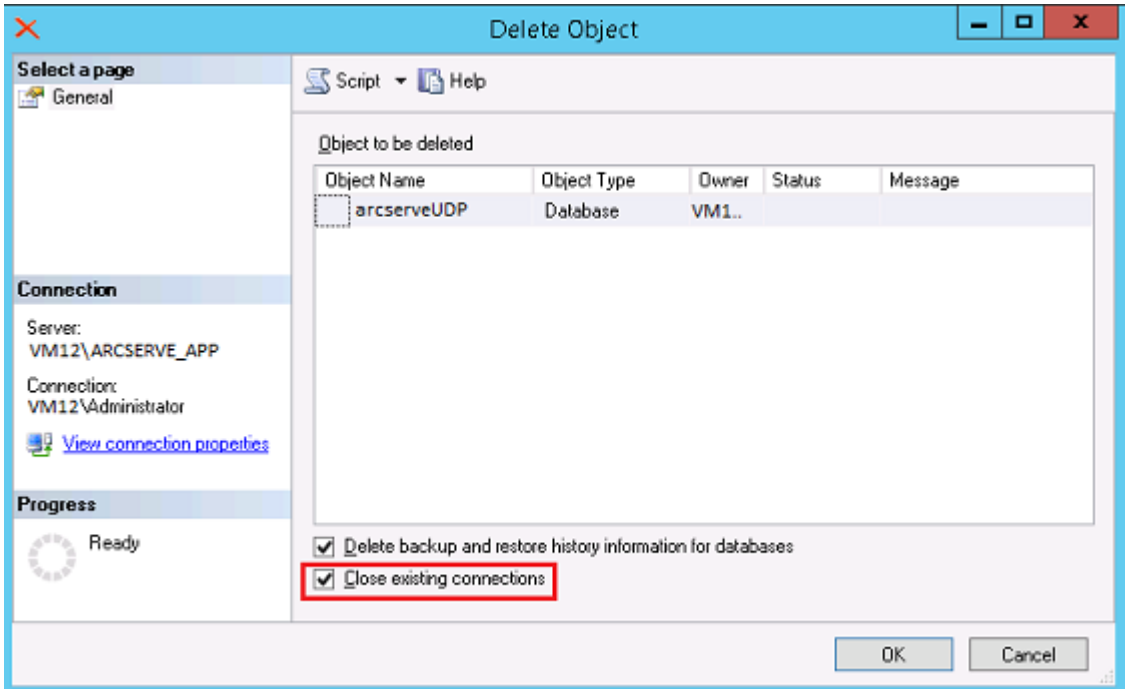
To re-create the Arcserve Unified Data Protection database

1. Open Microsoft SQL Server Management Studio Express and log into the ARCSERVE_APP instance.

   Note: If Microsoft SQL Server Management Studio Express is not installed on the Arcserve Unified Data Protection server, download the utility from the Microsoft Download Center.

2. Right-click arcserveUDP and click **Delete** on the pop-up dialog.

   The **Delete Object** dialog opens.



3. On the **Delete Object** dialog, click the **Close existing connections** option, and then click **OK**.

The existing Arcserve Unified Data Protection database is deleted.

4. Configure the new database. For more information, see Configure Database.

   The Arcserve Unified Data Protection solution re-creates the database. The name of the database instance is **ARCSERVE_APP**.

Configure Arcserve Backup Data Synchronization

You can configure the **Arcserve Backup Data Synchronization Schedule**.

Follow these steps:

1. From the Console, click the **settings** tab.
2. From the left pane, click **Arcserve Backup Data Synchronization Schedule**.
3. From the right pane, click **Enable**.

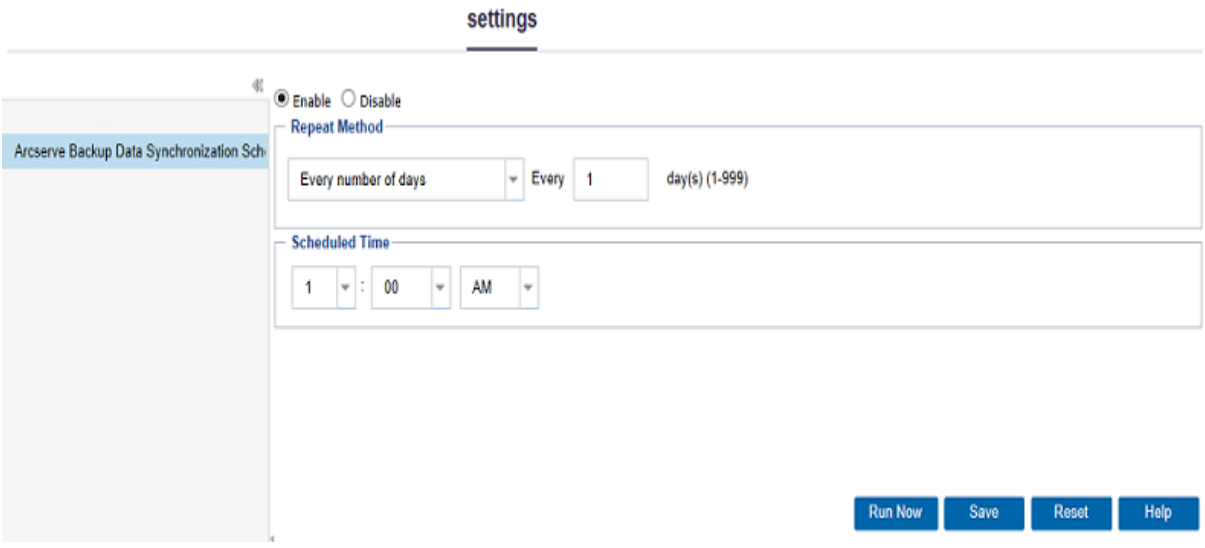
   By default, **Arcserve Backup Data Synchronization** configuration is enabled.

   Note: Clicking **Disable** stops scheduling.
4. Specify the following parameters to schedule Arcserve Backup Data Synchronization:
   - **Repeat Method**
   - **Scheduled Time**
5. Click **Save**.

   The schedule for Arcserve Backup Data Synchronization is applied.

   Note: Do not click **Save** if you want to run the synchronization immediately.
6. (Optional) To run the process immediately, click **Run Now**.

   The **Node** dialog is displayed with the list of nodes available for synchronization.



7. Select the nodes that you want to run for synchronization and click **OK**.

Configure SRM

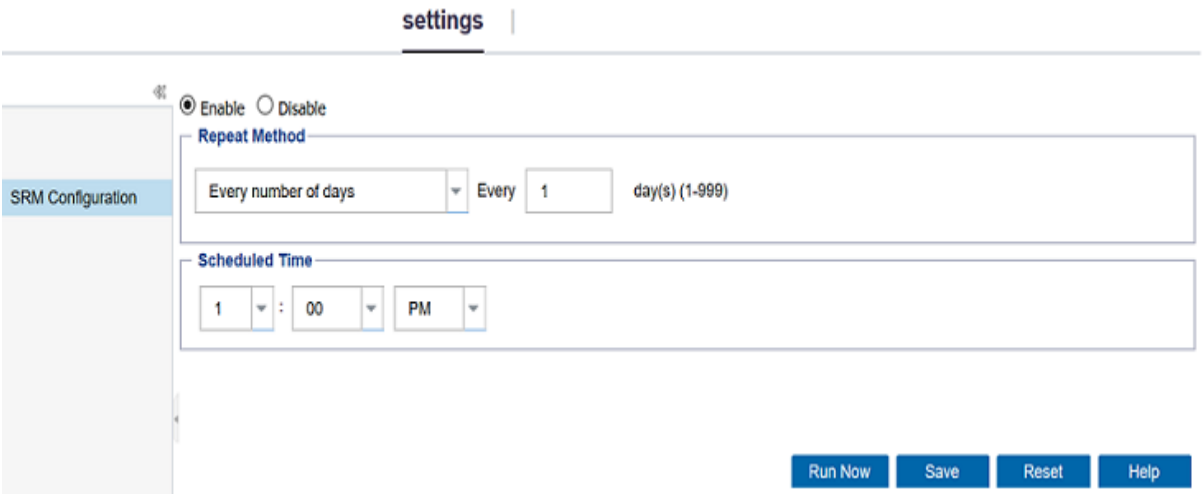
The SRM Configuration page lets you configure an SRM schedule for nodes that defines when and how often to collect SRM data. SRM (Storage Resource Management) is a functionality that collects information about the following data:

- Hardware, software, and application data for Microsoft SQL Server and Microsoft Exchange Server implementations.
- Performance Key Indicators (PKI) data from nodes.

Follow these steps:

1. From the Console, click the **settings** tab.
2. From the left pane, click **SRM Configuration**.
3. From the right pane, click **Enable**.

   By default, **SRM Configuration** is enabled.



   Note: Clicking **Disable** stops scheduling.

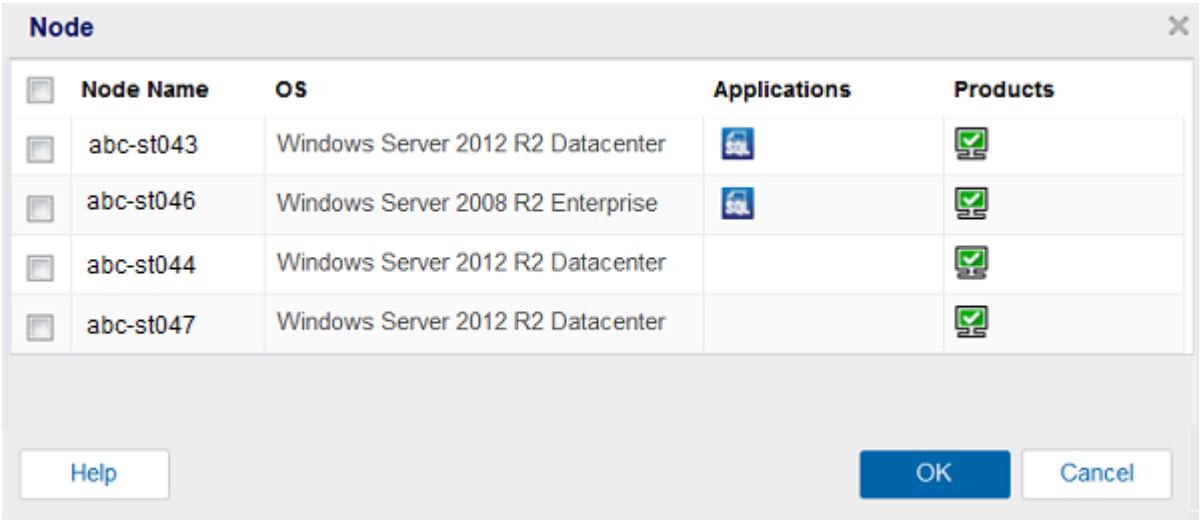
4. Specify the following parameters to schedule SRM:
   - **Repeat Method**
   - **Scheduled Time**
5. Click **Save**.

   The schedule for SRM is applied.

   Note: Do not click **Save**, if you want to collect the SRM data immediately.

6. (Optional) To run the process immediately, click **Run Now**.

   The **Node** dialog is displayed with the list of nodes available for synchronization.



7. Select the nodes that you want to run for synchronization, and click **OK**.

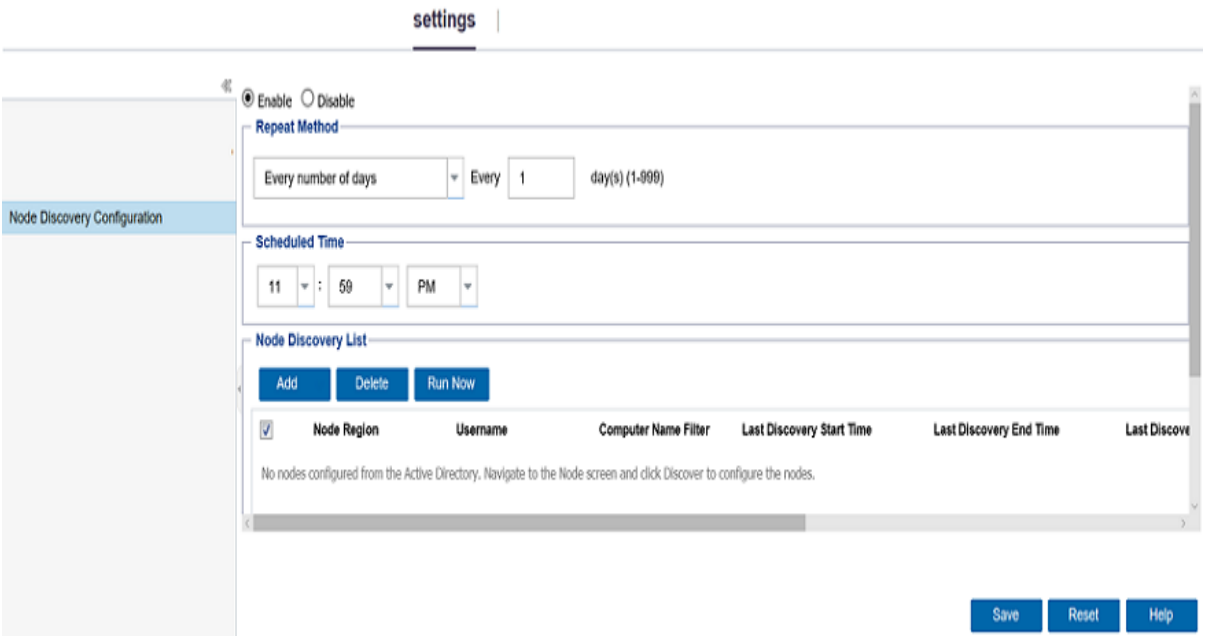
Node Discovery Configuration

The **Node Discovery Configuration** page lets you configure the Active Directory, VMware vSphere, and Microsoft Hyper-V node discovery schedule on a repeating basis and on a scheduled time. When new nodes are discovered, an email alert is sent to the administrator to manually add the new nodes. By default, **Discovery Configuration** is disabled.

Follow these steps:

1. From the Console, click the **settings** tab.
2. From the left pane, click **Node Discovery Configuration**.

   To enable the configuration, click the **Enable** option to specify the type of repeating method that you want and a scheduled time for the node discovery to begin.



You can specify the following parameters to configure your discovery schedule:

- **Every number of days:** Lets you repeat this method on the number of days that are specified. (Default)
- **Every selected day of the week:** Lets you repeat this method on the days that are specified. Monday, Tuesday, Wednesday, Thursday, and Friday are the default days of the week.
- **Every selected day of the month:** Lets you repeat this method on the specified day of the month. 1 is the default option for the day of the month.

- **Scheduled Time:** Lets you specify the time when the discovery runs according to the repeat schedule.
- **Node Discovery List>Add:** Select from where you want to add nodes from. Then, specify the credentials as required.

  Note: Optionally, click **Run Now** to run the discovery instantly.

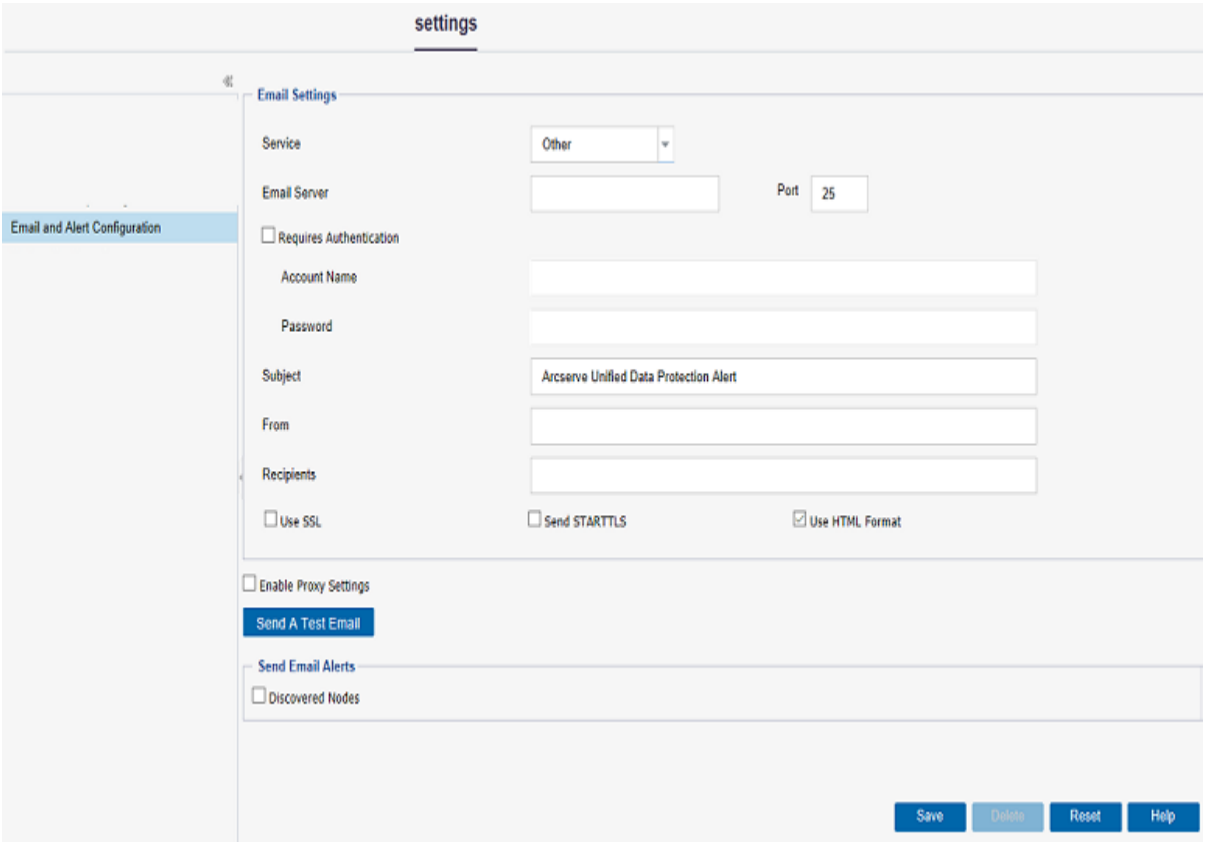
Configure Email and Alert

The **Email and Alert Configuration** page lets you provide email settings and email alerts configuration.

Note: As a prerequisite, install Microsoft .NET Framework (version 2.0 or higher) on the machine where you have installed the Console for the Report Chart export feature to export images in a report successfully.

Follow these steps:

1. From the Console, click the **settings** tab.
2. From the left pane, click **Email and Alert Configuration**.
3. Enter details to set default settings.



Service

Select email services from the available options.

Email Server

Specify the host name of the SMTP server that you can use to send email alerts.

Port

Specify the port number related to the Email server.

Requires Authentication

Select check box to enter credentials.

Use SSL/Send STARTTLS/Use HTML Format

Select the desired option to specify requirements.

Enable Proxy Settings

Select check box to enter **Proxy Server** and Authentication details.

Test Email

Click to verify the details entered in the Email Settings section.

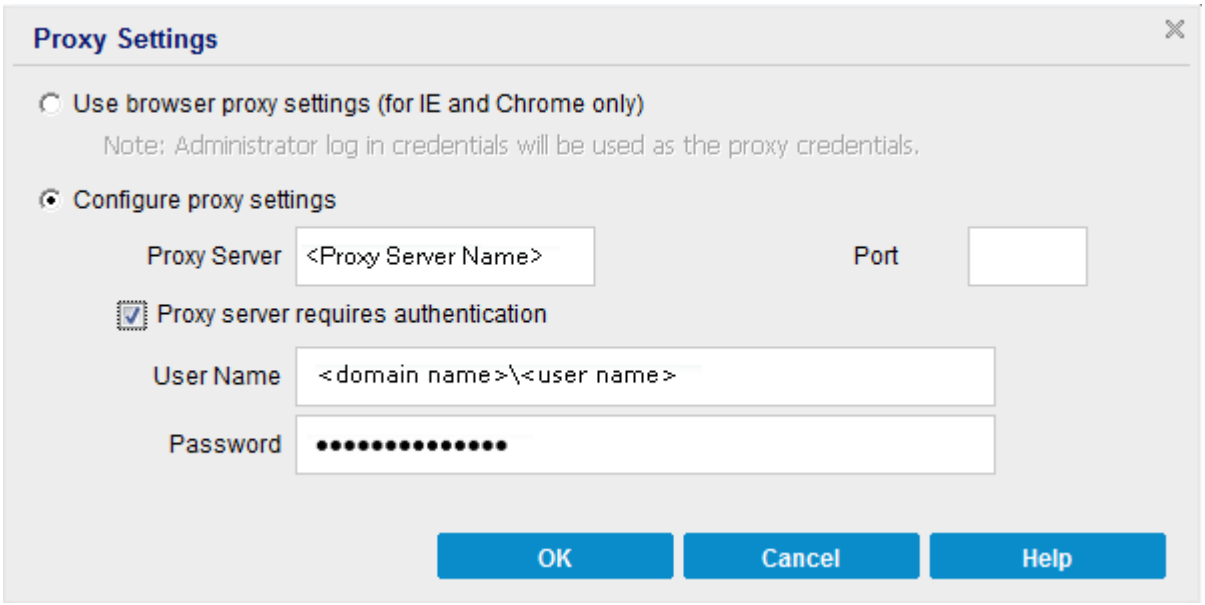
Send Email Alerts

Select **Discovered Nodes** to configure **Active Directory** nodes that you can find using the Discover feature available for Nodes under the **resources** tab.

Configure Proxy Settings

Select **Proxy Settings** to specify if you want the Arcserve UDP to communicate through proxy server. A proxy server acts as an intermediary between your server and the Arcserve server to ensure security, increased performance, and administrative control. This serves as the connection to the Arcserve server from which your download server gets the updates.

When you select the Arcserve Server as the download server, the **Proxy Settings** dialog opens.



- **Use browser proxy settings**

  This selection is only applicable to Windows Internet Explorer (IE) and Google Chrome.

  When selected, directs Arcserve UDP to automatically detect and use the same proxy settings that are applied to the browser to connect to the Arcserve server for Arcserve UDP update information.

- **Configure proxy settings**

  When selected, enables the specified proxy server to connect to the Arcserve server for Arcserve UDP update information. If you select this option, you must also include the IP address (or machine name) of the proxy server and the corresponding port number that is used by the proxy server for internet connections.

In addition, you can also specify if your proxy server will require authentication. When selected, specifies that authentication information (User ID and Password) are required to use the proxy server.

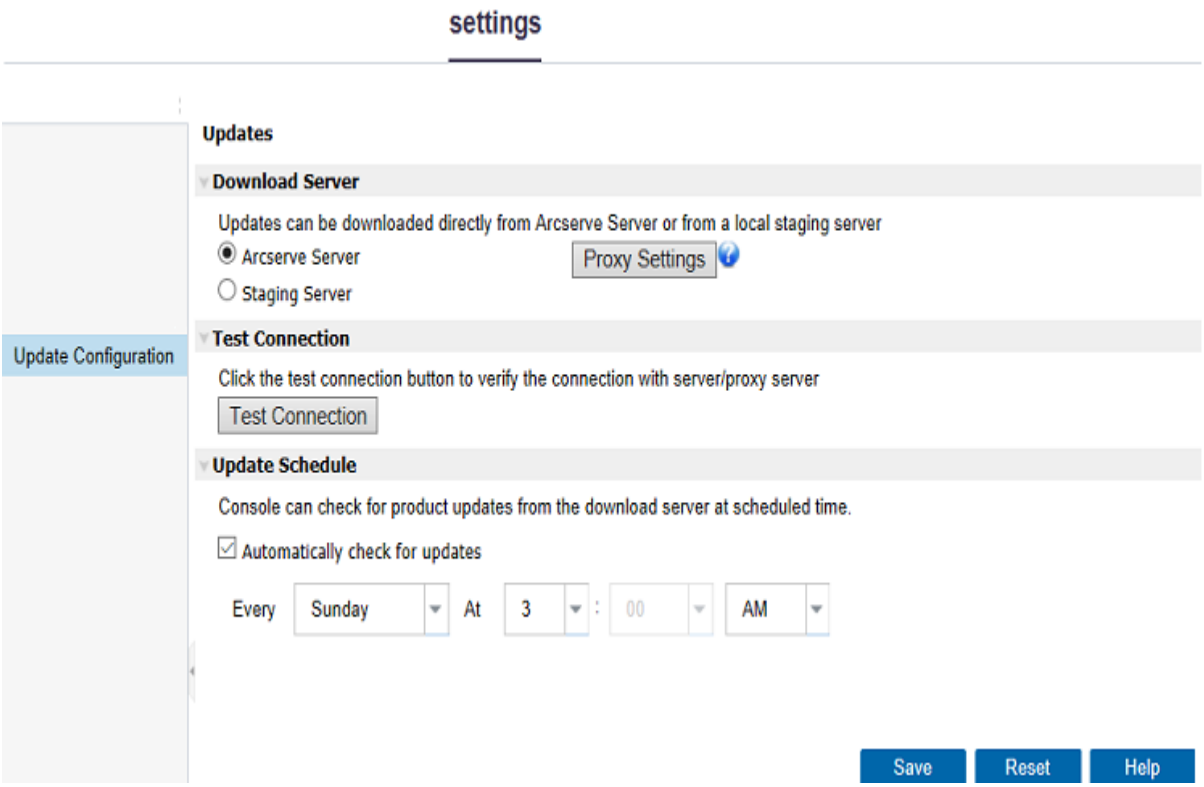
Note: The format for user name should be a fully qualified domain user name in the form of "<domain name>\<user name>".

Update Configuration

The Update Configuration page lets you set Download Server and Update schedule for configuring updates. You can provide details about Arcserve Server proxy settings or Staging server for Download Server.

Follow these steps:

1. From the Console, click the **settings** tab.
2. From the left pane, click **Update Configuration**.



Provide the details on the type of Update Server and Update Schedule. Update Server is either Arcserve Server or Staging Server.

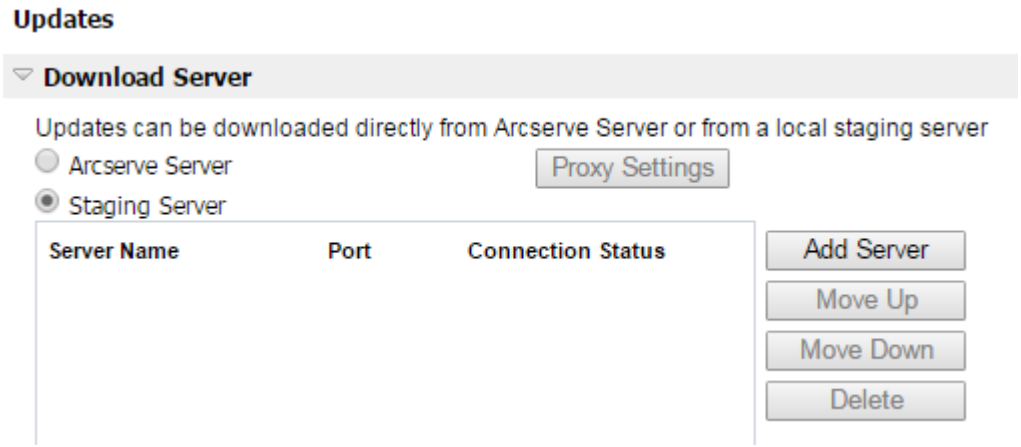
3. For **Download Server**, select one of the following options:
   - For **Arcserve Server**, click **Proxy Settings** to complete Proxy Setup.
   - For **Staging Server**, click **Add Server** to provide staging server details.

     To create a Staging Server, view How to Create a Staging Server.

     If the Staging server is behind a firewall, prepare one machine that can access internet, has UDP product installed and gets the latest updates from Arcserve Server. Then, copy **EngineUpdates** and **FullUpdates** folders from the fol-

lowing location to the staging server machine:

<UDP install path>\ Arcserve\Unified Data Protection\Update Manager



Notes:

- As Auto update feature does not need license, you do not need to activate license on staging server.
- For more information about ports, refer to Communication Ports Used by Arcserve UDP.

You can add multiple staging servers.

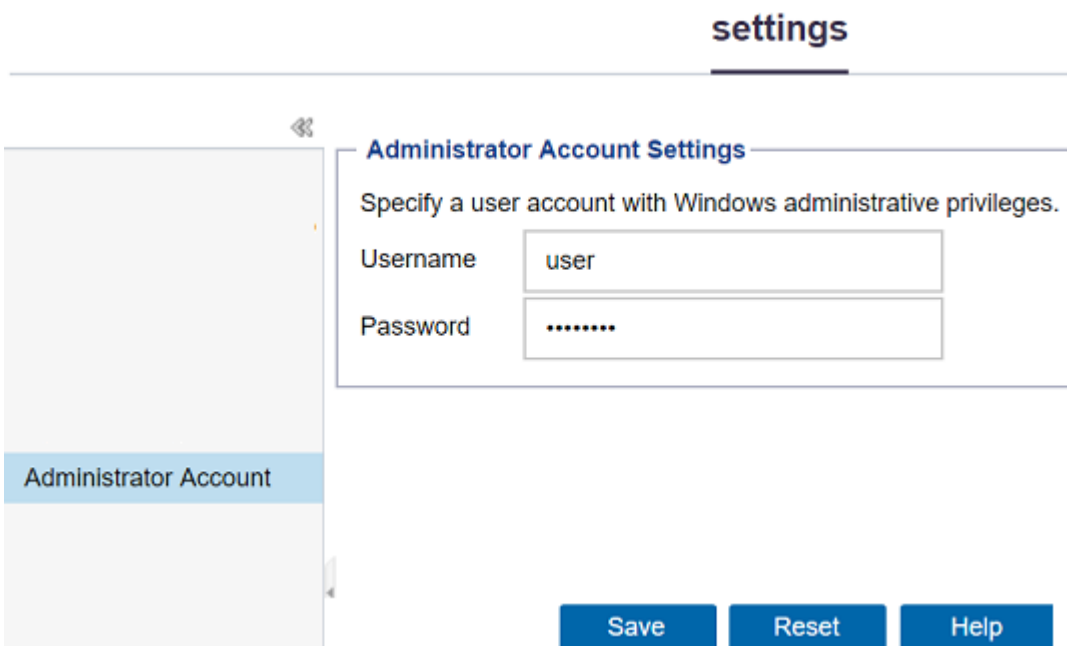
4. Click **Test Connection** to verify the **Download server** details.
5. Enter details for **Update Schedule**.
6. Select **Automatically check for updates**.
7. Click **Save** to complete the update.

Configure Administrator Account

The **Administrator Account** page lets you create a user account by providing a **user-name** and **password**.

Follow these steps:

1. From the Console, click the **settings** tab.
2. From the left pane, click **Administrator Account**.



3. Provide your administrator account credentials, and click **Save**.

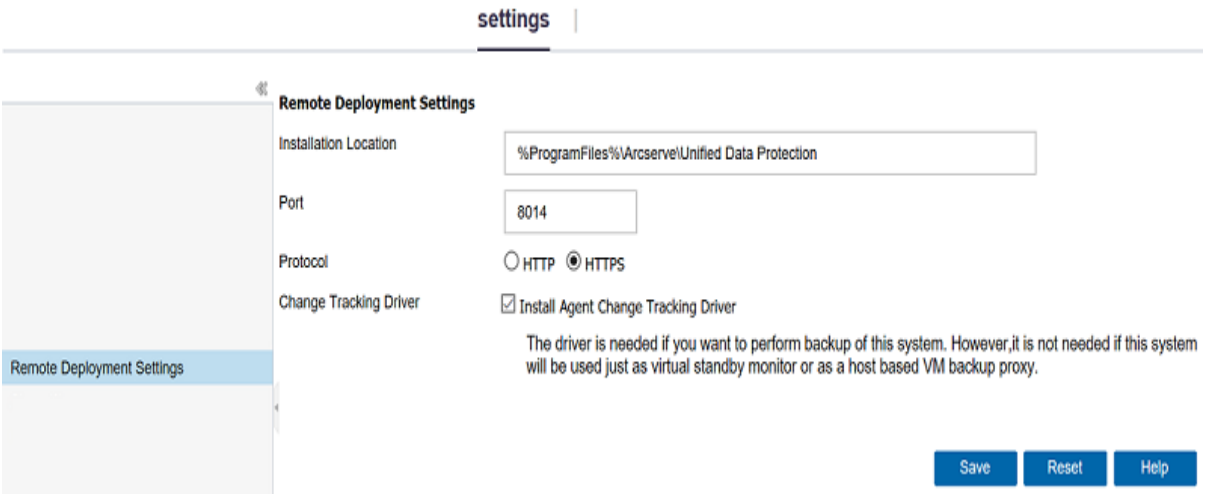
Remote Deployment Settings

The **Remote Deployment Settings** page lets you specify default settings for installing Arcserve Unified Data Protection Agent and Arcserve Unified Data Protection Recovery Point Server. Specify the default installation settings to provide location of installation.

Enter the details for **install path**, **protocol**, and **port**, and click **Save**.

Follow these steps:

1. From the Console, click the **settings** tab.
2. From the left pane, click **Remote Deployment Settings**.



3. Enter details as required, and click **Save**.

Map the Plan to the User Account

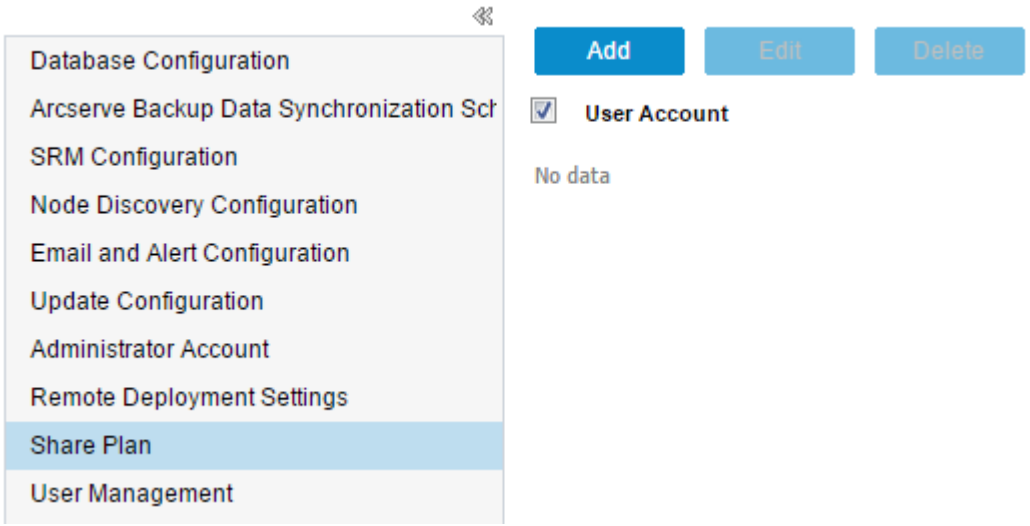
Destination Administrator

You have already created a user account and a plan for a source Console. To identify and manage replicated data, assign the plan to the user account.

Note: You can assign more than one plan to a user account but two different accounts cannot share a plan. However, we recommend assigning a single plan to a user account so that you can easily identify and manage the replicated data.

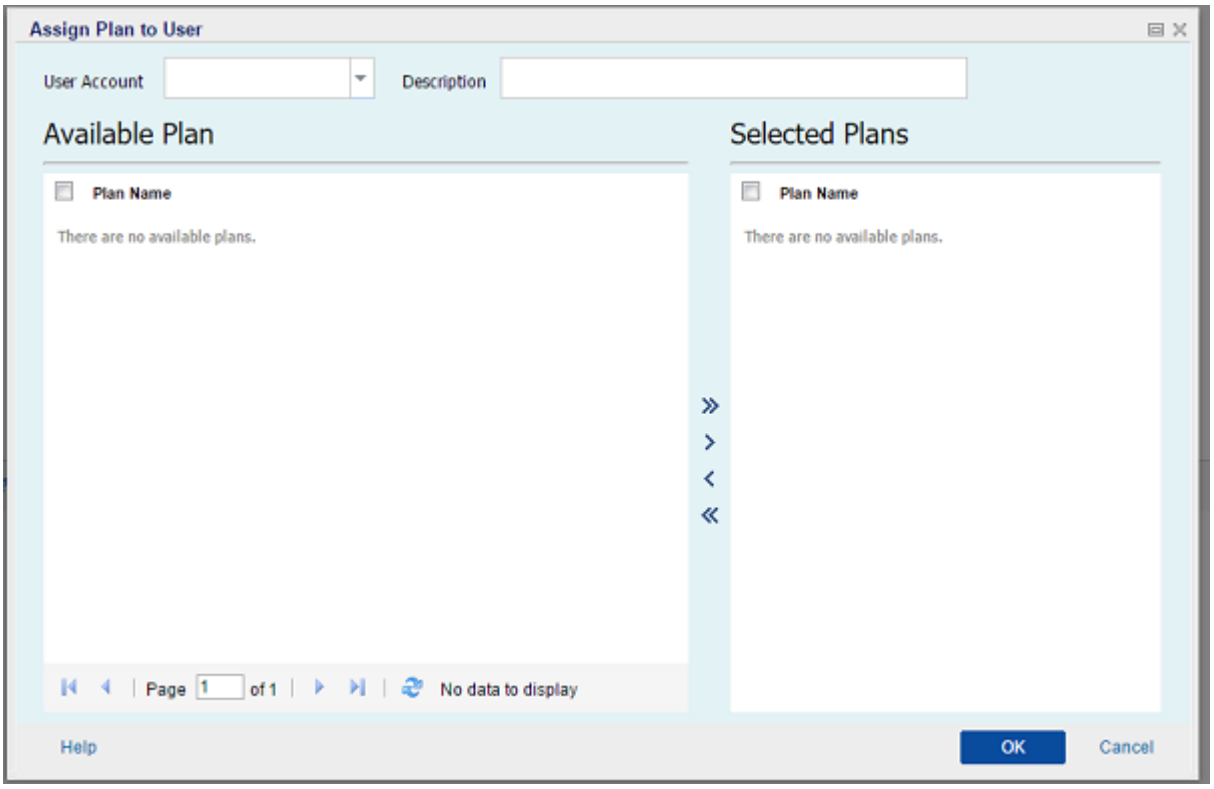
Follow these steps:

1. From the Console, click the **settings** tab.
2. From the left pane, click **Share Plan**.



3. From the center pane, click **Add**.

   The **Assign Plan to User** dialog opens.



4. Select the **User Account**.
5. Select a plan from the **Available Plan** column.

   Note: If a plan is already added to a user name, that plan is not displayed in the **Available Plan** column.
6. Click **Add all plans** or **Add selected plans** to add the plans in the **Selected Plans** column.
7. Click **OK**.

   The **Assign Plan to User** dialog closes. The user name and the associated plans are displayed on the **Share Plan** page.

   The user account is mapped to the plan created for the source Console.

   You can use **Edit** to modify the user configuration or **Delete** to remove the user account from the list.

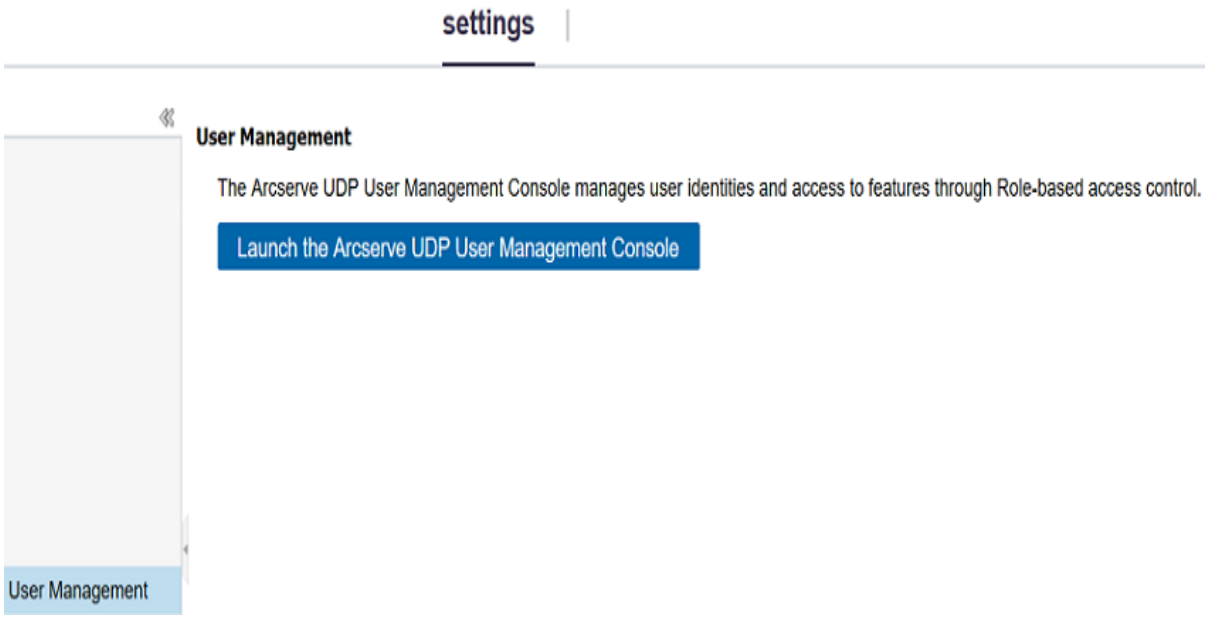
User Management

The User Management page lets you log in to the User Management Console (Identity Service Console) from Arcserve UDP Console. The Arcserve UDP User Management Console manages user identities and controls the access to features using the role-based access control.

Follow these steps:

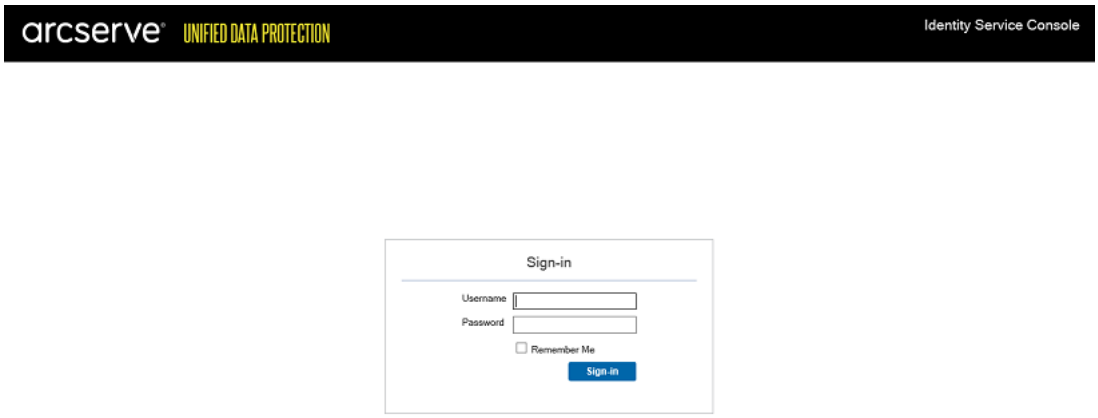
1. Log into the Arcserve UDP Console.
2. Click the **settings** tab.
3. Click **User Management** from the left pane.

   The **User Management** page opens on the center pane.



4. Click **Launch the Arcserve UDP User Management Console**.

   The **Identity Service Console** opens in a new window.



5. Specify the username and password and click **Sign-in**.

   The **Identity Service Console** home page opens.

   You have successfully accessed the Arcserve UDP User Management Console.

   Alternatively, you can also log into the Identity Service Console by entering the address in the following format in a new window:

   http(or https):///(IP address or hostname): (console port number)/carbon

Configure the Console Timeout Duration

If the Console is inactive for a certain duration, you are automatically logged out of the Console. You can change the default timeout value in the ConsoleConfiguration.xml file.

Follow these steps:

1. Log in to the machine where you have installed the Console.
2. Open the ConsoleConfiguration.xml file from the following location:

   <UDP_Home>\Management\Configuration\ConsoleConfiguration.xml
3. Change the value for consoleUISessionTimeout.

   The value is in seconds.

   For example:

   <consoleUISessionTimeout>3600</consoleUISessionTimeout> (default value is 1 hour, value is in seconds)

   3600 indicates that the timeout duration for the Console is 3600 seconds
4. Save the ConsoleConfiguration.xml file.

How to Migrate Arcserve r16.5 Recovery Points to Arcserve UDP

Important! If you are replicating from a shared folder to a data store selected on Recovery Point Server, see How to Perform an Offline Data Replication Using RPS Jumpstart.

To migrate Arcserve r16.5 recovery points to Arcserve UDP, perform the following steps:

1. Create a Data Store to Replicate Data from an Arcserve r16.5
2. Replicate Arcserve r16.5 Data to the UDP Data Store

Create a Data Store to Replicate Data from an Arcserve r16.5 Recovery Point

To replicate data from an existing Arcserve r16.5 D2D recovery point, you first create a data store from the Console where the data will be replicated.

Follow these steps:

1. Log into the Arcserve UDP Console.
2. Navigate to **Destinations**, **Recovery Point Server**.
3. Select the Recovery Point Server.
4. Right-click and select **Add a Data Store**.
5. Enter the details on the **Add a Data Store** page.
6. Save the data store.

   The data store is created.

Replicate Arcserve r16.5 Data to the UDP Data Store

After creating the data store, you can replicate the Arcserve r16.5 recovery point data using RPS Jumpstart.

Follow these steps:

1. Click **Actions** and then click **RPS Jumpstart**.

   The **RPS Jumpstart Wizard** opens.

2. Select From a shared folder to a data store on Selected Recovery Point Server.

3. Specify the source shared folder.

   The recovery point details are displayed.

   Note: If the session is not encrypted and the target data store is unencrypted, then session password is optional. If the session is not encrypted and then target data store is encrypted, then you have to provide a session password in the **Select Target Data Store** page.

4. Click **Next**.

   The **Select Target Data Store** page opens. If the source data is encrypted, only the encrypted data stores are displayed in the drop-down list.

5. (Optional) Specify the session password if the session is not encrypted in Step 3.

6. Click **Next**.

7. Click **Finish**.

   The recovery point data from Arcserve r16.5 is replicated to the Arcserve UDP data store.

Set up Configuration Wizard

Using the Configuration Wizard, configure your data protection environment as soon as you log in to the Console. If using Arcserve UDP for the first time, Configuration Wizard is an ideal way to create your first backup plan. The wizard creates plans to define how to protect data. A plan consists of single or multiple tasks to define source, destination, schedule, and advanced parameters. With the first login to the Console, the Configuration Wizard opens. You can hide the welcome page of the wizard by selecting **Do not show this page next time**.

You can create a plan to protect your physical nodes such as Windows and Linux nodes, and virtual machines such as VMware and Hyper-V.

The following steps describe a generic description on how to create a Windows Agent-based plan using the Configuration Wizard.

1. Log into the Console.

   If you are logging in for the first time, the Configuration Wizard opens.

2. If the wizard does not open, then click the **resources** tab and then click **Configuration Wizard** from the right pane.

   The welcome page of the wizard opens.

3. Click **Next**.

   The **Create a Plan** page opens.

4. Specify the plan name.

5. Select a node type to protect.

   For example, select **Backup: Agent-based Windows**.

6. Click **Next**.

   The **Add Nodes to Protect** page opens. The fields are optional in the **Add Nodes to Protect** page.

7. Select the method to add nodes from the drop-down list.

   The fields vary depending on your selection.

8. Provide the node details, click **Add to List**, and then click **Next**.

   The **Backup Destination** page opens.

9. Specify the destination type from the drop-down list.

   The remaining fields on the **Backup Destination** page vary depending on the destination type. You can also enable the session password.

   To create a data store, see Add a data store.

10. Specify the remaining destination details and click **Next**.

    The **Backup Schedule** page opens.

11. Specify the backup schedule and click **Next**.

    The **Plan Confirmation** page opens.

12. Verify the plan.

13. (Optional) Click **Create a Plan** to add another plan.

14. Click **Next**

    The resource configuration is complete.

15. Click **Finish**.

    The wizard closes and a new plan is created.

    You can see the plan in **resources**, **Plans**, **All Plans**.

Chapter 6: Working With Key Features of Arcserve UDP

This section contains the following topics:

Understanding Nutanix Feature in Arcserve UDP

To explore the functions available for Nutanix feature in Arcserve UDP, refer to How to Work With Nutanix Using Arcserve UDP.

Understanding OneDrive Feature in Arcserve UDP

Explore the functions available for OneDrive feature in Arcserve UDP.

- Manage OneDrive Nodes
- How to Create a Microsoft Office 365 OneDrive Backup Plan
- How to Restore OneDrive Data

Understanding Exchange Online Feature in Arcserve UDP

Exchange Online is an email application hosted on Microsoft cloud. To protect your Exchange Online mail items (Mails, Calendar items, Contacts, and so on) from Microsoft Cloud, you need to create a plan in Arcserve UDP. Explore the functions available for exchange online feature in Arcserve UDP.

- User Privileges for Exchange Online in Arcserve UDP
- Add an Exchange Online Node
- Manage Exchange Online Nodes
- Add the Required Role and Group to the Exchange Online Backup Account to Perform Backup and Restore
- How to Create an Exchange Online Backup Plan
- How to Restore Exchange Online Mailbox Data
- Applying Best Practices

Understanding Hardware Snapshot Feature in UDP

Arcserve UDP has the capability to utilize hardware storage snapshots for backup. You can specify whether you want to use the hardware snapshot while creating a backup task. If you select hardware snapshot, then Arcserve UDP first tries to create a hardware snapshot. If hardware snapshot fails, Arcserve UDP automatically reverts to the software snapshot without failing the backup job.

- How to Use Hardware Snapshot for Backup
- Supported Storage Array in Arcserve UDP
- Use Hardware Snapshot for VMware Agentless Backup
- Use Hardware Snapshot for Hyper-V Agentless Backup
- Use Hardware Snapshot for Agent-based Backup
- Verify that the Backup has Used Hardware Snapshot

Nimble

- Add a Storage Array
- Considerations for Nimble Storage When CHAP Authentication is enabled
- Use Hardware Snapshot for VMware Agentless Backup

HPE 3PAR storeserve

- Add a Storage Array
- Use Hardware Snapshot for VMware Agentless Backup

NetApp

- Add a Storage Array
- Considerations for NetApp iSCSI/FC Support for VMware
- Conditions applied to hardware snapshot for NetApp NFS VMware
- Use Hardware Snapshot for VMware Agentless Backup

Understanding Cloud Feature in UDP

Arcserve UDP Cloud feature ensures that you have protected data in Cloud in case of any accidental deletion. By using Arcserve UDP Cloud features, you can copy the specified files, Recovery Points, create instances using Recovery Points, create Virtual Standby Virtual Machines in Cloud, and so on.

View the links to use Cloud in Arcserve UDP:

- Add a Cloud Account
- Manage Nodes for Cloud
  - Download Recovery Point from Cloud
  - Upload Recovery Point to Cloud
  - Copy Recovery Point to local disk or network share
- How to configure an RPS in the cloud
  - Specify Cloud Configuration for Restoring from a file copy cloud location
  - Specify Cloud Configuration for Restoring from a file archive cloud location
- How to Backup data to the cloud
  - How to Create a Virtual Standby to AWS EC2 Plan
  - How to Create and Manage an Instant Virtual Machine on Amazon EC2
  - How to Create a Copy Recovery Points Plan
- Cloud for Linux
  - Install Arcserve UDP Agent (Linux) In AWS Cloud
  - How to Perform a Bare Metal Recovery (BMR) for Linux Machines in AWS Cloud
  - How to Perform a Migration BMR for Linux Machines from Amazon EC2 to local
  - How to Perform IVM migration from Cloud to Local
- Microsoft Office 365 Backup Plan
  - How to Create an Exchange Online Backup Plan
  - How to Restore Exchange Online Mailbox Data
  - How to Create a Microsoft Office 365 OneDrive Backup Plan

- How to Restore OneDrive Data
- Create a SharePoint Online Backup Plan (Watch Video)
- How to Restore SharePoint Online Site Collection Data

- Microsoft Azure
  - How to Create a Virtual Standby to Microsoft Azure Plan
  - How to Create and Manage an Instant Virtual Machine on Microsoft Azure

- Troubleshooting
  - How to Add Encryption Password for an Existing Encrypted Destination
  - Configure the Registry for Copy Recovery Point Job
  - Bandwidth Congestion with Copy Recovery Point to Cloud Jobs
  - Unable to Connect to Cloud

Understanding UNC/NFS Path Feature in UDP

The UNC/NFS Path is introduced as a new node type. To use the feature, you can refer to the following sections:

- User Privileges for UNC/NFS Path in Arcserve UDP
- How to Add and Manage UNC/NFS Path
- How to Create a UNC/NFS Path Backup Plan
- How to Restore From a UNC/NFS Path

Understanding SharePoint Online Feature in UDP

Arcserve UDP v6.5 Update 2 and higher supports working with Microsoft SharePoint Online environment. To use the feature, you can refer to the following sections:

- Prerequisites
- Add a SharePoint Node
- Manage SharePoint Online Nodes
- Create a SharePoint Online Backup Plan (Watch Video)
- How to Restore SharePoint Online Site Collection Data

Prerequisites

Arcserve UDP 8.0 has the following prerequisite to work with Microsoft SharePoint Online environment:

Ensure that you have appropriate user privileges for Arcserve UDP functions. For more information, view User Privileges for SharePoint Online in Arcserve UDP.

Chapter 7: Using Arcserve UDP Role-based Administration

Role-based Administration (RBAC) allows administrators to assign different roles and permission to different users for using the Arcserve UDP Console. Each role can have its own permissions. A super administrator role can create customized roles and permission for other users of the Arcserve UDP Console.

Using RBAC, you can assign varied level of security to each role.

This section contains the following topics:

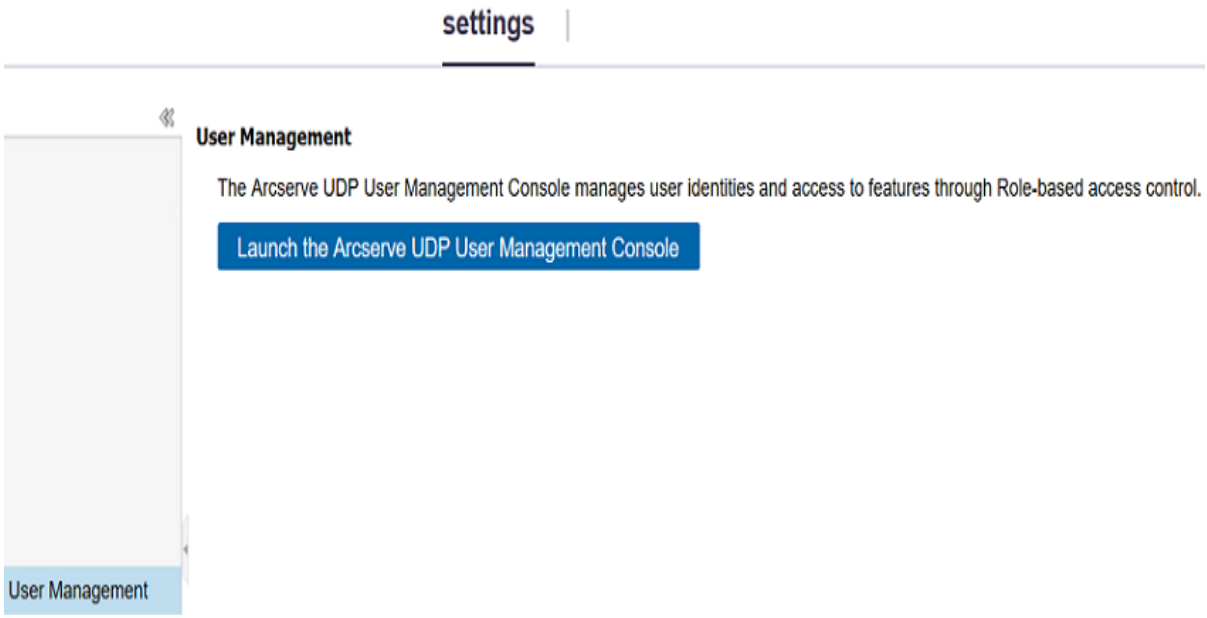
Access the User Management Console

The User Management page lets you log in to the User Management Console (Identity Service Console) from Arcserve UDP Console. The Arcserve UDP User Management Console manages user identities and controls the access to features using the role-based access control.

Follow these steps:

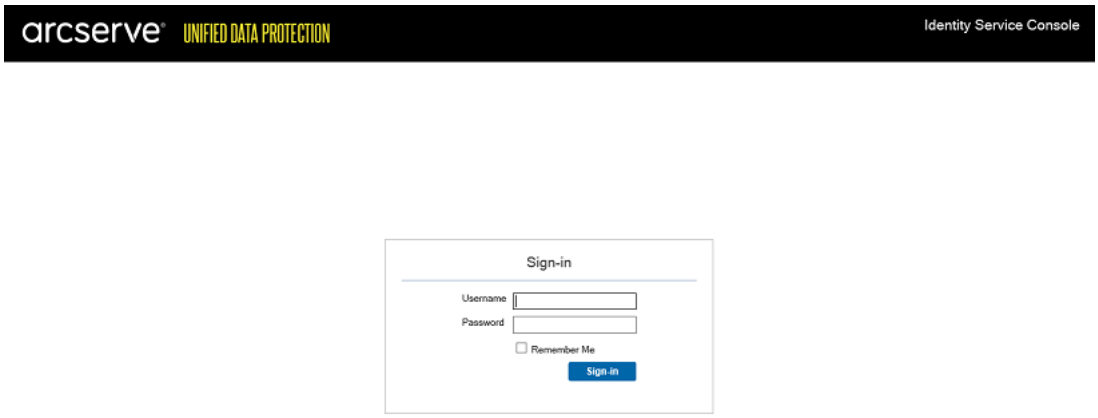
1. Log into the Arcserve UDP Console.
2. Click the **settings** tab.
3. Click **User Management** from the left pane.

   The **User Management** page opens on the center pane.



4. Click **Launch the Arcserve UDP User Management Console**.

   The **Identity Service Console** opens in a new window.



5. Specify the username and password and click **Sign-in**.

   The **Identity Service Console** home page opens.

   You have successfully accessed the Arcserve UDP User Management Console.

   Alternatively, you can also log into the Identity Service Console by entering the address in the following format in a new window:

   http(or https):///(IP address or hostname): (console port number)/carbon

Add User, Delete User, and Change Password

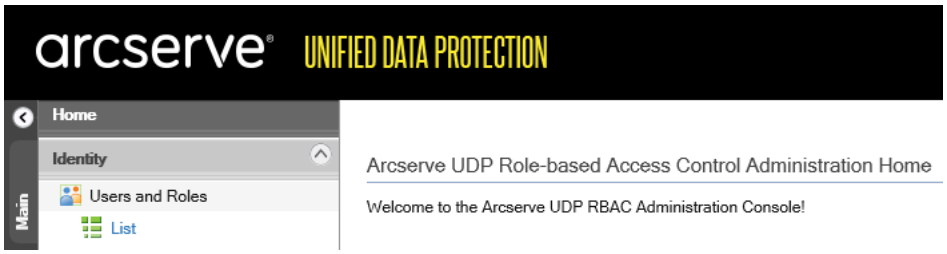
You can add or delete local or domain users from Windows User Control. The user management list on the Identity Service Console updates immediately.

Use Windows User Control to change the user password. When you update the password of a user, the user must log in to the Identity Service Console using the latest password. The role of the user is retained.

Configure User Management

You can assign different roles to different users and provide different permissions to different roles. Configuring user management helps you perform the following options:

- Pre-defined Roles
- Assign a Pre-defined Role
- View and Cancel a Role
- Add a New Role
- Search Users and Roles

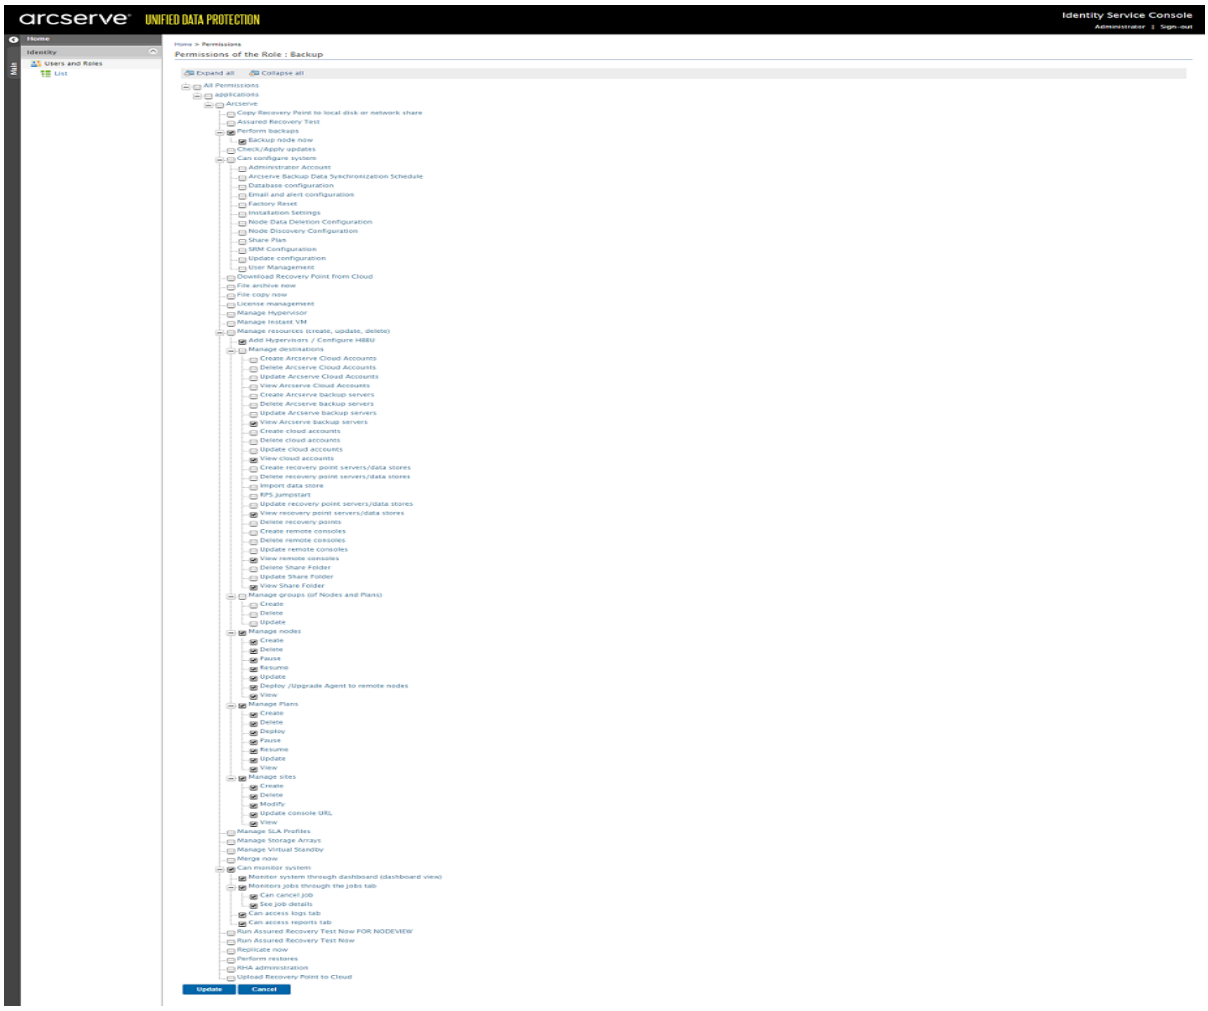


Pre-defined Roles

The function of a pre-defined role is to provide a reference for some typical role definition. Each role has a predefined set of permissions assigned.

For the admin role, all the options in the permission are selected. An admin role can access all the functions of Arcserve UDP.

Click the permission for the Backup role and the following selected permissions are pre-defined:

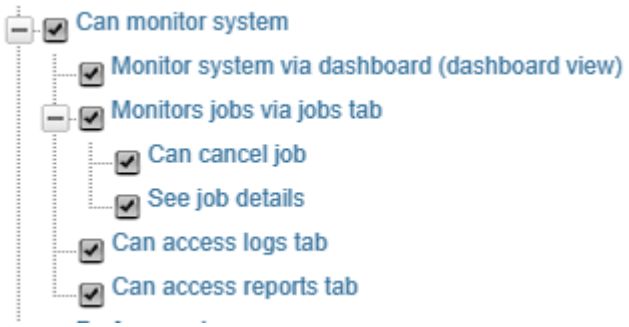

arcserve
UNIFIED DATA PROTECTION
Identity Service Console
Administrator | Sign-out
Home
Identity
Users and Roles
List
Main
Home > Permissions
Permissions of the Role : Backup
Expand all
Collapse all
All Permissions
applications
Arcserve
Copy Recovery Point to local disk or network share
Assured Recovery Test
Perform backups
Backup node now
Check/Apply updates
Can configure system
Administrator Account
Arcserve Backup Data Synchronization Schedule
Database configuration
Email and alert configuration
Factory Reset
Installation Settings
Node Data Deletion Configuration
Node Discovery Configuration
Share Plan
SRM Configuration
Update configuration
User Management
Download Recovery Point from Cloud
File archive now
File copy now
License management
Manage Hypervisor
Manage Instant VM
Manage resources (create, update, delete)
Add Hypervisors / Configure HBBU
Manage destinations
Create Arcserve Cloud Accounts
Delete Arcserve Cloud Accounts
Update Arcserve Cloud Accounts
View Arcserve Cloud Accounts
Create Arcserve backup servers
Delete Arcserve backup servers
Update Arcserve backup servers
View Arcserve backup servers
Create cloud accounts
Delete cloud accounts
Update cloud accounts
View cloud accounts
Create recovery point servers/data stores
Delete recovery point servers/data stores
Import data store
RPS jumpstart
Update recovery point servers/data stores
View recovery point servers/data stores
Delete recovery points
Create remote consoles
Delete remote consoles
Update remote consoles
View remote consoles
Delete Share Folder
Update Share Folder
View Share Folder
Manage groups (of Nodes and Plans)
Create
Delete
Update
Manage nodes
Create
Delete
Pause
Resume
Update
Deploy /Upgrade Agent to remote nodes
View
Manage Plans
Create
Delete
Deploy
Pause
Resume
Update
View
Manage sites
Create
Delete
Modify
Update console URL
View
Manage SLA Profiles
Manage Storage Arrays
Manage Virtual Standby
Merge now
Can monitor system
Monitor system through dashboard (dashboard view)
Monitors jobs through the jobs tab
Can cancel job
See job details
Can access logs tab
Can access reports tab
Run Assured Recovery Test Now FOR NODEVIEW
Run Assured Recovery Test Now
Replicate now
Perform restores
RHA administration
Upload Recovery Point to Cloud
Update
Cancel

The following permissions are pre-defined for the backup role:

- Perform backup
- View destination
- Manage nodes/plan/sites
- Monitor system function

The admin role has the complete flexibility to clear the selected permissions or select a new permission. When you click Update, the newly added permission becomes the default permission for the backup role. You can also rename the role.

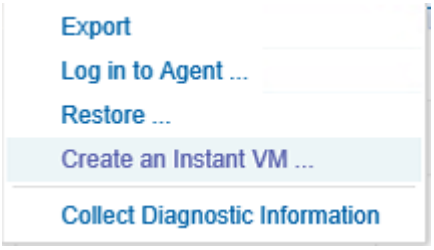
For the Monitor role, the dashboard job monitor and log/report permission are preselected.



For the Restore role, the following permissions are preselected:

- Manage instant VM
- View destination
- View node
- View plan
- Manage virtual standby
- Monitor jobs
- Access logs
- Perform restore

If you assign the Restore role to one user, the user can log in and will have the corresponding authority. For example, if a user has the Restore role, then if a node is backed up successfully, you can Create an Instant VM or Restore to go to the next activity.



For the RHA_Admin role, the administration permission has access to the high availability function.

Assign a Pre-defined Role

When a super administrator assigns a role to any user, only that user can log in to the Console. The Users button displays the complete users list, including domain users and local users.

You can assign the available (pre-defined) roles or self-defined roles to any local user or domain users.

Note: Only a Super-administrator (an administrator who installs Arcserve UDP) can assign the administrator role to other users. Administrators can assign only non-administrator roles to other users.

Follow these steps:

1. Click **Users and Roles** from the **Configure** pane.

   Users and Roles are displayed on the User Management page.

2. Click **Users** from the User Management page.

   The list of users is displayed.

3. Click **Assign Role** for a user.

   The **Role List of User** page opens.

4. Select one or more roles and click **Update**.

   The role for user is successfully updated.

5. Click **Finish** to go to the previous screen.

   You have successfully assigned a role to user.

View and Cancel a Role

You can view the current role that is assigned to a user. You can cancel an assigned role by clearing the check box for that role.

Follow these steps:

1. Click **Users and Roles** from the **Configure** pane.

   Users and Roles are displayed on the **User Management** page.

2. Click Users from the User Management page.

   The list of users is displayed.

3. Clear the check box to cancel a role and click **Update**.

   The roles are removed from the user.

   Note: If not assigned any role, a user cannot log into the Arcserve UDP Console.

4. Click **Finish** to go to the previous page.

   You have successfully viewed and canceled a role.

Add a New Role

You can create a customized role and select permission for that role.

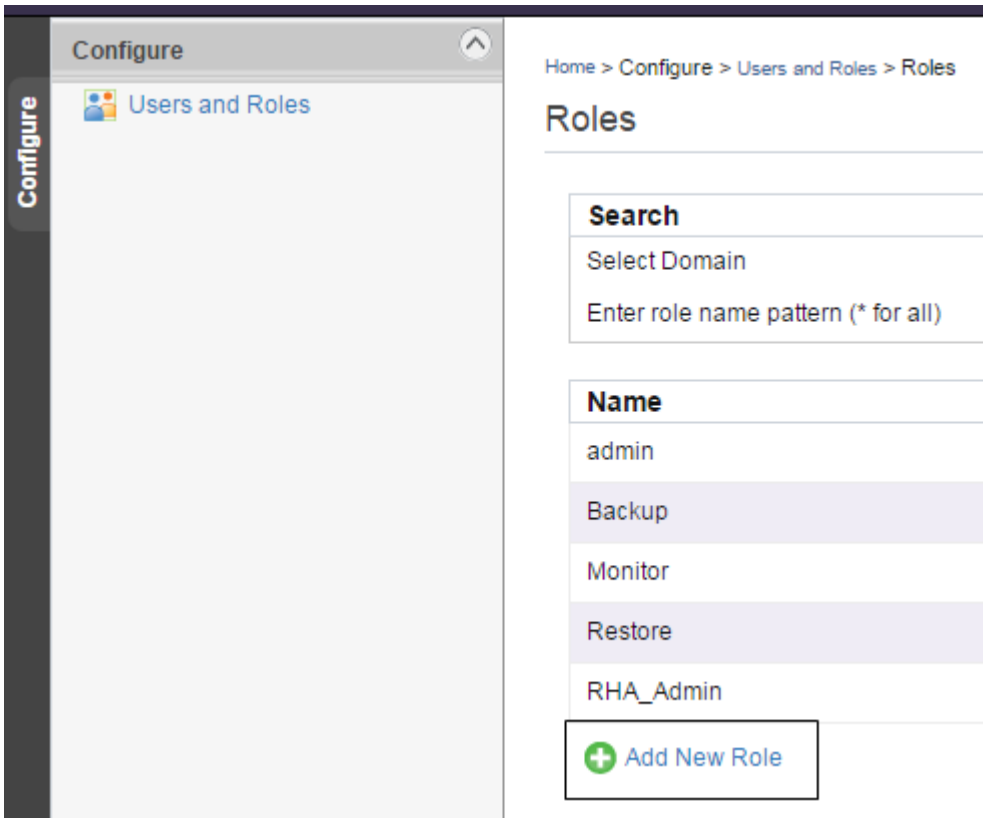
Follow these steps:

1. Click **Users and Roles** from the **Configure** page.

   Users and Roles are displayed on the **User Management** page.

2. Click **Roles**.

   The **Roles** page opens and lists all the available roles.



3. Click **Add New Role**.

   The **Add Role** page opens.

4. Provide a role name and click **Next**.

   Note: Do not use special characters such as **~!@#$%^&*** and others in the name of a role.

5. Select the required permission check boxes and click **Next**.

6. Select users for this role.

7. Click **Finish**.

   A new role is created and assigned permissions to this role.

Note: Some permissions work only when the related permissions are also selected. For example, to configure a role for managing Virtual standby, select the permission to manage Virtual standby and select the view node permission to ensure that the role can function normally.

Search Users and Roles

You can filter Users and Roles to find the required user or role. Enter * to search all users and roles.

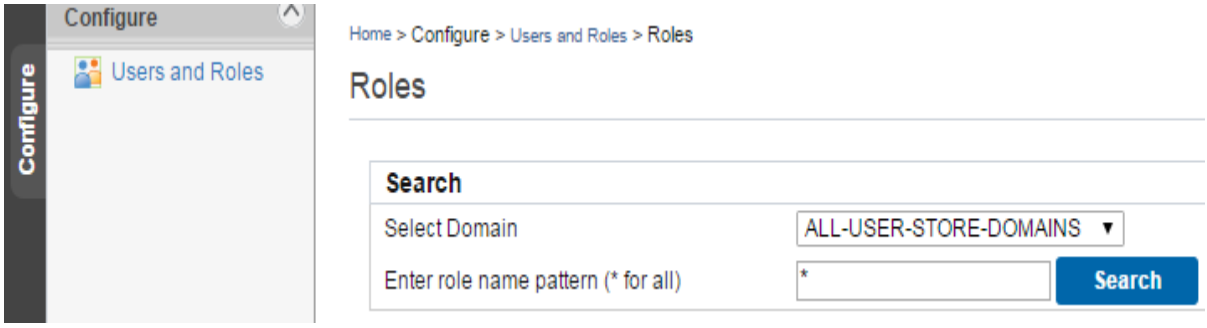
To search a role, follow these steps:

1. Click **Users and Roles** from the **Configure** page.

   Users and Roles are displayed on the **User Management** page.

2. Click **Roles**.

   The **Roles** page opens.



3. Specify the role name pattern and click **Search**.

   The filtered result is displayed.

To search a user, follow these steps:

1. Click **Users and Roles** from the **Configure** page.

   Users and Roles are displayed on the **User Management** page.

2. Click **Users**.

   The **Users** page opens.

3. Specify the user name pattern and click **Search**.

   The filtered result is displayed.

Integrate Arcserve UDP with Active Directory

Arcserve UDP supports the Active Directory (AD) integration using Windows and AD groups.

This section contains the following topics:

- How to Integrate Arcserve UDP 8.0 with Active Directory Using Windows Groups
- How to Integrate Arcserve UDP 8.0 with Active Directory Using Active Directory Groups

How to Integrate Arcserve UDP 8.0 with Active Directory Using Windows Groups

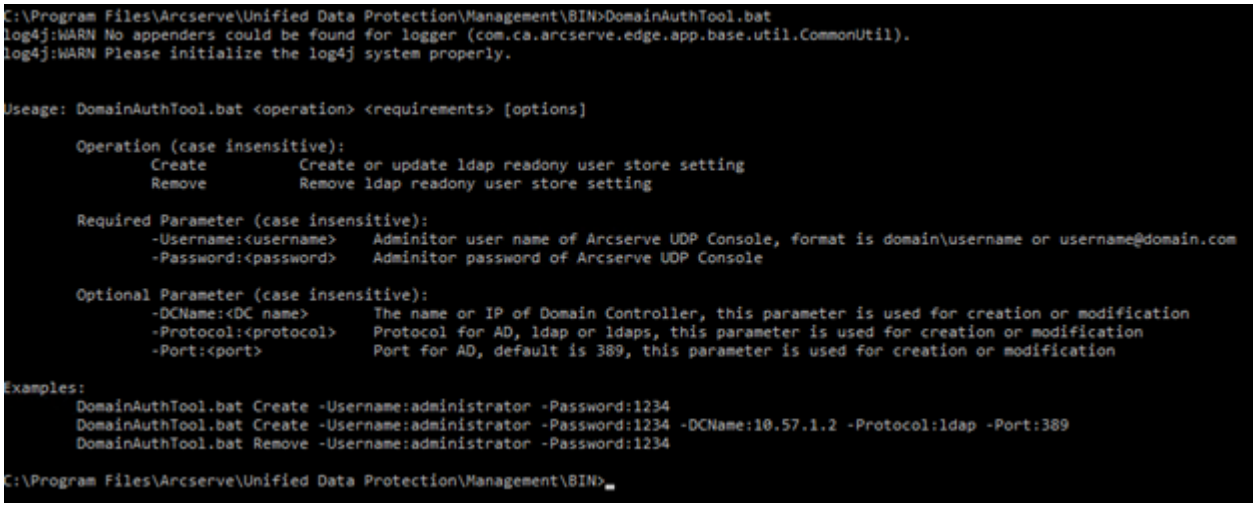
The existing feature of Role Based Administration (UDP-RBA) allows management of UDP permissions based on Active Directory (AD) groups. In the previous version of UDP-RBA, only individual AD user accounts were supported. Now, AD groups can serve as RBA roles.

The LDAP Read-only secondary user store is automatically added by running a utility.

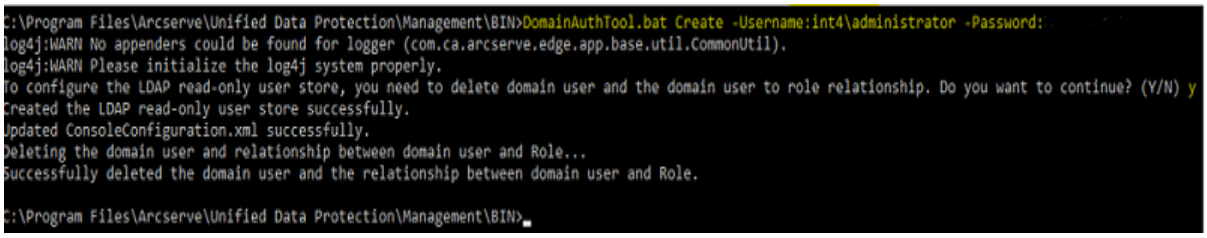
Follow these steps:

1. Launch the Command Prompt and run *C:\Program Files\Arcserve\Unified Data Protection\Management\BIN*.
2. Run the *DomainAuthTool.bat* utility.

   The following information appears on the screen:



3. Create an LDAP read-only secondary user store and at the same time modify the optional parameters. For example: Refer to the yellow marked command in the screenshot below.

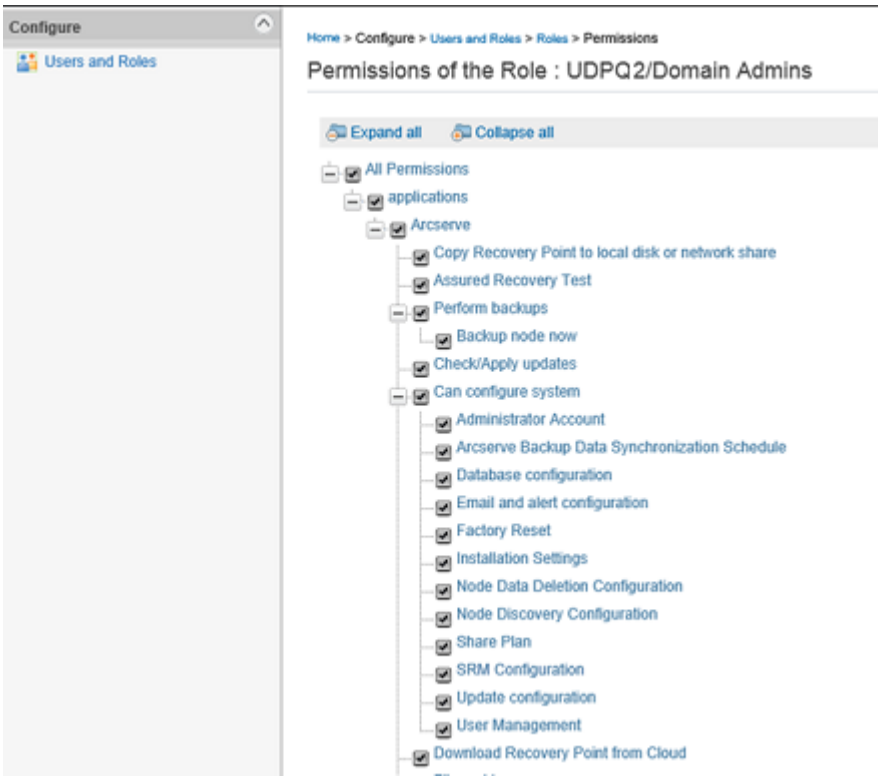


4. Assign the local user role.

For more information, see Assign a Pre-defined Role.

5. To add permission for the LDAP read-only domain users, perform the following steps:
   a. Join the group from Domain Controller.
   b. Add Arcserve UDP permission for the corresponding group as displayed in the screenshot below.

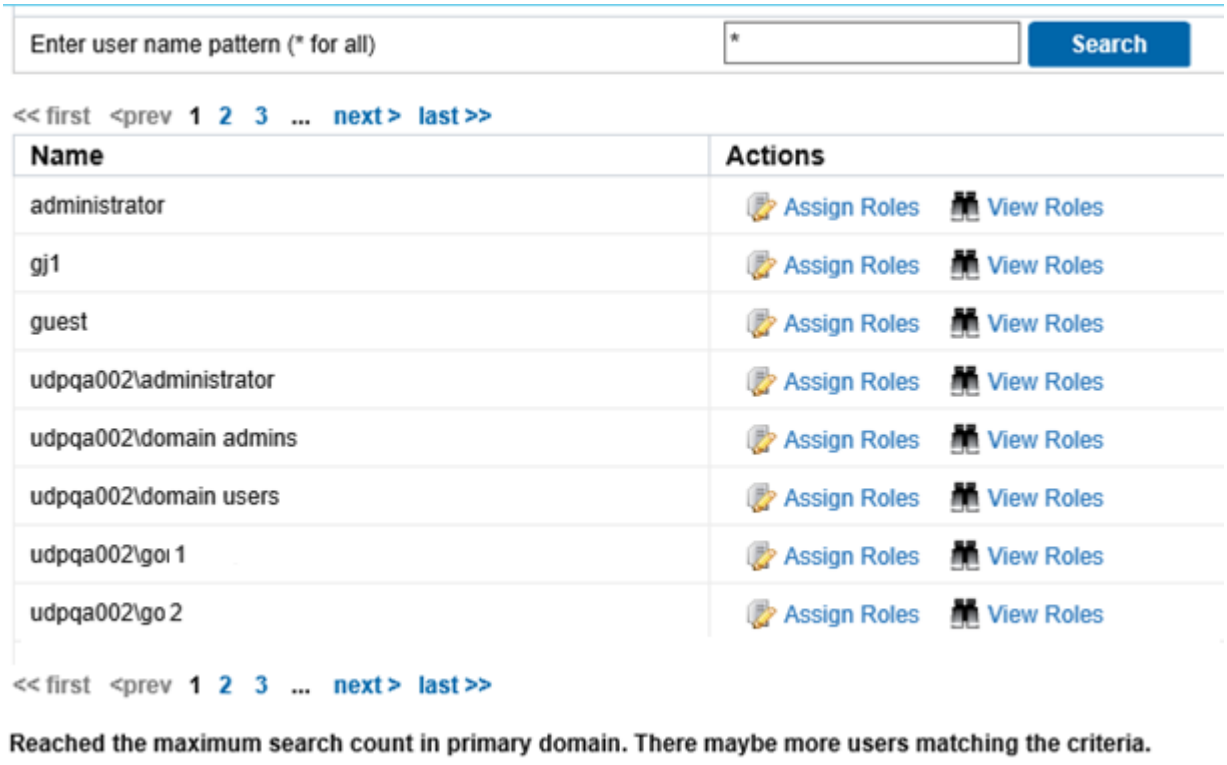
   Note: Assigning the local default role directly to the read-only domain users is not possible.



   The user can now log on to Arcserve UDP console with specific permission.

6. (Optional) Remove the LDAP user store with utility.

   Run the utility to remove the LDAP user store. Then, restart the management service and the previous domain user is listed as displayed below.


Enter user name pattern (* for all)
*
Search
<< first <prev 1 2 3 ... next > last >>
Name
Actions
administrator
Assign Roles
View Roles
gj1
Assign Roles
View Roles
guest
Assign Roles
View Roles
udpqa002\administrator
Assign Roles
View Roles
udpqa002\domain admins
Assign Roles
View Roles
udpqa002\domain users
Assign Roles
View Roles
udpqa002\go 1
Assign Roles
View Roles
udpqa002\go 2
Assign Roles
View Roles
<< first <prev 1 2 3 ... next > last >>
Reached the maximum search count in primary domain. There maybe more users matching the criteria.

How to Integrate Arcserve UDP 8.0 with Active Directory Using Active Directory Groups

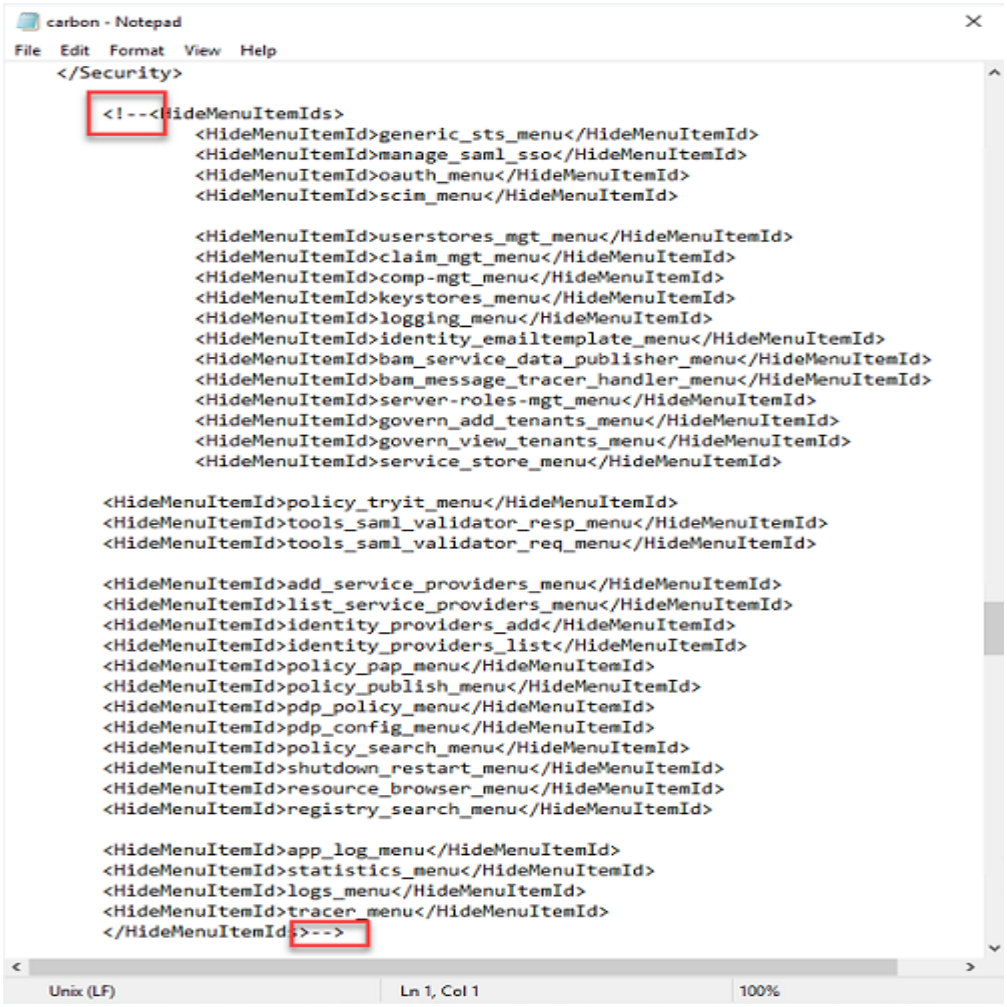
The role-based administration of Arcserve UDP 8.0 allows user-level permission where the Active Directory (AD) feature is not enabled by default. However, the WSO2 Carbon platform in Arcserve UDP 8.0 does not support AD groups that have secondary user store. You can enable the extension for Arcserve UDP 8.0 that configures the AD groups as Arcserve UDP roles and helps assign the permissions automatically for the members in the AD group.

Follow these steps:

1. Navigate to the following installation path of Arcserve UDP and open the *carbon.xml* file:

   …\Program Files\Arcserve\Unified Data Protection\Management\IdentityServer\repository\conf\carbon.xml

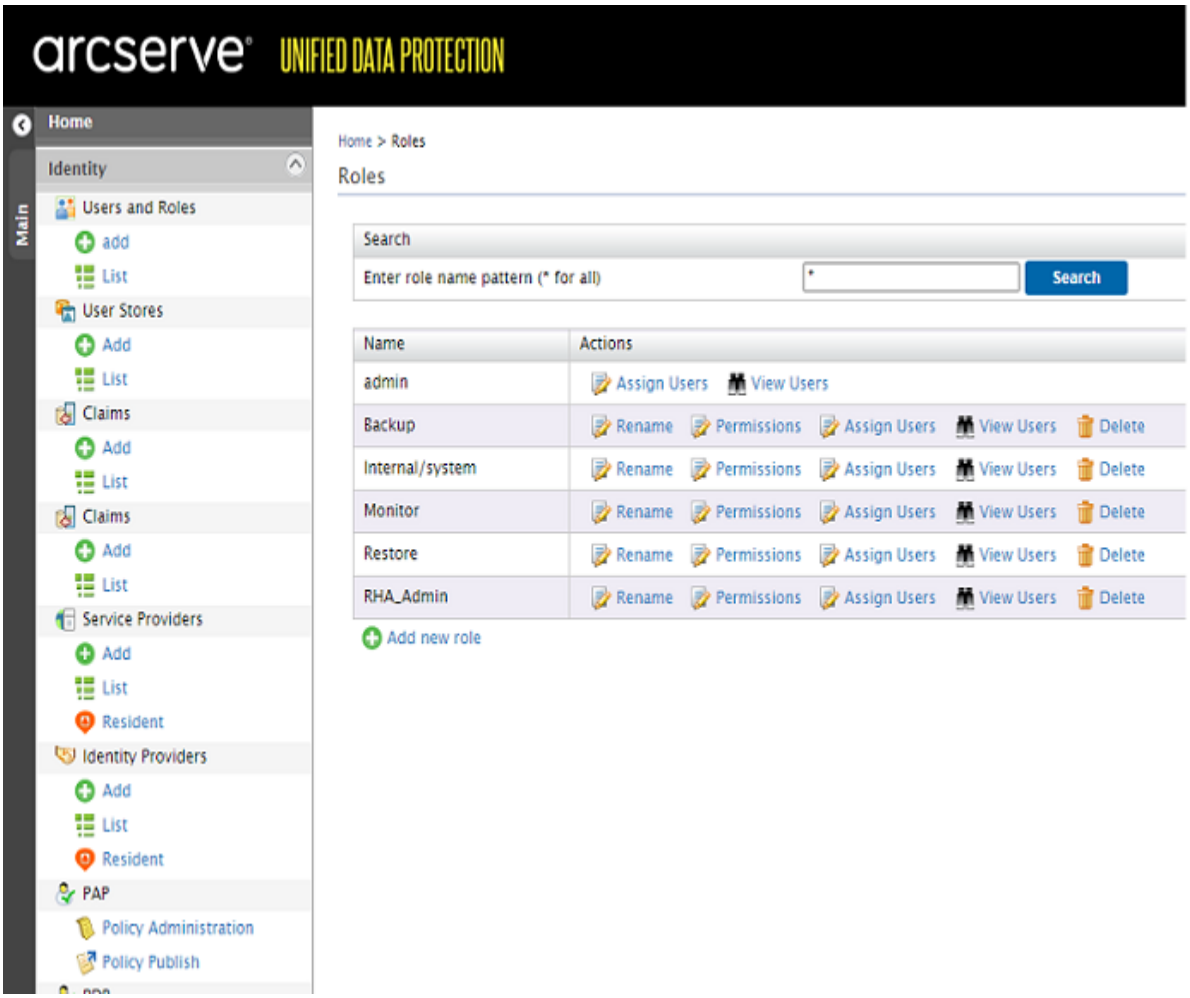
2. From the *carbon.xml* file, disable the contents of HideMenuItemIds using <!-- and --> as displayed in the screenshot below.



3. Save the *carbon.xml* file and restart the Arcserve UDP Management service.
4. Open the user management console using the following link:

   https://localhost:8015/carbon

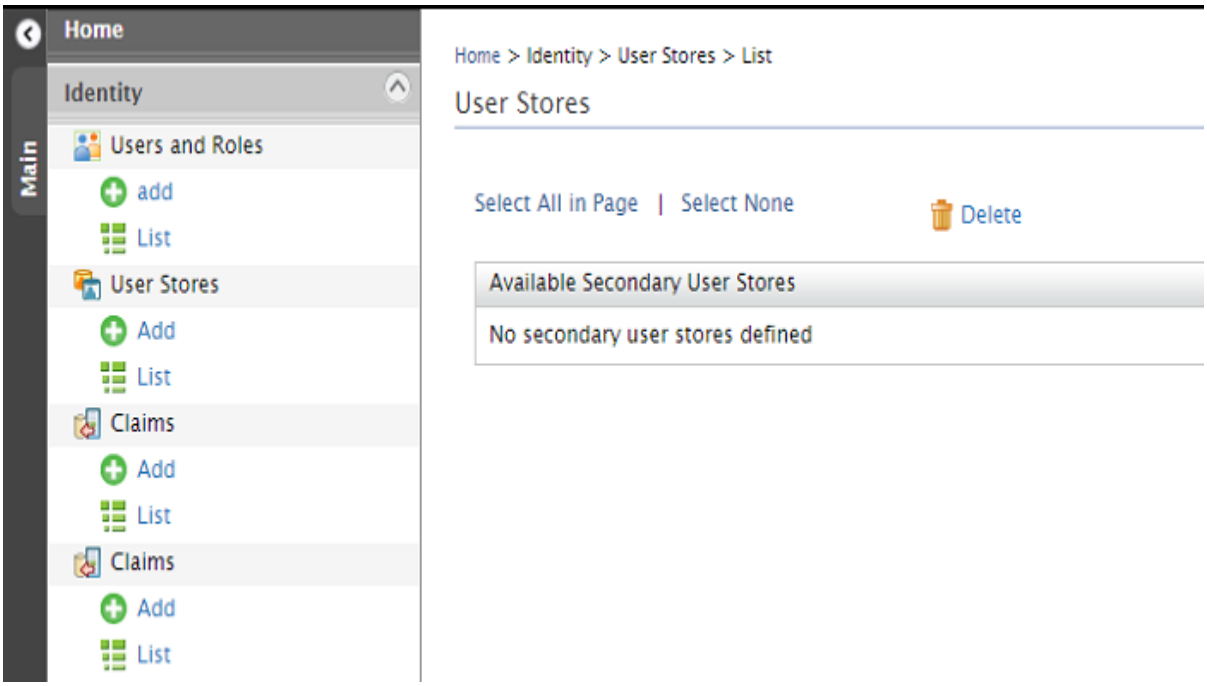
   The *Arcserve UDP Role-based Access Control Administration Home* page appears.



5. Click the **User Store Management** option available on the left pane.

   The **User Store Management** page appears.

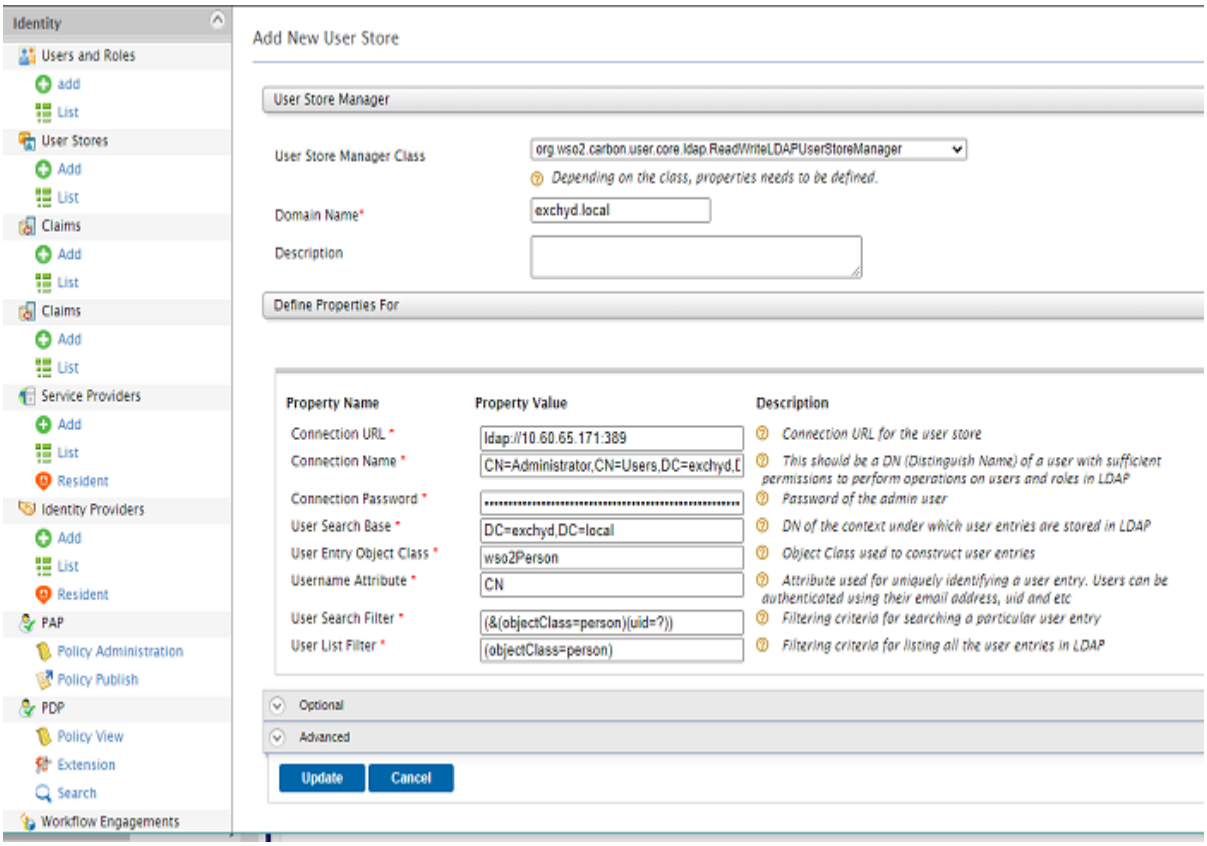
6. Click **Add Secondary User Store**.



The **User Store Manager** page appears.

7. Select the required **User Store Manager Class** option from the drop-down list and enter your domain name in the **Domain Name** field.
8. Enter the details in the fields as required under **Define Properties For** and **Optional** groups.

   The screenshot below is an example of the User Store Manager page after entering the details.

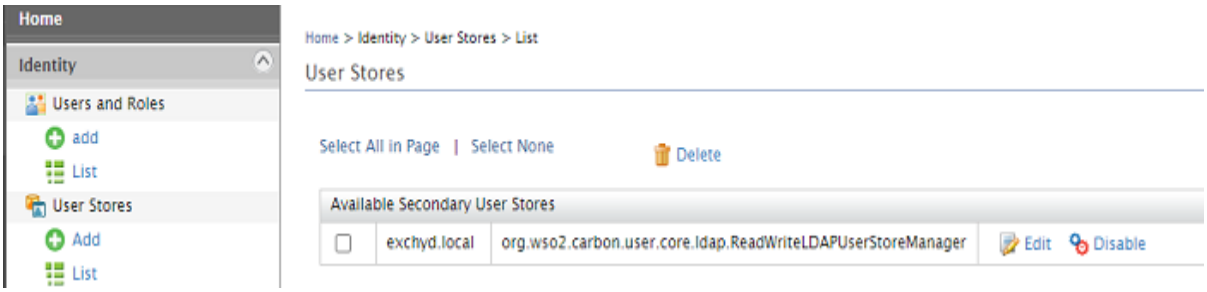


9. Click **Add**.

   The **UDP User Management** dialog appears.

10. Click **OK**.

    The User Store Management page appears and displays the added secondary user store.



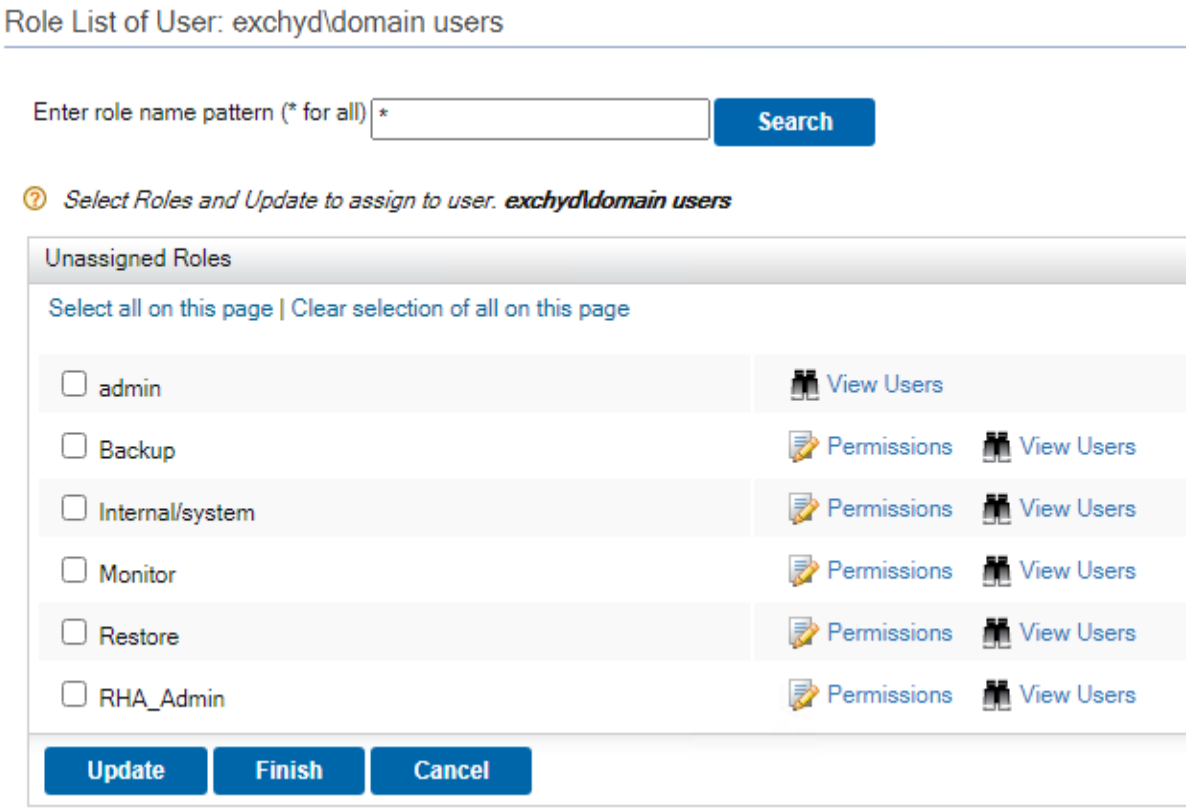
Note: If the secondary user store is not displayed, refresh the browser.

11. (Optional) Click the **Users and Roles** option from the left pane to view the list of users and roles.

    Note: You need to define roles in the domain using AD.

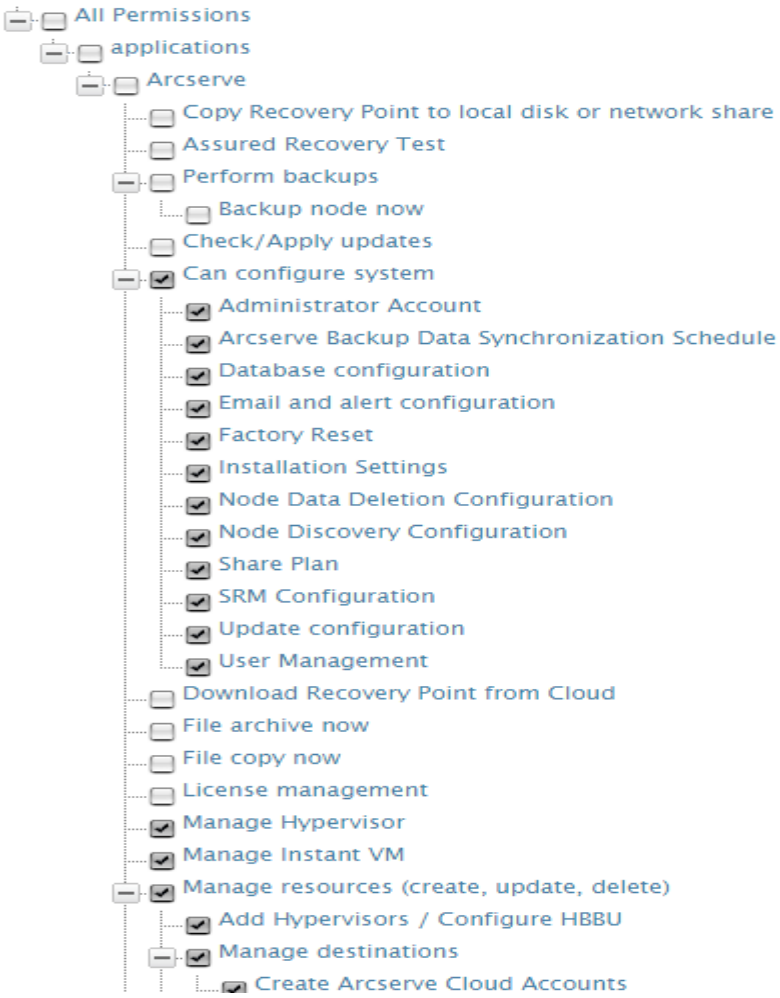
12. Now, perform the following steps to add the UDP role permissions:
    a. Select an AD user or AD group.
    b. Assign a role from the available list of roles.
    c. Click **View Role**.

       The **Role List of User** page appears.



    d. Click **Permissions**.

       The list of permissions appear.



e. Select the Permissions as required.

Now, the secondary user can log into the Arcserve UDP Console with the assigned permissions.

Access Arcserve UDP Using Integrated Windows authentication

Arcserve UDP users can now also login using Integrated Windows Authentication (IWA). IWA facilitates browser-based login. If authenticated once, IWA allows web browser to save the credentials of the user logged in using Windows. You would only need to enter the Protected web application's URL in the browser. The browser and server authenticate and automatically log in the user.

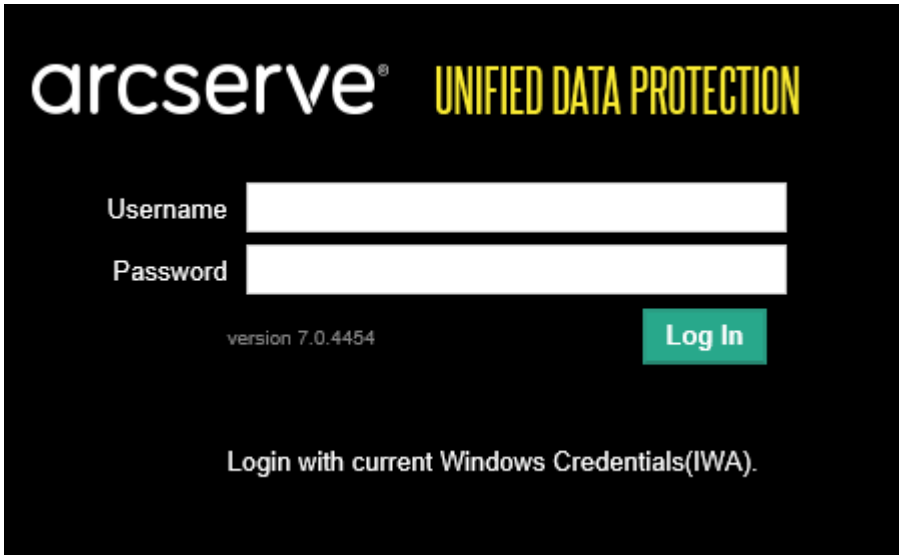
Note: Arcserve UDP Console does not support IWA if Console Database Connection uses Windows Authentication Mode. For a workaround, view troubleshooting link.

With introduction of IWA, you can use the same feature and access Arcserve UDP.

Follow these steps:

1. Open the Arcserve UDP Console login page.

   The IWA link is visible on the login page.



2. Click the link: **Login with current Windows Credentials (IWA)**.

   You are led to the Home page.

   To resolve any issue, see the Troubleshooting.

Troubleshooting for Integrated windows authentication (IWA)

If the IWA link does not open the Console page, you can use the following troubleshooting steps:

1. Verify if you are using IWA in your local machine instead of remote machine.
2. If you see a white screen after clicking the IWA link, verify if the user is assigned the role in your RBAC Administration Console.

   Note: If Console Database Connection uses Windows Authentication Mode, refer the link.
3. Configure using the steps given below if you encounter any issue during IWA login using Firefox to open the console URL.

   Follow these steps:

   a. In the browser's Location field, enter about:config.

   b. Click **I'll be careful, I promise!** to continue to the about:config page.

   c. Set values for the following options so that the browser trusts the ProxySG appliance and negotiates authentication:

   network.automatic-ntlm-auth.trusted-uris, network.negotiate.auth.delegation-uris, network.negotiate-auth.trusted-uris

   For each option, complete the following steps:

   - Locate the option that you want to set by scrolling or entering the option name in the Filter field.
   - Double-click the option to open the Enter string value dialog.
   - Enter the virtual URL (for transparent deployments).

     If you have more than one ProxySG appliance that will challenge for authentication credentials, separate the entries with commas. For example, if your opened URL for console is https://-localhost:8015; you can enter localhost as the string value, Or https://10.57.60.9:8015, then enter 10.57.60.9 as the string value.
   - Click **OK**.

4. While using Internet Explorer to open Console with IWA on the remote machine whose Windows credential is exactly same to the local machine that UDP installed, you may need to set using the following steps:
    a. Select Tools > Internet Options.
    b. Select the Security tab.
    c. Select the Local intranet zone and click Sites > **Advanced**.
    d. Enter the fully qualified domain name of the ProxySG appliance (for explicit deployments) or the virtual URL (for transparent deployments) in the Add this website to the zone field and then click Add > Close > OK.
    e. Select the Advanced tab and make sure the Security > Enable Integrated Windows Authentication option is selected.
    f. Click **OK** to save your changes and close the Internet Options dialog.

Arcserve UDP Console does not support IWA if Console Database Connection uses Windows Authentication Mode

To use Integrated Windows Authentication (IWA), you need to change the authentication mode of Console Database Configuration to SQL Server and Windows Authentication Mode.

Follow these steps:

1. From the SQL Server Management tool, verify if Console Database supports SQL Server and Windows Authentication Mode. If Console Database does not support, switch to SQL Server and Windows Authentication mode, and restart SQL Server.
2. Perform the following steps to create Arcserve UDP account in SQL Server:
    a. From the following folder, right-click **DBAccountUpdate.bat**, and click the **Run as administrator** option:

    <UDP Installation Folder>\Management\BIN

    The command interface opens.

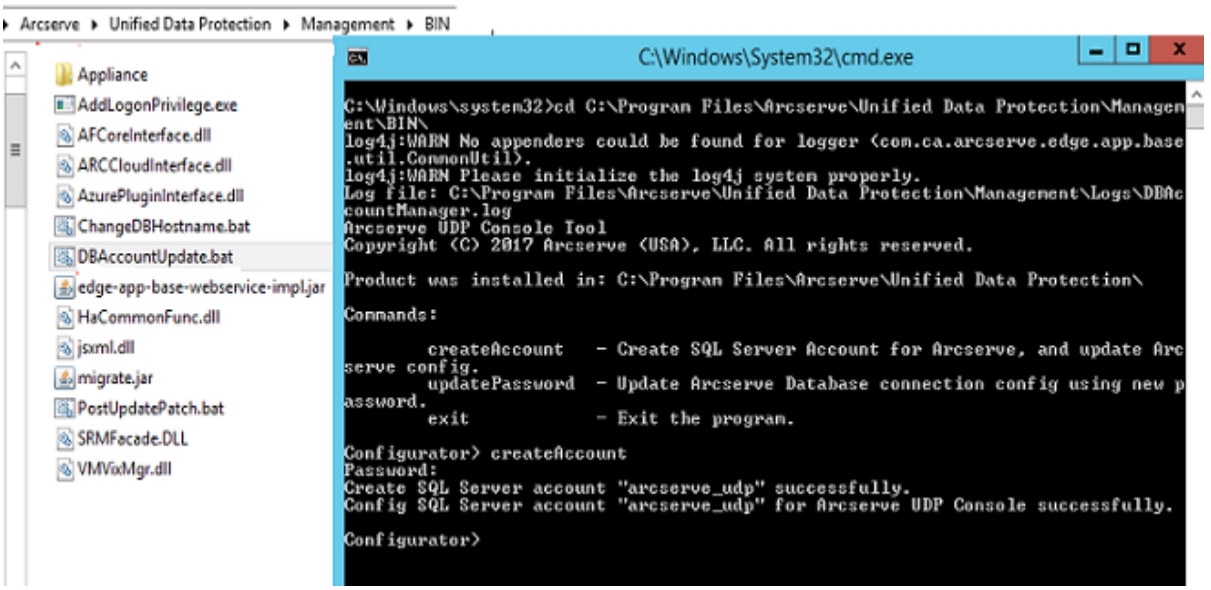
    b. In the command interface, type *createAccount* and press Enter.

    The command line interface prompts you to provide a Password.

    c. Specify the desired password for the account *arcserve_udp* and press Enter.

    Arcserve UDP Console creates SQL Server account *arcserve_udp*.

    d. Type *exit* to close the command line interface.



3. Restart the Arcserve UDP Management Service.

Troubleshooting

The following list provides the possible resolution for the errors that may appear with the User Management Console:

- **Symptom**

  Failed to log into the User Management Console

  Solution

  Verify that you have logged in as an administrator. Non-administrator users do not have permission to access the User Management Console. Verify that the username and password are correct.

- **Symptom**

  Authentication failure: User fails to log in to the Console

  Solution

  Verify whether the user is assigned any role. If not assigned any role, a user can-not log into the Arcserve UDP Console.

- **Symptom**

  User Management Console page timed out.

  Solution

  The login retaining time for the User Management page is 15 minutes. If the Console does not detect any operation on the page for 15 minutes, then the user is automatically logged out.

Chapter 8: Adding and Managing Source Nodes

This section contains the following topics:

How to Add Nodes to the Console

A node refers to a physical or virtual source machine on hypervisors that you want to protect. You can protect a node by backing up data to a destination. Arcserve Unified Data Protection lets you add the following types of nodes:

- Windows
- Linux
- Virtual machines in VMware ESX/vCenter and Microsoft Hyper-V servers

You can add nodes by manually specifying the node details, discovering from an active directory, or importing from a file and hypervisors.

Note: You can also add nodes while creating a plan.

What To Do Next?

- Review the Prerequisites
- Add Nodes
- Discover Nodes
- Import Nodes

Review the Prerequisites

Before you start adding a node, complete the following prerequisite tasks:

1. Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.
2. Log into the Console.
3. Click the **resources** tab.

   The **Nodes: All Nodes** page is displayed.
4. From the center pane, click **Add Nodes**.

   The **Add Nodes to Arcserve UDP Console** dialog opens.

   The dialog provides multiple options to add a node.

Add Nodes

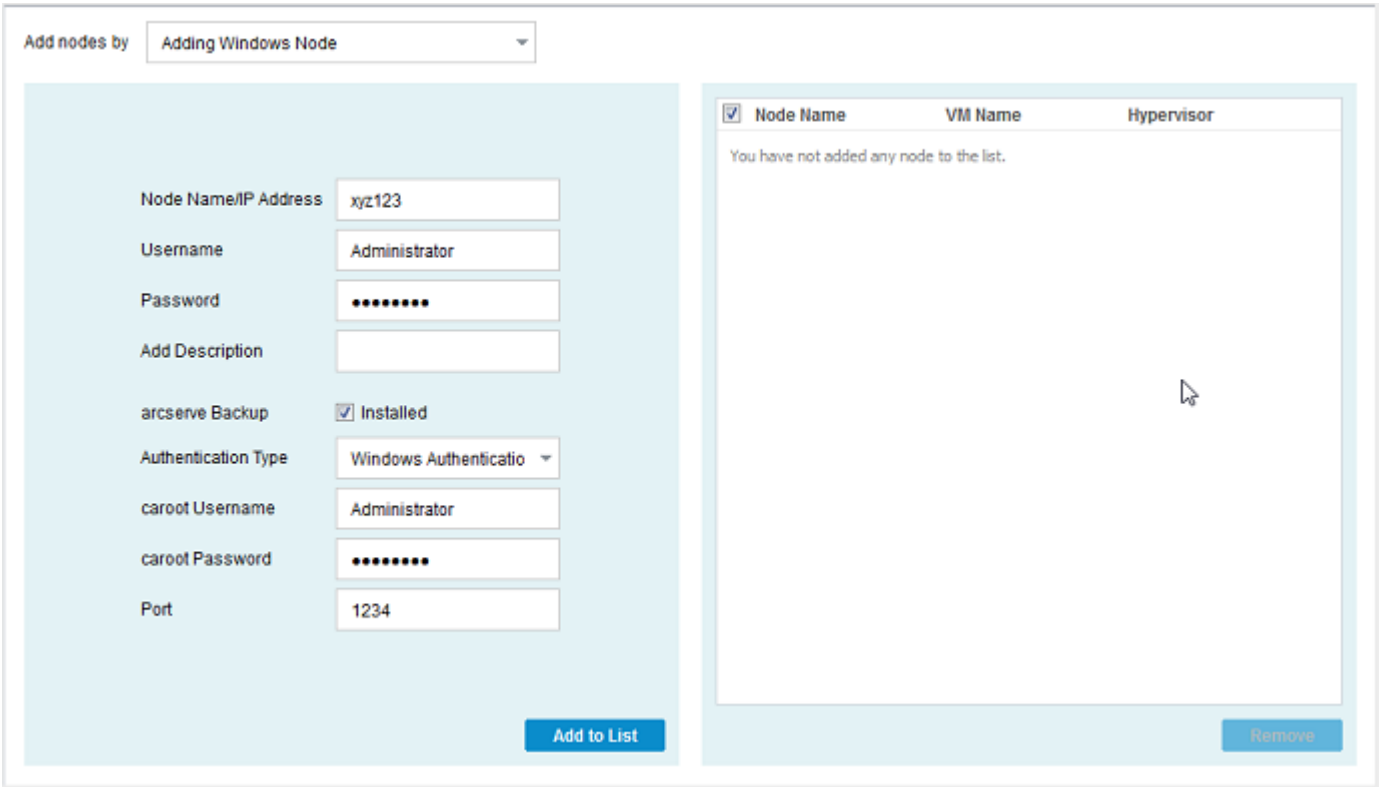
When you have the IP address or name of a node or set of nodes, you can add them to the Console by specifying their details manually. You can add the following types of nodes:

- **Windows:** Windows source nodes that you want to protect. When the Windows source node with SMB port is blocked, Arcserve UDP Agent(Windows) gets installed on this node.
- **Linux:** Linux source nodes that you want to protect. Arcserve Unified Data Protection Agent for Linux is installed on the Linux Backup Server and not on the Linux source nodes.
- **Linux Backup Server:** Linux server that manages Linux source nodes. Arcserve Unified Data Protection Agent for Linux is installed on this server.

Follow these steps:

1. From the **Add nodes by** drop-down list, select one of the following options:

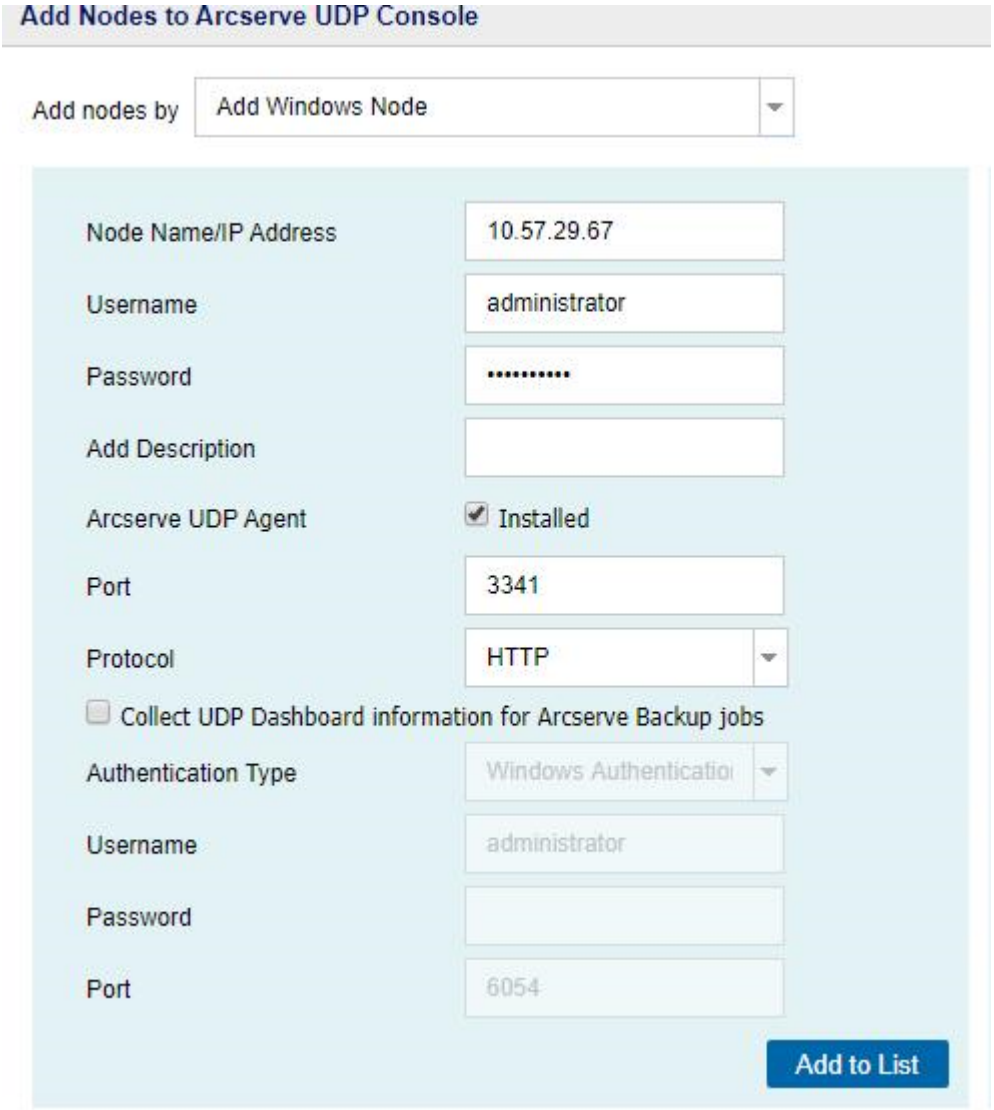
   Adding Windows Node



   Note: To enable the details for Arcserve Backup, select **Installed**.

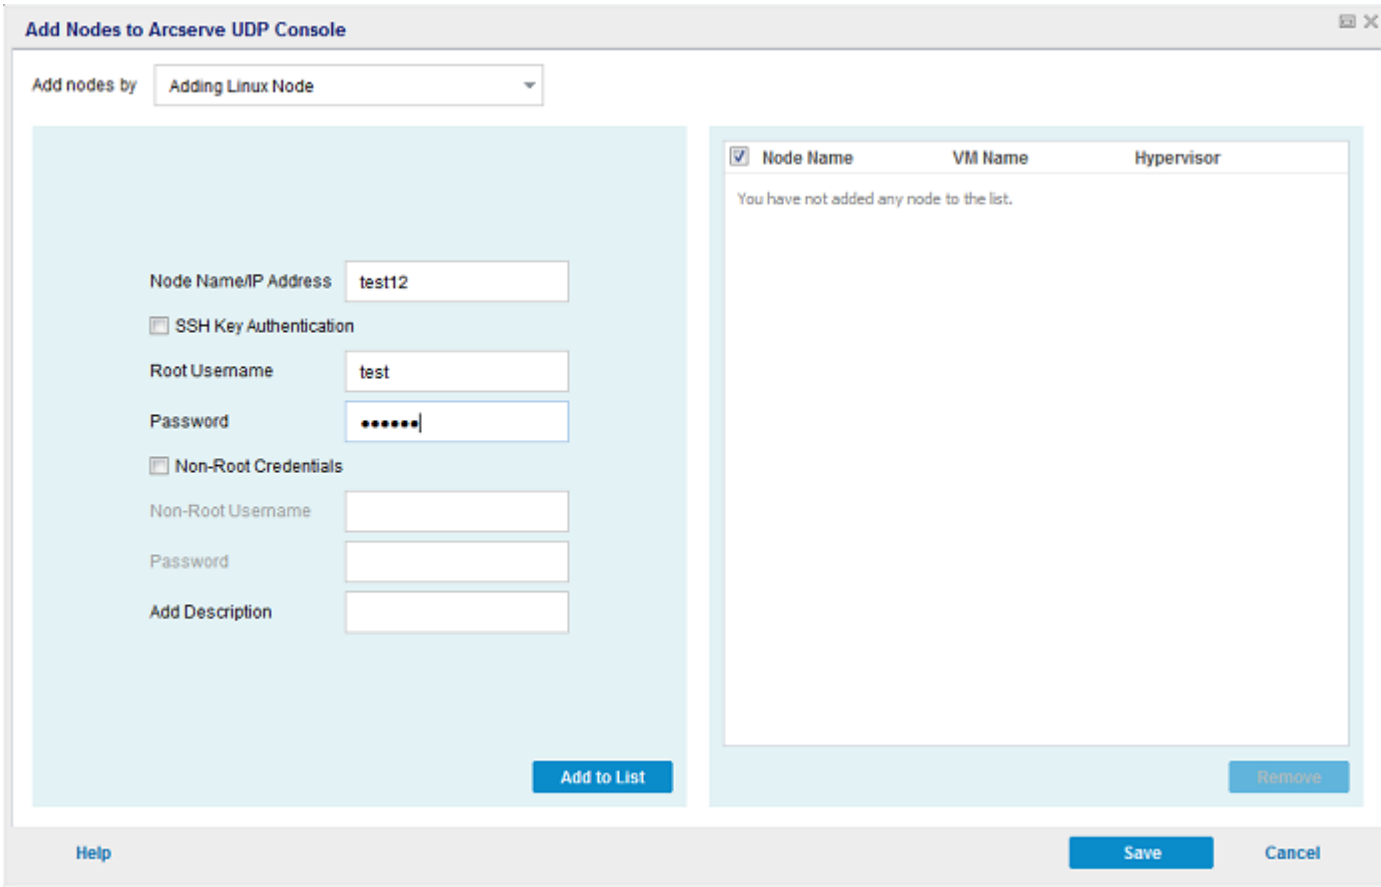
2. To add SMB port blocked node, go to the **Add nodes by** drop-down list, and select one of the following options:

Adding Windows Node



Note: Provide the customized port and protocol details, and then select installed.

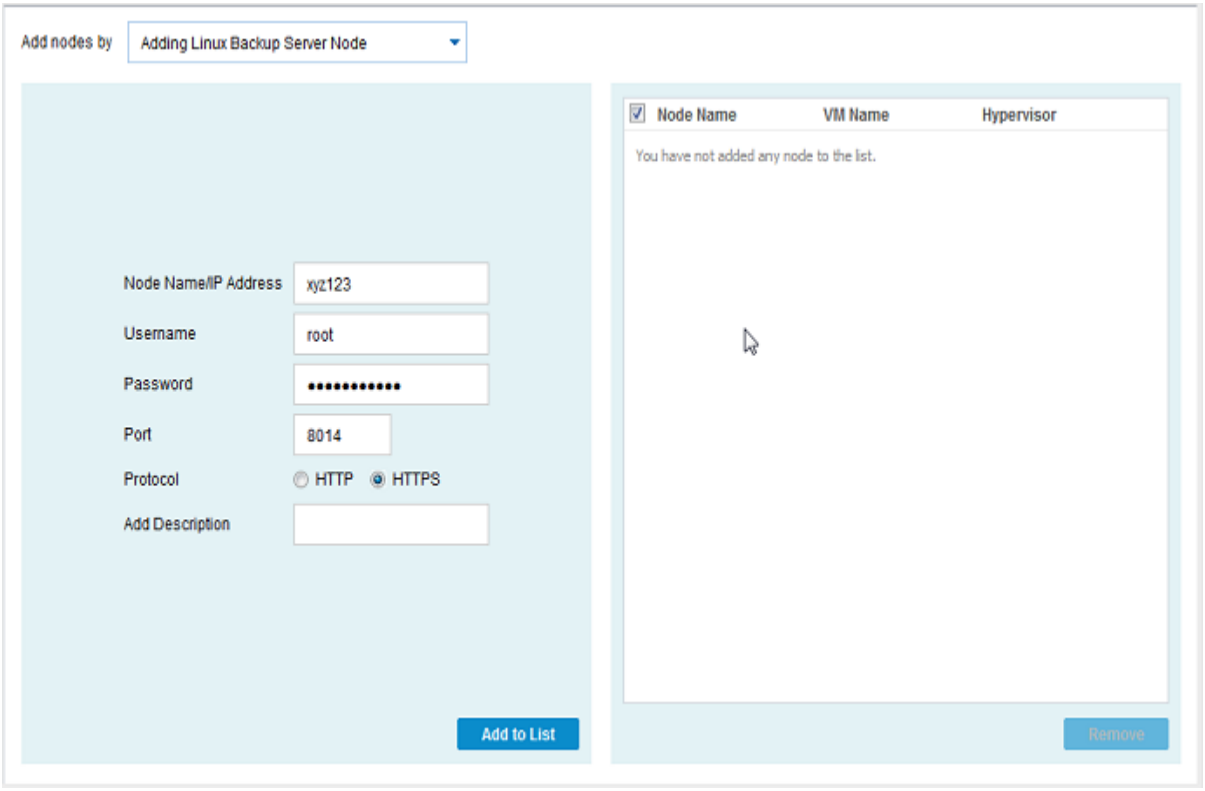
Adding Linux Node



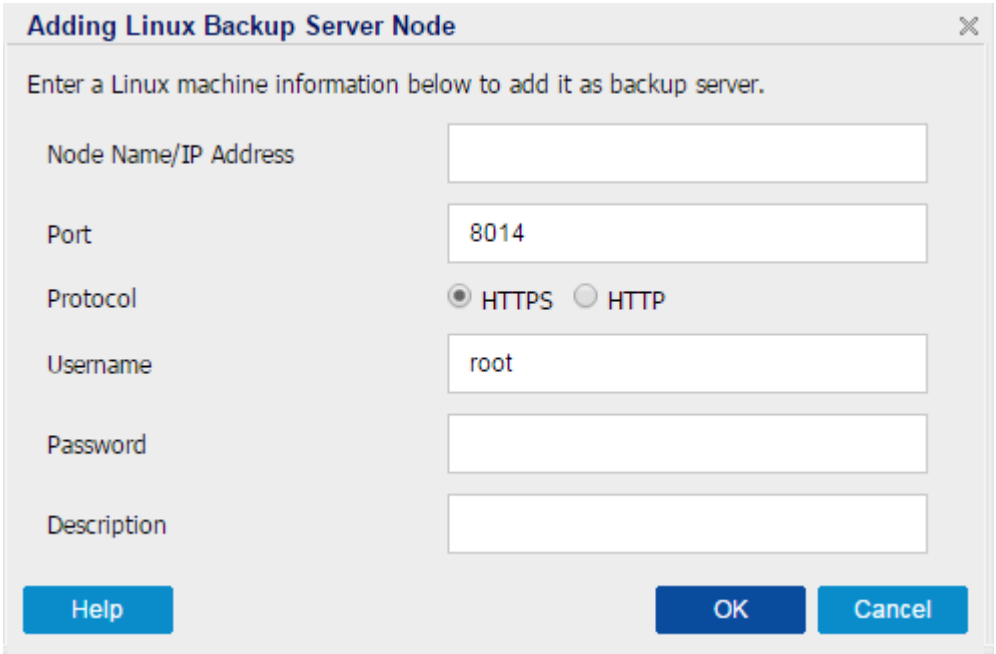
Notes:

- For Linux, on selecting **SSH Key Authentication**, you do not need to enter password details but the username is required. For more information about configuring the SSH Key, see Configure the Private Key and Public Key Authentication.
- Before adding a Linux node, you must add a **Linux Backup server** that manages the Linux nodes.
- For Debian Linux node, non-root credential is not supported. To add the Debian Linux with non-root user, use sudoers. For more information about configuring Sudo, see Configure sudo in Debian.
- You can log into the Linux Backup Server from the Arcserve UDP Console only when you perform a restore.

Adding Linux Backup Server Node



The following dialog appears when you add Linux Backup Server node from the **Backup: Agent-Based Linux** task:



The details of the selected option are displayed.

3. Enter the details of the node and click **Add to List**.

The node is added to the right pane. To add more nodes, follow the steps again. All the added nodes are listed on the right pane.

4. (Optional) To remove the added nodes from the list on the right pane, select the nodes and click **Remove**.
5. Select the nodes to add and click **Save**.

   The nodes are added and displayed at the **Nodes: All Nodes** page.

   If your Linux Backup server is in the NAT environment, perform the following steps before adding to the Arcserve UDP console:

   1. Create a file *server.cfg* on the following Linux backup server folder:

      /opt/Arcserve/d2dserver/configfiles/.

   2. Add the following line to the server.cfg file:

      nat_enable=true

   3. Restart Linux Agent using the following command line:

      /opt/Arcserve/d2dserver/bin/d2dserver restart

   Note: If this Linux backup server is already added, update the Linux Backup server on the UDP console.

Discover Nodes

To add nodes that are in an active directory, you can first discover the nodes by providing the active directory details and then adding the nodes to the Console.

Follow these steps:

1. From the **Add nodes by** drop-down list, select **Discovering Nodes from Active Directory**.
2. Specify the user credentials and click **Add.**

   Username

   Specifies the domain and user name in the domain\username format.

   Password

   Specifies the user password.

   Computer Name Filter

   Specifies the filter to discover node names.

   After validation, the user name is added to the list.
3. Select the added user name and click **Browse**.

   A successful node discovery opens the **Confirm** dialog that prompts you to add the nodes from the **Discovery** result.

   Note: The discovery process may take a while depending upon the factors such as the network and number of computers in the network.
4. Click **Yes**.

   The discovered nodes are listed.
5. Select the node, enter the user name and password, and then click **Apply**.

   Note: When you click Apply, the credentials are verified. You must verify each node before you add to the list.

   The green check marks are displayed for the verified nodes.
6. Click **Add to List**.

   The selected node is listed on the right pane.
7. To add the nodes to the Console, from the right pane select the node and click **Save**. To add all the nodes, select the **Node Name** check box.

   The verified nodes are added and available at the **Nodes: All Nodes** page.

Troubleshooting: The Specified Domain Either Does not Exist or Could not be Contacted

Symptom

When adding nodes by discovering from an Active Directory, I get the following error message:

"The specified domain either does not exist or could not be contacted. Verify that the Console server can access the domain controller through the network."

Solution

First verify the connectivity between the Arcserve UDP and domain controller. If the connectivity is okay, use the following command with "dsgetdc" argument to test if Windows can locate the domain controller from the domain name:

nltest.exe

For example, "nltest /dsgetdc:sample_domain", where sample_domain is the domain name.

If the command fails, then there may be a DNS problem in your environment.

Note: You should run the above command on the UDP machine. For more details, refer to the article from Microsoft.

Import Nodes

Arcserve Unified Data Protection lets you add multiple physical and virtual nodes by using the import method. Depending on the requirement, you can use one of the following import methods:

- Import Nodes from a CSV or TXT File
- Import Nodes from a vCenter/ESX Server
- Import Nodes from a Hyper-V Server

Import Nodes from a vCenter/ESX Server

Using this import method, you can import virtual machine nodes from the ESX or vCenter server. This option lists all the virtual machines that are detected on the specified server, even if they are already being managed in Arcserve Unified Data Protection.

Follow these steps:

1. From the **Add nodes by** drop-down list, select **Importing from vCenter/ESX**.
2. Specify the vCenter/ESX server details, and click **Connect**.

   In the left pane, a node tree is displayed.

   Note: VMware Virtual Disk Development Kit (VDDK) 6.x.x is bundled with Arcserve UDP Version 7.0, but VDDK 6.x.x does not support HTTP. Also, vCenter and ESX typically support HTTPS connection only by default. Select HTTPS, unless you manually replace the built-in VDDK 6.x.x by another version VDDK, and you manually configure vCenter/ESX to allow HTTP connection.
3. Expand the node tree.

   (Optional) You can type the node name in the filter field to locate the node in the tree.
4. Select the nodes that you want to add.

   Note: Arcserve UDP lets you perform the following options:

   - Add and protect container objects in vSphere infrastructure (such as data center and resource pool).
   - Add and protect VM template and protect VM by tag.

   But, currently you can perform this action only using the Plan wizard. For details, refer to *Add Nodes from vCenter/ESX* in the Specify the Source topic.
5. Select **Provide credentials for the selected nodes** check box and provide the user credentials.

   Note: User credentials are required for functions such as Pre-Flight Check (PFC), Application Log Truncation, Pre/Post Backup Commands.
6. Click **Add to List**.

   The selected nodes are added to the right pane.
7. Select the nodes and click **Save**.

   The nodes are added and displayed on the **Nodes: All Nodes** page.

Import Nodes from a Hyper-V Server

Using this import method you can import the virtual machine nodes from the Microsoft Hyper-V servers.

Follow these steps:

1. From the **Add nodes by** drop-down list, select **Importing from Hyper-V**.
2. Complete the following fields, and click **Connect**.

   Hyper-V

   Specifies the Hyper-V server name or the IP address. To import virtual machines that are in Hyper-V clusters, specify either the cluster node name or Hyper-V host name.

   Username

   Specifies Hyper-V user name having the administrator rights.

   Note: For Hyper-V clusters, use a domain account with administrative privilege of the cluster. For standalone Hyper-V hosts, we recommend using a domain account.

   Password

   Specifies the password of user name.

   The Arcserve Unified Data Protection solution searches and displays a node tree on the left pane.
3. Expand the node tree.

   (Optional) You can type the node name in the filter field to locate the node in the tree.

   Note: The virtual machines configured as cluster role are listed directly under the cluster node name on the tree. The virtual machines that are not part of the cluster are listed under the host name of individual Hyper-V host.
4. Select the nodes that you want to add.
5. Select **Provide credentials for the selected nodes** check box and provide the user credentials.

   Note: User credentials are required for functions such as Pre-Flight Check (PFC), Application Log Truncation, Pre/Post Backup Commands. Without user credentials, PFC fails for the selected nodes.
6. Click **Add to List**.

   The selected nodes are added to the right pane.

7. Select the nodes and click **Save**.

   The nodes are added and displayed on the **Nodes: All Nodes** page.

Import Virtual Machine Using Additional Administrative Account

Additional administrative account refers to those accounts that are not default administrators. Such accounts are also referred as non-built-in administrative accounts. To import virtual machine from a Hyper-V host, either use the built-in administrator account of the Hyper-V host or a domain account that is in the local administrators group of the Hyper-V host, or a non-built-in administrative user.

The user with additional administrative account can use the procedures explained below to disable UAC remote access.

Notes:

- This procedure is not similar to disabling UAC. Using this procedure you can disable some of the functions of UAC.
- Considering that remote Windows Management Instrumentation (WMI) technology is used for import, verify that WMI is not blocked by firewall.

Follow these steps:

1. Click Start, type regedit in the Search programs and files field, and then press Enter.

   The Windows Registry Editor opens.

   Note: You may need to provide administrative credentials to open Windows Registry Editor.

2. Locate and click the following registry key:

   HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System

3. From the Edit menu, click New and then click DWORD (32-bit) Value.
4. Specify LocalAccountTokenFilterPolicy as the name for the new entry and then press Enter.
5. Right-click LocalAccountTokenFilterPolicy and then click Modify.
6. Specify 1 in the Value data field and then click OK.
7. Exit the Registry Editor.

   For more information about Windows behavior, see Microsoft documentation.

Import Nodes from CSV or TXT File

When you have multiple physical nodes to add, instead of adding one node at a time, you can create a .csv or .txt file using the <NodeName>, <UserName>, <Password> format. Subsequently, you can select and import the .txt or .csv file to the Console using the browse and upload options.

Follow these steps:

1. From the **Actions** drop-down list, select **Import.**
2. Click **Browse** to select the saved file in .txt or .csv format.
3. Click **Upload**.

   The nodes are added and displayed on the **Nodes: All Nodes** page.

How to Manage Nodes

Using Arcserve UDP, you can perform multiple actions to manage a node such as update node and hypervisor, export nodes, delete, and perform preflight checks.

What To Do Next?

- Review the Prerequisites
- Update Hypervisor Information
- Specify the Hypervisor
- Update VM Information
- Update Nodes
- Export Node
- Pause Node
- Resume Node
- Synchronize Data
- Delete Nodes from the Console
- Deploy Agent to Nodes
- Perform Preflight Checks for Your Backup Jobs

Review the Prerequisites

Before starting to manage the nodes, complete the following prerequisites:

- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.
- Log into the Console.
- Add a node.

Update Hypervisor Information

After a VM node is added into Arcserve Unified Data Protection, the connection-related information, such as hostname or credentials of hypervisor of VM may change. In such cases, Arcserve Unified Data Protection lets you update the hypervisor information.

Follow these steps:

1. Click the **resources** tab.
2. Right-click the node group under **vCenter/ESX Groups** or **Hyper-V Groups**.
3. Click **Update vCenter/ESX** or **Update Hyper-V.**

   The **Update vCenter/ESX** or **Update Hyper-V** dialog is displayed.
4. Enter the new details in the dialog box and click **OK**.

   The **Update vCenter/ESX** or **Update Hyper-V** dialog closes.

   The hypervisor information is successfully updated.

Specify the Hypervisor

Specify the hypervisor details to avoid using extra license while protecting a VM. When you protect a virtual machine (VM) using a host-based agentless backup plan, the hypervisor host license is used to protect the VM. You do not have to install any agent on the VM. In certain cases, you may decide to install the agent on the VM and create an agent-based backup plan to protect the VM. In such cases, the VM uses another license, other than the hypervisor host license. Specify the hypervisor details in such cases and the VM uses the hypervisor host license instead of using another license.

Examples describing when to specify the hypervisor information:

- You have a Host-Based Agentless Backup plan to protect the VMs of ESX or Hyper-V Server. The plan uses the Hypervisor license to protect the VM. Now, you install the UDP Agent in a VM of the specified Hypervisor and create an Agent-Based plan to protect the VM. Typically the plan uses extra license to protect the VM. If you specify the hypervisor for the VM, the plan uses the license of the Hypervisor.
- You have an Agent-based Linux Plan to protect the Linux VM Agent nodes. If you specify the hypervisor for the VM, all the VMs on the same Hypervisor share the Hypervisor license.

Consider the following points before specifying the hypervisor:

- Cannot specify the hypervisor for a physical node.
- Cannot specify the hypervisor for a VM node that is imported from vCenter/ESX or Hyper-V.
- Can specify the hypervisor for multiple VMs that belong to the same hypervisor, at the same time.
- Verify that the latest VMware tools or Hyper-V integration service is installed, and the VM is powered on. Also, verify the Windows Management Instrumentation (WMI) is in the exception list of the firewall on the VM agent node.

Follow these steps:

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes** and click **All Nodes**.

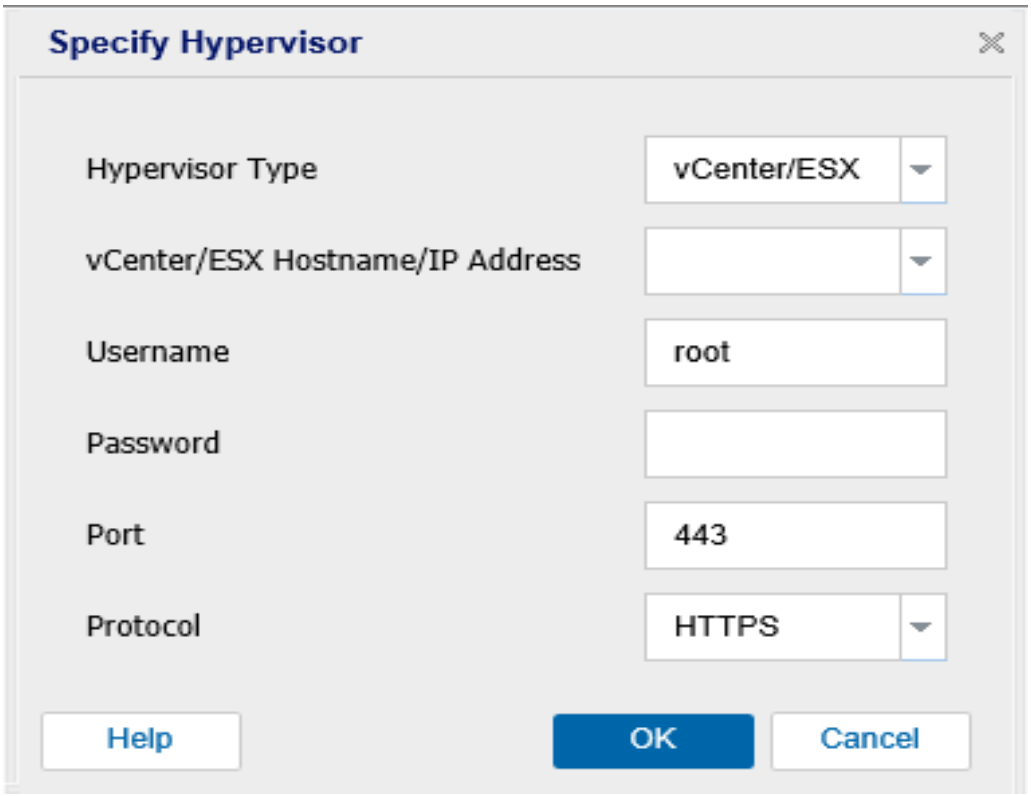
   The **Nodes: All Nodes** page is displayed.
3. Perform one of the following actions:

- Right-click the node name.
- Select the node name, and from the center pane click the **Actions** drop-down list.

A list of options is displayed.

4. Click **Specify Hypervisor**.

   The Specify Hypervisor dialog opens. The **Hypervisor Type** can be **Hyper-V**, **vCenter/ESX**, and **Other** (Xen, Kernel-based Virtual Machine, Red Hat Enterprise Virtualization).



5. Enter the hypervisor details and click **OK**.

   The hypervisor information is specified.

Update VM Information

Using Arcserve Unified Data Protection, you can update some of the properties of the VM nodes from their hypervisors. You can trigger the update manually or automatically. The following properties of the VM nodes are updated and synchronized with their corresponding VMs in the hypervisor:

- Node Name
- VM Name
- OS

To manually trigger the update, use the **Update VM Information** option.

Follow these steps:

1. Click the **resources** tab.

   The **Nodes: All Nodes** page is displayed.

2. From the center pane, click the **Actions** drop-down list, and then click **Update VM Information**.

   The **Update VM Information** dialog opens.

3. Click **OK**.

   You have triggered a manual discovery and that updates the virtual machine nodes.

   The auto update feature is triggered automatically when you perform the following actions:

   - Open the **resource** tab on the Console.
   - Send a scheduled report.

   Note: Even if you trigger multiple automatic updates, only one automatic update runs at a time. The remaining automatic updates are put in a queue.

Update Nodes

You can update information that is related to the existing nodes. You can update the node anytime. Some of the situations when you need to update a node are as follows:

- A new product is installed on the node after the node was registered with Arcserve Unified Data Protection.
- The user name or password of the node was updated after the node was registered with Arcserve Unified Data Protection.

Note: If a node acts as both recovery point server and agent, and you change the credentials or protocol of that node, then update the node from the **Destinations: Recovery Point Server** page. The plan will automatically deploy to the agent after you update the recovery point server. If you update the node from the **Nodes: All Nodes** page, then the plans involving those nodes are not deployed successfully. To deploy the plan, update the node from the **Destinations: Recovery Point Server** page again.

Follow these steps:

1. Click the **resources** tab.

   The **Nodes: All Nodes** page is displayed.

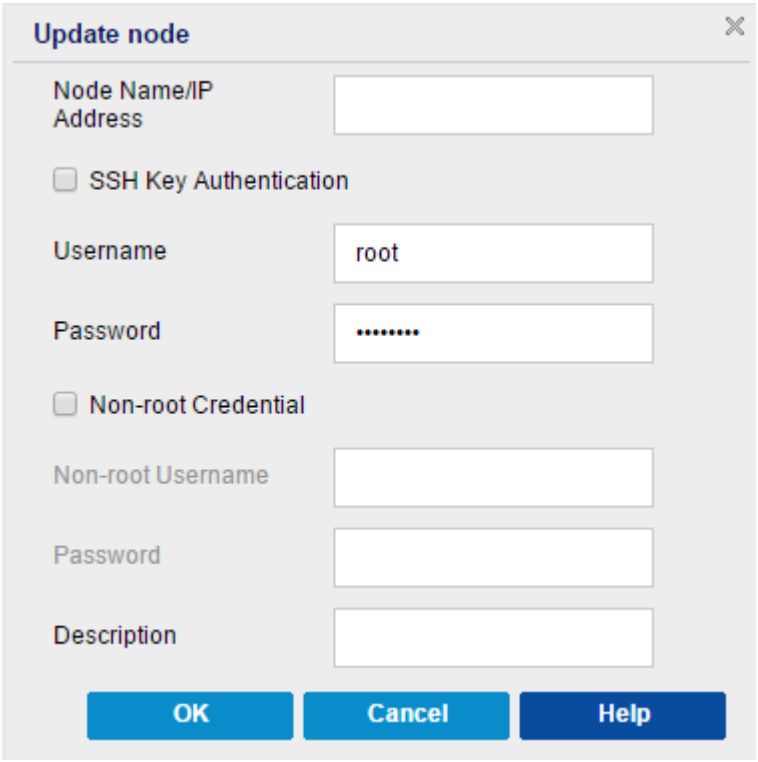
2. Perform one of the following actions:

   - Right-click the node name.
   - Select the node name, and from the center pane click the **Actions** drop-down list.

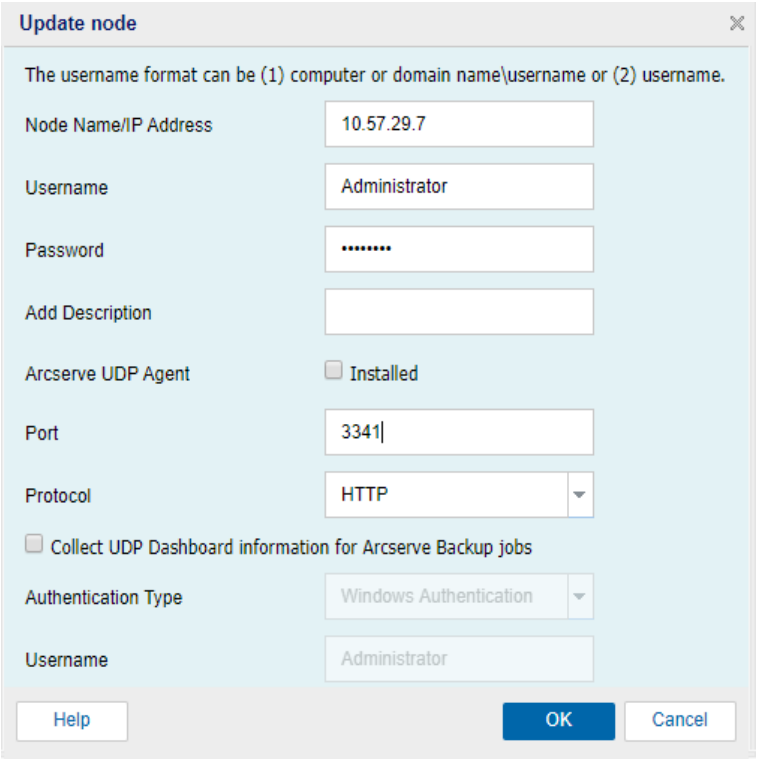
3. Click **Update**.

   The **Update node** dialog opens.

   The following dialog is for Linux nodes:



The following dialog for Windows node (When SMB port blocked) with customized port appears:



4. Update the details and click **OK**.

   The node information is updated.

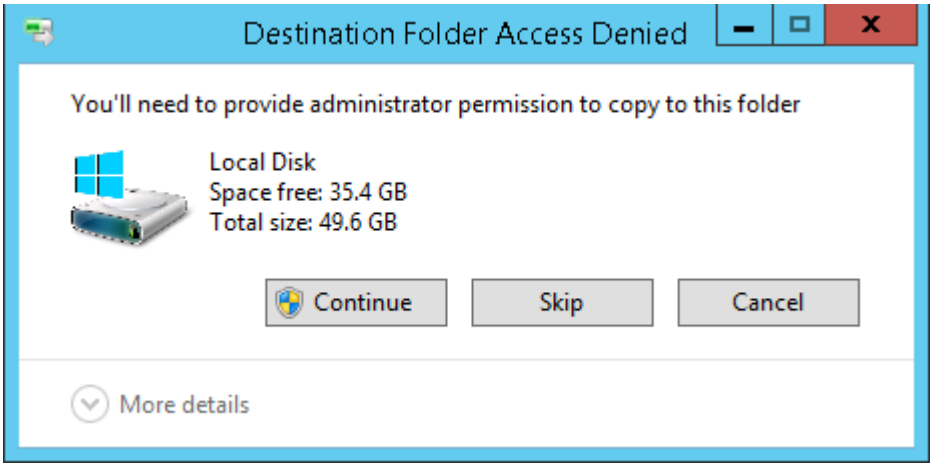
Update Nodes Using an Additional Administrative Account

An additional administrative account refers to those accounts that are not using default administrators. Such accounts are also referred as non-built-in administrative accounts. The Update Node and Preflight Check (PFC) functions use the account specified in Update Node to connect to a virtual machine and perform related checks.

Note: You should use either the built-in administrator or built-in domain administrator account when performing the Update Node function. If necessary, you can use a non-built-in administrator, but before doing so you should verify that the account you are using has the required administrator permissions.

Follow these steps:

1. Verify that you can access \\[VM host name]\ADMIN$ using the additional administrator account from another machine. If you have any problem, verify if the "File and Printer Sharing" is blocked by the firewall. If the firewall settings are good, you may need to disable the UAC remote access. To disable UAC remote access, see Import Virtual Machine Using Additional Administrative Account.
2. In VMware, when you update nodes, Arcserve UDP automatically installs some tools in the VM to perform PFC. To verify that the account has the required permissions, perform the following:
   a. Log in to the virtual machine using the non-built-in administrator account.
   b. Copy one file from C:\Windows into C:\ and ensure that the following message does not appear:



   c. If you experience any problem, you can modify the User Account Control (UAC) configurations. To modify, disable **Run all administrator in Admin**

Approval Mode in the Local Security Policy by changing the UAC settings at secpol.msc -> Local Policies -> Security Options. (Secpol.msc is Microsoft's security policy editor).

Note: Do not attempt to disable the UAC in the User Account Control Settings dialog box that opens from the control panel.

For more information about changing the UAC configuration settings, see the corresponding Microsoft documentation.

3. For Hyper-V VMs, the additional administrator account must have similar permissions as mentioned in Import Virtual Machine Using Additional Administrative Account.

Export Node

You can export the nodes as a ZIP (.zip) file. When required, you can import the ZIP file to retain the nodes. For example, exporting the nodes before upgrades or rebooting helps you import the same set of nodes.

You can export only such nodes that have valid credentials and are running the Windows operating system.

Follow these steps:

1. Click the **resources** tab.

   The **Nodes: All Nodes** page is displayed.
2. Select a node.
3. From the center pane, click the **Actions** drop-down list, click **Export**.

   A dialog opens requesting your action on the list.zip file.
4. Click **Open** or **Save**.

   The node list is exported.

Pause Node

Using Arcserve UDP, you can pause only selected node/s instead of the complete plan. To prevent a scheduled job from running, now you do not need to pause and resume the complete plan that stops all the associated nodes.

Important: Pause a node feature works only if that node is associated with either an Agent-Based Windows plan or a Host-Based Agentless plan. For other plans, you cannot pause a node.

Follow these steps:

1. Click the **resources** tab.

   The **Nodes: All Nodes** page is displayed.

2. Select desired node/s associated with a plan.

3. From the center pane, click the **Actions** drop-down list, click **Pause**.

   A confirmation dialog is displayed, explaining only manual jobs can run for paused nodes

4. Click **Yes**.

   The nodes are paused. If any node fails to pause, a message pops up appears with the reason.

Resume Node

Using Arcserve UDP, you can resume a paused node. Now, instead of pausing and resuming a plan, you can pause and resume specific nodes associated with a plan.

Important: Pause and Resume a node feature works only if the node is associated with either an Agent-Based Windows plan or a Host-Based Agentless plan. For other plans, you cannot pause or resume a node.

Follow these steps:

1. Click the **resources** tab.

   The **Nodes: All Nodes** page is displayed.

2. Select paused node/s.

3. From the center pane, click the **Actions** drop-down list, click **Resume**.

   A dialog opens requesting confirmation of your action.

4. Click **Yes**.

   The node starts running again.

Synchronize Data

Synchronizing data keeps the data that are in different databases consistent and up-to-date.

Follow these steps:

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and click **All Nodes**.

   The **Nodes: All Nodes** page is displayed.
3. (Group Level) Select one of the node groups that are displayed on the left pane, and right-click.
4. Click one of the following options:

   Note: You can view only those options that you have already added for the synchronization with Arcserve Unified Data Protection.

   - Full Synchronize Arcserve Backup
   - Incremental Synchronize Arcserve Backup

   The **Information** dialog explains that the selected synchronization method is submitted.

Delete Nodes from the Console

Using Arcserve Unified Data Protection, you get the option to delete a node. If you delete the nodes, the associated logs and job histories will also be deleted. You can add the deleted node later, if required.

Follow these steps:

1. Click the **resources** tab.

   The **Nodes: All Nodes** page is displayed.

2. Select a node that you want to delete.

3. Perform one of the following actions:

   - Right-click the node name.
   - Select the node name, and from the center pane click the **Actions** drop-down list.

4. Click **Delete**.

   A **Confirm** dialog opens.

5. Click **Yes**.

   The node is deleted from the Console.

Deploy Agent to Nodes

To upgrade or install Arcserve Unified Data Protection Agent for a node, use **Install/Upgrade Agent**. If the destination machine contains a prior version of Arcserve Unified Data Protection agent, then use the upgrade option to get the latest version. Otherwise, use the install option.

Note: You can deploy Arcserve Unified Data Protection agents to multiple nodes. At one time, you can run only 16 deploy tasks. If there are more than 16 tasks, other tasks remain in pending status and run only when some of the default 16 deploy tasks complete. To modify the maximum task count, update the following registry key:

deployMaxThreadCount

Follow these steps:

1. Click the **resources** tab.

   The **Nodes: All Nodes** page is displayed.

2. Select one or more nodes.

3. From the center pane, click the **Actions** drop-down list, and then click **Install/Upgrade Agent**.

   The details of Install or upgrade appear above the name of the node on the center pane.

4. Verify the details.

5. Specify the install/upgrade schedule and click **OK**.

   The node is installed or upgraded with the latest version of Arcserve Unified Data Protection Agent.

   Note: You can cancel an agent deployment if it is scheduled for a later time. To cancel an agent deployment, select the agent and click **Actions, Cancel Agent Deployment.**

Perform Preflight Checks for Your Backup Jobs

The Arcserve Unified Data Protection solution features a utility named Preflight Check (PFC) that enables you to run vital checks on specific nodes to detect conditions that can cause backup jobs to fail. PFC is only applicable to virtual machine nodes that are imported from vCenter/ESX or Hyper-V or Nutanix AHV. PFC does not work for VMware VM template. PFC runs automatically when you perform the following actions:

- Import virtual machines from a vCenter Server/ESX Server system or Hyper-V or Nutanix AHV.
- Let a backup job run.

Other References:

- Solutions for Preflight Check Items of VMware VMs
- Solutions for Preflight Check Items of Hyper-V VMs
- Solutions for Preflight Check Items of Nutanix AHV VMs

In addition, you can also perform a Preflight Check manually.

Follow these steps:

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and click **All Nodes**.

   The **All Nodes: Node** page is displayed in the center pane.
3. Right-click the name of a node, and click **Preflight Check**.

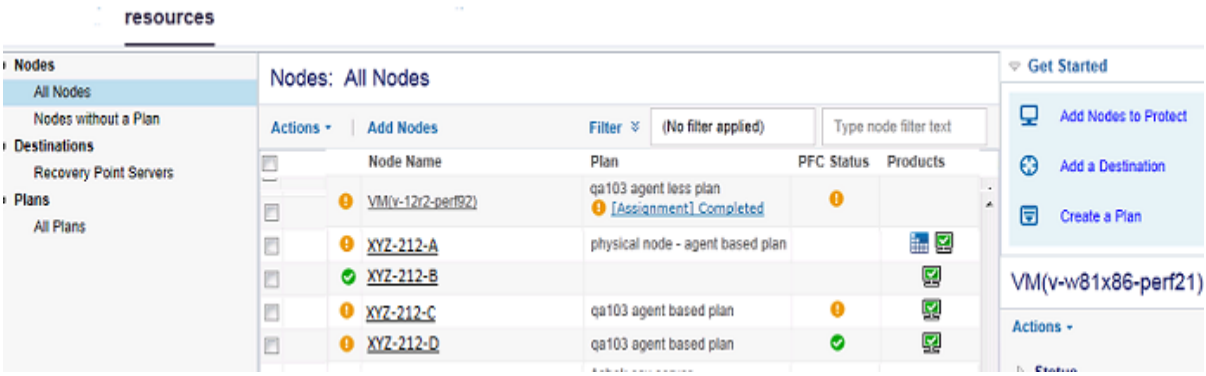
   Note: You can also perform **Preflight Check** using one of the following options:

   - (Node level) Click the check boxes of the nodes on which you want to run a preflight check, and then click **Actions** and select **Preflight Check**.
   - (Group level) Right-click the group containing the nodes, click **Preflight Check**.

   The following message is displayed: **Starting to preflight check the virtual machine**.
4. Navigate to the **PFC Status** column to view the status of the Preflight Check.

   Note: By default, the PFC Status column is not visible on the UI. You need to manually enable the PFC Status column on the UI.

   You can also view status of Preflight Check clicking **View Logs** from the right pane.



The following table describes the checks that PFC performs for VMware VM:

Item	Description
Changed Block Tracking (CBT)	A feature to tracks disk sectors that are on a virtual machine which has changed. This helps minimize the size of the backups. This item verifies that CBT is enabled.
VMware Tools	This item verifies that the VMware tools are installed on each virtual machine.
Disk	This item verifies the disks of the virtual machine.
Power State	This item verifies that the virtual machine is powered on.
Data Consistency	This item verifies if Application-consistent snapshot can be taken for the VM.
ESX Server	This item verifies if a server can detect the virtual machine. The option is visible only when the node is not available in ESX.

The following table describes the checks that PFC performs for Hyper-V VM:

Item	Description
Hyper-V Credentials	The product needs to deploy a backup utility and a Change Block Tracking utility to Hyper-V server through system share ADMIN$. The action helps in verifying if the product has necessary permission to the share. The Backup/restore job fails if the Hyper-V credentials are not correct or the administrator closed the ADMIN$ share.
Integration Services	This item verifies that the Hyper-V integration services are installed and enabled on each virtual machine. Without the integration services, Arcserve Unified Data Protection cannot complete the following actions: ■ Execute pre/post command and application log purge actions. ■ Perform application-consistent backup. Integration services contain several services. The Arcserve Unified Data Protection solution checks the statuses of the following two services: ■ Hyper-V Data Exchange Service: Required for collecting the VM info, executing the pre- or post-commands and the application log purge actions. ■ Hyper-V Volume Shadow Copy Requestor: Required for the application-con-

	sistent backup.
Power State	This item verifies that the virtual machine is powered on. A Suspended warning is shown when the VM is in the status other than powered on and power off, like the Saved status. The Arcserve Unified Data Protection solution cannot run the pre/post commands and the application log purge actions when the VM is not in the Powered On status. In addition, Arcserve Unified Data Protection cannot perform the application-consistent backup when VM is in the Suspended status.
Disk	This item verifies if unsupported disk is attached to the VM.
Data Consistency	This item verifies if Application consistent snapshot can be taken for the VM.

The following table describes the checks that PFC performs for Nutanix AHV VM:

Item	Description
Power State	This item verifies that the virtual machine is powered on.
Data Consistency	This item verifies if Application-consistent snapshot can be taken for the VM.
AHV Server	This item verifies if a server can detect the virtual machine. The option is visible only when the node is not available in Nutanix AHV.

Solutions for Preflight Check Items of VMware VMs

Other References:

- Solutions for Preflight Check Items of Hyper-V VMs
- Solutions for Preflight Check Items of Nutanix AHV VMs

The following tables describe the solutions to help you resolve errors and warnings from your Preflight Check results for VMware VMs:

Changed Block Tracking (CBT)

Status	Message	Solution
Error	Unable to enable changed block tracking.	If the virtual machine does not have hardware version 7 or higher, upgrade the hardware version of the virtual machine, or create an agent-based backup plan in Arcserve Unified Data Protection and use Arcserve Unified Data Protection Agent for Windows to back up the VM.
Warning	Changed Block Tracking is enabled with snapshots present. A full disk backup will be applied for Full and Verify backup jobs.	**Note:** This affects Full and Verify backup jobs only. For an Incremental backup job, only changed data are backed up. To apply the used block backup for Full and Verify backup jobs, perform the following steps: 1. Delete all the snapshots associated with the virtual machine. 2. Log in to the Backup proxy server. 3. Open the registry editor and locate the following key: *HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\ARCserve Unified Data Protection\Engine\AFBackupDll\<VM-InstanceUUID>* **Note:** Replace <VM-InstanceUUID> with the UUID value of the virtual machine where CBT is failing. You can find the value in the URL of the virtual machine that is used when connected to Arcserve UDP Agent (Windows). 4. Set registry key to "full disk backupForFullBackup"=0. 5. Create/set the registry to ResetCBT=1. 6. Submit the backup job.

VMware Tools

Status	Message	Solution
Warning	Out of date.	Install the latest version of VMware Tools.
Warning	Not installed or not running.	Install the latest version of VMware Tools and ensure that the tool is running.

Disk

Status	Message	Solution
Error	VM snapshots are not supported for the VM because it has a SCSI controller configured for bus-sharing configuration.	Create an agent-based backup plan in Arcserve Unified Data Protection or use Arcserve Unified Data Protection Agent for Windows to back up the VM.
Warning	The physical Raw Device Mapping (RDM) disk is not backed up.	Create an agent-based backup plan in Arcserve Unified Data Protection or use Arcserve Unified Data Protection Agent for Windows to back up the VM.
Warning	The virtual Raw Device Mapping (RDM) disk backs up as a full disk.	Create an agent-based backup plan in Arcserve Unified Data Protection or use Arcserve Unified Data Protection Agent for Windows to back up the VM.
Warning	The independent disk is not backed up.	Create an agent-based backup plan in Arcserve Unified Data Protection or use Arcserve Unified Data Protection Agent for Windows to back up the VM.
Warning	For Full or Verify backup jobs, the disk on the NFS data store is backed up as a full disk.	**Note**: This affects Full and Verify backup jobs only. For an Incremental backup job, only changed data are backed up. Move the virtual disk to a data store on a block storage device or create an agent-based backup plan in Arcserve Unified Data Protection or use Arcserve Unified Data Protection Agent for Windows to back up the VM.

Power State

Status	Message	Solution
Warning	Powered off	Power on the virtual machine.
Warning	Suspended	Power on the virtual machine.

Data Consistency

Status	Message	Solution
Warning	VMware does not support application-consistent quiescing for a VM that has IDE disks.	Create an agent-based backup plan in Arcserve Unified Data Protection or use Arcserve Unified Data Protection Agent for Windows to back up the Microsoft SQL Server and Exchange Server data.
Warning	VMware does not support application-consistent quiescing for a VM that has SATA disks.	Create an agent-based backup plan in Arcserve Unified Data Protection or use Arcserve Unified Data Protection Agent to back up the Microsoft SQL Server and Exchange Server data.
Warning	VMware does not support application-consistent quiescing because the version of the ESX server is prior to release 4.1.	Upgrade ESX Server to 4.1 or higher or create an agent-based backup plan in Arcserve Unified Data Protection or use Arcserve Unified Data Protection Agent for

		Windows to back up the Microsoft SQL Server and Exchange Server data.
Warning	VMware does not support application-consistent quiescing because there are not enough SCSI slots available	Add a SCSI Controller to the VM from the vSphere Web Client or create an agent-based backup plan in Arcserve Unified Data Protection or use Arcserve Unified Data Protection Agent for Windows to back up the Microsoft SQL Server and Exchange Server data.
Warning	VMware does not support application-consistent quiescing if the guest OS has dynamic disks.	Create an agent-based backup plan in Arcserve Unified Data Protection or use Arcserve Unified Data Protection Agent for Windows to back up the Microsoft SQL Server and Exchange Server data. **Note:** VMware does not support application-level quiescing on virtual machines that are Windows Server 2008 or later with dynamic disks running on ESX Server 4.1 or later.
Warning	Not verified because the application failed to access the virtual machine. Verify that the user credentials are correct and have administrative privileges.	Provide the built-in local administrator or domain administrator credentials to log in to the virtual machine guest operating system. Also, verify that the VMware Tools inside the virtual machine are update-to-date and running. Due to a VMware limitation, backup is supported only on VMs running on an ESX server that has a paid license. Backup is not supported on an ESXi server with a free license. **Note:** Data Consistency check is supported on Windows Server 2008 and later.
Warning	Unable to verify whether data-consistent backup is possible because the virtual machine is not powered on.	Refer to the Power State column
Warning	VMware does not support application-consistent quiescing if the guest OS has storage spaces enabled. File-level recovery is supported only for those volumes that do not have storage spaces enabled. (Full VM recovery is supported through Recover VM).	Create an agent-based backup plan in Arcserve Unified Data Protection or use Arcserve Unified Data Protection Agent for Windows to back up the Microsoft SQL Server and Microsoft Exchange Server data.

How to Create Application Consistent Snapshots for VMware

In some cases, the VMware VSS writer does not create application consistent snapshots on some virtual machines (VM). As a result, the backed up data may not be in an application-consistent state.

Verify Prerequisites:

Complete the following prerequisites to create application consistent snapshots:

- Latest VMware Tools must be installed and running inside the guest OS of the VM. Also, ensure that VMware Snapshot Provider service is installed (no need to be running).
- The VM must run on ESXi 4.1 or later.
- The VM must use only SCSI disks. The VM must have equal number of free SCSI slots to match the number of disks.
- Application consistent quiescing is not supported for VM that have IDE or SATA disks.
- All volumes in the VM are basic disks and there are no dynamic disks.
- The VM guest OS does not have storage spaces enabled.
- The disk.EnableUUID parameter of the VM must be enabled. VMs created on 4.1 or later have this parameter enabled by default. The following configurations are performed automatically by backup job to avoid data inconsistency and perform application-consistent backup. If backup job cannot enable disk.EnableUUID due to some reasons, configure the parameter manually using the following procedure:
  - If disk.EnableUUID exists and is FALSE, change it to TRUE.
  - If disk.EnableUUID does not exist, create it and set it to TRUE.
  - If disk.EnableUUID exists and is TRUE, keep it as it is.

  Note: For more information about creating application-consistent backup, see the VMware KB article.

Affected Features:

If any of the requirements are not met, the session data remains crash consistent. As a result, the following features are affected:

- Backed up data that includes application data of a VM, such as SQL, Exchange, and SharePoint, may remain in a crash consistent state.
- Catalog job may fail.

Solutions for Preflight Check Items of Hyper-V VMs

Other References:

- Solutions for Preflight Check Items of VMware VMs
- Solutions for Preflight Check Items of Nutanix AHV VMs

The following tables describe the solutions to help you resolve errors and warnings from your Preflight Check results for Hyper-V VMs:

Hyper-V Credentials

Status	Message	Solution
Error	Failed to access the ADMIN$ share of the Hyper-V server or does not have the proper credentials.	■ Verify if the Hyper-V server is running ■ Verify if the network of Hyper-V server is connectable. ■ Verify if the ADMIN$ share of Hyper-V Server is enabled. ■ Provide administrator rights of Hyper-V when importing VM from it.

Integration Services

Status	Message	Solution
Warning	Not installed, running, operational.	Install/Upgrade/Enable the integration services. **Notes:** ■ For Windows VM, if the integration services are installed, verify if the following two required services are running in the VM: Hyper-V Data Exchange Service and Hyper-V Volume Shadow Copy Requestor. Also verify, if there are errors of Hyper-V services in the event log of VM. ■ For Linux VM, verify the latest integration services are installed, and *Key-Value Pair* and *Live virtual machine backup* features are available on the specific Linux VM. For more information on Linux-integrated services on Hyper-V VM, see the Microsoft KB article.
Warning	Not responding	Restart the integration services in the guest OS of the VM.

Warning	The integration service inside the virtual machine is not compatible with the integration service in the Hyper-V server.	Create an agent-based backup plan in Arcserve Unified Data Protection or use Arcserve Unified Data Protection Agent to back up the VM.
Warning	Out of date.	Upgrade the integration services.

Power State

Status	Message	Solution
Warning	Powered off.	Power on the virtual machine.
Warning	Suspended.	Power on the virtual machine.

Disk

Status	Message	Solution
Warning	The physical hard disk that is attached to the virtual machine will not be backed up.	Create an agent-based backup plan in Arcserve Unified Data Protection or use Arcserve Unified Data Protection Agent to back up the virtual machine.
Warning	Failed to get the virtual machine by instance UUID.	Verify if the virtual machine exists on the Hyper-V server.

Data Consistency

Status	Message	Solution
Warning	An application-consistent snapshot is not supported. The virtual machine has a dynamic disk.	Create an agent-based backup plan in Arcserve Unified Data Protection or use Arcserve Unified Data Protection Agent to back up the virtual machine.
Warning	An application-consistent snapshot is not supported because the VM has storage spaces.	Create an agent-based backup plan in Arcserve UDP or use Arcserve UDP Agent to back up the virtual machine.
Warning	An application-consistent snapshot is not supported because the VM has a volume whose shadow copy storage is located on another volume.	Change the volume's shadow copy storage area to the volume itself. **Note:** For a VM in Hyper-V 2012 R2, if the latest update of Microsoft is applied on Hyper-V host, application-consistent snapshot is still supported in such cases.
Warning	An application-consistent snapshot is not supported. The virtual machine has different file systems other than NTFS/Refs.	If you want to back up the virtual machine but skip the File Systems other than NTFS/Refs, create an agent based backup plan in Arcserve Unified Data Protection or use Arcserve Unified Data Protection Agent to back up the virtual machine.

		Note: For a VM in Hyper-V 2012 R2, if the latest update of Microsoft is applied on Hyper-V host, application-consistent snapshot is still supported in such cases.
Warning	An application-consistent snapshot is not supported. The Scoped Snapshot feature is enabled in the virtual machine.	Disable the Scoped Snapshot inside VM by adding a DWORD registry key HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\SystemRestore\ScopeSnapshots with the value 0.
Warning	An application-consistent snapshot is not supported. The integration service is not operational (Failed state).	Refer to the Integration Services column.
Warning	Not verified because the application failed to get the virtual machine by instance UUID.	Verify if the virtual machine exists on the Hyper-V server.
Warning	Not verified because the virtual machine is not powered on.	Refer to the Power State column.
Warning	An application-consistent snapshot is not supported. Failed to log in to the virtual machine to check the reason.	Provide built-in local administrator or domain administrator credentials to log in to the virtual machine guest operating system. Also, verify that the virtual machine has network connectivity.
Warning	An application-consistent snapshot is not supported for an unknown reason.	To identify the reason of cannot take an application-consistent snapshot, check the event logs. The event log is located at the following location: Inside the VM: Event Viewer>Windows Logs>Application and System. In the log, look for errors that come from Disk, VSS, and VolSnap. On the Hyper-V server: Event Viewer>Windows Logs>Application and Services Logs> Microsoft>Windows>Hyper-V-*. In the log, look for errors for the respective VM.

How to Create Application Consistent Snapshots for Hyper-V

In some cases, the Hyper-V VSS writer does not create application consistent snapshots on some virtual machines (VM). As a result, the backed up data may not be in an application-consistent state.

Verify Prerequisites:

Complete the following prerequisites to create application consistent snapshots:

- In the child VM, the integration service named Hyper-V Volume Shadow Copy Requestor is installed and running.
- The child VM is in the running state.
- The Snapshot File Location for the VM is set to the same volume in the host operating system as the VHD files for the VM.
- All volumes in the child VM are basic disks and there are no dynamic disks.
- All disks in the child VM must use a file system that supports snapshots (for example, NTFS).

Verify Considerations:

Complete the following considerations to create application consistent snapshots:

- Integration Service installed in the child VM must be compatible with the Hyper-V host.
  - For example: Windows 8.1/2012R2 integration service inside VM is not compatible with Windows 2008R2 Hyper-V host.
- For Windows 8, 2012 and later, and the VM running in Windows 2008R2 Hyper-V host, the Scoped Snapshot feature in the VM must be disabled. To disable the Scoped Snapshot feature, follow these steps:
  1. Log into the VM.
  2. Navigate to the following location:

     HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion
  3. Open the SystemRestore key.

     Note: Create the key if it does not exist.
  4. Add a 32-bit DWORD registry value named "ScopeSnapshots" and set the value as 0.

Affected Features:

If any of the requirements are not met, the session data is crash consistent. As a result, the following features are affected:

- Backed up data that includes application data of a VM, such as SQL, Exchange, and SharePoint, may remain in a crash consistent state.
- Catalog job may fail.

Solutions for Preflight Check Items of Nutanix AHV VMs

Other References:

- Solutions for Preflight Check Items of VMware VMs
- Solutions for Preflight Check Items of Hyper-V VMs

The following tables describe the solutions to help you resolve errors and warnings from your Preflight Check results for Nutanix AHV VMs:

Nutanix Guest Tools (NGT)

Status	Message	Solution
Warning	Not running	Enable the Nutanix Guest Tools for the VM.
Warning	Tools Enabled	Guest Tools enabled but not running. Install the Guest Tools in the VM.

Power State

Status	Message	Solution
Warning	Powered off	Power on the virtual machine.
Warning	Suspended	Continue the virtual machine.

Data Consistency

Status	Message	Solution
Warning	Nutanix AHV does not support application-consistent quiescing for a VM that has IDE disks.	Create an agent-based backup plan in Arcserve UDP or use Arcserve UDP Agent (Windows) to back up the Microsoft SQL Server and Exchange Server data.
Warning	Nutanix AHV does not support application-consistent quiescing for a VM that has SATA disks.	Create an agent-based backup plan in Arcserve UDP or use Arcserve UDP Agent to back up the Microsoft SQL Server and Exchange Server data.
Warning	Nutanix AHV does not support application-consistent quiescing because enough SCSI slots are not available.	Create an agent-based backup plan in Arcserve UDP or use Arcserve UDP Agent (Windows) to back up the Microsoft SQL Server and Exchange Server data.
Warning	Not verified because the application failed to access the virtual machine. Verify that the user credentials are correct and have administrative privileges.	Provide the built-in local administrator or domain administrator credentials to log into the virtual machine guest operating system.
Warning	Unable to verify whether data-consistent backup is possible because the virtual machine is not powered on.	Refer to the Power State column.

How to Add and Manage UNC/NFS Path

The UNC/NFS Path is introduced as a node type in Arcserve UDP. Unlike other nodes, you cannot add the UNC/NFS Path node from the All Nodes section. Like Exchange Online node, you can add the UNC/NFS Path node only when you create a plan. After adding, you can manage the UNC/NFS Path node from the All Nodes section to update or delete.

What To Do Next?

- Add UNC/NFS Path
- Update UNC/NFS Path
- Delete UNC/NFS Path

Add UNC/NFS Path

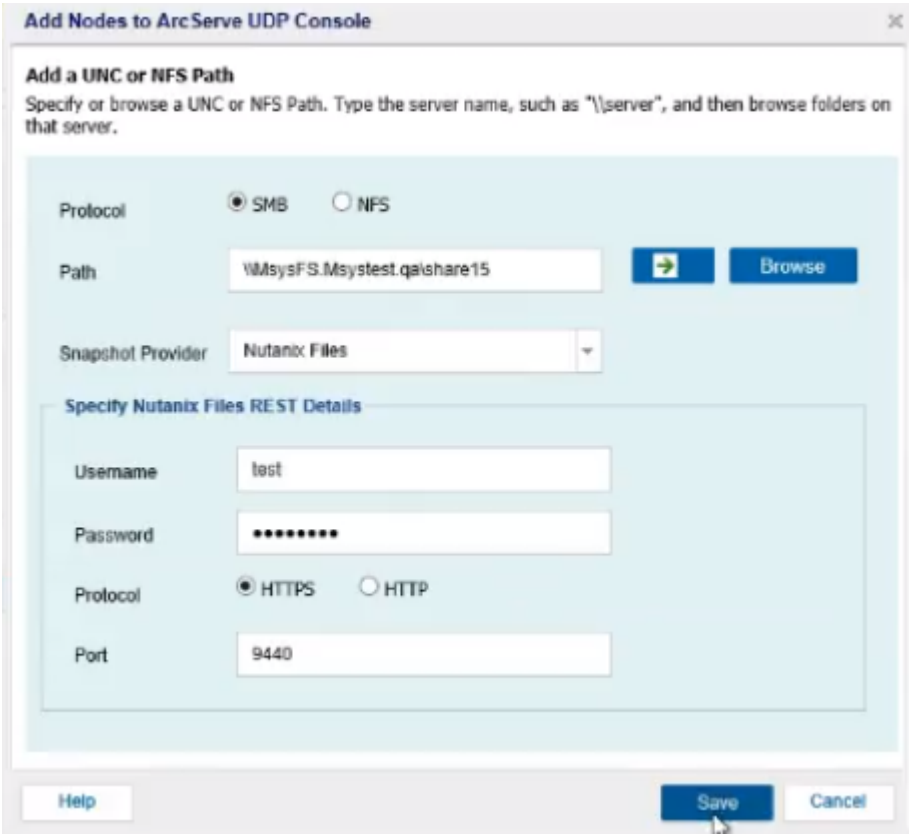
You can add a UNC/NFS path only when you create a UNC/NFS Path backup plan. While creating the plan, from the **Source** tab you can add a UNC/NFS Path node.

Follow these steps:

1. From the **Source** tab of the UNC Path backup plan, click the **Add UNC or NFS Path** option.

   The Add Nodes to Arcserve UDP Console dialog appears.

2. On the Add a UNC or NFS Path page, perform the following steps:



   a. Select **SMB** or **NFS** as a protocol.

      - If you select **SMB**, do the following:

        1. Enter the UNC Path in the following format:

           *Hostname**share*

           Note: To get the UNC path, do the following:

a. Select the newly created share in the center pane.

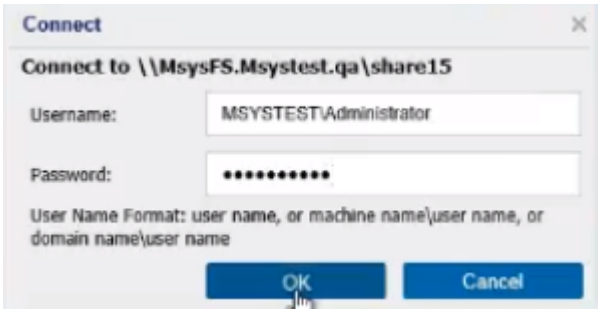
b. Under the SHARE/EXPORT DETAILS column, from the Summary pane, navigate to **Mount Path** and then copy the mount path.

   If the UNC path is valid, a right arrow and the Browse option gets enabled.

2. Click the right arrow (>) to validate the UNC Path.

   The Connect dialog appears.

3. Enter the username and password as needed, and then click **OK**.



   The Add Nodes to Arcserve UDP Console dialog displays the verified UNC Path.

- If you select NFS, do the following:

  1. Enter the UNC Path in the following format:

     *Hostname**share*

  2. Select Encoding from the available options - ANSI, GB2312-80, KSC5601, BIG5, SHIFT-JIS, EUC-KR, EUC-TW, EUC-JP.

     The default value is ANSI.

     Note: During the backup of an NFS Shared Folder, all such files/-folders whose name has an unsupported language encoding are skipped or their name appears as junk.

b. From the Snapshot Provider drop-down list, select **Nutanix Files**.

c. For the Specify Nutanix Files REST Details box, specify the following credentials:

Username

Refers to the user name created by the REST API user using the Nutanix prism.

Note: For a more secure communication, select the HTTPS protocol.

Password

Refers to password for the username.

Protocol

Specifies the protocol that you want to use to communicate with the destination server. The available options are HTTP and HTTPS.

Port

Refers to the port number. The default port value is 9440 to access the REST API calls and Prism.

3. Click **Save**.

You can update or delete the UNC/NFS Path node.

Update UNC/NFS Path

You can update information that is related to the existing nodes. When you modify the credentials of the UNC/NFS Path, you must update the UNC/NFS Path on Console.

Follow these steps:

1. Click the **resources** tab.

   The **Nodes: All Nodes** page is displayed.

2. From the left pane, navigate to **Nodes** and select **UNC/NFS Paths**.

   The existing UNC/NFS Path nodes are displayed on the center pane.

3. From the center pane, select a UNC/NFS Path and click **Actions**.

   A list of options appears.

4. From the list, click **Update**.

   The **Update node** dialog opens.

5. Update the details and click **OK**.

   The node information is updated.

Delete UNC/NFS Path

You can delete an existing UNC/NFS Path node from the **Resources** tab:

Follow these steps:

1. Click the **resources** tab.

   The **Nodes: All Nodes** page is displayed.

2. From the left pane, navigate to **Nodes** and select **UNC/NFS Paths**.

   The existing UNC/NFS Path nodes are displayed on the center pane.

3. Select the desired UNC/NFS Path and click **Actions**.

   A list of options appears.

4. Click **Delete**.

   A confirmation dialog appears.

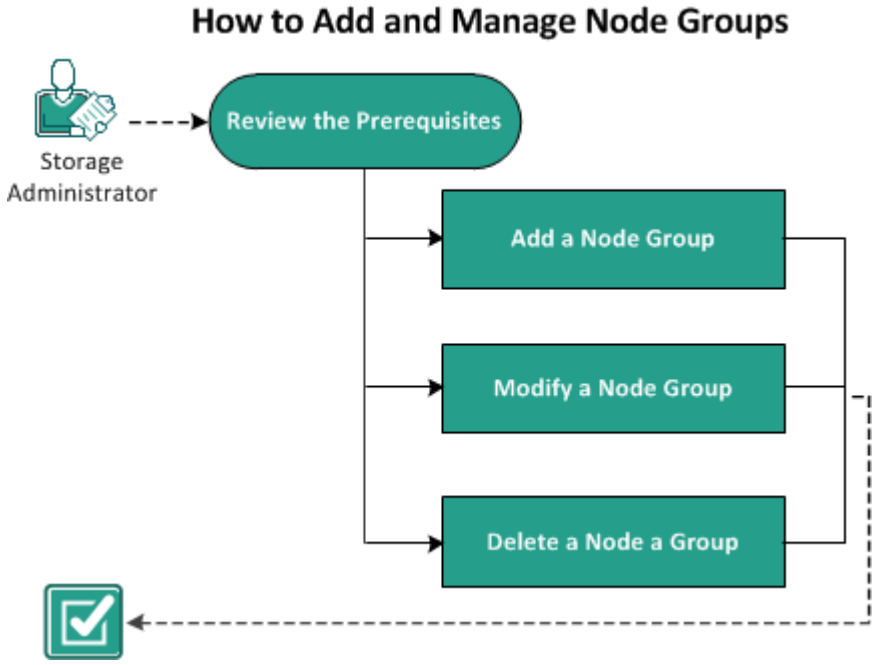
5. Click **OK**.

   The UNC/NFS Path is deleted successfully

How to Add and Manage Node Groups

Using Arcserve Unified Data Protection, you can add multiple nodes in to a group. You can add node groups to manage your physical and virtual machine environment.

The following diagram illustrates how you can add and manage the node groups:



The Arcserve Unified Data Protection solution contains following node groups:

- Default Groups:
  - **All Nodes**: Displays all the nodes that are added to the Console.
  - **Nodes without a Plan**: Displays the nodes that do not have any plan assigned.

  Note: You cannot modify or delete the default node groups.
- Groups that appear when you add child groups:
  - **Plan Groups**: Displays the list of plans that you have created. Select each plan under the group to view all the nodes associated with that plan.
  - **Custom Groups**: Displays the list of customized node groups that you have created. For example, the node group that you create by clicking **Actions**, **Node Group**, **Add** from the center pane.
  - **vCenter/ESX Groups**: Displays the nodes that you add using the **Importing from vCenter/ESX** option.
  - **Linux Backup Server Groups**: Displays the Linux Backup Server nodes.

- **Exchange Online Nodes**: Displays the Exchange Online nodes.
- **UNC Paths**: Displays UNC nodes.
- **SLA Profile Groups**: Displays Service Level Agreement (SLA) related nodes.
- **Hyper-V Groups**: Displays the nodes that you add using the **Importing from Hyper-V** option
- **Global Dashboard Groups**: Displays all the Arcserve Backup branch primary server under the GDB server. The Global Dashboard group is added when you add one Arcserve Backup Global Dashboard server into the Console and perform a Full Arcserve Backup synchronization for the added GDB server.

What To Do Next?

- Review the Prerequisites
- Add Node Groups
- Modify Node Groups
- Delete Node Groups

Review the Prerequisites

Before working on the node groups, complete the following prerequisites:

- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.
- Log in to the Console.
- Add a node.

Add Node Groups

To manage the list of nodes, you can create a group for selected nodes. For example, you can group nodes by business function or by installed application. You can also add nodes into any custom groups later after adding a blank group.

Follow these steps:

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and click **All Nodes**.

   The **Nodes: All Nodes** page is displayed.
3. From the center pane, click the **Actions** drop-down list.
4. Click **Create a Node Group**.

   The **Add Group** dialog opens displaying all the available nodes.
5. Complete the following actions to add nodes to the group, and click **OK**.
   - Select nodes that you want to add in a group.
   - Provide a name to the group.

     The **Information** dialog opens on the right pane to provide the message that the node group is created.

   The added group is placed below **Custom Groups** on the left pane.

   Note: The **Modify** and **Delete** options are enabled only when you have added a group.

Modify Node Groups

Using Arcserve Unified Data Protection solution, you can modify the node groups that you created. You can add and remove nodes from node groups and change the name of the node groups.

Follow these steps:

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and click **All Nodes**.

   The **Nodes: All Nodes** page is displayed.
3. From **Custom Groups** in the left pane, select a group.

   The details of selected group are displayed on the center pane.
4. Click the **Actions** drop-down list, and then click **Edit this node group**.

   The **Modify Group** dialog opens.
5. Update the details and click **OK**.

   The node group is updated.

Delete Node Groups

You can delete a group, if required. When you delete a group that was manually added, the virtual or physical machines are not removed from Arcserve Unified Data Protection. However, if you delete a group that was automatically created from an ESX or vCenter Server discovery, the group and all virtual machines are deleted from the Console.

Important! You cannot delete the default node groups.

Note: The process of deleting the node groups does not delete individual nodes from the Console.

Follow these steps:

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and click **All Nodes**.

   The **Nodes: All Nodes** page is displayed.
3. From **Custom Groups** in the left pane, select a group name.

   The details of selected group are displayed on the center pane.
4. Click the **Actions** drop-down list, and then click **Delete this Node Group**.

   The **Confirm** dialog opens.
5. Click **Yes**.

   The **Information** dialog opens on the right pane to provide the message that the node group is deleted.

How to Manage Nodes for Cloud

Using Arcserve UDP, you can perform multiple actions to manage a node for Cloud.

What To Do Next?

- Download Recovery Point from Cloud
- Upload Recovery Point to Cloud
- Copy Recovery Point to local disk or network share

Download Recovery Point from Cloud

Using this feature, you can download Recovery Points from Cloud to Local share or Network share.

Important!

You need a node that has a backup plan and Copy Recovery Point task with Cloud as destination configured.

You need a Recovery Point already copied to Cloud.

Follow these steps:

1. Click the **resources** tab.

   The **Nodes: All Nodes** page is displayed.

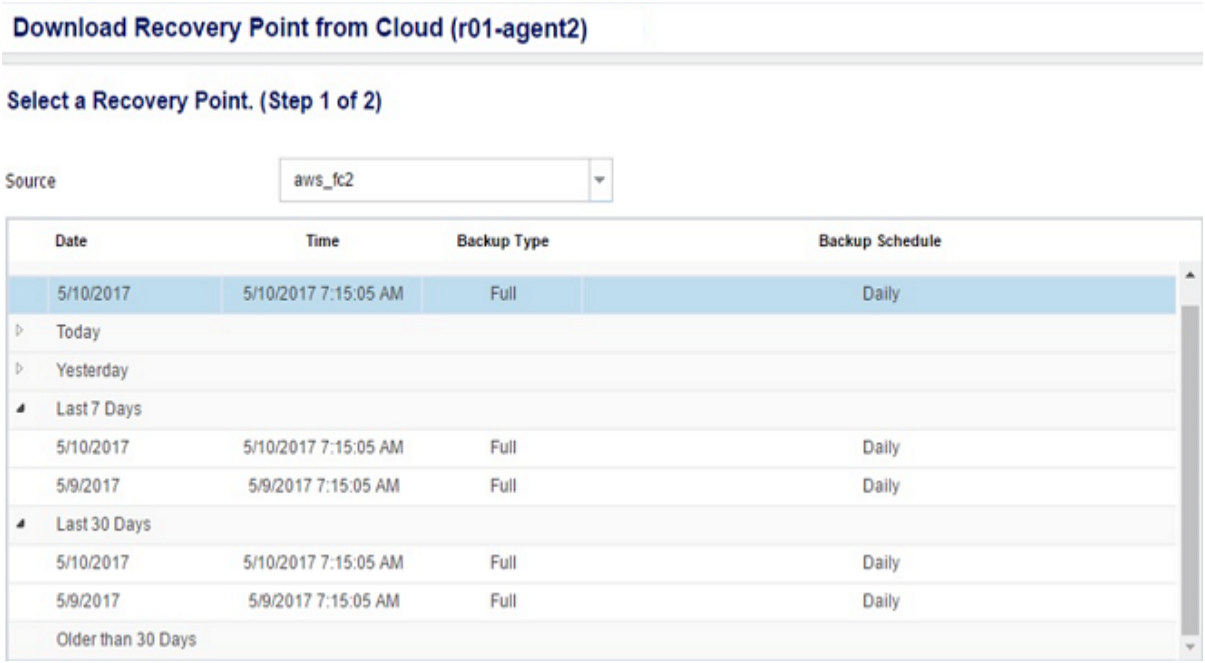
2. Perform one of the following actions:

   - Right-click the node name.
   - Select the node name, and from the center pane click the **Actions** drop-down list.

3. Click **Download Recovery Point from Cloud**.

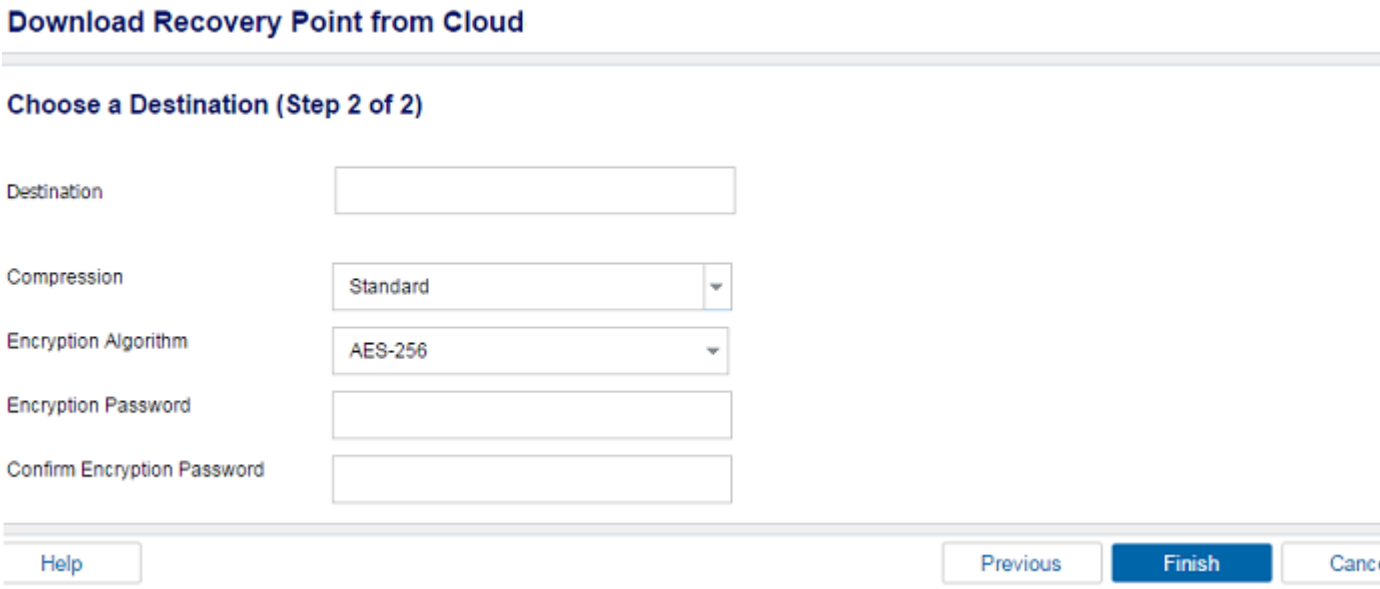
   The **Download Recovery Point from Cloud** dialog opens.

   You can download the recovery point from Cloud in two steps.



4. In Step 1, select a **Source** from the drop-down list and click **Next**.

The next screen of download appears.



5. In Step 2, provide the following details, and click **Finish**.
   a. Enter **Destination**.

      Note: If you have entered a remote destination that needs validation, after clicking **Finish** the **Connect** dialog appears. Enter details to get the destination validated.

   b. Select type of **Compression**.
   c. Select type of **Encryption Algorithm**.
   d. Enter **Encryption Password** twice.

The recovery point is downloaded from Cloud.

Upload Recovery Point to Cloud

Using this feature, you can upload Recovery Points from backup destinations like RPS\Local\Remote share to Cloud.

Important! You need a node that has a backup plan and a qualified backup session available to Copy to Cloud.

Follow these steps:

1. Click the **resources** tab.

   The **Nodes: All Nodes** page is displayed.

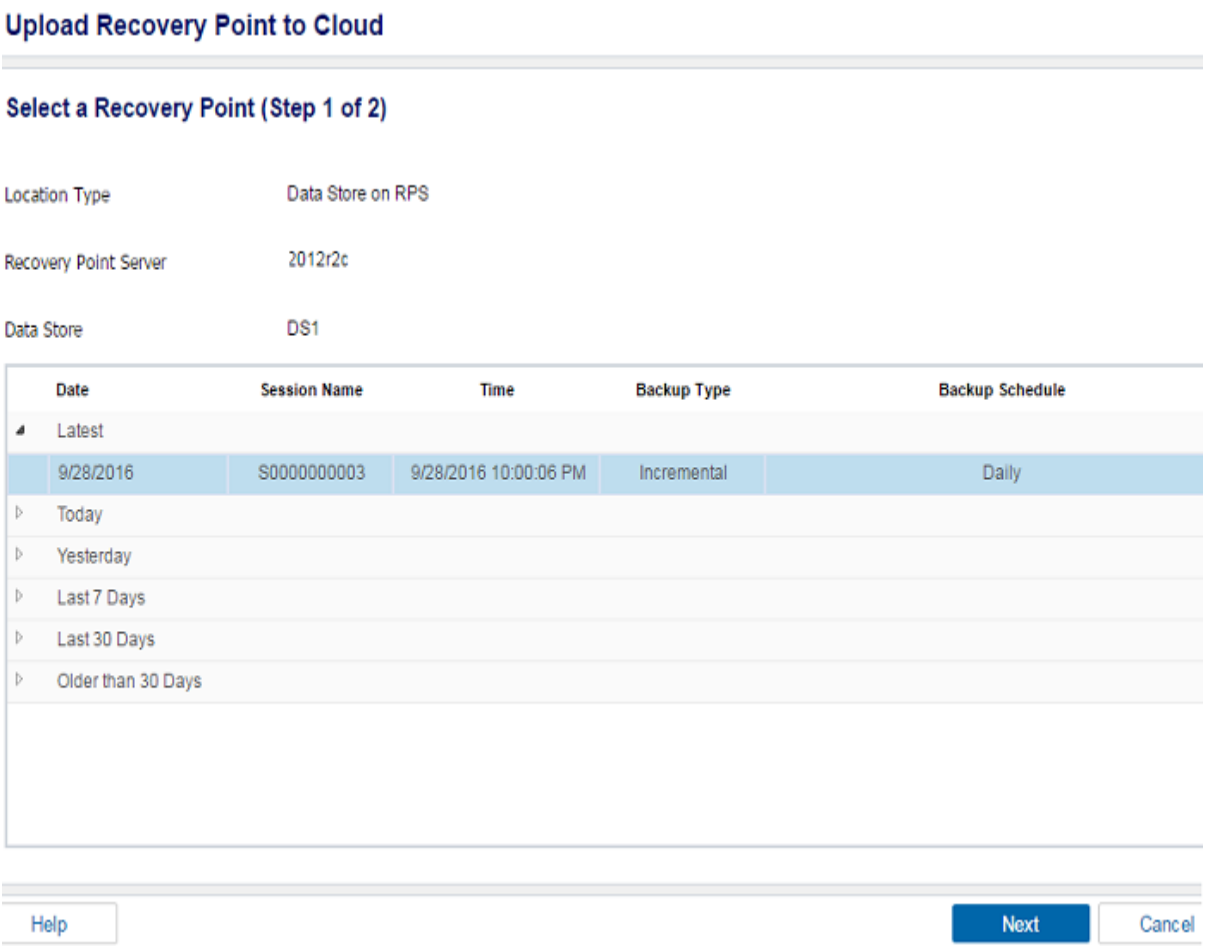
2. Perform one of the following actions:

   - Right-click the node name.
   - Select the node name, and from the center pane click the **Actions** drop-down list.

3. Click **Upload Recovery Point to Cloud**.

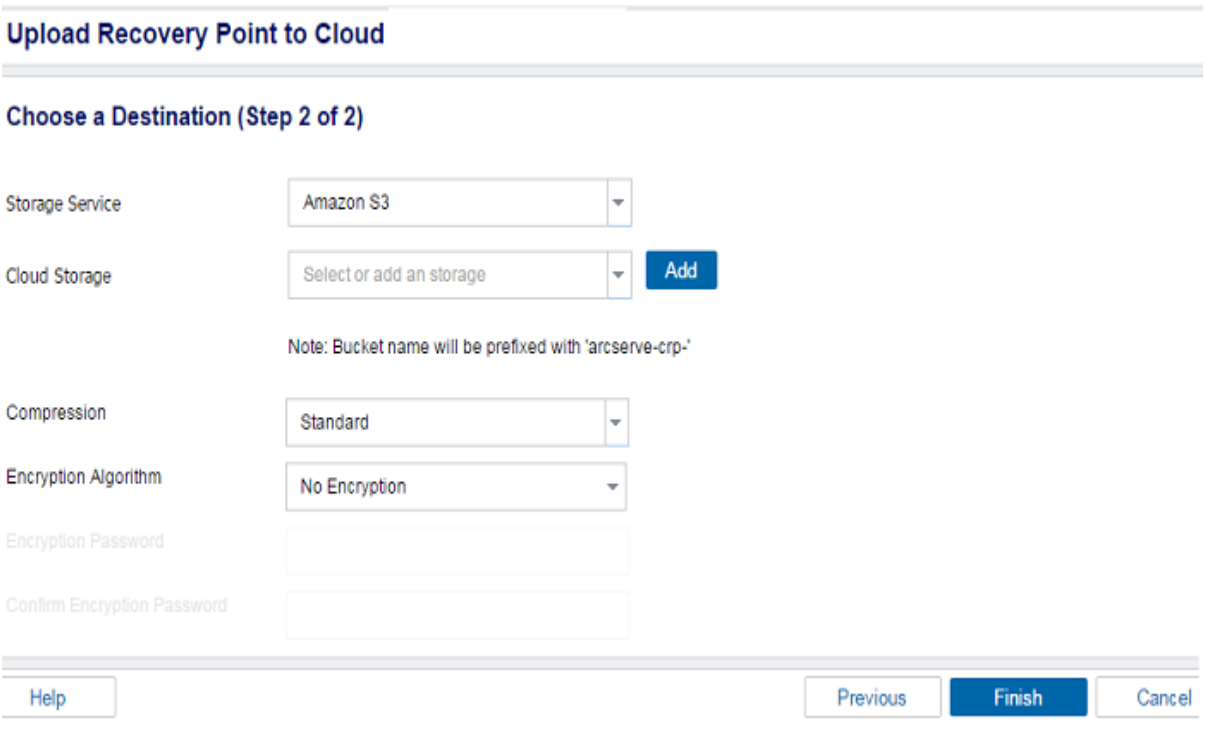
   The **Upload Recovery Point to Cloud** dialog opens.

   You can upload the recovery point to Cloud in two steps.



4. In Step 1, select a **Recovery Point** and click **Next**.

   The next screen of upload appears.



5. In Step 2, provide the following details, and click **Finish**.
   a. Select **Storage Service** from the drop-down list.
   b. Select **Cloud Storage** from the drop-down list.

      Note: If a cloud account is not added before, click **Add** to add a cloud account.
   c. Select type of Compression from the drop-down list.
   d. Select type of **Encryption Algorithm**.
   e. Enter **Encryption Password** twice.

   Note: A Bucket/Container is created with *arcserve-{random-generated-id}-<BucketName>* naming convention. However, each *Upload recovery point to cloud* job creates a new bucket/container with unique *{random-generated-id}* ID. For exact bucket/container name, refer to the job activity log.

   The recovery point is uploaded to Cloud.

Copy Recovery Point to Local Disk or Network Share

Using this feature, you can copy Recovery Points from backup destinations like RPS\Local\Remote share to Local or Remote.

Important! You need a node that has a backup plan configured and a qualified backup session available to copy to Local or Remote Share.

Follow these steps:

1. Click the **resources** tab.

   The **Nodes: All Nodes** page is displayed.

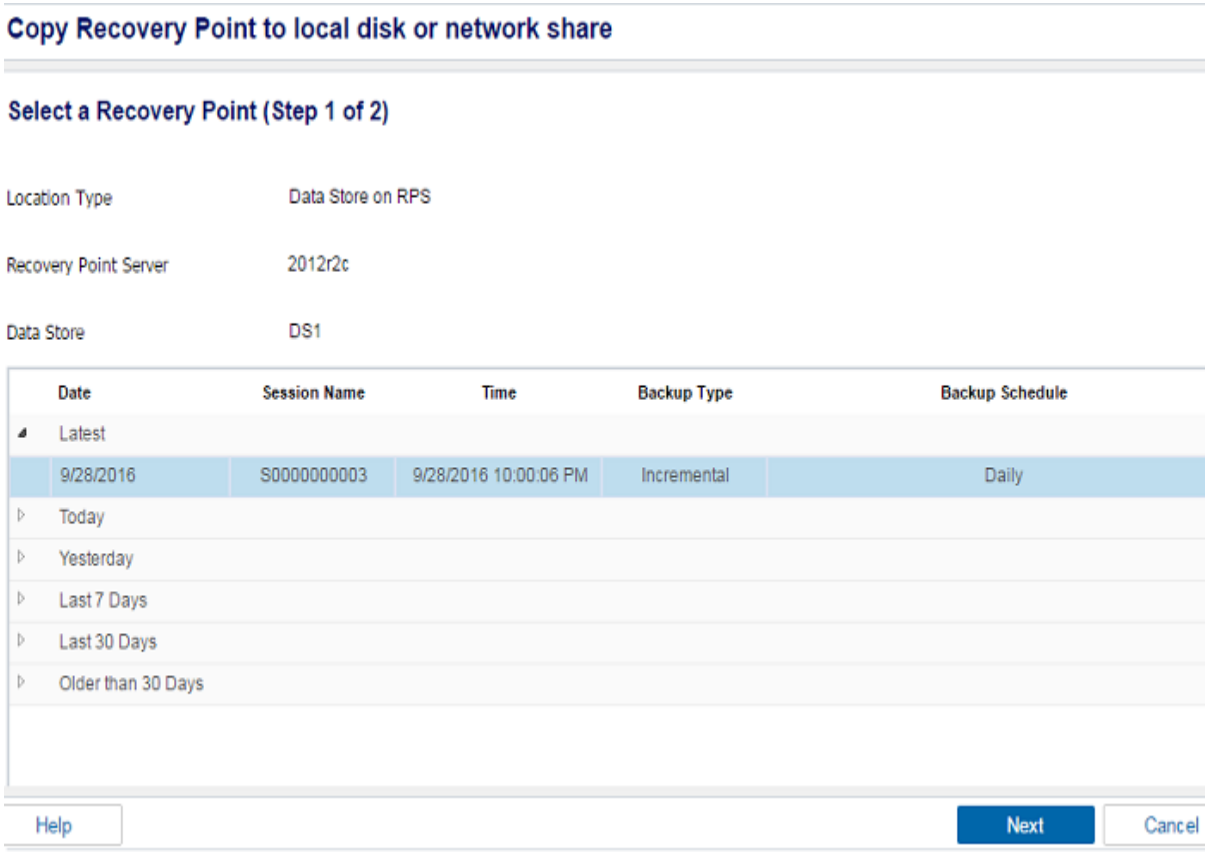
2. Perform one of the following actions:

   - Right-click the node name.
   - Select the node name, and from the center pane click the **Actions** drop-down list.

3. Click **Copy Recovery Point to local disk or network share**.

   The **Copy Recovery Point to local disk or network share** dialog opens.

   You can Copy Recovery Point to local disk or network share in two steps.



4. In Step 1, select a **Recovery Point** and click **Next**.

   The next screen of copy recovery point appears.



5. In Step 2, provide the following details, and click **Finish**.

   a. Enter **Destination**.

      Note: If you have entered a remote destination that needs validation, after clicking **Finish** the **Connect** dialog appears. Enter details to get the destination validated.

   b. Select type of **Compression**.

   c. Select type of **Encryption Algorithm**.

   d. Enter **Encryption Password** twice.

   The recovery point is copied from Cloud.

Add a Storage Array

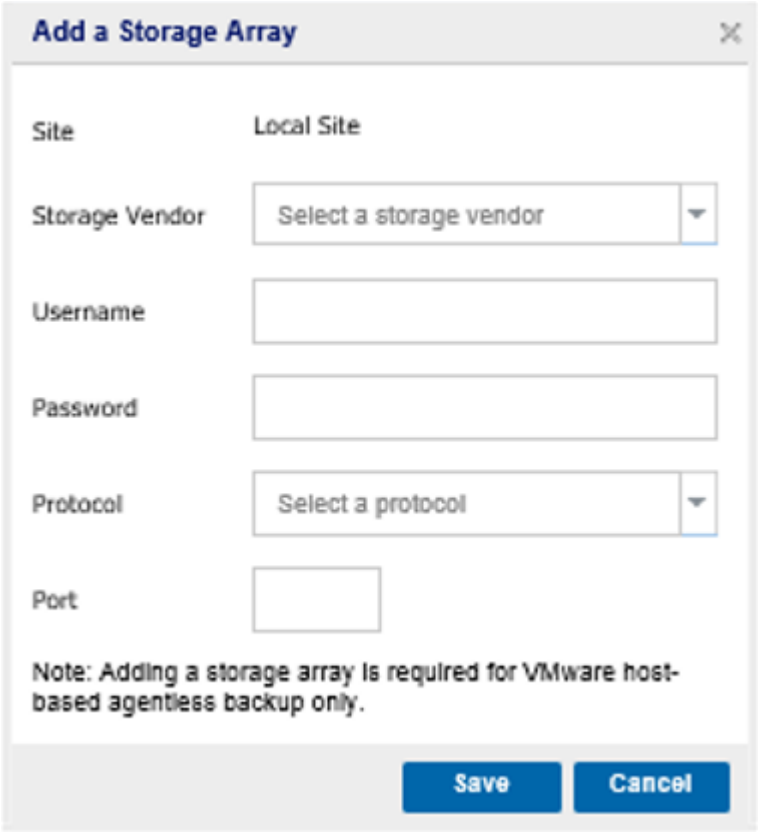
If you are using hardware snapshots, you have to add the storage array details to the Console. If you do not add a storage array and you submit a backup job using hardware snapshot, the backup job will first look for the storage array details in the Console. When the backup job does not find the storage array details, the job uses a software snapshot to create backup sessions.

Adding a storage array is required for VMware host-based agentless backup only.

Follow these steps:

1. From the Console, click **resources**.
2. Navigate to **Infrastructure** on the left pane and click **Storage Arrays**.

   The **Add a Storage Array** dialog opens.



You can add storage array for the following options:

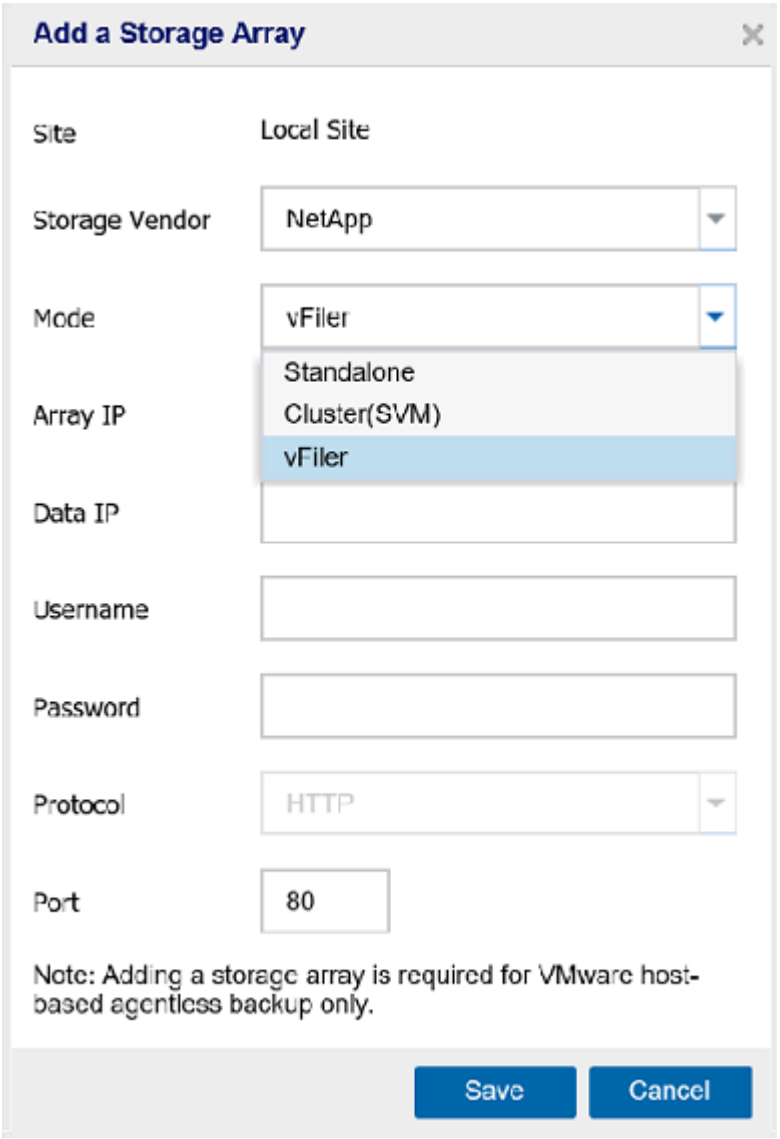
- NetApp
- HPE 3PAR
- Nimble

Add a Storage Array for NetApp

You can add a NetApp storage array using three options.

Follow these steps:

1. From the **Add a Storage Array** dialog, select NetApp from the **Storage Vendor** drop-down option.



2. From Mode, select one of the following options for the type of array:

Standalone

Specifies that the storage array is a standalone appliance.

Cluster(SVM)

Specifies that the storage array belongs to a cluster of arrays.

vFiler

Specifies that the storage array is a vFiler.

3. Provide the following details for the specific type of array:

Site

Specifies the name of the site.

Array IP

Specifies the IP address of the array. For 7-mode, provide the IP address that can access the iSCSI interface and NFS. VMware(ESXi) uses the IP address to access the storage array for read and write operations. For cluster and vFiler, provide the management IP address of the Storage Array or management IP address of the Storage Virtual Machine (SVM).

To find the Array IP address, refer to the topic Find Array and Data IP of the SVM Storage Array for the Arcserve UDP Console.

Data IP

Specify the Data access IP (Logical Interface) of the storage virtual machine (SVM), where the target machines (ESXi) can access this storage array for read and write operations. This option is applicable for Cluster and vFiler only. If the Storage Virtual Machine supports only FC protocol, then enter NA.

To find the Data IP address, refer to the topic Find Array and Data IP of the SVM Storage Array for the Arcserve UDP Console.

Username

Specifies the user name to connect to the storage array.

Password

Specifies the password for the user name.

Protocol

Specifies the protocol to connect to the array.

Port

Specifies the port number of the array.

Default: 443 (HTTPS), 80 (HTTP)

4. Click **Save**.

The storage array is added to the Console.

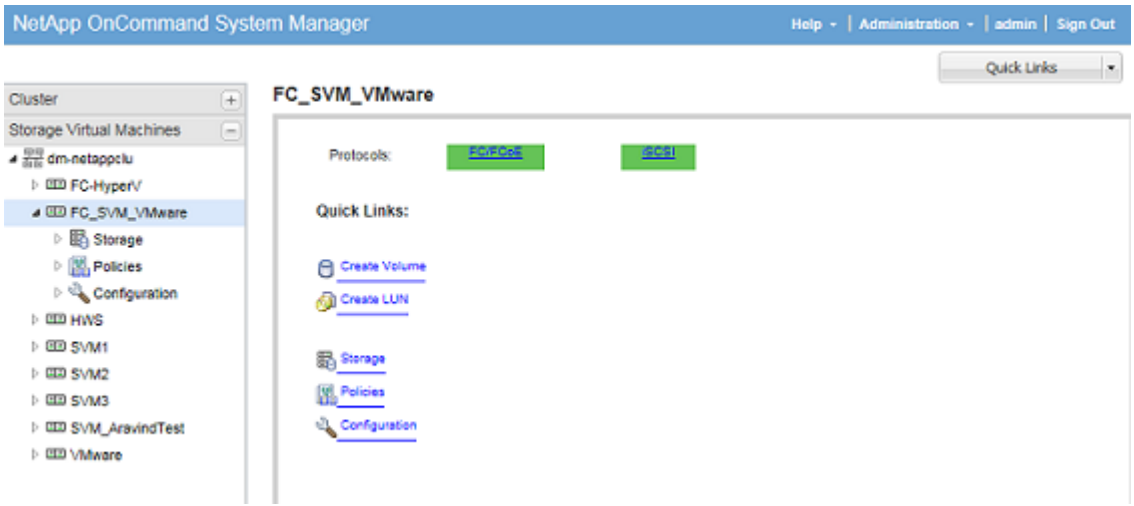
Find Array and Data IP of the SVM Storage Array for the Arcserve UDP Console

This section describes how to find Array and Data IP when the NetApp Storage array is running in the cDOT(clustered) mode.

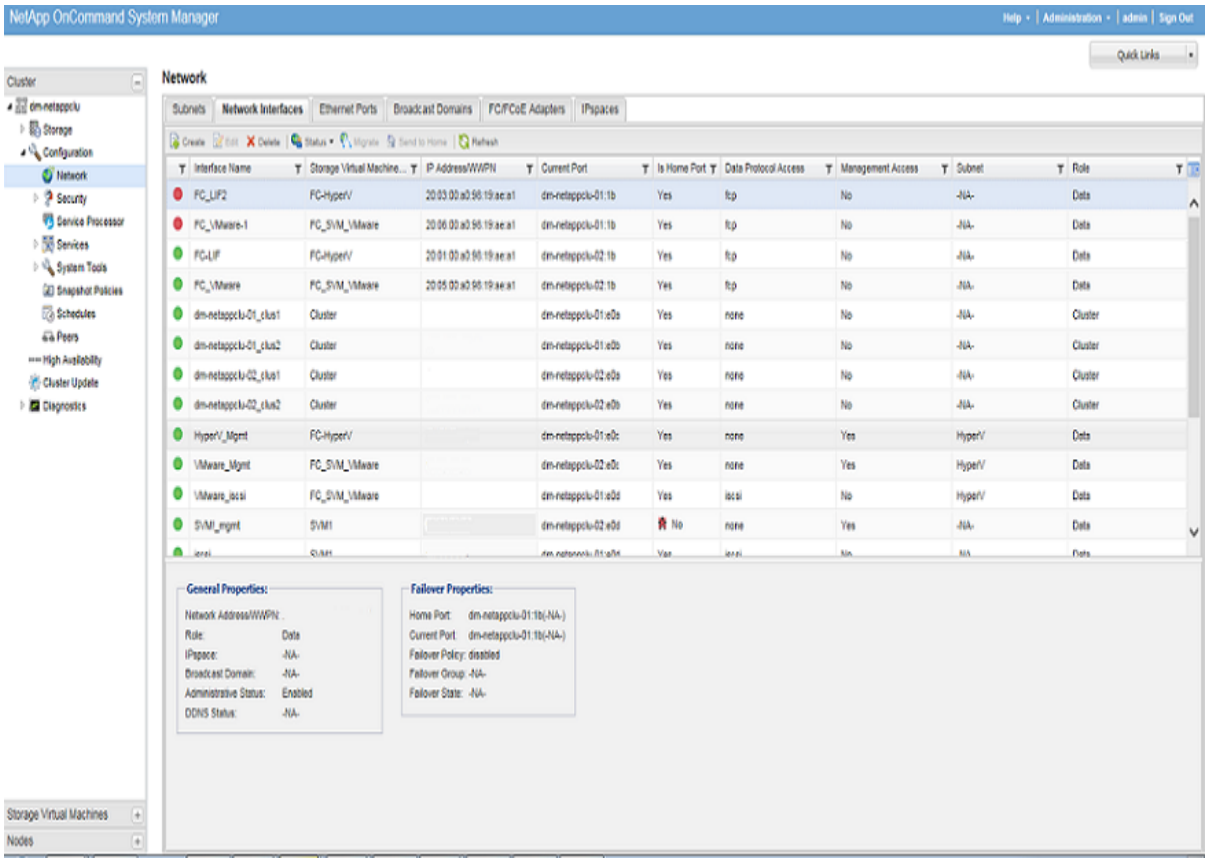
Note: Only applicable to the Cluster (SVM).

Follow these steps:

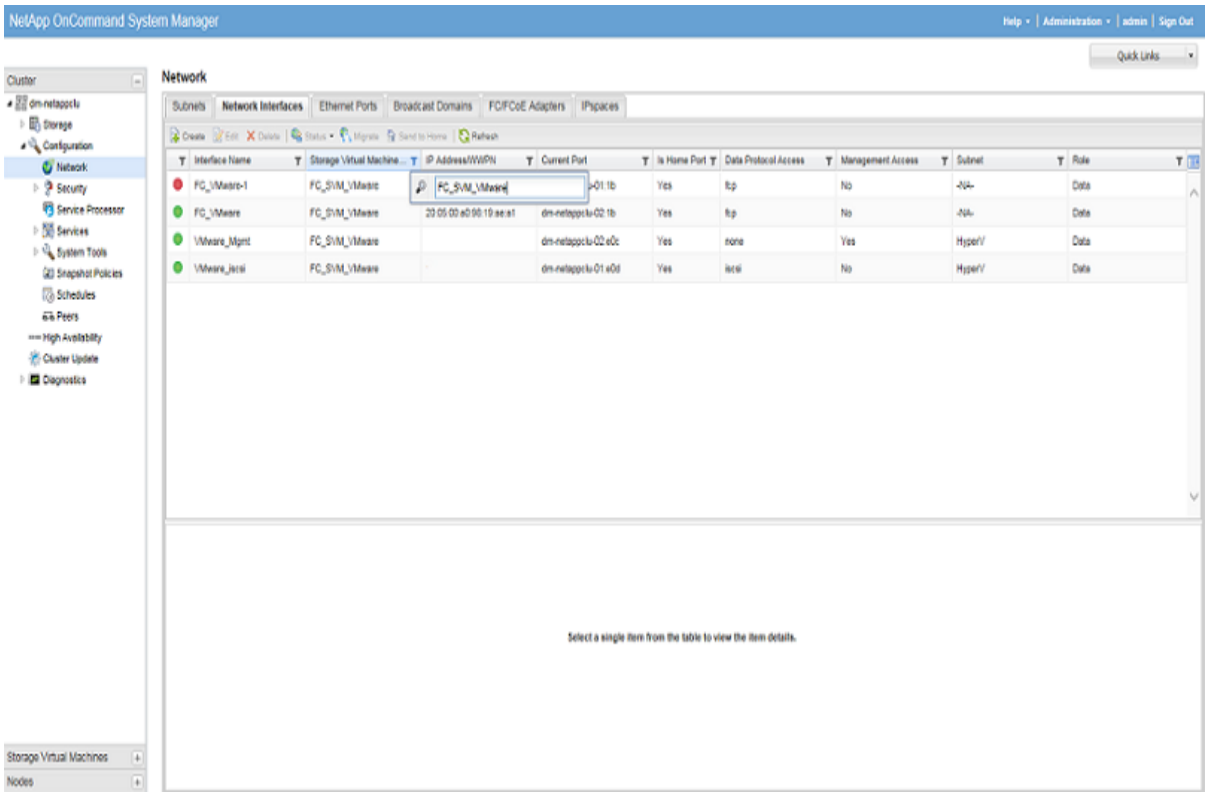
1. Log into the NetApp system manager and connect to the cluster with valid credentials.
2. Identify the SVM to add to the Arcserve UDP console.



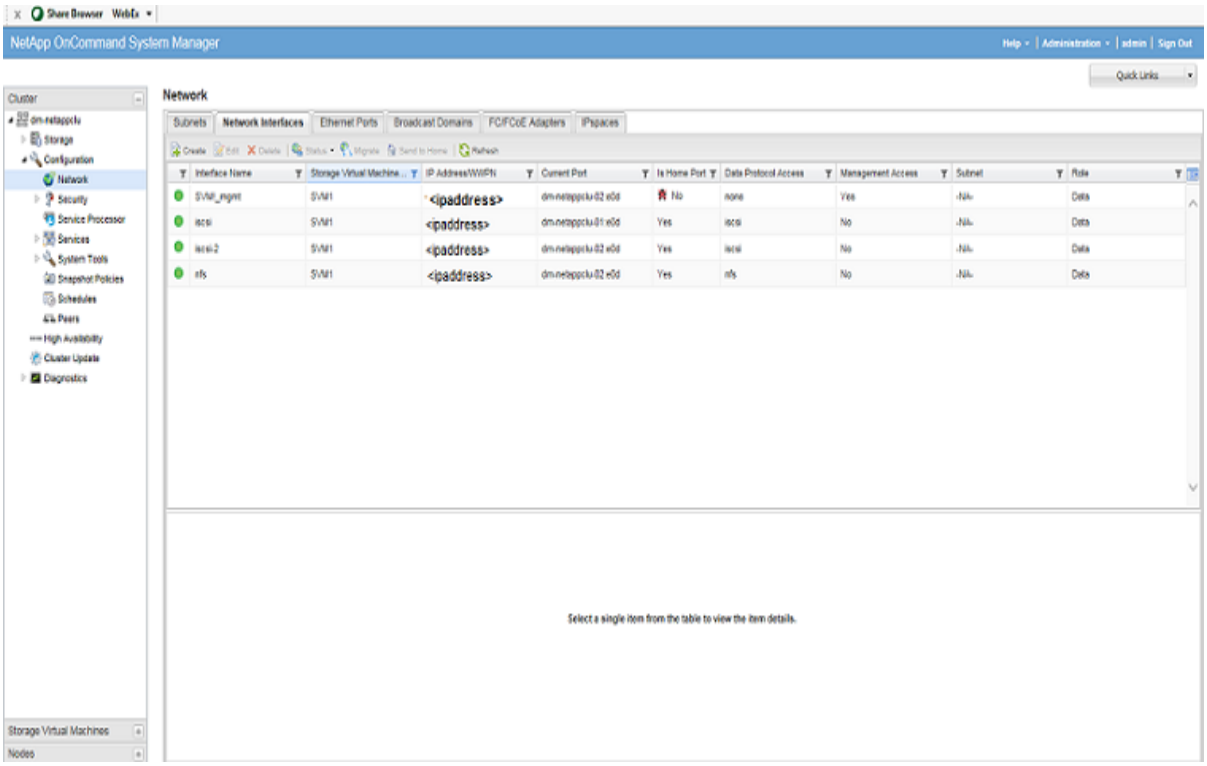
3. Change the View to Cluster.
4. From Cluster>Configuration, click Network, and then open Network Interfaces.

5. 

6. Filter the storage virtual machine by the SVM that you have identified.

7. 

8. Find out which interface has the Management Access IP address and the Data protocol access IP address, and provide them at Add Storage Array in the UDP Console.
   - Get Management Access IP address:
     - Locate the Interface that has the option **Yes**under the **Management Access** column.
     - From the interface, select the IP address available under **IP Address/WWPN**.
   - Get Data Protocol Access IP address:
     - Locate the Interface that has the option **iSCSI** or **NFS** under the **Data Protocol** column.
     - For the ISCSI or NFS interface, select the respective IP address available under **IP Address/WWPN.**



Note: The ESXi server must use the same Data protocol access IP address for read and write to the NetApp storage array for iSCSI or NFS.

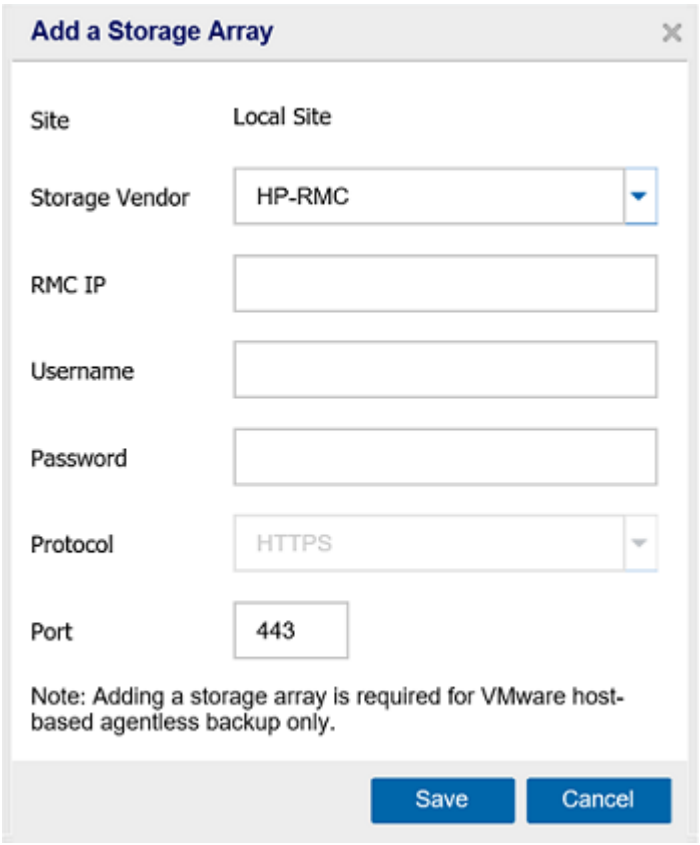
Add Details of HPE RMC Managing HPE 3PAR Store-serv storage array

You can add details of HPE RMC that manages an HPE 3PAR storeserve array.

Important! HPE 3PAR storeserve must be managed by HP RMC..

Follow these steps:

1. From the **Add a Storage Array** dialog, select HP-RMC in the **Storage Vendor** drop-down.



2. Provide the following details:

RMC IP

Specifies the HPE RMC management IP address.

Username

Specifies the user name to connect to the storage array.

Password

Specifies the password for the user name.

Protocol

Specifies the protocol to connect to the array.

Note: Only HTTPS is supported.

Port

Specifies the port number of the array.

Default: 443

3. Click **Save**.

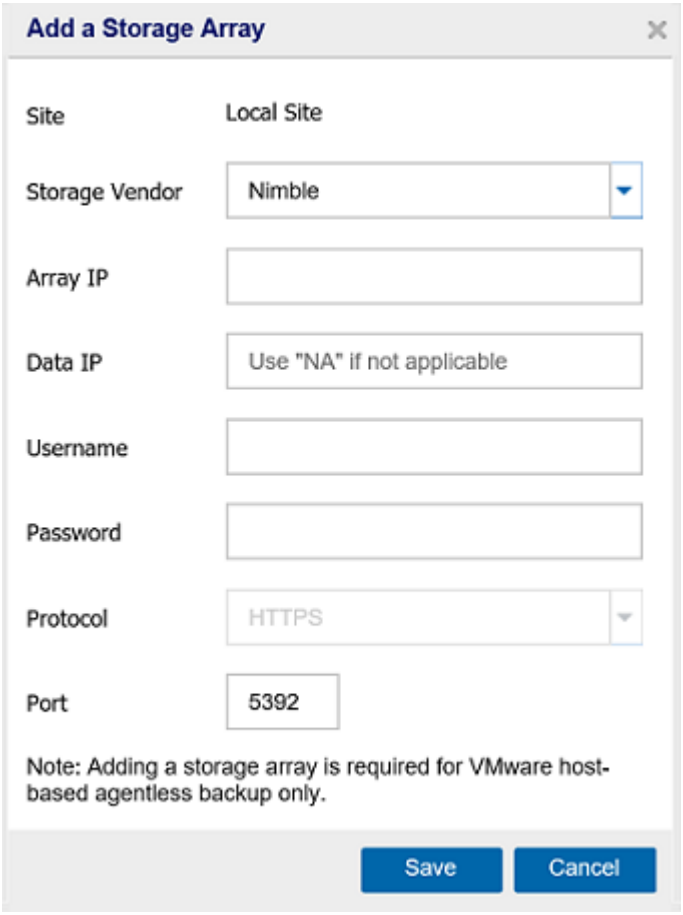
The storage array is added to the Console.

Add a Nimble Storage Array

You can add a Nimble storage array.

Follow these steps:

1. From the **Add a Storage Array** dialog, select Nimble from the **Storage Vendor** drop-down option.



2. Provide the following details for the specific type of array:

Array IP

Specifies the IP address of the array. Enter the Nimble storage array management IP address.

Data IP

Specifies the Data access IP (Logical Interface) of the nimble storage array, where the target machines (ESXi) can access this storage array for read and write operations. If the nimble storage array supports only the FC protocol, then enter NA.

Username

Specifies the user name to connect to the storage array.

Password

Specifies the password for the user name.

Protocol

Specifies the protocol to connect to the array.

Note: Only HTTPS is supported.

Port

Specifies the port number of the array.

Default: 5392

3. Click **Save**.

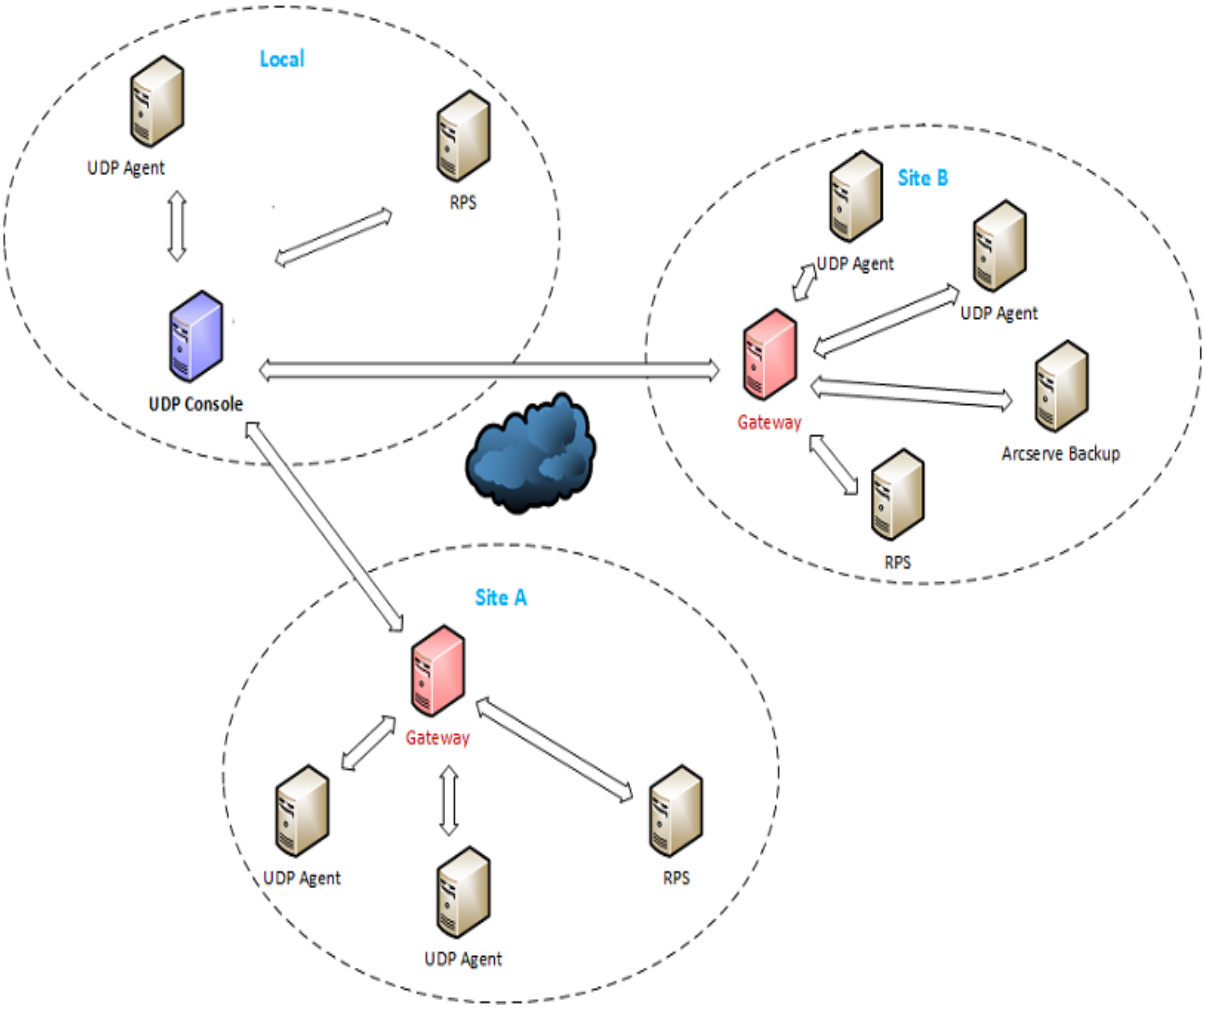
The storage array is added to the Console.

How to Add and Manage a Site

Arcserve UDP Console can manage remote nodes and recovery point servers from another subnet across a WAN. The remote nodes and server in a Site interact with the Console using a gateway. The gateway is installed in Sites. Although the Console cannot connect to the remote nodes directly, Arcserve UDP uses the gateway to establish a connection between the nodes and Console.

Important! If the recovery points are in a remote site, then to restore data, the Console must connect to the remote site using a VPN connection.

The following diagram shows the connection between local and sites.



After adding a site, you can modify, update, or delete sites from the Console. Also, you can manage the remote nodes from the Console.

Specify the Site Name

The Site Name page lets you specify a name for the site and select a heartbeat interval. This site name is displayed on the Console.

Follow these steps:

1. Click the **resources** tab on the Console.
2. From the left pane, navigate to **Infrastructures** and click **Sites**.

   The Infrastructure: Remote Sites page opens on the center pane.
3. Click **Add a Site**.

   The **Add a Site** wizard opens.
4. Provide the following details on the Site Page:

   Site Name

   Provide a name for the site.

   Heartbeat Interval

   Select a time interval for the heartbeat from the drop-down list. The heartbeat will check the connection between the Console and the site at specified interval.

   UDP Console Connection URL

   Specifies the URL of the UDP Console. The gateway server connects to this URL.

   Remember UDP Console Connection URL

   Select this check box to remember the Console URL.
5. Click **Next**.

   The Arcserve Remote Management Gateway Installation Instructions page opens.

   You have specified the name for the remote site.

Share the Registration Instruction

The Share the Registration Instruction includes the information to download and install the gateway. The download instructions are sent to the remote administrator. The remote administrator needs to use the download information to download and install the gateway.

Follow these steps:

1. Copy the instructions and save it for your reference.

   If you have not configured Arcserve UDP email, send the copied instructions to the remote administrator from another email server.

2. Select **Send instructions using Arcserve UDP email** and provide the email address of the remote administrator.

3. Click **Next**

   The **Confirmation** page opens.

   You have successfully shared the registration instruction.

Verify and Add the Site

Verify the details before adding the site. You can click **Previous** to visit the previous pages.

Follow these steps:

1. Verify the details on the confirmation page.

   You can modify any information on the previous pages if required.

2. Click **Finish**.

   The wizard closes.

   The remote site is created on the **Infrastructure: Sites** page.

   When the remote administrator installs Arcserve Remote Management Gateway and successfully provides the gateway authorization code, a green check mark is displayed beside the site name on the Console.

   You have successfully added a site.

Modify the Console URL

When you change the Console URL, the gateway must be registered again to manage the site from the Console. The gateway registration is performed by the remote administrator. When you update the Console URL, Arcserve UDP sends an email to the remote administrator with detailed instructions on how to register the gateway to the Console.

Follow these steps:

1. From the Console, click **resources**.
2. Navigate to **Infrastructures** on the left pane and click **Sites**.

   The **Infrastructure: Site** page opens on the center pane.
3. Select the site and click **Actions**, **Update Console URL**.

   The **Update Console URL** dialog opens.
4. Specify the new URL of the Console.
5. Click **Send**.

   An email is sent to the remote administrator. When the remote administrator updates the new URL on the gateway server, the Console displays a green check mark beside the updated site.

   You have successfully modified the Console URL.

Modify the Site

You can modify the site to change any of the parameters of the added site. For example, you can rename the site or change the heartbeat interval.

Follow these steps:

1. From the Console, click **resources**.
2. Navigate to **Infrastructures** on the left pane and click **Sites**.

   The **Infrastructure: Site** page opens on the center pane.
3. Select the site and click **Actions**, **Modify**.

   The **Modify a Site** wizard opens.
4. Use **Previous** and **Next** to visit any pages and modify any parameter.
5. Click **Finish**.

   The **Modify a Site** wizard closes.

   You have successfully modified the site.

Delete a Site

You can delete a site that you do not want to manage. Before deleting a site, all nodes and node discovery filters related to this site must be deleted first.

Follow these steps:

1. From the Console, click **resources**.
2. Navigate to **Infrastructures** on the left pane and click **Sites**.

   The **Infrastructure: Site** page opens on the center pane.
3. Select the site and click **Actions**, Delete.

   The **Confirm** dialog opens.
4. Click **Yes**.

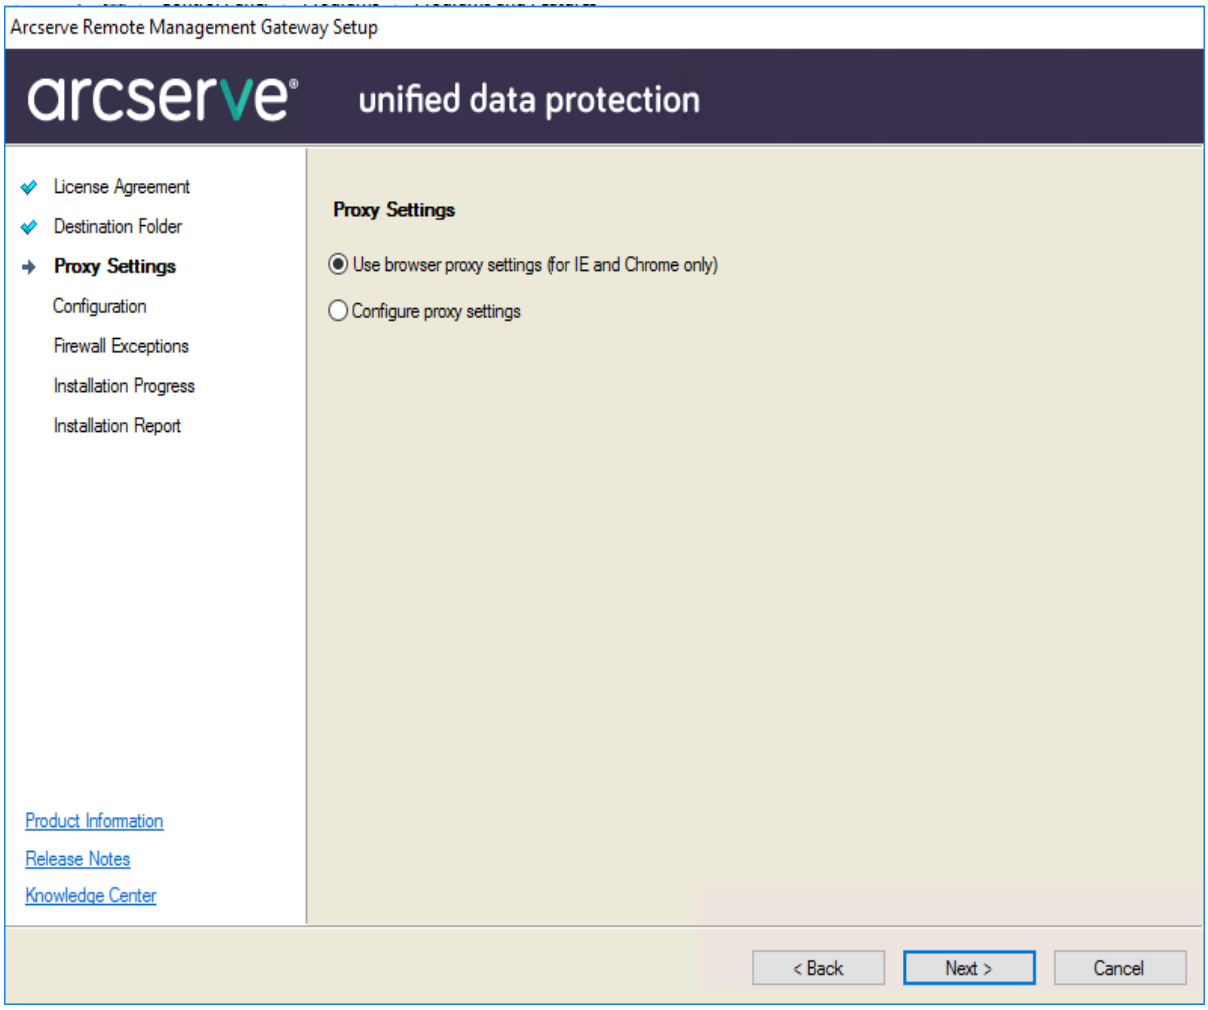
   The site is deleted.

Set up a Proxy Server for the Gateway

Arcserve UDP supports installing a proxy server on the Gateway machine. Gateway uses this proxy setting to communicate with its registered Console.

Follow these steps:

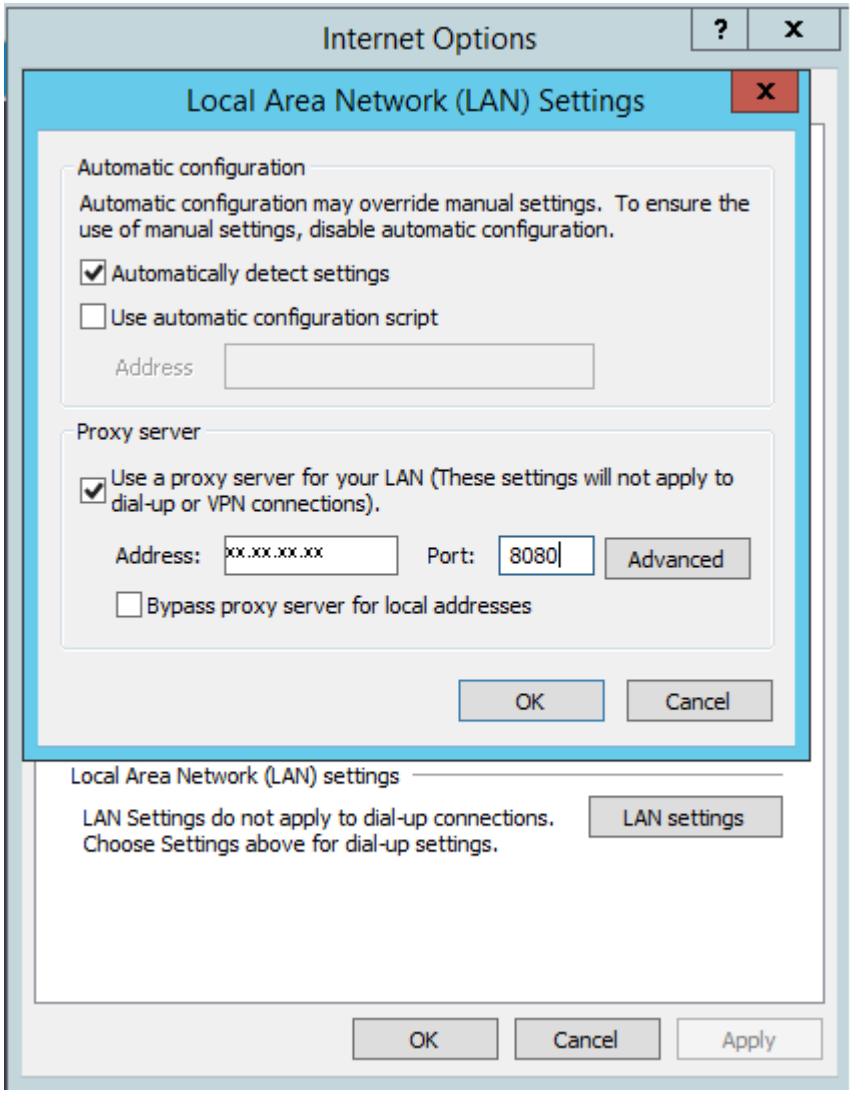
1. Open the Arcserve Remote Management Gateway Setup wizard.



2. On the Proxy Settings dialog, select one of the two options:

   Use browser proxy settings (for IE and Chrome only)

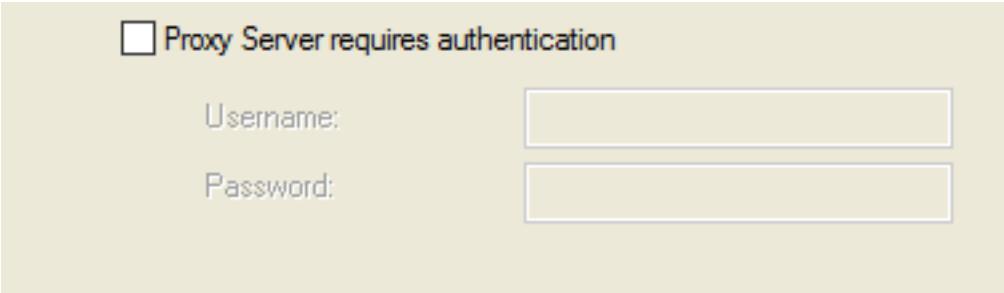
   Specifies to use the proxy settings of the browser. You must update the browser proxy settings. Open the browser and click **Options**, **Connection**, **LAN setup**.



Configure proxy settings

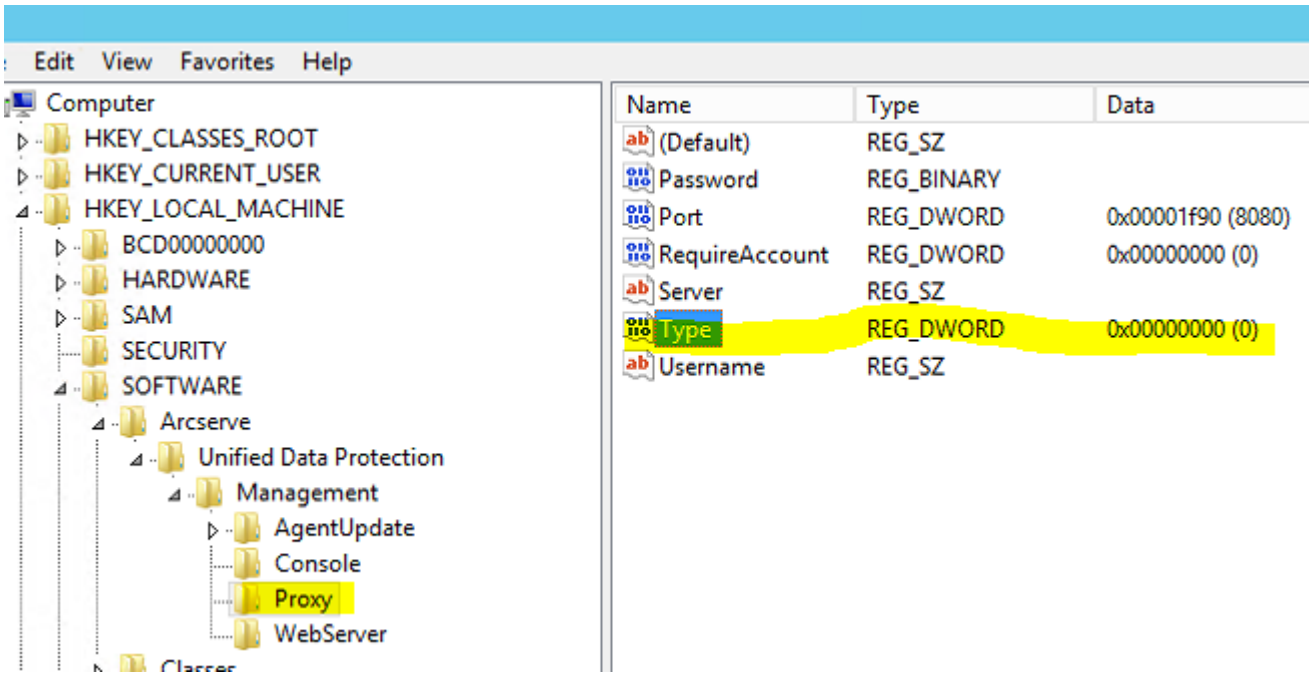
Specifies that you have to provide the proxy server details on the wizard page.

3. Clear the authentication check box if the proxy does not support credentials.



The proxy server is set up for the gateway.

4. To verify the proxy setting, open regedit, and navigate to Proxy.



On the Type field, 0 indicates browser settings and 1 indicates other settings.

Note: Use regedit to modify any of the proxy settings such as port, server IP address, or type.

Managing Exchange Online Nodes

You cannot add the Exchange Online nodes to all the nodes view directly. Once you add to the backup plan, then the Exchange Online node is added to all the nodes view directly.

You can perform the following tasks:

1. Add an Exchange Online Node
2. Update an Exchange Online Node
3. Delete an Exchange Online Node
4. Public folder Mailbox Support for Exchange Online Protection

Add an Exchange Online Node

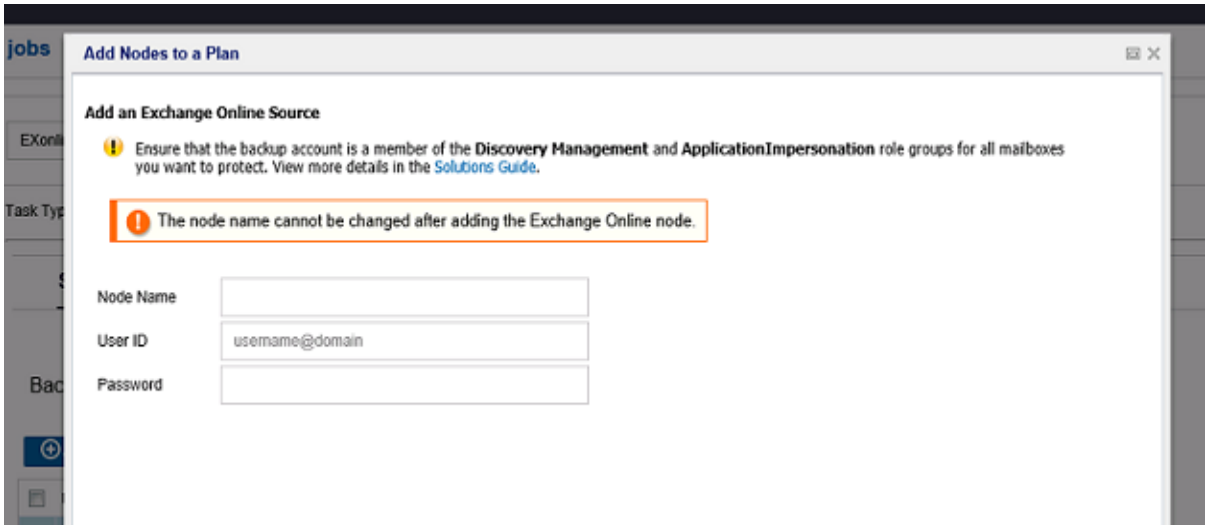
Note: Unlike other nodes, you cannot add the Exchange Online node from All Nodes page. An Exchange Online node is added only in a plan or while modifying a plan.

For information about how to create a plan, see How to Create an Exchange Online Plan.

Follow these steps:

1. Enter the user name of Exchange Online backup account that meets the pre-requisites in Username or user account for the exchange online node.

   Multiple Exchange online nodes can use the same user account (service account) of Exchange online. To add Exchange node by plan, specify the node name, user name, and password. Once created, you cannot change the node name of Exchange online node.



   Note: Updating / changing the user account may change the number of protected mailboxes. Verify that the new / updated service account have impersonation rights for the mailboxes to be protected.

2. Enter password and click **Connect**.
3. Select Exchange Online accounts to protect and click the right arrow (>) to move them to the protected list.

   Note: Select the check box to protect all Office 365 Exchange sources to protect all the Exchange Online accounts across all the pages. To add all the exchange online accounts listed on the page to the protected list, click the right (>) arrow.

4. Click **Save**.

You may select the folder(s) you do not want to backup from the Select folders to exclude from Backup option on the Source tab.

The selected Exchange Online accounts are added.

Update an Exchange Online Node

When you modify the credentials of the Exchange Online node, update the Exchange Online node on Console.

Follow these steps:

1. Click the **resources** tab on the Console.
2. From the left pane, navigate to **Nodes** and click **Exchange Online Nodes**.

   The created Exchange Online nodes are displayed on the center pane.
3. Select the Exchange Online node and click **Actions**.
4. Click **Update**.

   The **Update Node** dialog appears.
5. Modify the details as desired, and click **OK**.

   The Exchange Online node is updated.

Delete an Exchange Online Node

Delete any Exchange Online node that you do not require anymore.

Follow these steps:

1. Click the **resources** tab on the Console.
2. Navigate to **Nodes** and click **Exchange Online Nodes**.

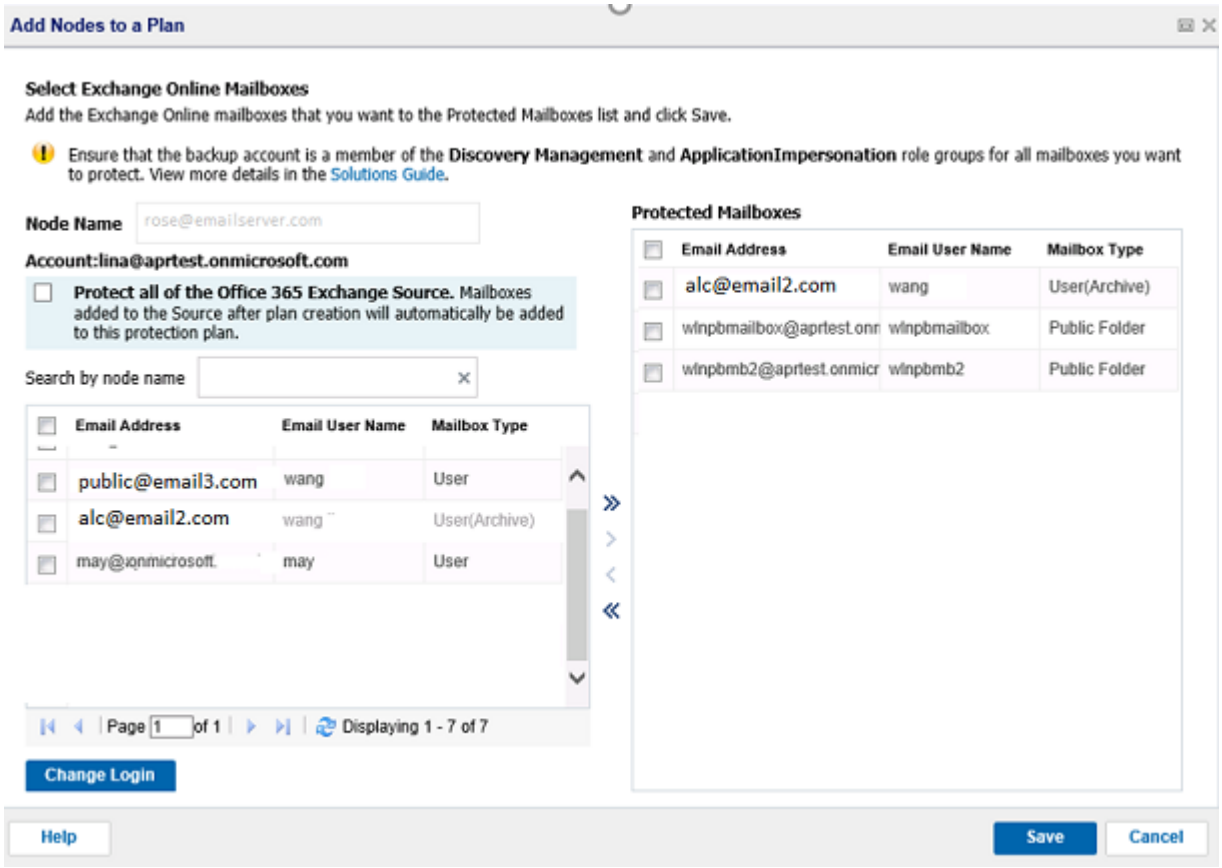
   The created Exchange Online nodes are displayed on the center pane.
3. Select the Exchange Online node and click **Actions**.
4. Click **Delete**.

   A confirmation dialog appears.
5. Click **OK**.

   The Exchange Online node is deleted successfully.

Public Folder Mailbox Support for Exchange Online Protection

When adding the Exchange online node, you can see mailbox Type column in **Add Nodes to a Plan** window. Public folder mailbox appears in this window and in the column appears as **public folder**. To protect public folder, add the corresponding public mailbox into **Protected Mailboxes**, and then save the plan.



Note: For more information about Exchange Online public folder mailbox and permission, refer to link 1 and link 2.

Manage SharePoint Online Nodes

You cannot add the SharePoint Online nodes to all the nodes view directly. Once you add to the backup plan, then the SharePoint Online node is added to all the nodes view directly.

You can perform the following tasks:

1. Add a SharePoint Online Node
2. Update a SharePoint Online Node
3. Delete a SharePoint Online Node

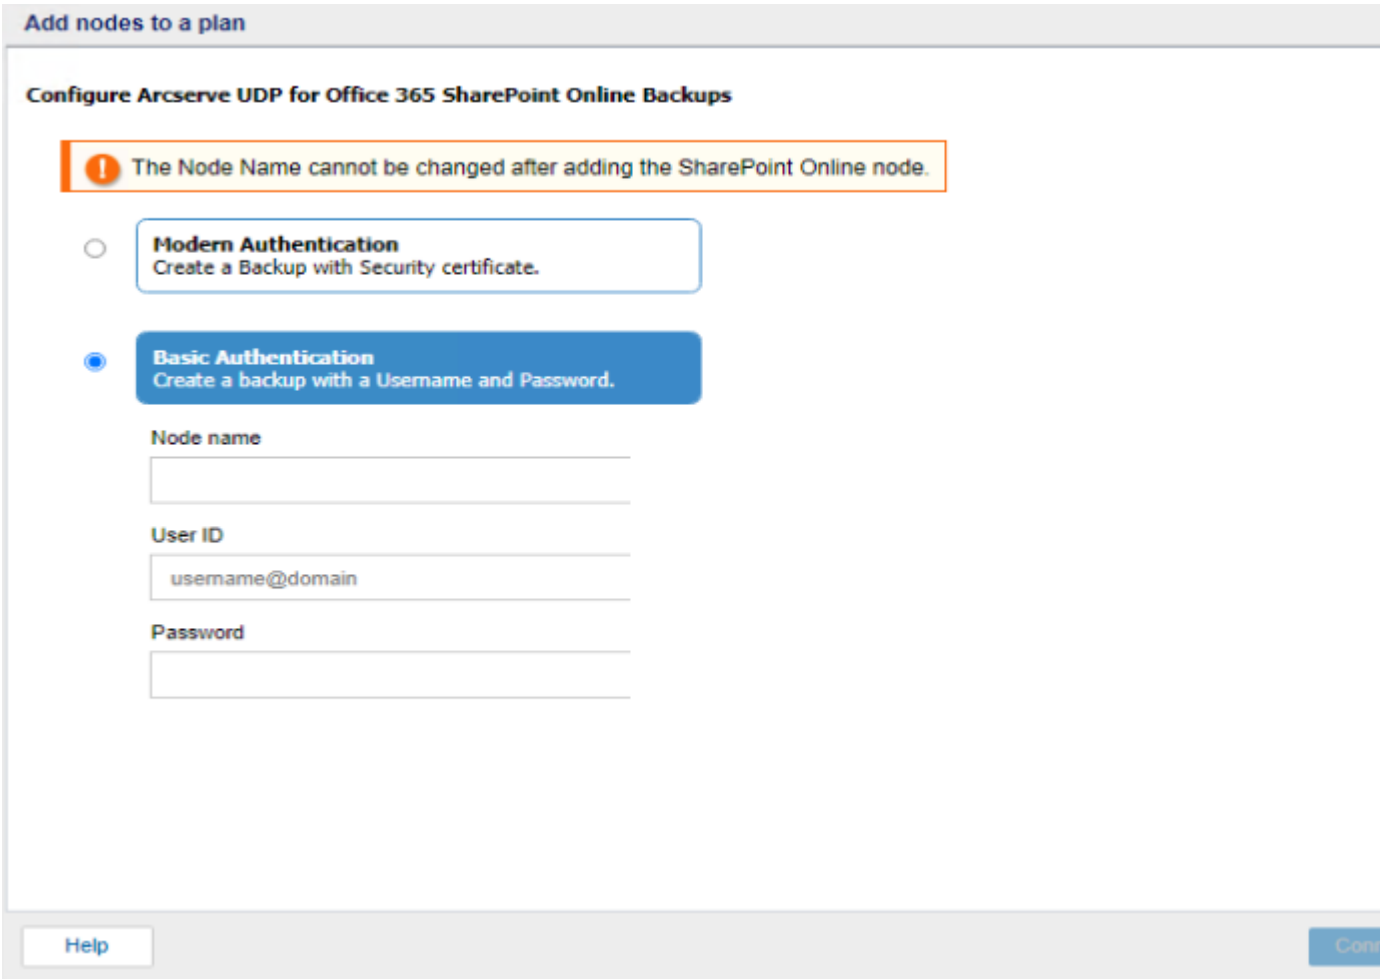
Add a SharePoint Online Node

Note: Unlike other nodes, you cannot add the SharePoint Online node from the All Nodes page. A SharePoint Online node is added only in a plan while creating or modifying a plan.

For information about how to create a plan, see How to Create a Sharepoint Online Plan.

Follow these steps:

1. Enter the user name of SharePoint node name, and enter the Site collection URL, site owner user name, and password.



2. Enter password and click **Connect**.
3. Select SharePoint list/Library or list items to protect.
4. Click **Save**.

   The selected SharePoint Online accounts are added.

Update a SharePoint Online Node

When you modify the credentials of the SharePoint Online node, update the SharePoint Online node on Console.

Follow these steps:

1. Click the **resources** tab on the Console.
2. From the left pane, navigate to **Nodes** and click **SharePoint Online Nodes**.

   The created SharePoint Online nodes are displayed on the center pane.
3. Select the SharePoint Online node and click **Actions**.
4. Click **Update**.

   The **Update Node** dialog appears.
5. Modify the details as desired, and click **OK**.

   The SharePoint Online node is updated.

Delete a SharePoint Online Node

Delete any SharePoint Online node that you do not require anymore.

Follow these steps:

1. Click the **resources** tab on the Console.
2. Navigate to **Nodes** and click **SharePoint Online Nodes**.

   The created SharePoint Online nodes are displayed on the center pane.
3. Select the SharePoint Online node and click **Actions**.
4. Click **Delete**.

   A confirmation dialog appears.
5. Click **OK**.

   The SharePoint Online node is deleted successfully.

Manage OneDrive Node

You cannot add the OneDrive nodes to all the nodes view directly. Once you add to the backup plan, then the OneDrive node is added to all the nodes view directly.

You can perform the following tasks:

1. Add a OneDrive Node
2. **Delete the Node:** From the option of Actions on OneDrive node, click **Delete** and provide confirmation on pop up to delete the node from Console.
3. **Export the node:** From the option of Actions on OneDrive node, click **Export** and receive Node as zip file.
4. **Modify Plan:** From the option of Actions on OneDrive node, click **Modify Plan** to open related plan and update as required.
5. **Log into Agent:** From the option of Actions on OneDrive node, click **Log into Agent** to perform agent specific tasks, such as Restore using Mount Volume Option.
6. **Backup Now:** From the option of Actions on OneDrive node, click **Backup now** to perform manual backup.
7. **Restore:** From the option of Actions on OneDrive node, click **Restore** to login agent and perform restore.
8. Download Recovery Point from Cloud
9. Update Recovery Point to Cloud
10. Copy Recovery Point to local disk or network share

Add a OneDrive Node

Note: Unlike other nodes, you cannot add the OneDrive node from the All Nodes page. A OneDrive node is added only in a plan while creating or modifying a plan.

For information about how to add the node, view "Specify the Source" on page 992 while creating OneDrive Backup plan.

Delete a OneDrive Online Node

Delete any OneDrive Online node that you do not require anymore.

Follow these steps:

1. Click **Delete** from the option of **Actions** on OneDrive node.
2. Provide confirmation on pop up to delete the node from Console.

   The OneDrive Online node is deleted successfully.

Chapter 9: Adding and Managing Destinations

This section contains the following topics:

How to Add a Destination

A destination is a location where you store your backup data. Arcserve UDP lets you add multiple type of destinations.

What To Do Next:

1. Add a Destination Using RPS
2. Add Arcserve Backup Servers
3. Add a Remote Console
4. Add a Cloud Account
5. Add Arcserve Cloud Account

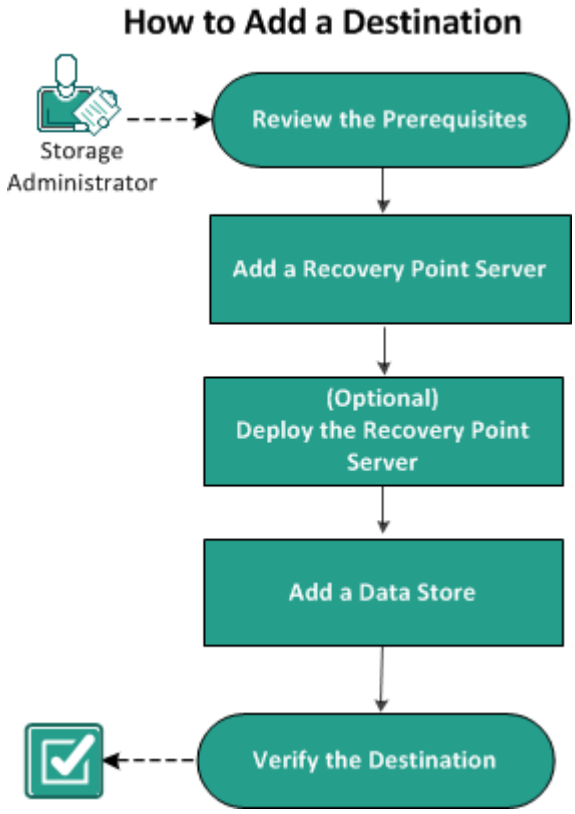
How to Add a Destination

For Arcserve UDP, you can assign a recovery point server (RPS) as a central destination. You can store data from multiple nodes in a recovery point server and then recover data when necessary. Adding a destination primarily involves two steps:

a. Adding a recovery point server to the Console.

b. Adding a data store to the recovery point server.

The following diagram illustrates how to add a destination:



What To Do Next?

1. Review the Prerequisites
2. Add a Recovery Point Server
3. (Optional) Deploy the Recovery Point Server
4. Add a Data Store
5. Verify the Destination

Review the Prerequisites

Before you set up a recovery point server, complete the following prerequisites:

- Review the Release Notes for a description of system requirements, supported operating systems, and a list of issues that are known to exist with this release of Arcserve Unified Data Protection.
- Verify that you have administrator privileges to install Arcserve Unified Data Protection.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Add a Recovery Point Server

Adding a destination starts with addition of a recovery point server to the Console. Later, you add data stores to the RPS.

Follow these steps:

1. Log into Arcserve Unified Data Protection, and click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Recovery Point Servers**.

   The **Destinations: Recovery Point Servers** page is displayed in the center pane.
3. Click **Add a Recovery Point Server**.

   The **Add a Recovery Point Server** page is displayed.
4. Enter the following details:

   Node Name/IP Address

   Defines the node name of the recovery point server that you want to add to the Console.

   Username and Password

   Defines the user name and password that helps you log in to the node.

   Note: Use one of the following formats for the user name: Computer name, domain name/username, or username.

   Description

   (Optional) Defines any additional information about the node.
5. Enter the following fields for the **Installation Settings**:

   Note: If the node already has Recovery Point Server installed, ignore these installation settings.

   Installation Location

   Specify the location where you want to install the recovery point server. You can accept the default path or can specify an alternative path.

   Port

   Specifies the port number that connects to the web-based UI.

   Default: 8014.

   Protocol

   Specify the protocol that you want to use to communicate with the destination server. The available selections are HTTP and HTTPS.

   Note: For a more secure communication, select the HTTPS protocol.

Change Tracking Driver

Specify if you want to **Install Agent Change Tracking Driver**.

6. Schedule the installation or upgrade by selecting one of the options from **Start Time to Install or Upgrade**.

   Note: If the server already has Recovery Point Server installed, ignore these settings.

7. Click **Save**.

   The deployment progress is displayed in the right pane. The recovery point server is added.

   Now, the recovery point server is deployed. You can add data stores after the recovery point server is added.

(Optional) Deploy the Recovery Point Server

Using Arcserve Unified Data Protection, you can discover and deploy the latest version of the RPS component to recovery point servers. After you deploy the RPS component, the node is ready to store the backup sessions and serve as a recovery point server.

Note: The RPS components are installed with the Arcserve Unified Data Protection installation.

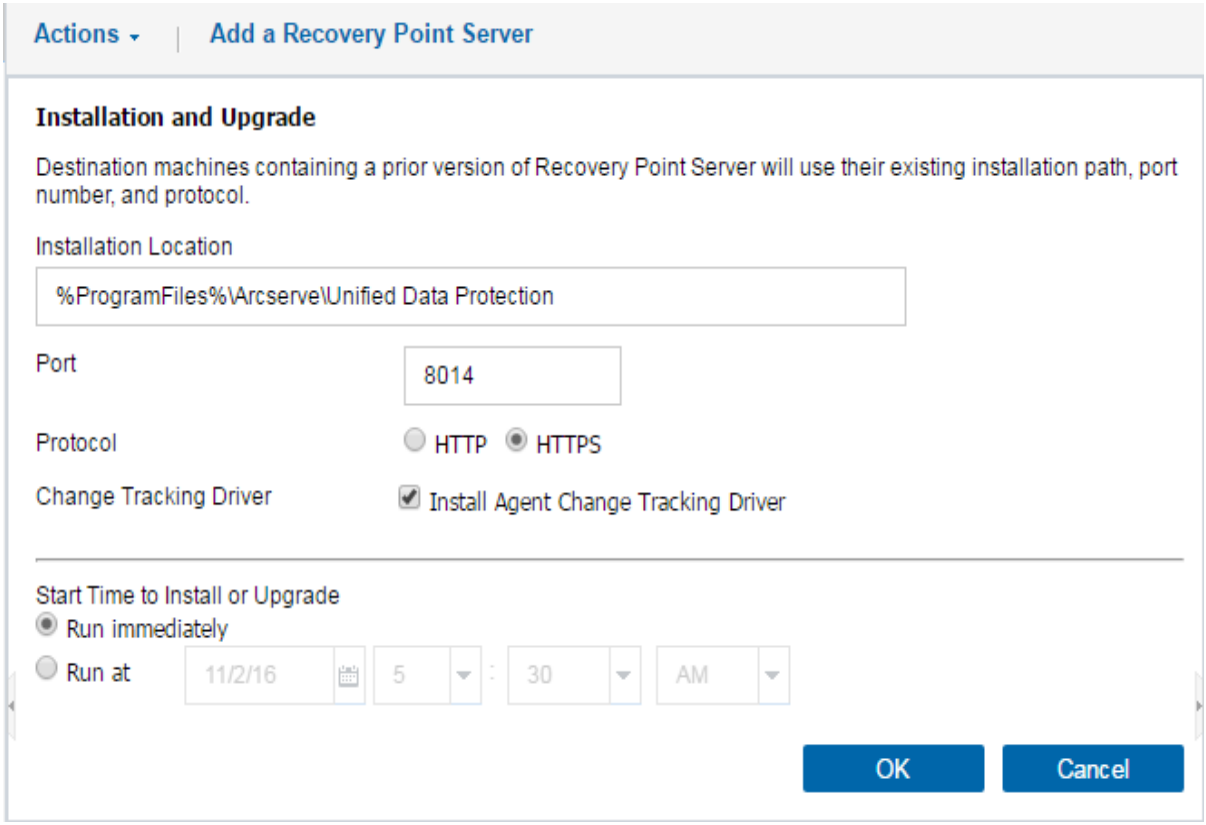
Follow these steps:

1. Click the **resources** tab.
2. From the left pane navigate to **Destinations**, and click **Recovery Point Servers**.

   The **Destination: Recovery Point Server** page is displayed.
3. Perform one of the following actions:
   - Right-click a recovery point server.
   - Select a recovery point server, and from the center pane click the **Actions** drop-down list.

   A list of options is displayed.
4. Click **Install/Upgrade Recovery Point Server**.

   The **Installation and Upgrade** page is displayed.



5. Modify the deployment settings, and click **OK** to deploy the recovery point server on the selected node.

   The recovery point server deployment starts. You can view the deployment progress on the right pane.

Add a Data Store

To create the destination, the recovery point server needs data stores. The data store specifies where the backup data is stored. You can add multiple data stores to an RPS.

Follow these steps:

1. Click the **resources** tab.
2. From the left pane navigate to **Destinations**, and click **Recovery Point Servers**.

   The **Destinations: Recovery Point Servers** page is displayed.
3. Perform one of the following actions:
   - Right-click a recovery point server.
   - Select a recovery point server, and from the center pane click the **Actions** drop-down list.

   A list of options is displayed.
4. Click **Add a Data Store**.

   The **Create a Data Store** page is displayed with the name of the specified recovery point server.
5. Specify the following fields and click **Save**.

   Recovery Point Server

   Defines the recovery point server where the data store is created. The recovery point server is already added by default.

   Data Store Name

   Defines the name of the data store.

   Data Store Folder

   Defines the location of the folder where the data store is created. Click Browse to select the destination folder.

   Notes:

   - You can now enable and view hidden volumes. For more information, see How to Enable and View Hidden Volumes while Creating a Data Store.
   - For non-deduplication and deduplication data store, the destination path should be an empty folder.

   Concurrent Active Nodes Limit to

Specifies the maximum concurrent jobs on the data store.

Default Value: 4

Refers to a value from 1 to 9999. The value indicates the number of jobs that can concurrently run. If the running jobs meet the number, another job is placed in to the queue and job can only start when one of the running job completes. The completed job could mean a finished, canceled, or a failed job.

The number applies to the Job Types but not to the Server nodes. For example, number 5 indicates that five backup jobs are running. Any job scheduled after five backup jobs waits in the queue, but you can submit another job such as File System Catalog.

If the value is more than 16 or 32, messages are displayed to warn about the increased demand on hardware.

Note: Limit to number only impacts the replication outbound job, not the replication inbound job. Limit to number does not impact the Restore or BMR jobs. Such jobs are not placed in a queue.

Enable Deduplication

Specifies that deduplication is enabled for this data store. Arcserve UDP supports both types of deduplication: Source-side deduplication and Global deduplication. Source-side deduplication prevents duplicate data blocks to move across network from a particular agent. Global deduplication eliminates duplicate data across all client machines based on the volume cluster level.

Deduplication Block Size

Defines the deduplication block size. The options are 4 KB, 8 KB, 16 KB, 32 KB, and 64 KB. The deduplication block size also impacts the Deduplication capacity estimation. For example, if you change the default 16 KB to 32 KB, the Deduplication capacity estimations double. Increasing the deduplication block size can decrease the deduplication percentage.

Hash Memory Allocation

Specifies the amount of physical memory that you allocate to keep hashes. This field is pre-filled with a default value. The default value is based on the following calculation:

If the physical memory of the RPS is smaller than 4 GB (or is identical to 4 GB), the default value of **Hash Memory Allocation** is identical to the physical memory of the RPS.

If the physical memory of the RPS is greater than 4 GB, Arcserve UDP calculates the available free memory at this time. Assume that the available

free memory is X GB at present. Arcserve UDP further checks the following conditions:

- If (X * 80%) > = 4 GB, the default value of **Hash Memory Allocation** is (X * 80%).
- If (X * 80%) < 4 GB, the default value of **Hash Memory Allocation** is 4 GB.

  Example: Consider the RPS has 32 GB of physical memory. Assume that operating system and other applications use 4 GB memory while creating the data store. So, the available free memory at this time is 28 GB. Then, the default value of **Hash Memory Allocation** is 22.4 GB (22.4 GB = 28 GB * 80%).

Hash Destination is on a Solid State Drive (SSD)

Specifies if the hash folder is on a solid state drive.

Note: Configure the hash destination on local SSD, if the Hash destination is on a Solid State Drive(SSD) option is enabled

Data Destination

Defines the data destination folder to save the actual unique data blocks. Use the largest disk to store data as that contains the original data blocks of source.

Note: The **Data Destination** path should be a blank folder.

Index Destination

Defines the index destination folder to store the index files. Choose a different disk to improve the deduplication processing.

Note: The **Index Destination** path should be a blank folder.

Hash Destination

Defines the path to store the hash database. Arcserve UDP uses the SHA1 algorithm to generate the hash for source data. The hash values are managed by the hash database. Selecting a high speed Solid State Drive (SSD) increases the deduplication capacity and requires a lower memory allocation. For better hash performance, we recommend to format the SSD volume as NTFS file system with 4KB volume cluster size.

Note: The **Hash Destination** path should be an empty folder.

Note: You cannot specify the same path for the following four folders: **Data Store folder, Data Destination, Index Destination,** and **Hash Destination**.

Enable Compression

Specifies that the data compression settings are enabled.

Compression Type

Specifies whether to use the standard or maximum compression type.

Compression is often selected to decrease the usage of the disk space, but also has an inverse impact on your backup speed due to the increased CPU usage. Based on your requirement, you can select one of the three available options.

Note: For more information, see Compression Type.

Enable Encryption

Specifies that encryption settings are enabled. When you select this option, you must specify and confirm the encryption password.

Data encryption is the translation of data into a form that is unintelligible without a deciphering mechanism. The Arcserve Unified Data Protection solution uses secure, AES (Advanced Encryption Standard) encryption algorithms to achieve maximum security and privacy of your data. For data stores, encryption or No encryption is supported. For Encryption, only AES-256 is available.

A password is not required when you are attempting to restore to the computer from which the backup was performed. However, when you attempt to restore to a different computer, a password is required. By default, only for the first login password is required. To enter password even after the first login, the administrator needs to manually stop Arcserve UDP Agent Explorer Extension Service.

Send an email alert when a destination is nearing full capacity

Selecting this option configures the data store to send email alert. RPS sends out email alerts to recipients when data store destination folder is nearing full capacity.

Configure Email

This button appears only when you enable the option of *Send an email alert when a destination is nearing full capacity*. The button helps you provide as email ID to receive alerts. Click the **Configure Email** button to load global email alert settings from Console>Settings>Email and Alert configuration. If global Email settings is not available, clicking the **Configure Email** button opens **Email settings** dialog to set email details.

The data store is created and gets displayed on the center pane. Click the data store to view the details in the right pane.

Various States of Data Store

The data store displays different status depending on the task performed by the data store. When you select a data store from the **resources** tab, the data store status is displayed on the right pane

- **Stopped:** The data store is inactive. You cannot submit any job in this state.
- **Starting:** The data store is starting. When the data store is getting started, the progress is displayed on the Console.
- **Running:** The data store is active. You can submit jobs in this state.
- **Stopping:** The data store is stopping. When the data store is stopping, the progress is displayed on the Console.
- **Modifying:** The data store is getting updated with the new data. When the data store is getting modified, the progress is displayed on the Console.
- **Deleting:** The data store is getting deleted. When the data store is getting deleted, the progress is displayed on the Console.
- **Out of Service:** The data store is not functioning properly. You cannot submit any jobs in this state. Stop the data store and verify the reason for this behavior. The following cases can result in the Out of Service status of a data store:
  - The data store backup destination cannot be accessed.
  - The configurations in registry or file are corrupted.
  - The GDD index or data role has internal errors.
  - The GDD index or data role process is manually stopped.
- **Restore Only:** In the Restore Only state, any jobs that require to write data to the data store does not run. Jobs such as backup, replication (in) job, jumpstart (in), data migration job. All others jobs run, which require to read data from the data store. The data store status changes to Restore Only in the following conditions:
  - When the hash role process is manually stopped.
  - When the backup destination/data/index/hash path volume capacity/assigned hash memory reaches its maximum limit.

Important! When the status of the data store is Restore only (Degraded State) or Out of service (Bad State), the data store does not function properly. You must stop the data store and verify the root cause for the status. For example, the problem may be the data deduplication volume has reached its maximum. After you resolve the root cause, start the data store and resubmit the backup job.

Verify the Destination

After completing all the procedures involved in adding an RPS, verify if the RPS is added successfully.

Follow these steps:

1. Click the **resources** tab.
2. From the left pane navigate to **Destinations**, and click **Recovery Point Servers**.

   The **Destinations: Recovery Point Servers** page is displayed.
3. Verify the following details:
   - The RPS that you created is displayed.
   - The data stores are displayed under the RPS.

Add Arcserve Backup Servers

Add an Arcserve Backup Server to archive data to a tape. When you create a plan to archive data to a tape device, you can use this destination.

Follow these steps:

1. Log in to Arcserve Unified Data Protection, and click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Arcserve Backup Servers**.

   The **Destinations: Arcserve Backup Servers** page is displayed in the center pane.
3. Click **Add an Arcserve Backup Server**.

   The **Add an Arcserve Backup Server** page is displayed.
4. Enter the following details:

   Node Name/IP Address

   Specifies the node name or the IP address of the Arcserve Backup server.

   Authentication Type

   Specifies the type of authentication used to log in to the Arcserve Backup server. The following two options are available:

   Windows Authentication

   Specifies that the Windows authentication is used to log in to the Arcserve Backup server.

   Note: The Windows user must be registered in Arcserve Backup first using Arcserve Backup User Profile Manager.

   Arcserve Backup Authentication

   Specifies that the Arcserve Backup authentication is used to log in to the Arcserve Backup server.

   Username and Password

   Specifies the user name and its password that helps you log in to the node.

   Note: Use one of the following formats for the user name: Computer name, domain name/username, or username.

   Port

   Specifies the port number that is used to connect to the Arcserve Backup server.

Note: Arcserve UDP uses the port number to connect to both the servers, the Arcserve Backup Primary server and the Member server in the Arcserve Backup domain.

5. Click **Save**.

   The Arcserve Backup Server is added to the Console.

   After adding the Arcserve Backup Server to the Console, you can navigate to **resources, Destination, Arcserve Backup Servers** and check detailed information of the tape media.

Add a Remote Console

Add a remote console to replicate recovery points to a remotely managed recovery point server. You can use the remote console to create a Replication Plan to Send Data to the Destination Console. You can also manage the added remote console.

Follow these steps:

1. Log into Arcserve Unified Data Protection, and click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Remote Console**.

   The **Destinations: Remote Console** page is displayed in the center pane.
3. Click **Add a Remote Console**.

   The **Add a Remote Console** page is displayed.
4. Enter the following details:

   Remote Console

   Refers to the URL of the remote console account that you want to add to the Console.

   Username

   Refers to the user name that helps you log into the Remote Console.

   Note: Use one of the following formats for the user name: Computer name, domain name/username, or username.

   Password

   Refers to password for the user name.

   Port

   Refers to the port number port number that connects to the web-based UI.

   Default: 8015.

   Protocol

   Specify the protocol that you want to use to communicate with the destination server. The available selections are HTTP and HTTPS.

   Note: For a more secure communication, select the HTTPS protocol.

   Proxy Settings

   Specifies the proxy server settings. Select the check box of **Connect using a proxy server** if you want to enable this option. If you select this option, you must also include the IP address (or machine name) of the proxy server and the corresponding port number that is used by the proxy server for internet

connections. You can also select this option if your proxy server requires authentication. Then, you must provide the corresponding authentication information (Domain Name\Username and Password) that is required to use the proxy server.

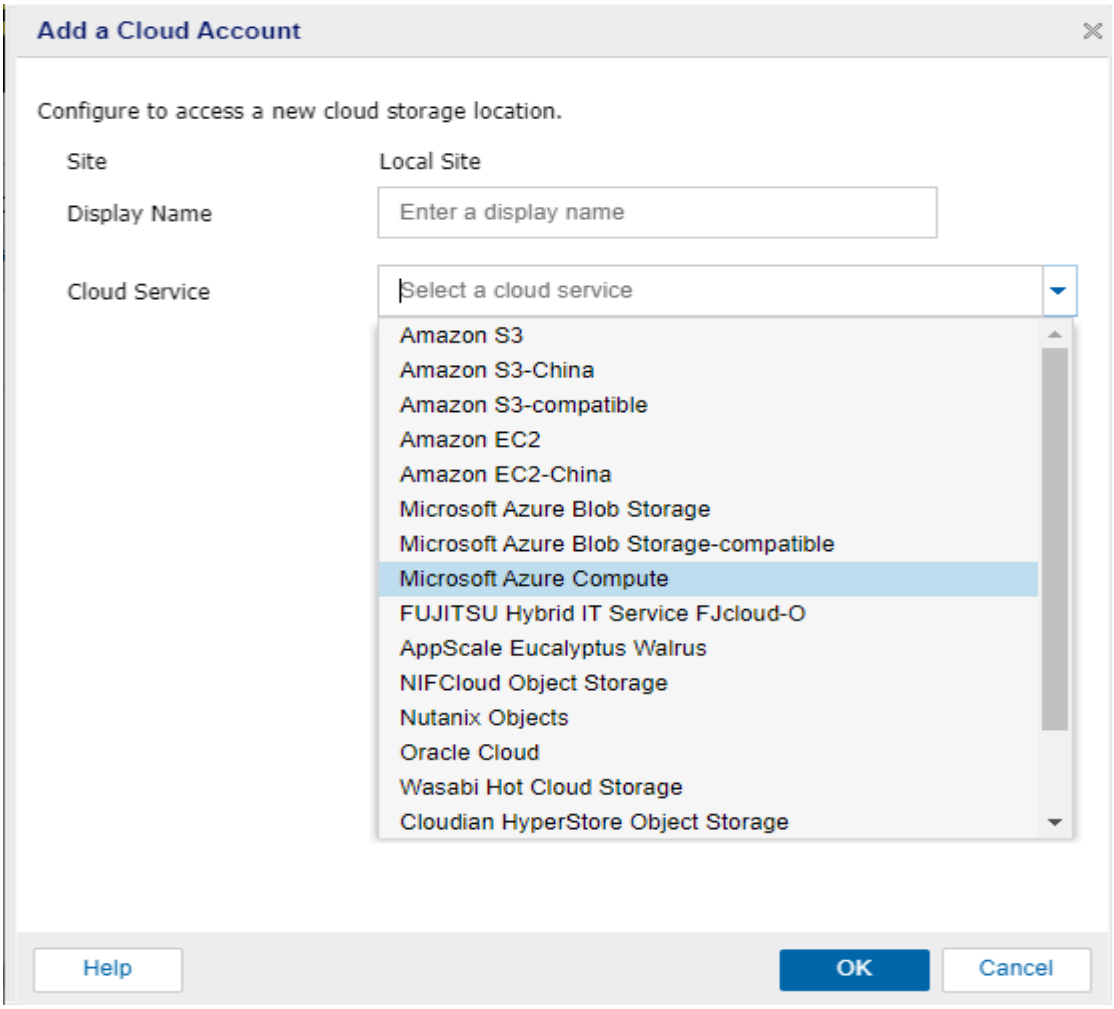
5. Click **OK**.

   Now, the remote console is added.

Add a Cloud Account

Add a cloud account to copy Files or recovery points to a cloud storage. You can use related accounts while creating tasks or one or more plans on Copy Recovery Point / File Copy /File Archive /Virtual Standby to Cloud / Instant Virtual Machine on Amazon EC2. Provide a unique storage name and select required storage service from the multiple options displayed in the drop-down list.

Note: If you have configured File Copy and File Archive in versions Arcserve UDP v6.0 or previous versions and now you upgrade to the latest version of Arcserve UDP, then for File Archive Arcserve UDP creates a new cloud bucket suffixed with **-fa**. File Copy uses the same bucket that was used in the previous version.



The configuration options for each cloud vendor are similar (with some different terminology), and any differences are described. Add a cloud account of your choice.

Follow these steps:

1. Log into Arcserve Unified Data Protection, and click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Cloud Accounts**.

   The **Destinations: Cloud Accounts** page is displayed in the center pane.
3. Click **Add a Cloud Account**.

   The **Add a Cloud Account** page is displayed.
4. For **Display Name**, provide a unique name.

   Display Name specifies the name of the cloud storage. This name will be added to Console for identifying the cloud account. Each cloud account must have a unique display name.
5. Select the option from the **Cloud Service** drop-down list.

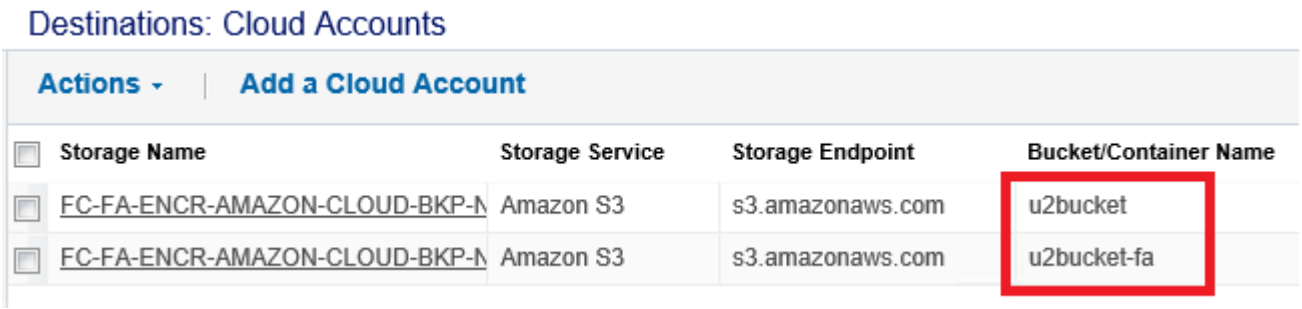
   Multiple fields appear for configuration.
6. Enter details in the following fields to configure and click **OK**:

   The available Storage Service options are Amazon S3, Amazon S3-compatible, Microsoft Azure Blob Storage, Microsoft Azure Blob Storage-Compatible, FUJITSU Hybrid IT Service FJcloud-O, AppScale Eucalyptus Walrus, Amazon EC2, Amazon EC2-China, Microsoft Azure Compute, Nutanix Objects, Wasabi Hot Cloud Storage and Oracle Cloud.

   Note: Click storage service names to view how to add a cloud account for that storage service.

   The configuration option varies depending on the storage service that is selected.

   The selected Storage Service cloud account is added to the Arcserve UDP Console and displayed on the **Destinations: Cloud Accounts** screen. For example, Amazon S3 accounts are displayed below.



What To Do Next?

Add a Cloud Account for:

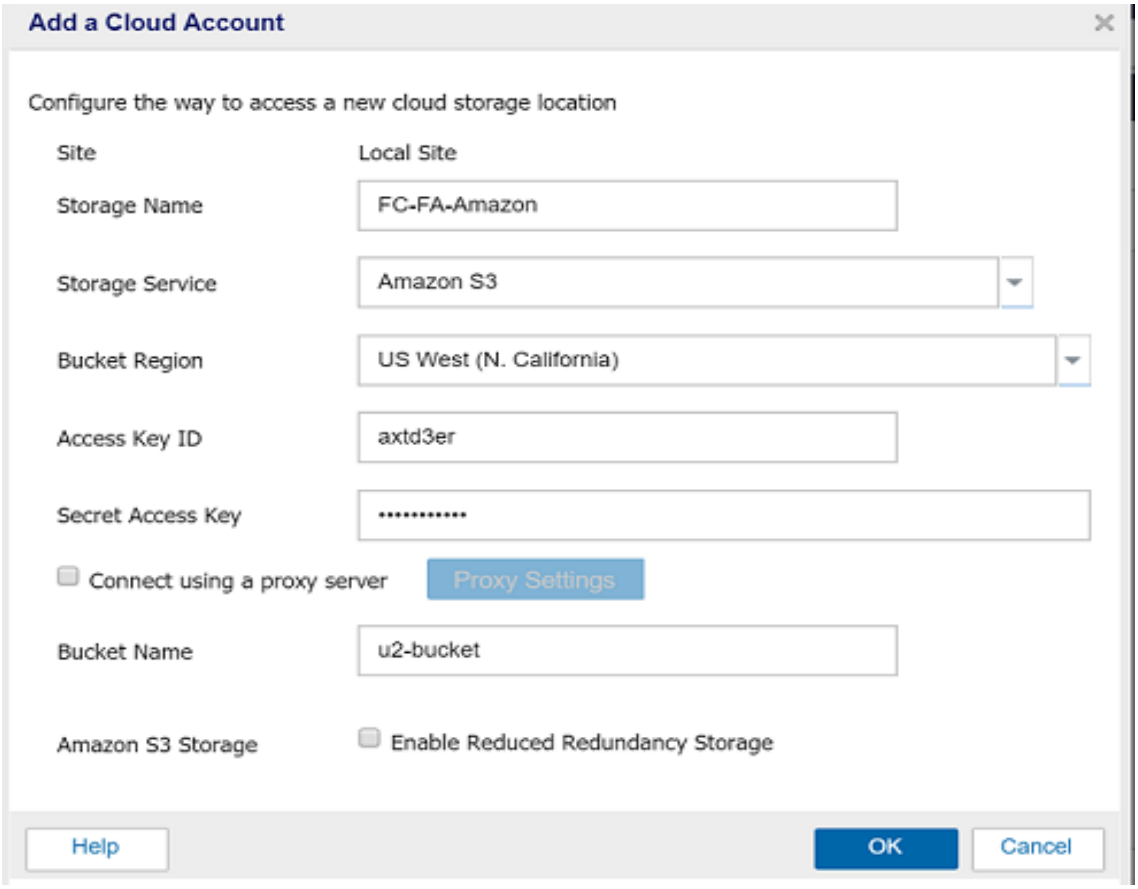
- Amazon S3
- Amazon S3-compatible

- Microsoft Azure Blob Storage
- Microsoft Azure Blob Storage-Compatible
- FUJITSU Hybrid IT Service FJcloud-O
- AppScale Eucalyptus Walrus
- Amazon EC2
- Amazon EC2-China
- Microsoft Azure Compute
- Nutanix Objects
- Wasabi Hot Cloud Storage
- Oracle Cloud

Add a Cloud Account for Amazon S3

Add an Amazon S3 cloud account to copy Files or recovery points to cloud storage.

You can use this account while creating Copy Recovery Point / File Copy /File Archive task.



Follow these steps:

1. Log into Arcserve Unified Data Protection, and click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Cloud Accounts**.

   The **Destinations: Cloud Accounts** page is displayed in the center pane.
3. Click **Add a Cloud Account**.

   The **Add a Cloud Account** page is displayed.
4. For **Display Name**, provide a unique name.

   Display Name specifies the name of the cloud storage. This name will be added to Console for identifying the cloud account. Each cloud account must have a unique display name.
5. Select the option from the **Cloud Service** drop-down list.

Multiple fields appear for configuration.

6. Enter details in the following fields to configure and click **OK**:

Bucket Region

Refers to the region of bucket in Amazon.

Access Key ID

Identifies the user who is requesting access to this location.

Secret Access Key

Refers to a password that is used to verify the authenticity of the request to access this location because your Access Key is not encrypted.

Important! This Secret Access Key is crucial for maintaining the security of your accounts. You should keep your keys and your account credentials in a secure location. Do not embed your Secret Access Key in a web page or other publicly accessible source code and do not transmit it over insecure channels.

Proxy Settings

Specifies the proxy server settings. Select the check box of **Connect using a proxy server** if you want to enable this option. If you select this option, you must also include the IP address (or machine name) of the proxy server and the corresponding port number that is used by the proxy server for internet connections. You can also select this option if your proxy server requires authentication. Then, you must provide the corresponding authentication information (Domain Name\Username and Password) that is required to use the proxy server.

Bucket Name

All files and folders moved or copied to the cloud vendor are stored and organized in your buckets. Buckets are like a container for your files and are used to group and organize objects together. Every object stored at the cloud vendor is placed in a bucket.

Enable Reduced Redundancy Storage

For Amazon S3 only, this option lets you select to enable Reduced Redundancy Storage (RRS). RRS is a storage option within Amazon S3 that helps you reduce cost by storing non-critical, reproducible data at lower levels of redundancy than Amazon S3's standard storage. Both the standard and reduced redundancy storage options store data in multiple facilities and on multiple devices, but with RRS the data is replicated fewer times, so the cost is less. You should expect the same latency and throughput using either the Amazon S3 standard

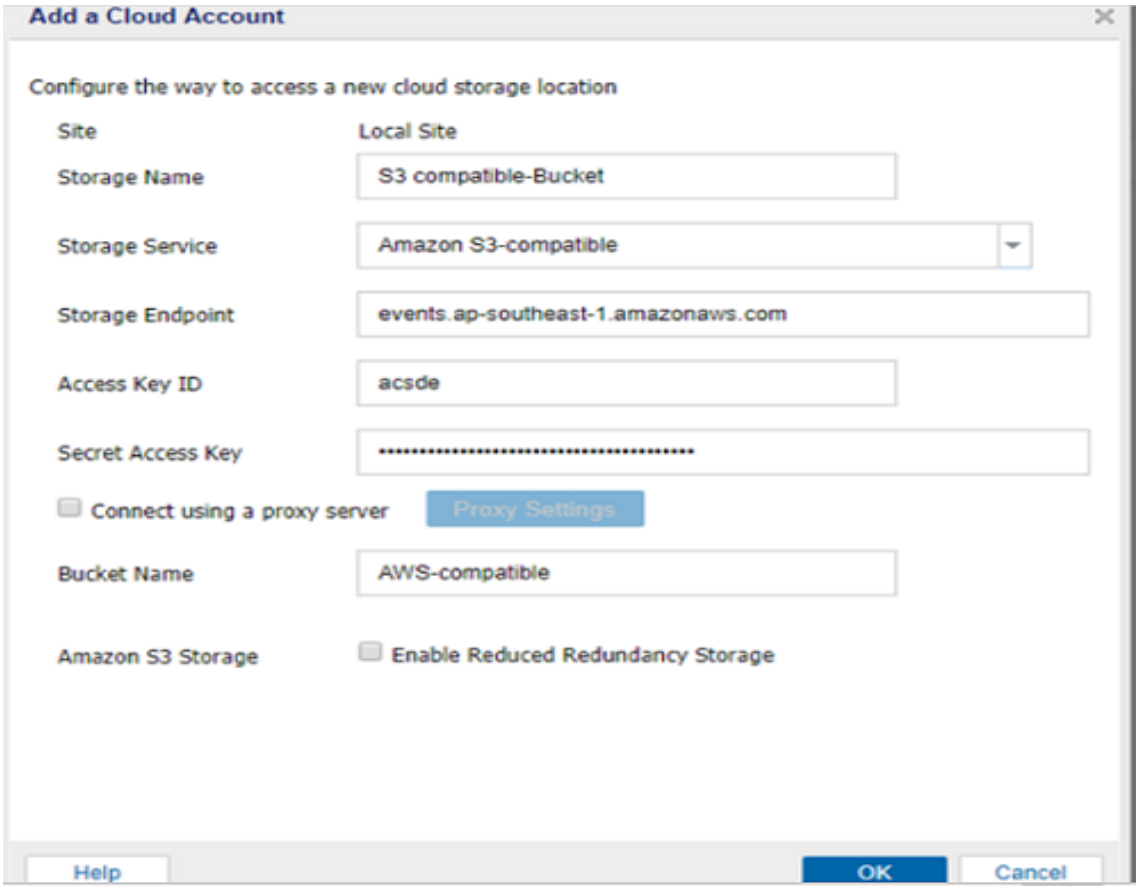
storage or RRS. By default this option is not selected (Amazon S3 uses the standard storage option).

The cloud account is added to the Console.

Note: If you have configured File Copy and File Archive in versions Arcserve UDP v6.0 or previous versions and now you upgrade to the latest version of Arcserve UDP, then for File Archive Arcserve UDP creates a new cloud bucket suffixed with **-fa**. File Copy uses the same bucket that was used in the previous version.

Add a Cloud Account for Amazon S3-Compatible

Add an Amazon S3-Compatible cloud account to copy Files or recovery points to cloud storage. You can use this account while creating Copy Recovery Point / File Copy /File Archive task.



Notes:

- For S3 Compatible/S3 sub-vendors, who use the V3/V2 authentication type to get certified with Arcserve CCI, it works fine as the default Signer type override flag is set to **True**. But, sub-vendors using the V4 authentication need to change the flag to **False** in the AmazonPlugin.properties file and restart services.

  SIGNER_OVERRIDE=false

- To support HGST cloud for Amazon-S3 Compatible, you need to modify the following AmazonPlugin.property:

  SET_STORAGECLASS_HEADER=false

  This property helps to skip the storage header. As a result, using this property if you add a File Copy / File Archive task with Amazon as the Cloud destination, by default the storage header is skipped.

The AmazonPlugin.properties file is located at the following location:

C:\Program Files\Arcserve\Unified Data Protection\Engine\BIN\CCI\Config

Follow these steps:

1. Log into Arcserve Unified Data Protection, and click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Cloud Accounts**.

   The **Destinations: Cloud Accounts** page is displayed in the center pane.
3. Click **Add a Cloud Account**.

   The **Add a Cloud Account** page is displayed.
4. For **Display Name**, provide a unique name.

   Display Name specifies the name of the cloud storage. This name will be added to Console for identifying the cloud account. Each cloud account must have a unique display name.
5. Select the option from the **Cloud Service** drop-down list.

   Multiple fields appear for configuration.
6. Enter details in the following fields to configure and click **OK**:

   Storage Endpoint

   Specifies the Vendor service URL. For example, *http://[server name]:Port No*

   Access Key ID

   Identifies the user who is requesting access to this location.

   Secret Access Key

   Refers to a password that is used to verify the authenticity of the request to access this location because your Access Key is not encrypted.

   Important! This Secret Access Key is crucial for maintaining the security of your accounts. You should keep your keys and your account credentials in a secure location. Do not embed your Secret Access Key in a web page or other publicly accessible source code and do not transmit it over insecure channels.

   Proxy Settings

   Specifies the proxy server settings. Select the check box of **Connect using a proxy server** if you want to enable this option. If you select this option, you must also include the IP address (or machine name) of the proxy server and the corresponding port number that is used by the proxy server for internet connections. You can also select this option if your proxy server requires authentication. Then, you must provide the corresponding authentication information

(Domain Name\Username and Password) that is required to use the proxy server.

Bucket Name

All files and folders moved or copied to the cloud vendor are stored and organized in your buckets. Buckets are like a container for your files and are used to group and organize objects together. Every object stored at the cloud vendor is placed in a bucket.

Enable Reduced Redundancy Storage

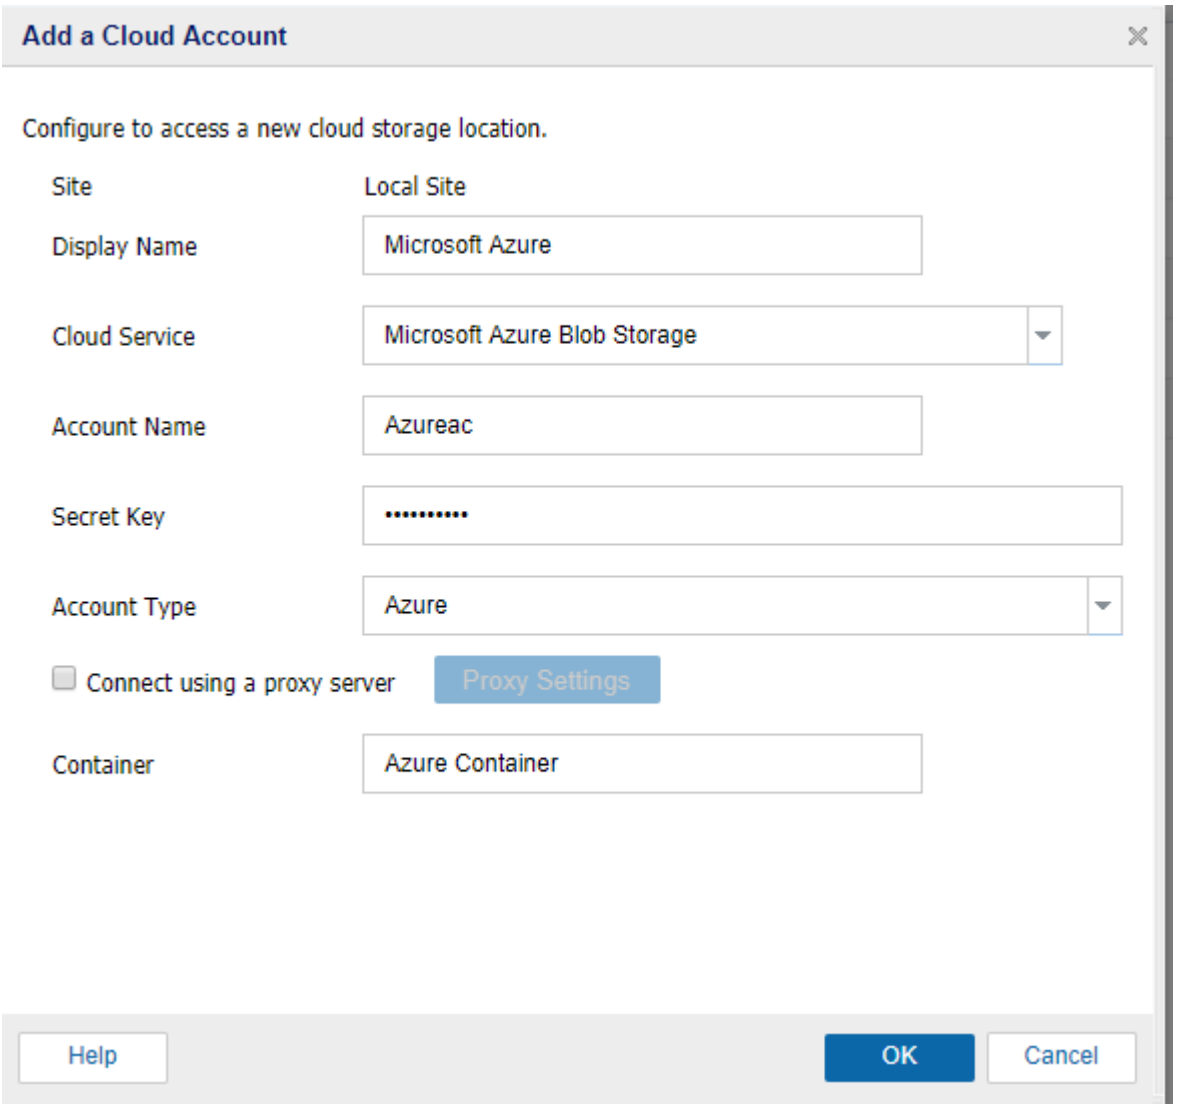
For Amazon S3 only, this option lets you select to enable Reduced Redundancy Storage (RRS). RRS is a storage option within Amazon S3 that helps you reduce cost by storing non-critical, reproducible data at lower levels of redundancy than Amazon S3's standard storage. Both the standard and reduced redundancy storage options store data in multiple facilities and on multiple devices, but with RRS the data is replicated fewer times, so the cost is less. You should expect the same latency and throughput using either the Amazon S3 standard storage or RRS. By default this option is not selected (Amazon S3 uses the standard storage option).

The cloud account is added to the Console.

Note: If you have configured File Copy and File Archive in versions Arcserve UDP v6.0 or previous versions and now you upgrade to the latest version of Arcserve UDP, then for File Archive Arcserve UDP creates a new cloud bucket suffixed with **-fa**. File Copy uses the same bucket that was used in the previous version.

Add a Cloud Account for Microsoft Azure Blob Storage

Add a Microsoft Azure Blob Storage cloud account to copy Files or recovery points to cloud storage. You can use this account while creating Copy Recovery Point / File Copy /File Archive task.



Follow these steps:

1. Log into Arcserve Unified Data Protection, and click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Cloud Accounts**.

   The **Destinations: Cloud Accounts** page is displayed in the center pane.
3. Click **Add a Cloud Account**.

   The **Add a Cloud Account** page is displayed.
4. For **Display Name**, provide a unique name.

Display Name specifies the name of the cloud storage. This name will be added to Console for identifying the cloud account. Each cloud account must have a unique display name.

5. Select the option from the **Cloud Service** drop-down list.

   Multiple fields appear for configuration.

6. Enter details in the following fields to configure and click **OK**:

   Account Name

   Identifies the user who is requesting access to this location.

   Secret Key

   Refers to a password that is used to verify the authenticity of the request to access this location because your Access Key is not encrypted.

   Important! This Secret Key is crucial for maintaining the security of your accounts. You should keep your keys and your account credentials in a secure location. Do not embed your Secret Key in a web page or other publicly accessible source code and do not transmit it over insecure channels.

   Proxy Settings

   Specifies the proxy server settings. Select the check box of **Connect using a proxy server** if you want to enable this option. If you select this option, you must also include the IP address (or machine name) of the proxy server and the corresponding port number that is used by the proxy server for internet connections. You can also select this option if your proxy server requires authentication. Then, you must provide the corresponding authentication information (Domain Name\Username and Password) that is required to use the proxy server.

   Container

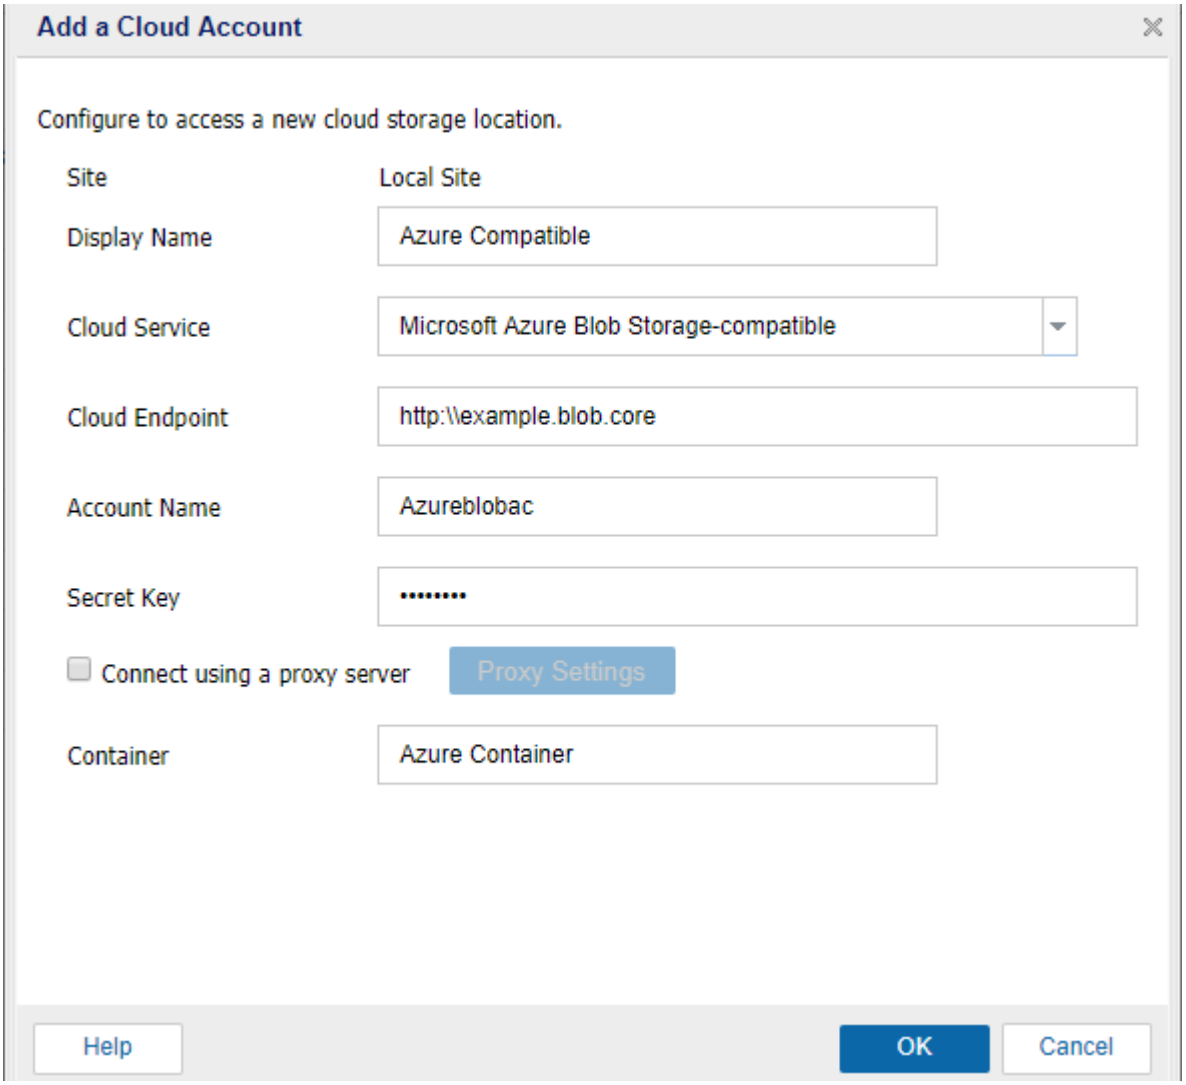
   All files and folders moved or copied to the cloud vendor are stored and organized in your containers. Using containers, you can group and organize objects together. Every object stored at the cloud vendor is placed in a bucket.

   The cloud account is added to the Console.

   Note: If you have configured File Copy and File Archive in versions Arcserve UDP v6.0 or previous versions and now you upgrade to the latest version of Arcserve UDP, then for File Archive Arcserve UDP creates a new cloud bucket suffixed with **-fa**. File Copy uses the same bucket that was used in the previous version.

Add a Cloud Account for Microsoft Azure Blob Storage-compatible

Add a Microsoft Azure Blob Storage-compatible cloud account to copy Files or recovery points to cloud storage. You can use this account while creating Copy Recovery Point / File Copy /File Archive task.



Follow these steps:

1. Log into Arcserve Unified Data Protection, and click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Cloud Accounts**.

   The **Destinations: Cloud Accounts** page is displayed in the center pane.
3. Click **Add a Cloud Account**.

   The **Add a Cloud Account** page is displayed.
4. For **Display Name**, provide a unique name.

Display Name specifies the name of the cloud storage. This name will be added to Console for identifying the cloud account. Each cloud account must have a unique display name.

5. Select the option from the **Cloud Service** drop-down list.

   Multiple fields appear for configuration.

6. Enter details in the following fields to configure and click **OK**:

   Storage Endpoint

   Specifies the Vendor service URL. For example, *http://[server name]:Port No*

   Account Name

   Identifies the user who is requesting access to this location.

   Secret Key

   Refers to a password that is used to verify the authenticity of the request to access this location because your Access Key is not encrypted.

   Important! This Secret Key is crucial for maintaining the security of your accounts. You should keep your keys and your account credentials in a secure location. Do not embed your Secret Key in a web page or other publicly accessible source code and do not transmit it over insecure channels.

   Proxy Settings

   Specifies the proxy server settings. Select the check box of **Connect using a proxy server** if you want to enable this option. If you select this option, you must also include the IP address (or machine name) of the proxy server and the corresponding port number that is used by the proxy server for internet connections. You can also select this option if your proxy server requires authentication. Then, you must provide the corresponding authentication information (Domain Name\Username and Password) that is required to use the proxy server.

   Container

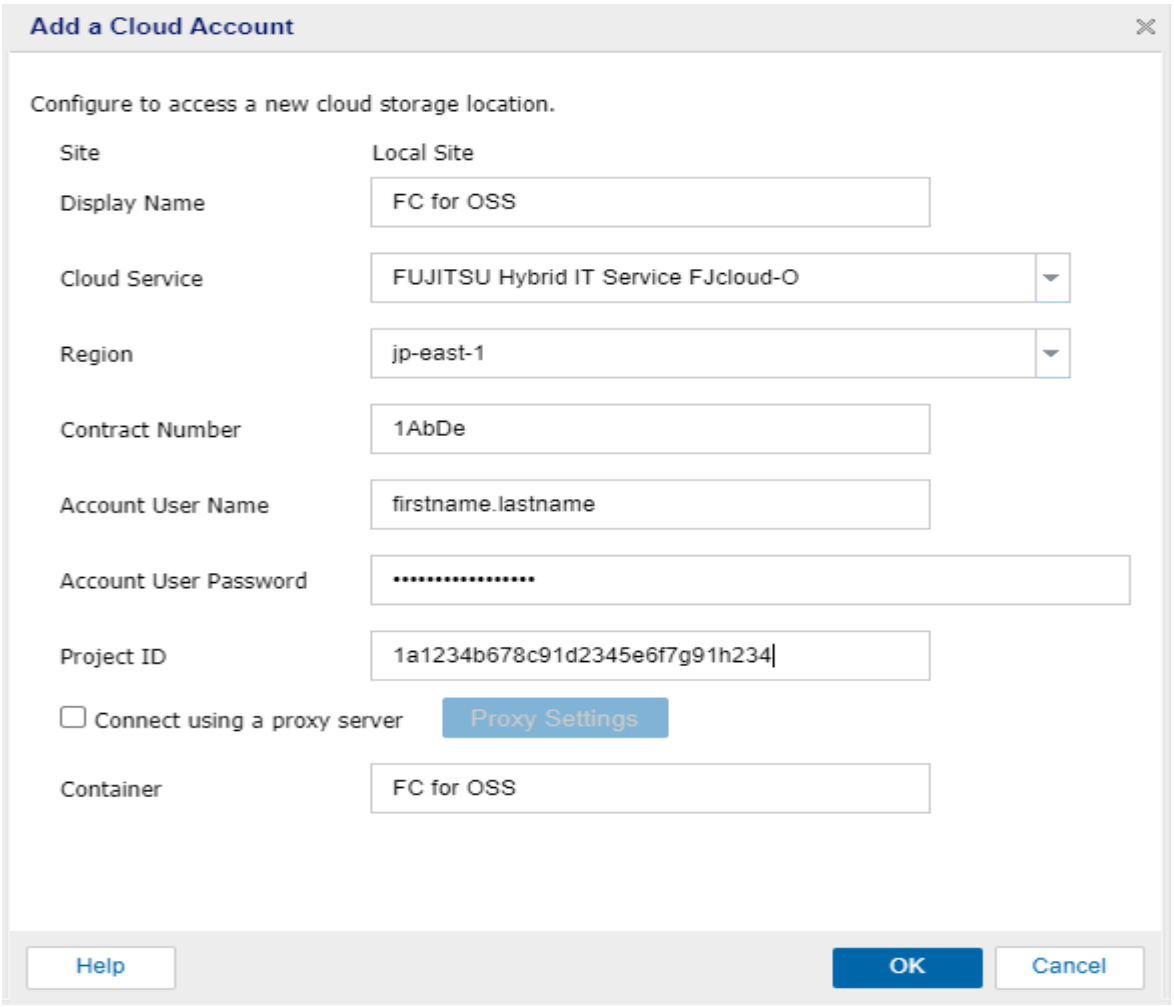
   All files and folders moved or copied to the cloud vendor are stored and organized in your containers. Using containers, you can group and organize objects together. Every object stored at the cloud vendor is placed in a bucket.

   The cloud account is added to the Console.

   Note: If you have configured File Copy and File Archive in versions Arcserve UDP v6.0 or previous versions and now you upgrade to the latest version of Arcserve UDP, then for File Archive Arcserve UDP creates a new cloud bucket suffixed with **-fa**. File Copy uses the same bucket that was used in the previous version.

Add a Cloud Account for FUJITSU Hybrid IT Service FJcloud-O

Add a FUJITSU Hybrid IT Service FJcloud-O cloud account to copy files or recovery points to cloud storage. You can use this account while creating File Copy /File Archive task.



Follow these steps:

1. Log into Arcserve Unified Data Protection, and click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Cloud Accounts**.

   The **Destinations: Cloud Accounts** page is displayed in the center pane.
3. Click **Add a Cloud Account**.

   The **Add a Cloud Account** page is displayed.
4. For **Display Name**, provide a unique name.

Display Name specifies the name of the cloud storage. This name will be added to Console for identifying the cloud account. Each cloud account must have a unique display name.

5. Select the option from the **Cloud Service** drop-down list.

   Multiple fields appear for configuration.

6. Enter details in the following fields to configure and click **OK**:

   Bucket Region

   Refers to the region of bucket in FUJITSU Hybrid IT Service FJcloud-O.

   Account User Name

   Identifies the user who is requesting access to this location.

   Account User Password

   Refers to a password that is used to verify the authenticity of the request to access this location because your password is not encrypted.

   Important! This password is crucial for maintaining the security of your accounts. You should keep your keys and your account credentials in a secure location. Do not embed your password in a web page or other publicly accessible source code and do not transmit it over insecure channels.

   Proxy Settings

   Specifies the proxy server settings. Select the check box of **Connect using a proxy server** if you want to enable this option. If you select this option, you must also include the IP address (or machine name) of the proxy server and the corresponding port number that is used by the proxy server for internet connections. You can also select this option if your proxy server requires authentication. Then, you must provide the corresponding authentication information (Domain Name\Username and Password) that is required to use the proxy server.

   Contract Number

   Refers to the number of contract that FUJITSU Hybrid IT Service FJcloud-O provides.

   Project ID

   Refers to the ID of project that FUJITSU Hybrid IT Service FJcloud-O generates.

   Container

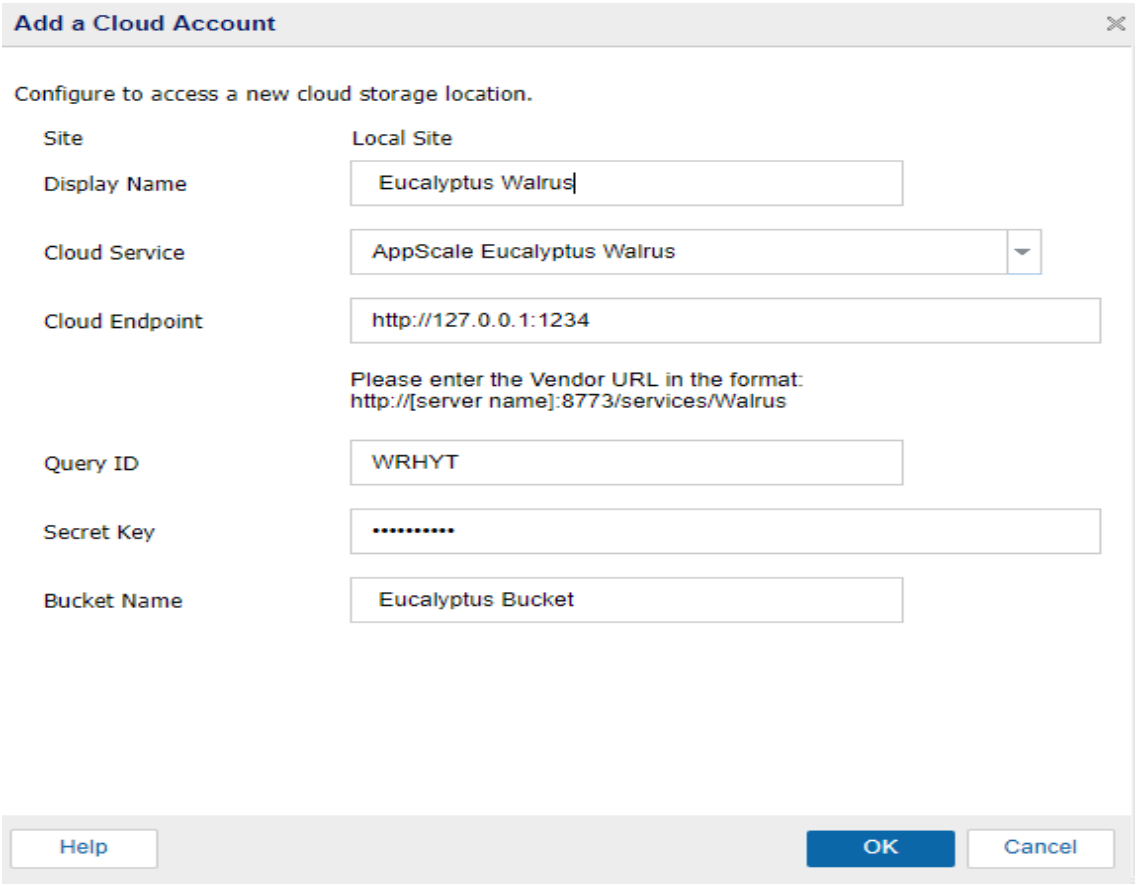
   All files and folders moved or copied to the cloud vendor are stored and organized in your containers. Using containers, you can group and organize objects together. Every object stored at the cloud vendor is placed in a bucket.

The cloud account is added to the Console.

Note: If you have configured File Copy and File Archive in versions Arcserve UDP v6.0 or previous versions and now you upgrade to the latest version of Arcserve UDP, then for File Archive Arcserve UDP creates a new cloud bucket suffixed with **-fa**. File Copy uses the same bucket that was used in the previous version.

Add a Cloud Account for AppScale Eucalyptus Walrus

Add an AppScale Eucalyptus Walrus cloud account to copy Files or recovery points to cloud storage. You can use this account while creating Copy Recovery Point / File Copy /File Archive task.



Note: Using AppScale Eucalyptus Walrus as your file copy cloud vendor, you cannot copy files whose path length is greater than 170 characters.

Follow these steps:

1. Log into Arcserve Unified Data Protection, and click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Cloud Accounts**.

   The **Destinations: Cloud Accounts** page is displayed in the center pane.
3. Click **Add a Cloud Account**.

   The **Add a Cloud Account** page is displayed.
4. For **Display Name**, provide a unique name.

Display Name specifies the name of the cloud storage. This name will be added to Console for identifying the cloud account. Each cloud account must have a unique display name.

5. Select the option from the **Cloud Service** drop-down list.

   Multiple fields appear for configuration.

6. Enter details in the following fields to configure and click **OK**:

   Storage Endpoint

   Specifies the Vendor service URL. For example, *http://[server name]:Port No*

   Query ID

   Identifies the user who is requesting access to this location.

   Secret Key

   Refers to a password that is used to verify the authenticity of the request to access this location because your Access Key is not encrypted.

   Important! This Secret Key is crucial for maintaining the security of your accounts. You should keep your keys and your account credentials in a secure location. Do not embed your Secret Key in a web page or other publicly accessible source code and do not transmit it over insecure channels.

   Bucket Name

   All files and folders moved or copied to the cloud vendor are stored and organized in your buckets. Buckets are like a container for your files and are used to group and organize objects together. Every object stored at the cloud vendor is placed in a bucket.

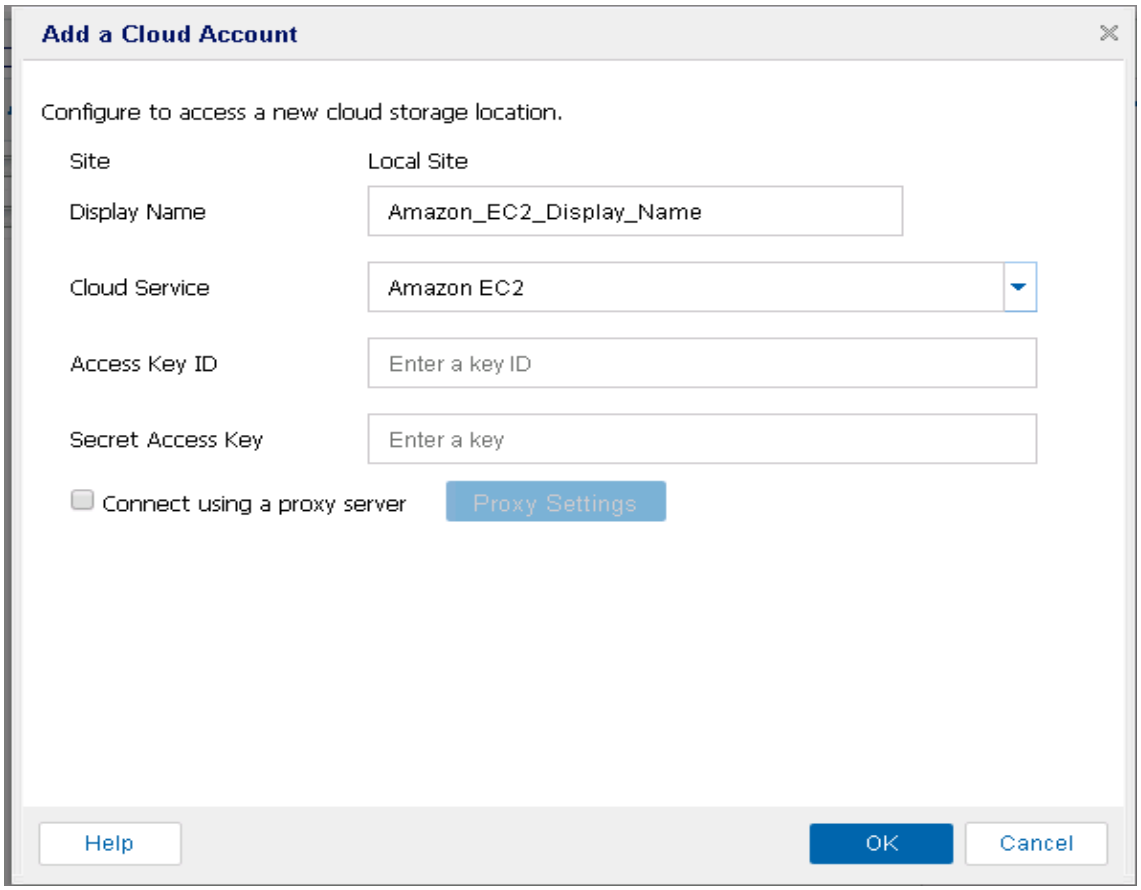
   The cloud account is added to the Console.

   Note: If you have configured File Copy and File Archive in versions Arcserve UDP v6.0 or previous versions and now you upgrade to the latest version of Arcserve UDP, then for File Archive Arcserve UDP creates a new cloud bucket suffixed with **-fa**. File Copy uses the same bucket that was used in the previous version.

Add a Cloud Account for Amazon EC2

Add an Amazon EC2 cloud account to copy Files or recovery points to cloud storage.

You can use this account while creating tasks for Virtual Standby to Cloud or Instant Virtual Machine on Amazon EC2 plans.



Follow these steps:

1. Log into Arcserve Unified Data Protection, and click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Cloud Accounts**.

   The **Destinations: Cloud Accounts** page is displayed in the center pane.
3. Click **Add a Cloud Account**.

   The **Add a Cloud Account** page is displayed.
4. For **Display Name**, provide a unique name.

   Display Name specifies the name of the cloud storage. This name will be added to Console for identifying the cloud account. Each cloud account must have a unique display name.

5. Select the option from the **Cloud Service** drop-down list.

   Multiple fields appear for configuration.

6. Enter details in the following fields to configure and click **OK**:

   Access Key ID

   Identifies the user who is requesting access to this location.

   Secret Access Key

   Refers to a password that is used to verify the authenticity of the request to access this location because your Access Key is not encrypted.

   Important! This Secret Access Key is crucial for maintaining the security of your accounts. You should keep your keys and your account credentials in a secure location. Do not embed your Secret Access Key in a web page or other publicly accessible source code and do not transmit it over insecure channels.

   Proxy Settings

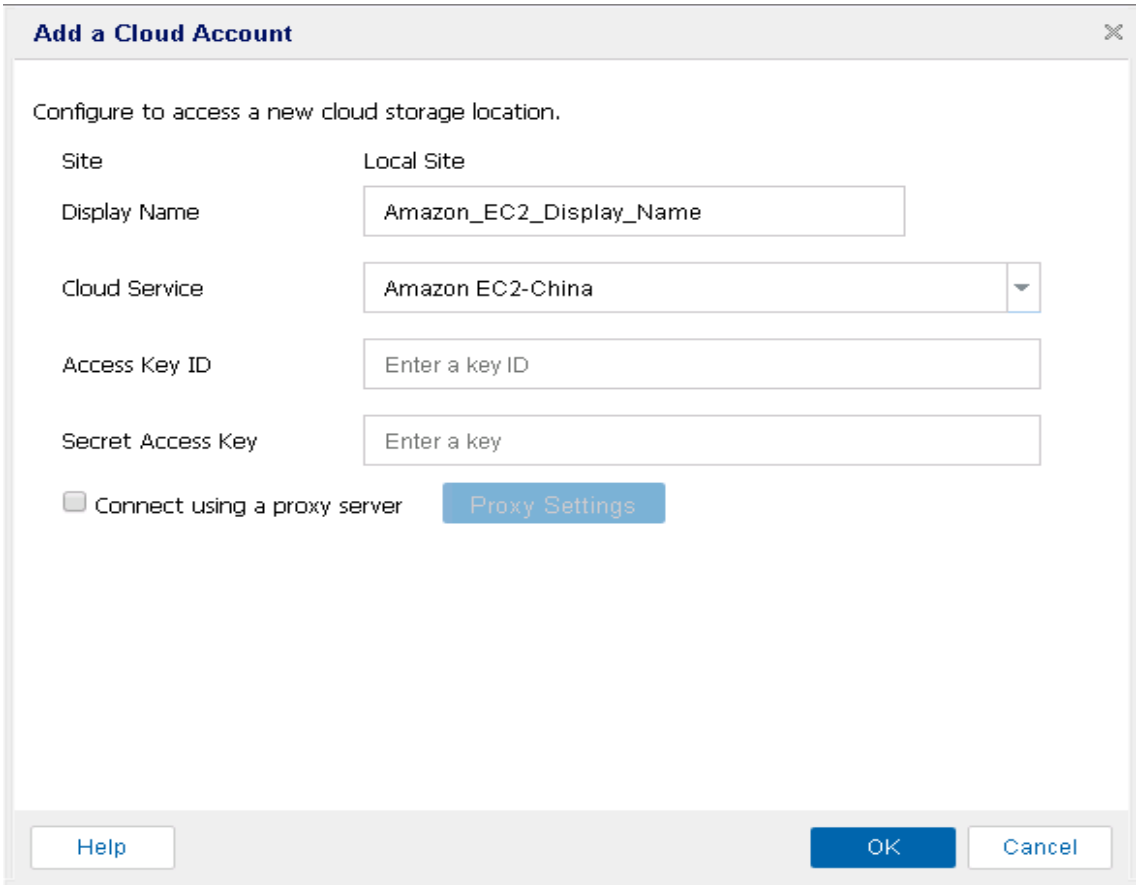
   Specifies the proxy server settings. Select the check box of **Connect using a proxy server** if you want to enable this option. If you select this option, you must also include the IP address (or machine name) of the proxy server and the corresponding port number that is used by the proxy server for internet connections. You can also select this option if your proxy server requires authentication. Then, you must provide the corresponding authentication information (Domain Name\Username and Password) that is required to use the proxy server.

   The cloud account is added to the Console.

   Note: If you have configured File Copy and File Archive in versions Arcserve UDP v6.0 or previous versions and now you upgrade to the latest version of Arcserve UDP, then for File Archive Arcserve UDP creates a new cloud bucket suffixed with **-fa**. File Copy uses the same bucket that was used in the previous version.

Add a Cloud Account for Amazon EC2-China

Add an Amazon EC2-China cloud account to copy Files or recovery points to cloud storage. You can use this account while creating tasks for Virtual Standby to Cloud or Instant Virtual Machine on Amazon EC2 plans.



Follow these steps:

1. Log into Arcserve Unified Data Protection, and click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Cloud Accounts**.

   The **Destinations: Cloud Accounts** page is displayed in the center pane.
3. Click **Add a Cloud Account**.

   The **Add a Cloud Account** page is displayed.
4. For **Display Name**, provide a unique name.

   Display Name specifies the name of the cloud storage. This name will be added to Console for identifying the cloud account. Each cloud account must have a unique display name.
5. Select the option from the **Cloud Service** drop-down list.

Multiple fields appear for configuration.

6. Enter details in the following fields to configure and click **OK**:

Access Key ID

Identifies the user who is requesting access to this location.

Secret Access Key

Refers to a password that is used to verify the authenticity of the request to access this location because your Access Key is not encrypted.

Important! This Secret Access Key is crucial for maintaining the security of your accounts. You should keep your keys and your account credentials in a secure location. Do not embed your Secret Access Key in a web page or other publicly accessible source code and do not transmit it over insecure channels.

Proxy Settings

Specifies the proxy server settings. Select the check box of **Connect using a proxy server** if you want to enable this option. If you select this option, you must also include the IP address (or machine name) of the proxy server and the corresponding port number that is used by the proxy server for internet connections. You can also select this option if your proxy server requires authentication. Then, you must provide the corresponding authentication information (Domain Name\Username and Password) that is required to use the proxy server.

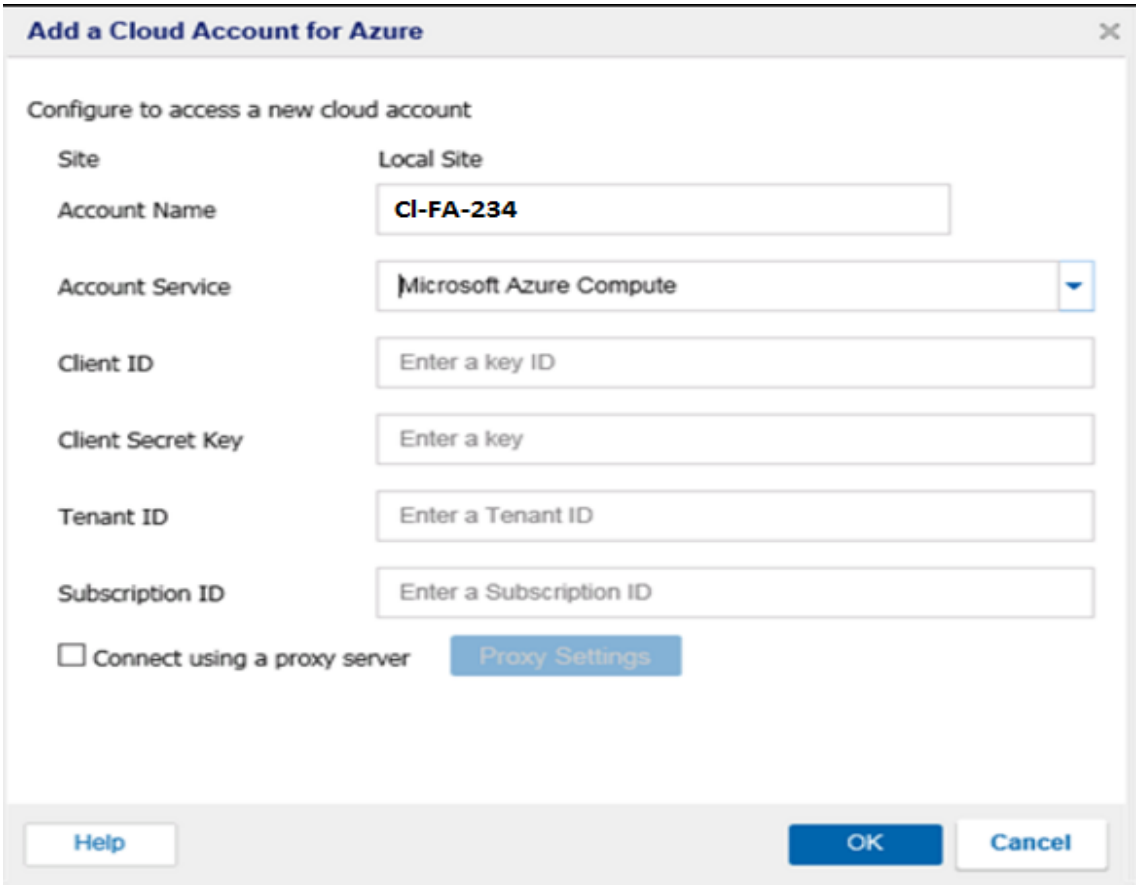
The cloud account is added to the Console.

Note: If you have configured File Copy and File Archive in versions Arcserve UDP v6.0 or previous versions and now you upgrade to the latest version of Arcserve UDP, then for File Archive Arcserve UDP creates a new cloud bucket suffixed with **-fa**. File Copy uses the same bucket that was used in the previous version.

Add a Cloud Account for Microsoft Azure Compute

Add a Microsoft Azure Compute cloud account to copy Files or recovery points to cloud storage. You can use this account while creating tasks for Virtual Standby to Cloud or Instant Virtual Machine on Microsoft Azure plans.

Note: To add a Cloud Account for Microsoft Azure, you must meet the pre-requisites. For details, view Prerequisites.



Follow these steps:

1. Log into Arcserve Unified Data Protection, and click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Cloud Accounts**.

   The **Destinations: Cloud Accounts** page is displayed in the center pane.
3. Click **Add a Cloud Account**.

   The **Add a Cloud Account** page is displayed.
4. For **Account Name**, provide a unique name.

Account Name specifies the name of the cloud storage. This name will be added to Console for identifying the cloud account. Each cloud account must have a unique storage name.

5. Select the option from the **Account Service** drop-down list.

   Multiple fields appear for configuration.

6. Enter details in the following fields to configure and click **OK**:

   Client ID

   Refers to the Application ID of the Azure Active Directory application. Copy your Client ID prepared in the text editor.

   Client Secret Key

   Refers to the authentication key generated for the Azure Active Directory application that you enter as Client ID. Copy your Client Secret Key prepared in the text editor.

   Important! This Secret Key is crucial for maintaining the security of your accounts. You should keep your keys and your account credentials in a secure location. Do not embed your Secret Key in a web page or other publicly accessible source code and do not transmit it over insecure channels.

   Tenant ID

   Refers to the ID of the Azure Active Directory where you created the Azure Active Directory application. Copy your Tenant ID prepared in the text editor.

   Subscription ID

   Refers to a Globally Unique Identifier (GUID) that uniquely identifies your subscription to use Azure services. Copy your Subscription ID prepared in the text editor.

   Proxy Settings

   Specifies the proxy server settings. Select **Connect using a proxy server** to enable this option. If you select this option, you must also include the IP address (or machine name) of the proxy server and the corresponding port number that is used by the proxy server for internet connections. You can also select this option if your proxy server requires authentication. You then must provide the corresponding authentication information that is required to use the proxy server.

The cloud account is added to the Console.

Prerequisites to Add a Cloud Account for Microsoft Azure

Before you can create a cloud account for Microsoft Azure, you must complete the following mandatory prerequisite tasks in the given order:

1. Prepare servers to deploy as Arcserve UDP Console, Recovery Point Server, and those servers must meet the system requirements for each component.
2. Ensure that you have the required permissions to create an application in Azure Active Directory.

   For instructions, see Check Azure Active Directory permissions in the Microsoft documentation.
3. Create an Azure Active Directory application. In a text editor (such as Notepad), copy the application ID of the application and label as Client ID.

   For instructions, see Check Azure Active Directory permissions in the Microsoft documentation.
4. Get the Application ID and generate an authentication key for this application. Copy the authentication key string to the text editor (such as Notepad), and label the string as Client Secret Key.

   For instructions, see Get application ID and authentication key in the Microsoft documentation.
5. Get the Tenant ID, which is the ID of the Azure Active Directory in which you created the application. In a text editor (such as Notepad), copy the ID and label it as Tenant ID.

   For instructions, see Get tenant ID in the Microsoft documentation.
6. Perform the following steps to assign Contributor role to the application.
   a. From the left pane of the Microsoft Azure portal menu, select **Subscriptions**.
   b. Select your subscription.
   c. Select the Access Control (IAM) tab.
   d. Add your application.
   e. Assign the Contributor role to the application.

   For details, see Assign application to role in the Microsoft documentation.
7. Get your Azure subscription ID.

Note: The subscription ID is a GUID that uniquely identifies your subscription to use Azure services.

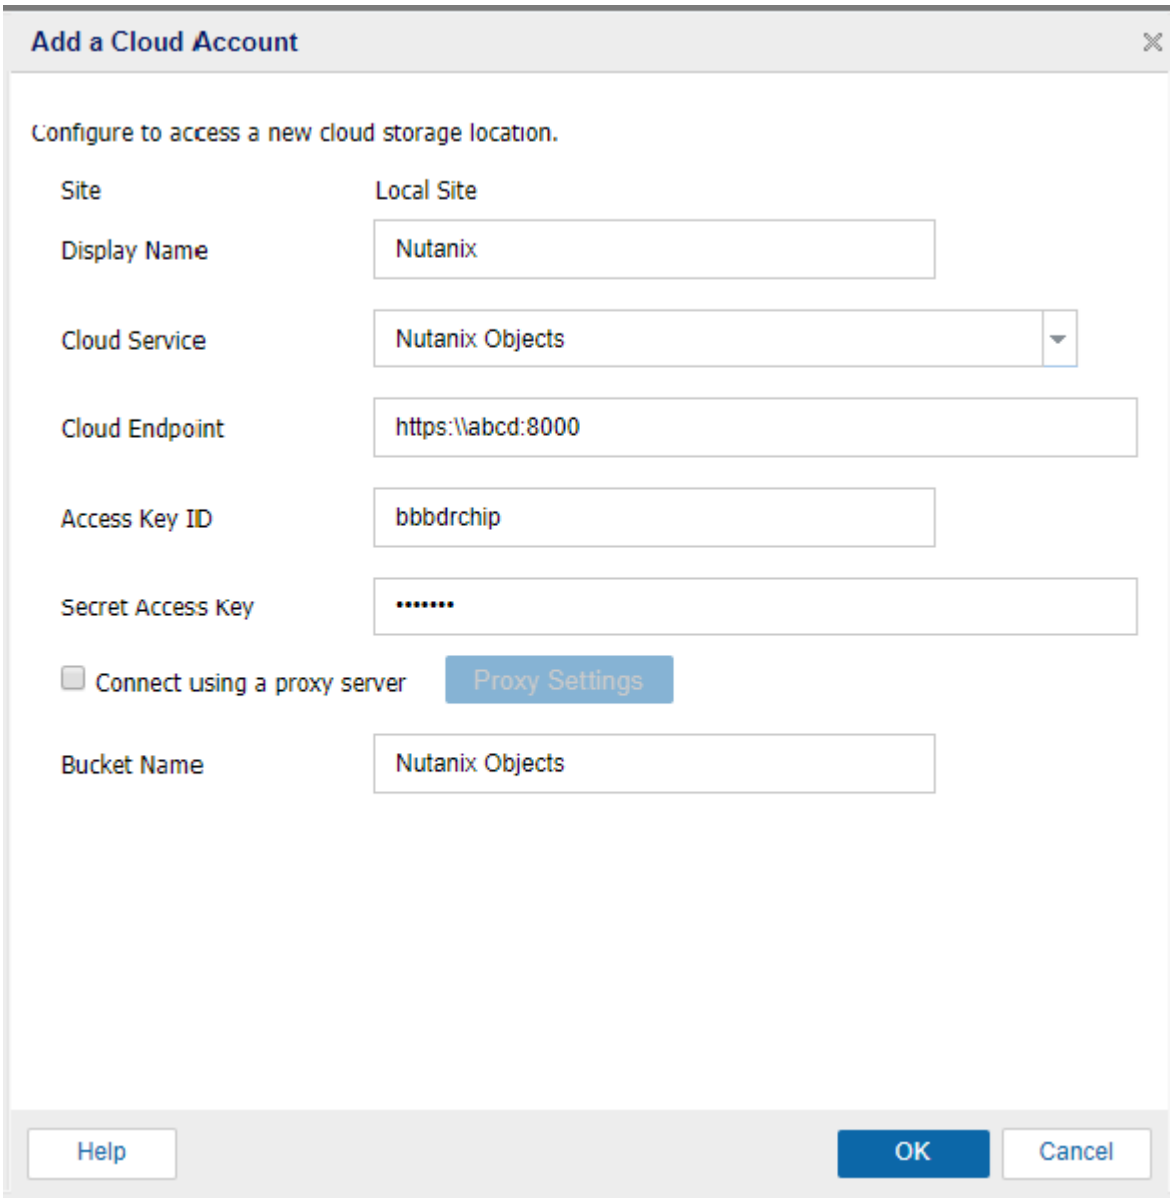
a. Log on to the Microsoft Azure portal.

b. In the left navigation panel, click Subscriptions.

The list of your subscriptions is displayed along with the subscription ID.

Add a Cloud Account for Nutanix Objects

Add an Nutanix Objects cloud account to copy Files or recovery points to cloud storage. You can use this account while creating Copy Recovery Point / File Copy /File Archive task.



Follow these steps:

1. Log into Arcserve Unified Data Protection, and click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Cloud Accounts**.

   The **Destinations: Cloud Accounts** page is displayed in the center pane.
3. Click **Add a Cloud Account**.

   The **Add a Cloud Account** page is displayed.
4. For **Display Name**, provide a unique name.

   Display Name specifies the name of the cloud storage. This name will be added to Console for identifying the cloud account. Each cloud account must have a unique display name.
5. Select the option from the **Cloud Service** drop-down list.

   Multiple fields appear for configuration.
6. Enter details in the following fields to configure and click **OK**:

   Storage Endpoint

   Specifies the Vendor service URL. For example, *http://[server name]:Port No*

   Access Key ID

   Identifies the user who is requesting access to this location.

   Secret Access Key

   Refers to a password that is used to verify the authenticity of the request to access this location because your Access Key is not encrypted.

   Important! This Secret Access Key is crucial for maintaining the security of your accounts. You should keep your keys and your account credentials in a secure location. Do not embed your Secret Access Key in a web page or other publicly accessible source code and do not transmit it over insecure channels.

   Proxy Settings

   Specifies the proxy server settings. Select the check box of **Connect using a proxy server** if you want to enable this option. If you select this option, you must also include the IP address (or machine name) of the proxy server and the corresponding port number that is used by the proxy server for internet connections. You can also select this option if your proxy server requires authentication. Then, you must provide the corresponding authentication information (Domain Name\Username and Password) that is required to use the proxy server.

   Bucket Name

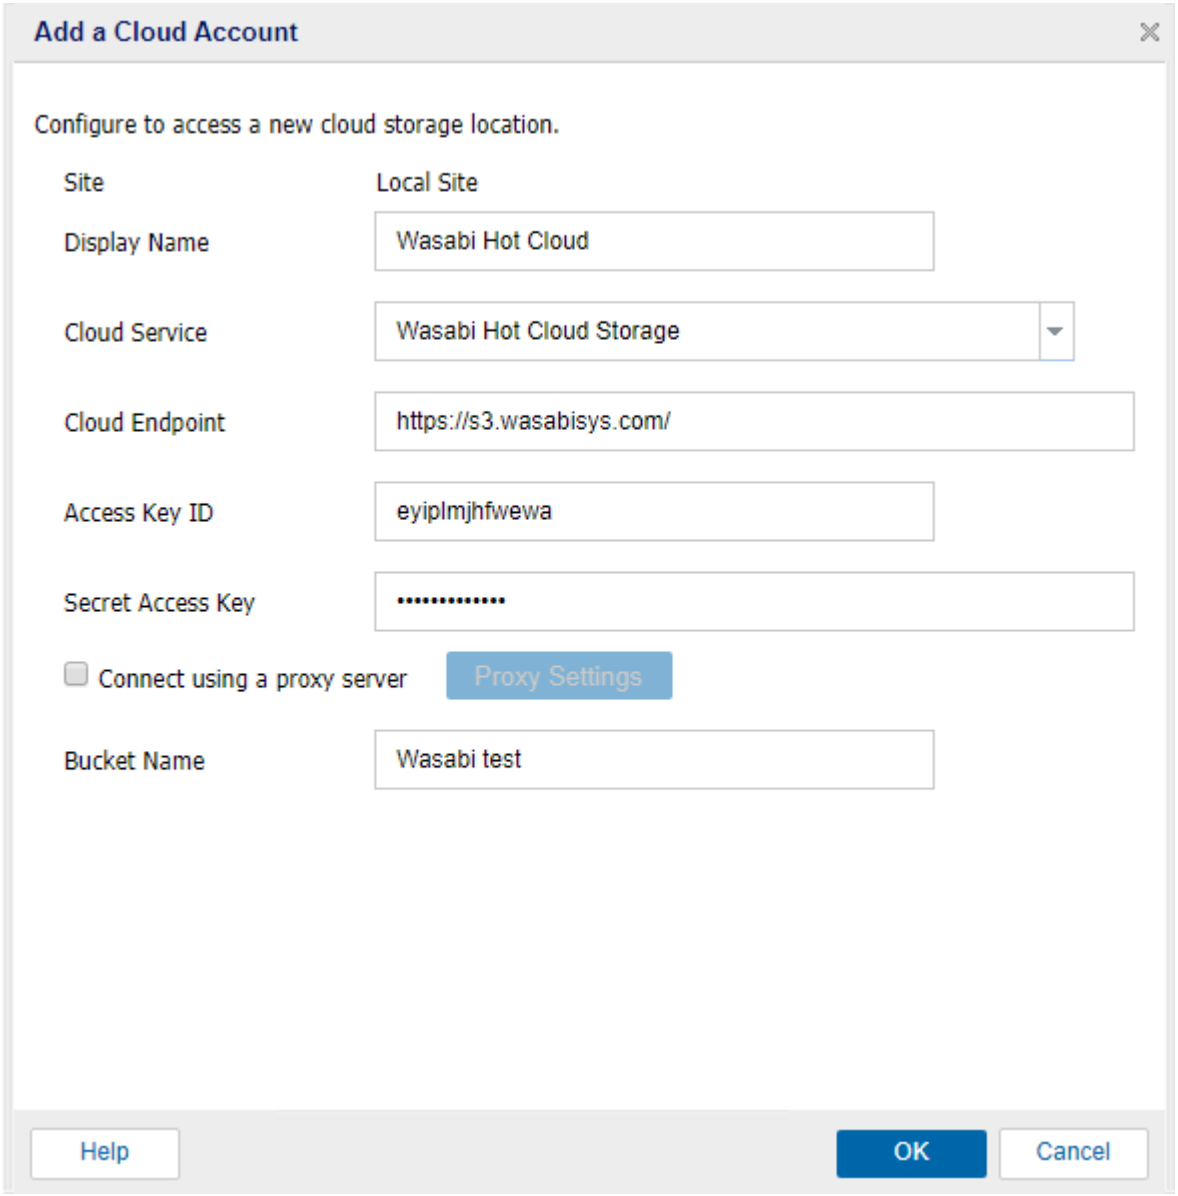
All files and folders moved or copied to the cloud vendor are stored and organized in your buckets. Buckets are like a container for your files and are used to group and organize objects together. Every object stored at the cloud vendor is placed in a bucket.

The cloud account is added to the Console.

Note: If you have configured File Copy and File Archive in versions Arcserve UDP v6.0 or previous versions and now you upgrade to the latest version of Arcserve UDP, then for File Archive Arcserve UDP creates a new cloud bucket suffixed with **-fa**. File Copy uses the same bucket that was used in the previous version.

Add a Cloud Account for Wasabi Hot Cloud Storage

Add a Wasabi Hot Cloud Storage cloud account to copy Files or recovery points to cloud storage. You can use this account while creating Copy Recovery Point / File Copy /File Archive task.



Follow these steps:

1. Log into Arcserve Unified Data Protection, and click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Cloud Accounts**.

   The **Destinations: Cloud Accounts** page is displayed in the center pane.
3. Click **Add a Cloud Account**.

   The **Add a Cloud Account** page is displayed.
4. For **Display Name**, provide a unique name.

   Display Name specifies the name of the cloud storage. This name will be added to Console for identifying the cloud account. Each cloud account must have a unique display name.
5. Select the option from the **Cloud Service** drop-down list.

Multiple fields appear for configuration.

6. Enter details in the following fields to configure and click **OK**:

Storage Endpoint

The following three regions are supported for Wasabi Hot Cloud Storage:

- For region us-east: The Endpoint is s3.wasabisys.com.
- For region us-west: The Endpoint is s3.us-west-1.wasabisys.com
- For region eu-central: The Endpoint is s3.eu-central-1.wasabisys.com

Access Key ID

Identifies the user who is requesting access to this location.

Secret Access Key

Refers to a password that is used to verify the authenticity of the request to access this location because your Access Key is not encrypted.

Important! This Secret Access Key is crucial for maintaining the security of your accounts. You should keep your keys and your account credentials in a secure location. Do not embed your Secret Access Key in a web page or other publicly accessible source code and do not transmit it over insecure channels.

Proxy Settings

Specifies the proxy server settings. Select the check box of **Connect using a proxy server** if you want to enable this option. If you select this option, you must also include the IP address (or machine name) of the proxy server and the corresponding port number that is used by the proxy server for internet connections. You can also select this option if your proxy server requires authentication. Then, you must provide the corresponding authentication information (Domain Name\Username and Password) that is required to use the proxy server.

Bucket Name

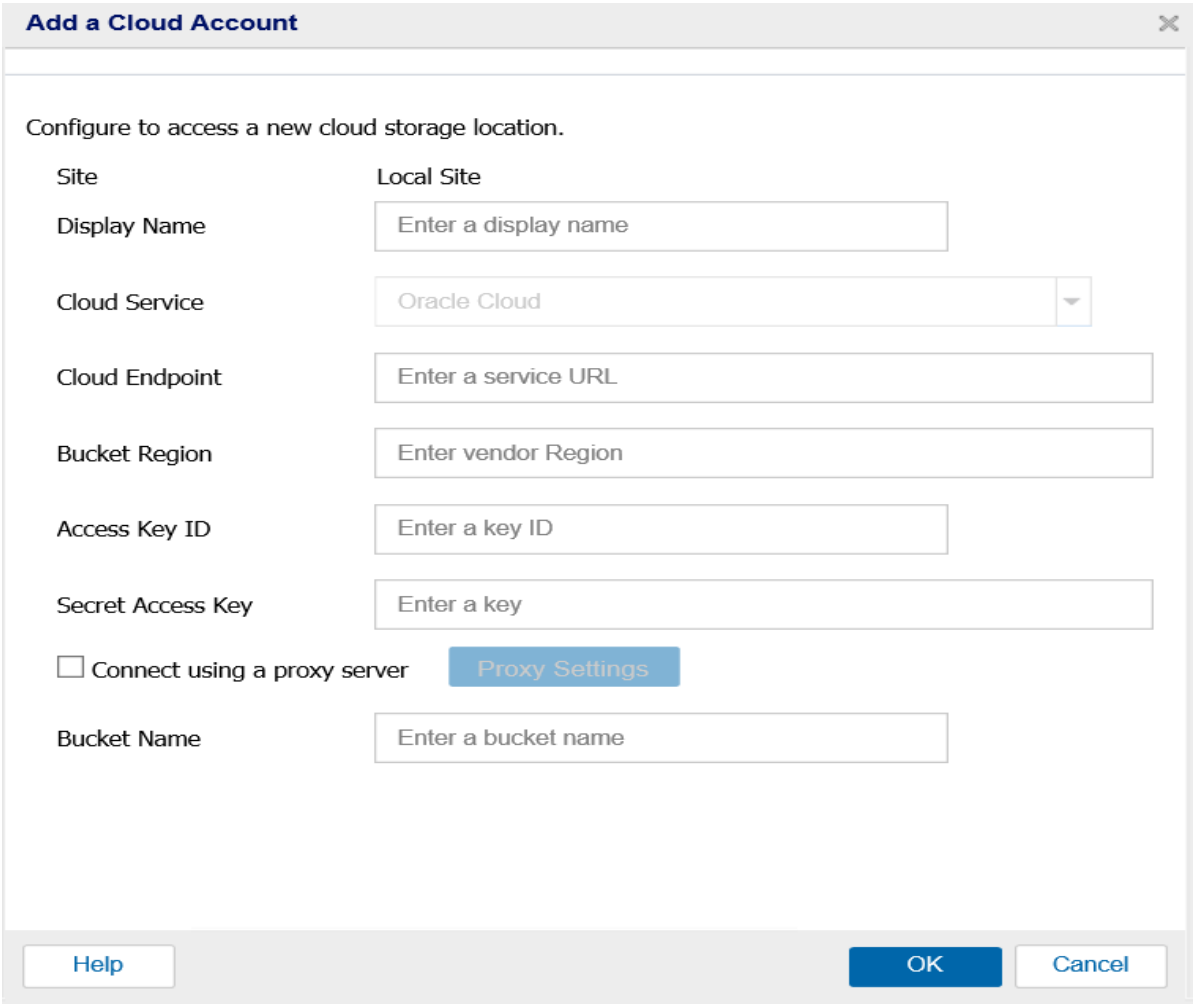
All files and folders moved or copied to the cloud vendor are stored and organized in your buckets. Buckets are like a container for your files and are used to group and organize objects together. Every object stored at the cloud vendor is placed in a bucket.

The cloud account is added to the Console.

Note: If you have configured File Copy and File Archive in versions Arcserve UDP v6.0 or previous versions and now you upgrade to the latest version of Arcserve UDP, then for File Archive Arcserve UDP creates a new cloud bucket suffixed with **-fa**. File Copy uses the same bucket that was used in the previous version.

Add a Cloud Account for Oracle Cloud

Add an Oracle cloud account to copy Files or recovery points to cloud storage. You can use this account while creating Copy Recovery Point / File Copy /File Archive task.



Follow these steps:

1. Log into Arcserve Unified Data Protection, and click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Cloud Accounts**.

   The **Destinations: Cloud Accounts** page is displayed in the center pane.

3. Click **Add a Cloud Account**.

   The **Add a Cloud Account** page is displayed.

4. For **Display Name**, provide a unique name.

   Display Name specifies the name of the cloud storage. This name will be added to Console for identifying the cloud account. Each cloud account must have a unique display name.

5. Select the option from the **Cloud Service** drop-down list.

   Multiple fields appear for configuration.

6. Enter details in the following fields to configure and click **OK**:

   Cloud Endpoint

   Specifies the Vendor service URL. For example, *https://-namespace.compat.objectstorage.region.oraclecloud.com*

   Bucket Region

   Refers to the vendor bucket region.

   Access Key ID

   Identifies the user who is requesting access to this location.

   Secret Access Key

   Refers to a password that is used to verify the authenticity of the request to access this location because your Access Key is not encrypted.

   Important! This Secret Access Key is crucial for maintaining the security of your accounts. You should keep your keys and your account credentials in a secure location. Do not embed your Secret Access Key in a web page or other publicly accessible source code and do not transmit it over insecure channels.

   Proxy Settings

   Specifies the proxy server settings. Select the check box of **Connect using a proxy server** if you want to enable this option. If you select this option, you must also include the IP address (or machine name) of the proxy server and the corresponding port number that is used by the proxy server for internet connections. You can also select this option if your proxy server requires authentication. Then, you must provide the corresponding authentication information (Domain Name\Username and Password) that is required to use the proxy server.

   Bucket Name

All files and folders moved or copied to the cloud vendor are stored and organized in your buckets. Buckets are like a container for your files and are used to group and organize objects together. Every object stored at the cloud vendor is placed in a bucket.

The cloud account is added to the Console.

Note: If you have configured File Copy and File Archive in versions Arcserve UDP v6.0 or previous versions and now you upgrade to the latest version of Arcserve UDP, then for File Archive Arcserve UDP creates a new cloud bucket suffixed with **-fa**. File Copy uses the same bucket that was used in the previous version.

Add an Arcserve Cloud Account

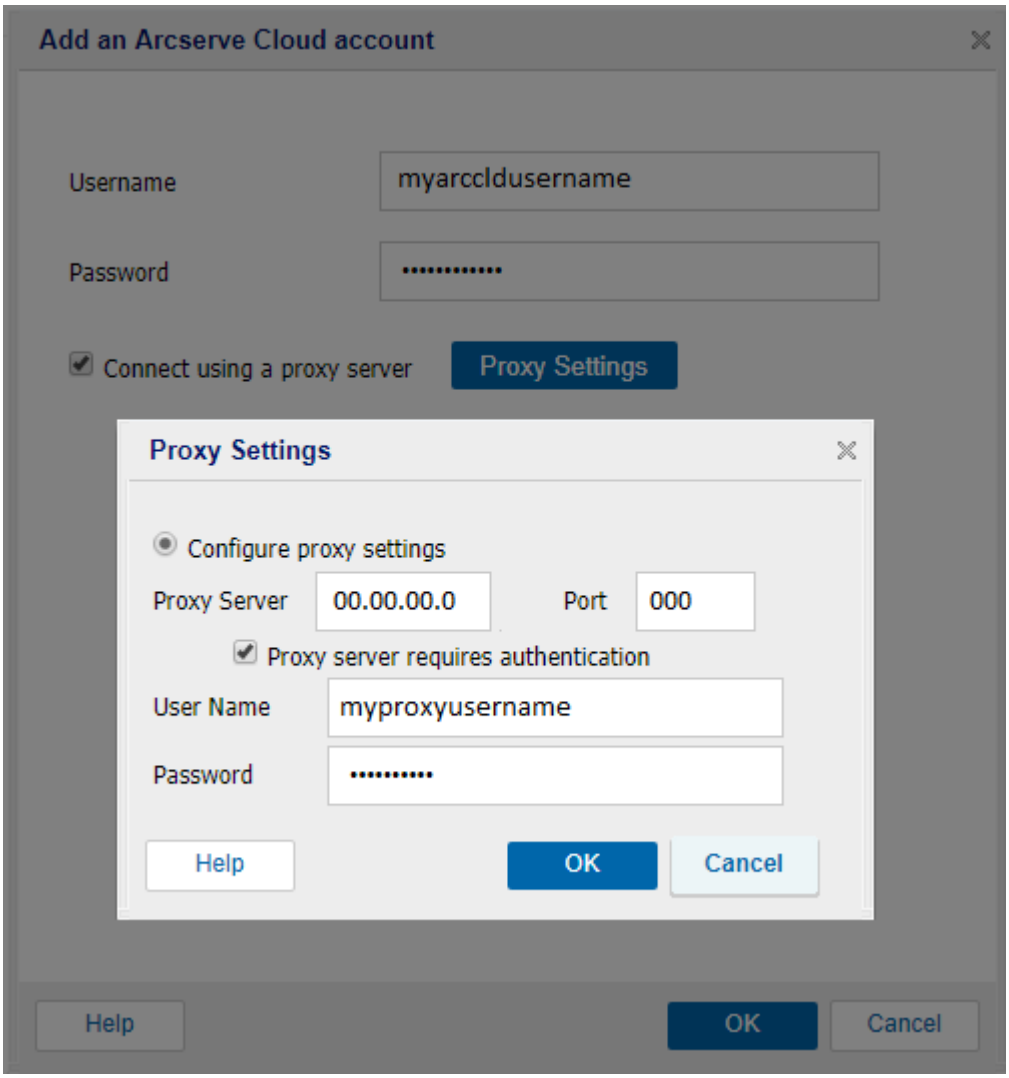
Add an Arcserve cloud account to copy files or recovery points to Arcserve cloud storage.

Follow these steps:

1. Log into Arcserve UDP, and click the **Resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Arcserve Cloud**.

   The Destinations: Arcserve Cloud page screen appears.
3. Click **Add an Arcserve Cloud account**.

   The Add an Arcserve Cloud account screen appears.



4. Enter the following details:

   Username

   Specifies user name of the Arcserve Cloud account that is registered with Arcserve.

   Password

   Specifies password of the corresponding user's Arcserve Cloud account that is registered with Arcserve.

   Connect using a proxy server

   Select this option to connect the Arcserve Cloud using a proxy server that has internet connection.

   Proxy Settings

   Click Proxy Settings if Connect using a proxy server is selected.

   Proxy Settings window appears. Enter the following details and click **OK**.

Proxy Server

Specifies the IP address of the proxy server.

Port

Specifies the port number that is open in the proxy server.

Proxy Server Requires Authentication

Select this option if you want the access to proxy server must require authentication.

User Name

Specifies user name that has access to proxy server.

Password

Specifies password of the corresponding user's account that has access to proxy server.

5. Click **OK**.

   The Arcserve Cloud Account is successfully added.

How to Manage a Data Store

After you create a data store, you may need to perform various operations such as modify, delete, stop, and start a data store.

You can also run on-demand merge jobs for multiple nodes to create more space on a data store.

What To Do Next?

- Review Prerequisites
- Modify a Data Store
- Delete a Data Store from the Console
- Stop a Data Store
- Start a Data Store
- Monitor Data Store Space Capacity
- Browse Recovery Points in a Data Store
- Delete Node Data from a Data Store
- Modify Concurrent Active Nodes Limit for Manual Backup
- Run a Manual or On-demand Merge Job
- Troubleshooting: How to Use a Data Store When the Backup Destination Folder is Full

Review the Prerequisites

To manage a data store, complete the following prerequisites:

- You have already added a data store.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Modify a Data Store

You can modify an already existing data store, however, there are some restrictions and you cannot modify the following details of a data store:

- Compression details
- Non-deduplication data store to a deduplication data store or deduplication data store to a non-deduplication data store.
- Deduplication options: Deduplicate Data and Deduplication Block Size.

Considerations before you modify a data store:

- If you change the path of the data store or the encryption password, all jobs running in that data store, including the jobs waiting in queue are canceled. Any change in the data store name, hash memory size, or concurrent active nodes number does not impact the running jobs.
- For non-deduplication data store: To change the data store path, keep the backup destination folder empty.
- For deduplication data store: To change the data store path, keep the following folders empty:
  - Backup Destination folder
  - Data Destination
  - Index Destination
  - Hash Destination
- The **Encryption Password** options are editable only if you selected the **Encrypt Data** option while creating the data store.

Follow these steps:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Recovery Point Servers**.

   The **Destinations: Recovery Point Servers** page displays the list of available recovery point servers.
3. Expand a recovery point server.

   You can see the list of data stores associated with the recovery point server.
4. Perform one of the following actions:

- Right-click the data store name.
- Select the data store, and from the center pane click the **Actions** drop-down list.

A list of options is displayed.

5. Click **Modify**.

   The **Modify a Data Store** page is displayed.

6. Update the required fields and click **Save**.

   Recovery Point Server

   Defines the recovery point server where the data store is created. The recovery point server is already added by default.

   Data Store Name

   Defines the name of the data store.

   Data Store Folder

   Defines the location of the folder where the data store is created. Click Browse to select the destination folder.

   Notes:

   - You can now enable and view hidden volumes. For more information, see How to Enable and View Hidden Volumes while Creating a Data Store.
   - For non-deduplication and deduplication data store, the destination path should be an empty folder.

   Concurrent Active Nodes Limit to

   Specifies the maximum concurrent jobs on the data store.

   Default Value: 4

   Refers to a value from 1 to 9999. The value indicates the number of jobs that can concurrently run. If the running jobs meet the number, another job is placed in to the queue and job can only start when one of the running job completes. The completed job could mean a finished, canceled, or a failed job.

   The number applies to the Job Types but not to the Server nodes. For example, number 5 indicates that five backup jobs are running. Any job scheduled after five backup jobs waits in the queue, but you can submit another job such as File System Catalog.

   If the value is more than 16 or 32, messages are displayed to warn about the increased demand on hardware.

Note: Limit to number only impacts the replication outbound job, not the replication inbound job. Limit to number does not impact the Restore or BMR jobs. Such jobs are not placed in a queue.

Enable Deduplication

Specifies that deduplication is enabled for this data store. Arcserve UDP supports both types of deduplication: Source-side deduplication and Global deduplication. Source-side deduplication prevents duplicate data blocks to move across network from a particular agent. Global deduplication eliminates duplicate data across all client machines based on the volume cluster level.

Deduplication Block Size

Defines the deduplication block size. The options are 4 KB, 8 KB, 16 KB, 32 KB, and 64 KB. The deduplication block size also impacts the Deduplication capacity estimation. For example, if you change the default 16 KB to 32 KB, the Deduplication capacity estimations double. Increasing the deduplication block size can decrease the deduplication percentage.

Hash Memory Allocation

Specifies the amount of physical memory that you allocate to keep hashes. This field is pre-filled with a default value. The default value is based on the following calculation:

If the physical memory of the RPS is smaller than 4 GB (or is identical to 4 GB), the default value of **Hash Memory Allocation** is identical to the physical memory of the RPS.

If the physical memory of the RPS is greater than 4 GB, Arcserve UDP calculates the available free memory at this time. Assume that the available free memory is X GB at present. Arcserve UDP further checks the following conditions:

- If (X * 80%) > = 4 GB, the default value of **Hash Memory Allocation** is (X * 80%).
- If (X * 80%) < 4 GB, the default value of **Hash Memory Allocation** is 4 GB.

Example: Consider the RPS has 32 GB of physical memory. Assume that operating system and other applications use 4 GB memory while creating the data store. So, the available free memory at this time is 28 GB. Then, the default value of **Hash Memory Allocation** is 22.4 GB (22.4 GB = 28 GB * 80%).

Hash Destination is on a Solid State Drive (SSD)

Specifies if the hash folder is on a solid state drive.

Note: Configure the hash destination on local SSD, if the Hash destination is on a Solid State Drive(SSD) option is enabled

Data Destination

Defines the data destination folder to save the actual unique data blocks. Use the largest disk to store data as that contains the original data blocks of source.

Note: The **Data Destination** path should be a blank folder.

Index Destination

Defines the index destination folder to store the index files. Choose a different disk to improve the deduplication processing.

Note: The **Index Destination** path should be a blank folder.

Hash Destination

Defines the path to store the hash database. Arcserve UDP uses the SHA1 algorithm to generate the hash for source data. The hash values are managed by the hash database. Selecting a high speed Solid State Drive (SSD) increases the deduplication capacity and requires a lower memory allocation. For better hash performance, we recommend to format the SSD volume as NTFS file system with 4KB volume cluster size.

Note: The **Hash Destination** path should be an empty folder.

Note: You cannot specify the same path for the following four folders: **Data Store folder, Data Destination, Index Destination,** and **Hash Destination**.

Enable Compression

Specifies that the data compression settings are enabled.

Compression Type

Specifies whether to use the standard or maximum compression type.

Compression is often selected to decrease the usage of the disk space, but also has an inverse impact on your backup speed due to the increased CPU usage. Based on your requirement, you can select one of the three available options.

Note: For more information, see Compression Type.

Enable Encryption

Specifies that encryption settings are enabled. When you select this option, you must specify and confirm the encryption password.

Data encryption is the translation of data into a form that is unintelligible without a deciphering mechanism. The Arcserve Unified Data Protection

solution uses secure, AES (Advanced Encryption Standard) encryption algorithms to achieve maximum security and privacy of your data. For data stores, encryption or No encryption is supported. For Encryption, only AES-256 is available.

A password is not required when you are attempting to restore to the computer from which the backup was performed. However, when you attempt to restore to a different computer, a password is required. By default, only for the first login password is required. To enter password even after the first login, the administrator needs to manually stop Arcserve UDP Agent Explorer Extension Service.

Send an email alert when a destination is nearing full capacity

Selecting this option configures the data store to send email alert. RPS sends out email alerts to recipients when data store destination folder is nearing full capacity.

Configure Email

This button appears only when you enable the option of *Send an email alert when a destination is nearing full capacity*. The button helps you provide as email ID to receive alerts. Click the **Configure Email** button to load global email alert settings from Console>Settings>Email and Alert configuration. If global Email settings is not available, clicking the **Configure Email** button opens **Email settings** dialog to set email details.

The data store is updated.

Modify the Data Store Threshold

This topic provides information about threshold registry locations, threshold key name and default values, when does a threshold error or warning message appears, and how to modify the data store threshold.

A data store has space capacity threshold configuration that helps to monitor the free space usage information of the data store destination. In a deduplication data store, the threshold monitors the memory allocated to the hash destination and the disk space allocated for the backup destination folder, index destination, and data destination. For a non-deduplication data store, the threshold monitors the storage space only for the backup destination folder. All the five items for the thresholds monitors have two types of values:

- Error Threshold
- Warning Threshold

The threshold value is saved in the system registry. You can manually modify the default value.

Thresholds registry locations

1. Data Store Folder: [HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\DataStore\XXXXXXX\CommStore]
2. Deduplication data destination:[HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\DataStore\XXXXXXX\GDD\DataRole]
3. Deduplication hash destination and memory: [HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\DataStore\XXXXXXX\GDD\HashRole]

   Note: Hash role monitors the memory and disk usage both. Path represents the disk usage and Mem represents the memory.
4. Deduplication index destination: [HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\DataStore\XXXXXXX\GDD\IndexRole]

Threshold key names and Default values

Path Type	Threshold Type	Registry Key Name	Default Value	Default Value (Before v6.5 Update 1)
Data store Folder	Warning	WarnPathThreshold	0.05	0.03
	Error	ErrorPathThreshold	2048	100

Deduplication Index Path	Warning	WarnPathThreshold	0.05	0.03
	Error	ErrorPathThreshold	2048	100
Deduplication Hash Path	Warning	WarnPathThreshold	0.05	0.03
	Error	ErrorPathThreshold	2048	100
Deduplication Data Role Path	Warning	WarnPathThreshold	0.05	0.03
	Error	ErrorPathThreshold	2048	100
Memory	Warning	WarnMemThreshold	0.05	0.03
	Error	ErrorMemThreshold	30	10

- Value less than 1 indicates free space percentage, For example 0.05 means 5%. If the free space is less than 5% of total space size, warning threshold is reached.
- Value larger than or equal to 1 indicates actual free space size and the unit is MB. For example, default 2048 means 2048MB. If the free space is less than 2048MB, error threshold is reached.

The default threshold value is designed for better performance of the data store. You can change the threshold values. We do not recommend changing these values, unless extra space is required.

Follow these steps to modify the threshold:

1. Navigate to the respective registry location.
2. Manually modify the default value of threshold.

Modify Only the Hash Destination

When a deduplication data store is changed, only the hash path destination can be changed to an empty folder. Arcserve UDP regenerates the hash path for the new data store. This process is useful when the hash folder runs out of space and all the jobs get canceled. You can modify the data store and provide a new hash destination folder.

How to Switch the Hash Destination Modes

When you create a deduplication data store, you specify whether the hash destination is on a Solid State Drive (SSD mode) or the hard disk drive (RAM mode). If you configured hard disk as the hash destination, you need more memory to process hash keys. As a result when your backup size grows, all your memory may get exhausted. In that case, you can add an SSD to back up more data. Similarly if you had configured an SSD as the hash destination, you need less memory to process hash keys. However, if you are moving to a higher memory machine, you might want to switch to the RAM mode for a faster hash processing.

To switch the hash destination from a RAM to SSD or SSD to a RAM, Arcserve UDP lets you modify an existing data store and change the mode as required.

You can modify an existing data store even when it is running but the data store restarts after you save the change.

Changing from the RAM to SSD Mode

When you switch from the RAM to SSD mode, you would need less memory. So, Arcserve UDP automatically decreases the minimum value of "Hash Memory Allocation". However, you can manually change Hash Memory Allocation. For this case, you change the hash destination folders to SSD. When you change the hash destination, Arcserve UDP automatically copies the hash files to the new location on SSD.

Changing from the SSD to RAM Mode

When you switch from the SSD to RAM mode, the RAM should be large enough to accommodate the current hash database. For example, before the change, the data store created 30 GB of hash files on SSD. Now after the change, you should allocate at least a 30 GB memory for hash files. If the RAM is not enough, the switch fails. In this case, Arcserve UDP automatically increases the following two parameters:

- Minimum value of Hash Memory Allocation
- Hash Memory Allocation

This ensures that data store starts after the modification.

For this case, you change the hash destination folders to the hard disk drive. When you change the hash destination, Arcserve UDP automatically copies the hash files to the new location on hard disk drive.

Start a Data Store

If you have stopped a running data store for any routine maintenance check, then you can start the data store again after the maintenance check is over. When you start the data store, the pending jobs will start from the point they were paused.

Note: To start a deduplication data store, depending on the Hash size, the hash data takes time to load from the hard disk to memory. On the right pane, the progress of the data store is displayed in percentage.

Follow these steps:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Recovery Point Servers**.

   The **Destinations: Recovery Point Servers** page displays the list of available recovery point servers.
3. Expand a recovery point server.

   You can see the list of data stores associated with the recovery point server.
4. Perform one of the following actions:
   - Right-click the data store name.
   - Select the data store, and from the center pane click the **Actions** drop-down list.

   A list of options is displayed.
5. Click **Start**.

   The right pane displays the information that the data store is starting. The status icon of the selected data store changes from **Stopped** to **Running**.

Stop a Data Store

If you do not want a data store to run, use the stop option. Stopping the data store ensures that no job is running on it.

Notes:

- If you stop a data store, all the jobs running, including the jobs waiting in queue, on that data store are canceled.
- If you stop a data store while a replication job is in progress, then on restarting the data store, the replication job starts from the same point at which you stopped the data store.
- If you stop the data store while a replication job (for example, Job-10) is in progress, and by that time two more backup jobs complete (for example, Job-11, Job-12), then when you restart the data store, the replication jobs complete in a sequence (Job-10, Job-11, Job-12, respectively).

Follow these steps:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Recovery Point Servers**.

   The **Destinations: Recovery Point Servers** page displays the list of available recovery point servers.
3. Expand a recovery point server.

   You can see the list of data stores associated with the recovery point server.
4. Perform one of the following actions:
   - Right-click the data store name.
   - Select the data store, and from the center pane click the **Actions** drop-down list.

   A list of options is displayed.
5. Click **Stop**.

   The **Confirm** dialog opens.
6. Select **Yes** to stop.

   The right pane displays the information that the data store is stopping.

   The data store stops and the status icon for the selected data store changes from **Running** to **Stopped**.

Delete a Data Store from the Console

If you no longer want to use a data store, you can delete the data store. When deleted, the data store is removed from the Console. However, the deleted data store exists in the recovery point server.

Notes:

- You can import the deleted data store, when required.
- To delete a data store that is linked to plans, first delete the plan that is linked to the data store.

Follow these steps:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Recovery Point Servers**.

   The **Destinations: Recovery Point Servers** page displays the list of available recovery point servers.
3. Expand a recovery point server.

   You can see the list of data stores associated with the recovery point server.
4. Perform one of the following actions:
   - Right-click the data store name.
   - Select the data store, and from the center pane click the **Actions** drop-down list.

   A list of options is displayed.
5. Click **Delete**.

   A **Confirm** dialog opens.

   Note: If the data store is linked to a plan, instead of the Confirm dialog you get a **Warning** dialog.
6. Click **Yes**.

   The data store is removed.

Delete Node Data from a Data Store

As a storage administrator, you may want to delete backed-up node data from a data store to free up space and effectively manage your storage space. Arcserve UDP lets you select the node data in a data store and delete. You can select multiple nodes in a data store. You can delete any type of node data including the encrypted and deduplicated data. The data store should be in the running state when you start this job, called as Purge job.

Note: You can delete only when the node is not running any job (For example, Backup / Restore / Merge job).

Follow these steps:

1. From the Console, click the **resources** tab.
2. Click the data store that contains the node data that you want to delete.
3. The **Recovery Points Summary** page appears.
4. Select the node that you want to delete.

   Warning! Only select, do not click the node name. Clicking the node name leads to another page from where you cannot perform the delete action.
5. Click **Actions**, **Delete**.

   Note: If you are using deduplication data store, after deleting the recovery points of a node, the disk space may not be released immediately. The delay happens because for deduplication data store:

   - The backend purge and reclamation process is scheduled at the idle time of data store.
   - Part of the data blocks of the deleted recovery points of that node may still be used by recovery points of other nodes.

     The volume of reclaiming the disk space depends on the deduplication ratio of the node that you want to delete.
6. Confirm that you want to delete the node data.

   The purge job is initiated and the node data is deleted from the data source. You can see the status of purge job from **Recent Events** and logs.

Monitor Data Store Space Capacity

Arcserve UDP data stores are created on RPS volume or remote shared disks. The data store space capacity monitor mechanism helps to monitor the destination folder usage, when destination is nearing full capacity. The free space in destination folder is not enough for running data store. An email alert can be sent out if the data store has enabled email alert configuration, and activity log provides the out-put. Meanwhile, running data stores is changed to restore only status. As a result, new backup / replication jobs are blocked from backing up the new backup data.

Monitoring includes the following options:

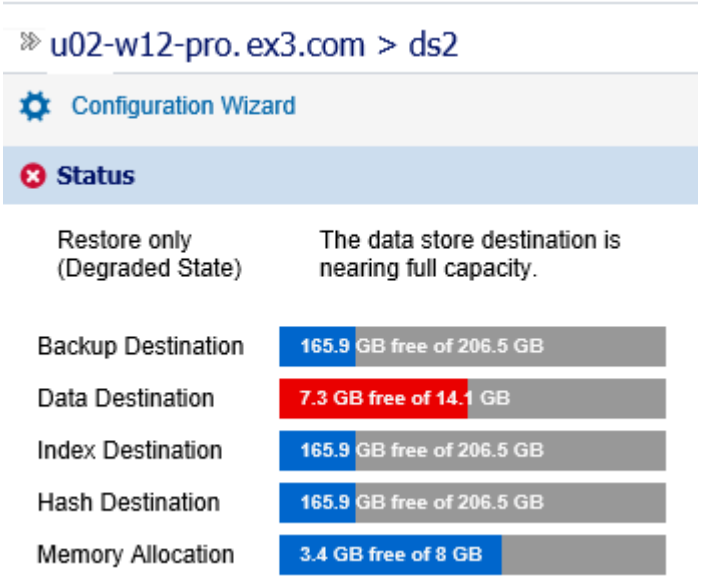
- Modify the Data Store Threshold
- Error Threshold usage
- Warning Threshold Usage
- Enabling Email Alert and Understanding Email Format

Error Threshold usage

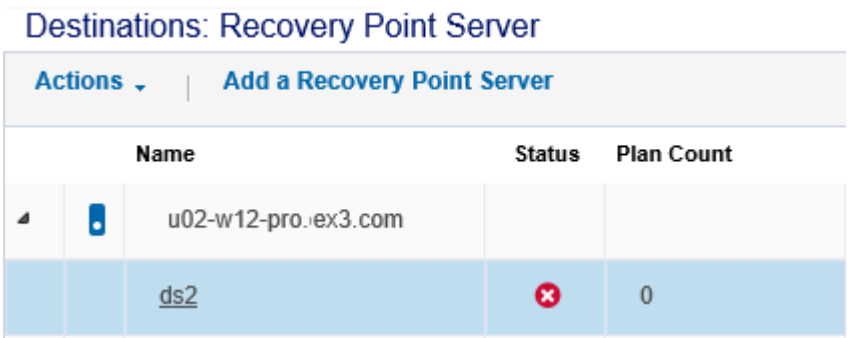
Error threshold indicates absence of required free space for the running data store. Also, some kernel components cannot work properly. In this situation, the data store is changed to restore only status automatically if the data store is not stopped manually.

When data store free space reaches error threshold:

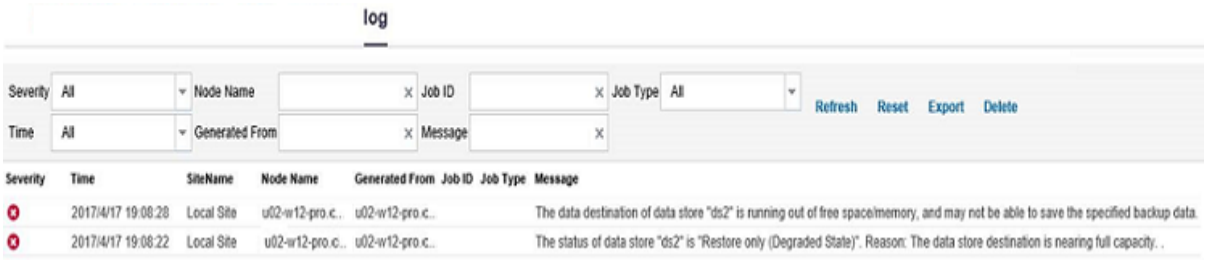
- The data store status is automatically changed to the “Restore Only (Degraded Status)” status. Meanwhile, any new backup / replication jobs cannot be triggered to this data store. Any other type of jobs that do not generate new data could still be triggered. For example, restore jobs and merge jobs.
- The data store space bar is marked with red color.



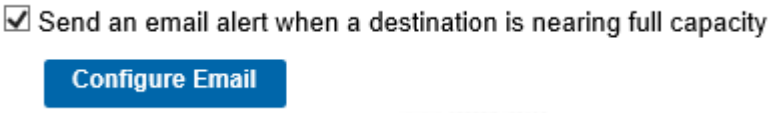
- The data store status is marked with the red icon.



- Activity logs appears to inform that the data store destination folder has run out of free space, and may not be able to save new backup data and reaches the error level.



- If the data store is configured to send email alert when destination folder is nearing full capacity, RPS sends out email alerts to recipients.

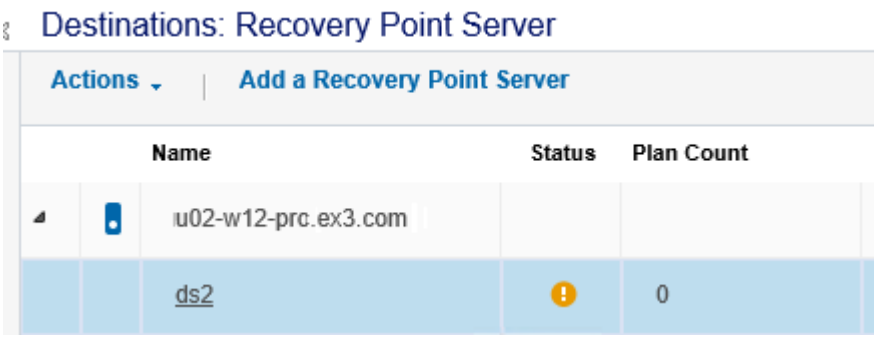


Warning Threshold Usage

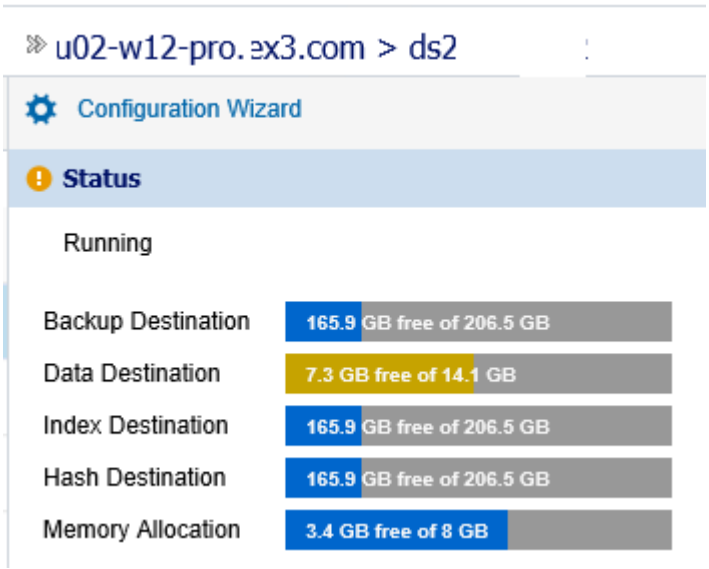
Warning threshold indicates that some free space is left in the destination. But, the space may exhaust in short time if the new backup / replication jobs are still working.

When data store free space reaches warning threshold:

- The data store status is not impacted and still keeps the original status. For example, the “running” status..
- The job schedule is not impacted.
- The data store status is marked with the yellow exclamation.



- The data store space bar is marked with yellow color.



- An activity log appears to inform that the data store destination folder is close to the maximum capacity and reached the warning level.



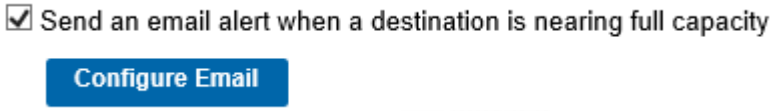
- If the data store is configured to send email alert when destination folder is nearing full capacity, RPS sends out email alerts to recipients.

Enabling Email Alert and Understanding Email Format

When data store destination is nearing full capacity or the free space of destination folder reaches the warning threshold or error threshold, email alert is sent to special recipients to notify about the situation. To receive email alert, you need to configure email settings and enable the option to receive email alerts.

To enable the configuration, follow these steps:

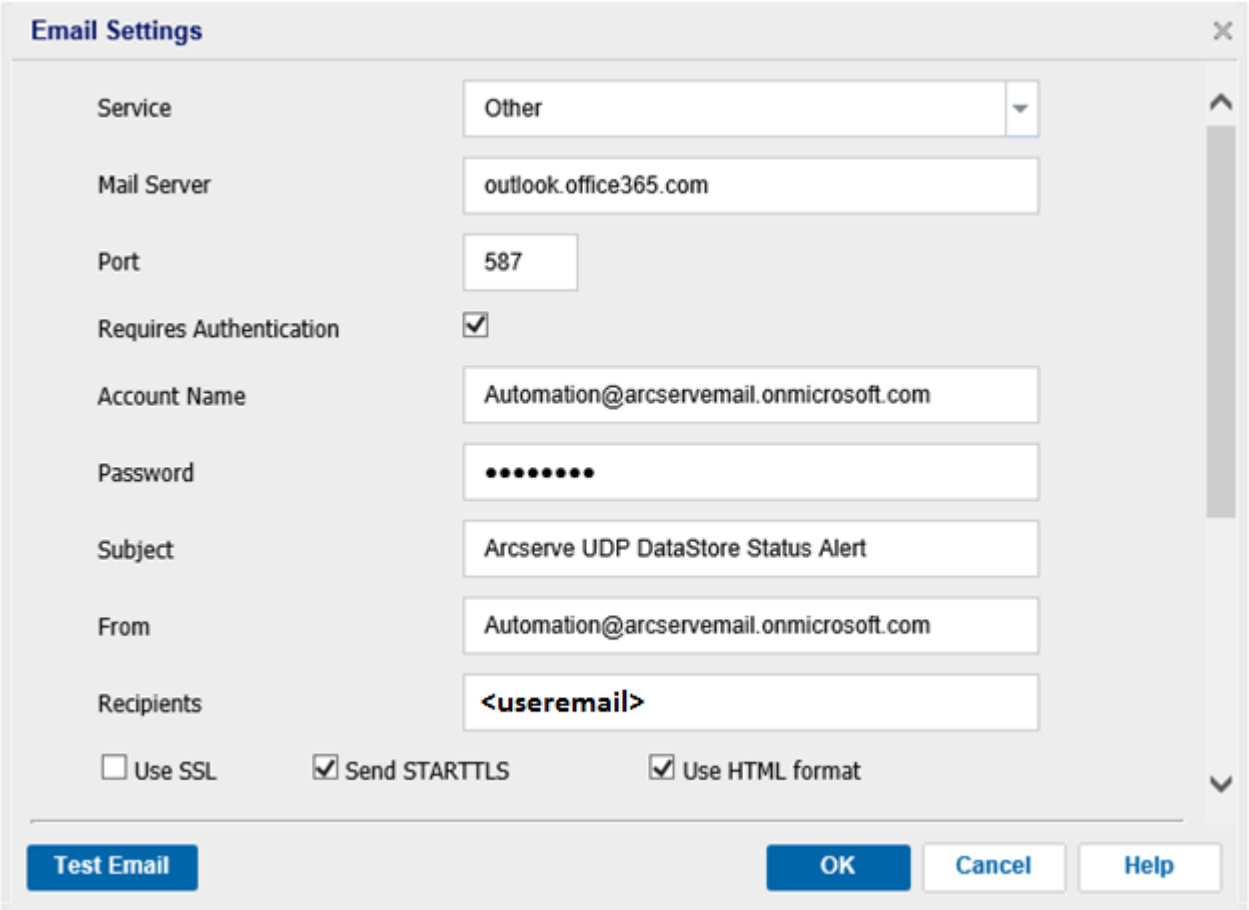
1. Select the check box of **Send an email alert when a destination is nearing full capacity** while adding / modifying / importing the data store.

   Configure Email button appears.

2. Click **Configure Email**.

   Email Settings dialog appears.

3. Enter details in Email Settings, and click **OK**.

Email is configured to receive all alerts for data store.

Note: Test Email verifies if the Recipients email ID is receiving email. We recommend to click and test before closing this dialog.

Email alert is sent to configured recipients when:

- Data store path capacity is changed from healthy to warning range (reaches warning threshold value).
- Data Store path capacity is changed from warning range to error range (reaches error threshold value).
- Data store destination folder releases some space (such as, after data store merge job, enlarging destination volume space, and so on) and leaves warning / error threshold to healthy range.

Format of Alert Received in Email

Arcserve UDP Data Store Capacity Error Alert

Arcserve UDP Recovery Point Server	u02-w12-pro	
Data Store Name:	ds2	
Data Store Path:	e:\ds2\c	
Level:	Error	
Backup Destination:	Total Size:	211454 MB
	Free Size:	169924 MB (80.36%)
Data Destination:	Total Size:	14399 MB
	Free Size:	7435 MB (51.64%)
Index Destination:	Total Size:	211454 MB
	Free Size:	169924 MB (80.36%)
Hash Destination:	Total Size:	211454 MB
	Free Size:	169924 MB (80.36%)
Memory Allocation:	Total Size:	8191 MB
	Free Size:	3441 MB (42.01%)
Occurrence Time:	2017/4/17 19:38:01	

To view current data store capacity, access address: https://u02-w12-pro.ex3.com

See more about data store capacity configuration, click here.

Browse Recovery Points in a Data Store

You can use the **Browse Recovery Points** option to view the details related to recovery points and the plans associated with that data store. For example, you can view the details related to data store settings and recent events.

You can also browse the recovery points from the Shared Folders view.

To delete a node from a data store, see Delete Node Data From a Data Store.

Follow these steps:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Recovery Point Servers**.

   The **Destinations: Recovery Point Servers** page displays the list of available recovery point servers.
3. Expand a recovery point server.

   You can see the list of data stores associated with the recovery point server.
4. Perform one of the following actions:
   - Right-click the data store name.
   - Select the data store, and from the center pane click the **Actions** drop-down list.

   A list of options is displayed.

   Note: You can also click the name of a Data store to browse the data store.
5. Click **Browse Recovery Points** from the options displayed after selecting a data store.

   The page for the selected data store appears with the summary displaying information about **Recovery Points**. For example, the page displays information related to **Datastore Settings** and **Recent Events**.
6. To update information about the plan or data store, select the plan or that data store, and click **Actions**, **Refresh**.
7. To restore, select the Agent node, and click **Actions**, **Restore**.

   You can see the **Restore** dialog box where you can opt for the restore option that you want to perform for the data store.

FQDN Support for a Data Store

When creating a data store in a local disk, Arcserve UDP creates a shared folder for the data store, so that the jobs running on other servers can access the data.

In Arcserve UDP Version 5.0, the folder was shared with the hostname (for example, <hostname>\sharename). As a result, when the remote server could access only the RPS with FQDN (Fully qualified domain name), it could not access the RPS with hostname, and the jobs failed.

Now, the data store folder can be shared with FQDN or IP address. This allows the remote server to access the data store.

To achieve this, update the Recovery Point Server on the Console to FQDN or IP address, and create new data store.

Note: For an existing database, stop the data store and import it again. You can overwrite the data store instead of deleting it. Then, re-deploy the plans that use the data store.

Run a Manual or On-demand Merge Job

As a storage administrator, you can run an on-demand merge job, and specify the number of recovery points to retain. The values can be different than the Plan settings. The merge job deletes the selected backup sessions from a data store to free up the space and effectively manages your storage space. You can select multiple nodes and run the on-demand merge job.

Note: For merge job, when a replication task is configured and you run an on-demand merge job from the source data store, the job does not check whether the sessions are replicated or not. As a result, the merged sessions cannot be replicated to target data store and you end up replicating more data. For example, consider there are five sessions, s1, s2, s3, s4, and s5 respectively. s1 and s2 are replicated. Now, you run an on-demand merge job on the source side and retain two sessions. s4 and s5 are retained. s4 is a full session. So, when the next replication job starts, the job needs to replicate a full session.

Follow these steps:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Destinations** and click **Recovery Point Servers**.
3. From the center pane, click the data store that contains the node data that you want to merge.

   The **Recovery Points Summary** page is displayed.
4. Select the node that you want to merge.
5. Click **Actions**, **Merge Now**.

   Note: To merge multiple nodes from a data store, press the Ctrl or Shift key and select the nodes, then click **Action**, **Merge Now**.

   The **Run a Merge Now** dialog is opened
6. Specify the number of recovery points that you want to retain and click **OK**.

   The On-demand merge job is initiated and the node data is deleted from the data source. You can see the status of merge job from Recent Events and logs.

Modify Concurrent Active Nodes Limit for Manual Backup

To submit manual backups concurrently for multiple nodes, not controlled by the *Concurrent Active Nodes Limit* option, add one registry key on Recovery Point Server. Limits for concurrent active nodes is configured in the Data Store user interface.

Follow these steps:

1. Open the Windows Registry Editor opens. To open, you can click **Start**, type *regedit* in the Search programs and files field, and press **Enter**.

   Note: You may need to provide administrative credentials to open Windows Registry Editor.

2. Locate and click the following registry key:

   HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine

3. From the Edit menu, click **New**, and then click **DWORD (32-bit) Value**.

   You can see the list of data stores associated with the recovery point server.

4. Specify ManualJobsIgnoreConcurrentJobLimit as the name for the new entry and then press **Enter**.

5. Right-click ManualJobsIgnoreConcurrentJobLimit and then click **Modify**.

6. Specify 1 in the **Value** data field and then click **OK**.

7. Exit the Registry Editor.

   Note: After being enabled with 1, the registry key ManualJobsIgnoreConcurrentJobLimit will work for all Data Stores on the RPS server.

Troubleshooting: How to Use a Data Store When One or Multiple folders are Full

Symptom:

How do I continue to use the data store when one of the following folders is full:

- Data store backup destination
- Deduplication index
- Hash
- Data

Solution:

Change Backup Destination Path

1. Stop data store.
2. Modify the plans that are using this data store. In the Destination tab, point the destination to a different data store or a Shared Folder, select the Check box of **Pause this plan** and **Save** the plan.
3. Delete the Data Store.
4. Copy the corresponding Backup Destination folder into a large volume. Verify that size and number of files on the Source and Target folders are same.

   Note: Verify that before copying the folder, you stop the data store. During copy, ensure that all files are copied to the destination folders, without skipping any files.

   Warning! Skipping any files may result in data corruption. If the number of files or size between Source and Target folders does not match, do not proceed to next step.
5. Import data store. Specify the new Data / Index path while importing the data store.
6. From the UDP Console, navigate to Resources, Destinations, Recovery Point Servers.
7. Select the Data Store under the Recovery Point Server.
8. From the Right Pane, under settings verify that the Data / Index Destination Path points to correct location.

   Important! While performing the above mentioned steps, DO NOT start the data store until Step 8 is complete.

9. Start the Data Store.
10. Modify the plans that you paused and reconfigure them to use the imported Data Store (as per the requirement). Cleary selection of the option **Pause this plan** and **Save** the plan.

    Note: After importing data store to link to the new location and starting the data store, **DO NOT switch back** to the original path. Such action may cause data corruption.

Change Data Or Index Destination Path

1. Stop data store.
2. Copy the corresponding folder (Data/Index as per the requirement) into a large volume. Verify that the size and number of files on the Source and Target folders are same.

   Note: Verify that before copying the folder, you stop the data store. During copy, ensure that all files are copied to the destination folders, without skipping any files.

   Warning! Skipping any files may result in data corruption. If the number of files or size between Source and Target folders does not match, do not proceed to next step.
3. Import data store. Specify the new Data / Index path while importing the data store.
4. From the UDP Console, navigate to Resources, Destinations, Recovery Point Servers.
5. Select the Data Store under the Recovery Point Server.
6. From the Right Pane, under settings verify that the Data/Index Destination Path points to correct location.

   Important! While performing the above mentioned steps, DO NOT start the data store until Step 6 is complete.
7. Start the Data Store.

   Note: After importing data store to link to the new location and starting the data store, **DO NOT switch back** to the original path. Such action may cause data corruption.

Change Hash Destination Path

For details, refer to Modify Hash Path.

How to Manage a Recovery Point Server

Using Arcserve UDP, you can perform various operations on the existing recovery point server such as update, delete, import and upgrade.

The recovery point server is displayed under **Name** on the **Destinations: Recovery Point Server** page. Click the **Actions** tab or the name of the recovery point server on the **Destinations: Recovery Point Server** page to receive all the options to manage your recovery point server.

What To Do Next?

- Review the Prerequisites
- Update a Recovery Point Server
- Delete a Recovery Point Server from the Console
- Import a Data Store
- Install/Upgrade Recovery Point Server

Review the Prerequisites

To manage a recover point server, complete the following prerequisites:

- Log into the Console.
- Add a recovery point store.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Update a Recovery Point Server

When the credentials or protocol is changed for the recovery point server, you must update the recovery point server. Otherwise, the recovery point server fails to function properly.

Note: If a node acts as both recovery point server and agent, and you change the credentials or protocol of that node, then update the node from the **Destinations: Recovery Point Server** page. The plan will automatically deploy to the agent after you update the recovery point server. If you update the node from the **Nodes: All Nodes** page, then the plans involving those nodes are not deployed successfully. To deploy the plan, update the node from the **Destinations: Recovery Point Server** page again.

Follow these steps:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Recovery Point Servers**.

   The **Destinations: Recovery Point Servers** page is displayed.
3. Perform one of the following actions:
   - Right click a recovery point server.
   - Select a recovery point server, and from the center menu click the **Actions** drop-down list.
4. Click **Update**.

   The **Update Node** dialog opens.
5. Modify the details as desired, and click **OK**.

   The recovery point server is updated.

Delete a Recovery Point Server from the Console

To remove a recovery point server from the Console, use the **Delete** option.

Note: When you remove a recovery point server, the associated data stores are not deleted. A recovery point server that is used in any plan cannot be deleted.

Follow these steps:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Recovery Point Servers**.

   The **Destinations: Recovery Point Servers** page is displayed.
3. Perform one of the following actions:
   - Right click a recovery point server.
   - Select a recovery point server, and from the center menu click the **Actions** drop-down list.
4. Click **Delete**.

   The **Confirm** dialog opens.
5. Click **Yes**.

   The recovery point server is deleted.

Import a Data Store

The **Import Data Store** feature lets you add a data store to the recovery point server. You can import any existing data store to a recovery point server. The data stores that you have deleted earlier from a recovery point server are available to import.

Note: When the hash data of a deduplication data store is missed or corrupted, you can still import the data store. Provide an empty folder as hash folder. The data store, then, starts in the **restore only** status and an error message shows that hash role is not working for the empty hash folder. Using this method, you can use the data store only for restore jobs. Rebuild the hash if you want to run the backup job.

Follow these steps:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Recovery Point Servers**.

   The **Destinations: Recovery Point Servers** page is displayed.
3. Perform one of the following actions:
   - Right click a recovery point server.
   - Select a recovery point server, and from the center menu click the **Actions** drop-down list.
   - To receive email alerts, select the checkbox of *Send an email alert when a destination is nearing full capacity*.

     Note: If you have not configured before, then click **Configure email** to provide your Email Settings.
4. Click **Import Data Store**.

   The **Import a Data Store** page is displayed.
5. Perform the following actions, and click **Next**:
   - **Browse** to select the **Backup Destination Folder** from where you want to import the data store.
   - Enter **Encryption Password**.

   Note: Leave it empty if the data store is not encrypted.

   After authenticating the **Backup Destination folder**, the **Import a Data Store** page displays the details of the data store.
6. Modify the details, if necessary, and click **Save**.

If you have copied folder of Data Destination, Index Destination, and Hash Destination for Deduplication data store, change the folder path.

Note: You cannot enable or disable the encryption option for an existing data store.

The data store is added to the recovery point server and displayed at the **Destinations: Recovery Point Servers** dialog.

Install/Upgrade Recovery Point Server

Use the **Install/Upgrade Recovery Point Server** option for the following reasons:

- When the installation fails.
- When you want to upgrade the product.

Follow these steps:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Recovery Point Servers**.

   The **Destinations: Recovery Point Servers** page is displayed.
3. Perform one of the following actions:
   - Right click a recovery point server.
   - Select a recovery point server, and from the center menu click the **Actions** drop-down list.
4. Click **Install/Upgrade Recovery Point Server**.

   The install path details appear on the same page above the list of the added recovery point server.
5. Update the details as required.
6. Specify the install/upgrade schedule and click **OK**.

   The install or upgrade starts per the schedule. You can view the install or upgrade progress on the right pane.

   Note: You can cancel a recovery point server deployment if it is scheduled for a later time. To cancel a recovery point server deployment, select the agent and click Actions, Cancel Agent Deployment.

How to Manage Arcserve Backup Servers

You can manage Arcserve Backup Servers from the Arcserve UDP Console. You can also update and delete the Arcserve Backup Server from the Console.

Update an Arcserve Backup Server

When the credentials or the Arcserve web service port of the Arcserve Backup Server is changed, you must update the same on the Console.

Follow these steps:

1. Log in to the Console and click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Arcserve Backup Servers**.

   The **Destinations: Arcserve Backup Servers** page is displayed on the center pane.
3. Select an Arcserve Backup Server, right-click, and select **Update**.

   The **Update an Arcserve Backup Server** page opens.
4. Update the required fields and click **Save**.

   The **Update an Arcserve Backup Server** page closes.

   You have successfully updated the Arcserve Backup Server.

Delete an Arcserve Backup Server

If you no longer need an Arcserve Backup Server, you can delete that server from the Console. Before deleting an Arcserve Backup Server, ensure that the Arcserve Backup Server is not included in any task in a plan. If the server is included in any task, either delete the task or change the backup destination.

Follow these steps:

1. Log in to the Console and click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Arcserve Backup Servers**.

   The **Destinations: Arcserve Backup Servers** page is displayed on the center pane.
3. Select an Arcserve Backup Server, right-click, and select **Delete**.

   The confirmation dialog opens.
4. Click **Yes**.

   You have successfully deleted the Arcserve Backup Server from the Console.

How to Manage Shared Folder

You can manage Shared Folders from the Arcserve UDP Console. You can also modify and delete the Shared Folder from the Console.

Modify a Shared Folder

Arcserve UDP lets you modify details of an added shared folder. If the information related to shared folder changed, you need to modify the shared folder added in Arcserve UDP for plans to function.

Follow these steps:

1. Log into Arcserve Unified Data Protection, and click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Shared Folders**.

   The **Destinations: Shared Folders** page displays added shared folders in the center pane.
3. To manage, right click one of the added shared folders or select a shared folder and click **Actions**.
4. From the displayed options, click **Update**.

   The **Update** dialog is displayed.
5. Modify the information and click **OK**.

   Now, the shared folder is updated.

Delete a Shared Folder

Arcserve UDP lets you delete an added shared folder.

Follow these steps:

1. Log into Arcserve Unified Data Protection, and click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Shared Folders**.

   The **Destinations: Shared Folders** page displays added shared folders in the center pane.
3. To manage, right click one of the added shared folders or select a shared folder and click **Actions**.
4. From the displayed options, click **Delete**.

   A **Confirmation** or **Error** dialog is displayed.
5. For confirmation message, click **OK** to delete.
6. (Optional) For Error message, resolve the error and try deleting again.

   Now, the shared folder is updated.

How to Manage a Remote Console

After you create a remote console, you may need to perform various operations such as modify, delete, or test connection.

What To Do Next?

- Modify a Remote Console
- Delete a Remote Console
- Test Remote Console Connection

Modify a Remote Console

Arcserve UDP lets you modify details of an added remote console. If the connection information changed for the remote console, you need to modify the remote console account added in Arcserve UDP.

Follow these steps:

1. Log into Arcserve Unified Data Protection, and click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Remote Console**.

   The **Destinations: Remote Console** page displays added remote console in the center pane.
3. To manage, right click one of the added remote consoles or select a remote console and click **Actions**.

   The **Modify a Remote Console** page is displayed.
4. Modify the information and click **OK**.

   Now, the remote console is updated.

Delete a Remote Console

Arcserve UDP lets you remove a remote console when required.

Follow these steps:

1. Log into Arcserve Unified Data Protection, and click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Remote Console**.

   The **Destinations: Remote Console** page displays added remote console in the center pane.
3. To manage, right click one of the added remote consoles or select a remote console and click **Actions**.

   A **Confirm** or **Error** message is displayed.

   Confirm message

   Appears when the remote console is not added to any plan.

   Error message

   Appears when the remote console is part of an existing plan.
4. (Optional) When the Error message appears:
   a. Click **OK** to close the Error message.
   b. Open the plan with which the remote console is associated.
   c. Delete the related plan or modify the plan to change the remote console in the destination tab of **Replicate to a remotely-managed RPS** plan.

   After removing the remote Console from the plan, return to the Remote Console page and try deleting the remote console again.
5. From the Confirm message, click **Yes**.

   Now, the remote console is removed.

Test Remote Console Connection

Arcserve UDP lets you test the connection of an added remote console. A successful connection for the remote console connection is required for the Replication plan to work as the remote console account is associated with the plan.

Remote console connection may fail due to multiple reasons:

- When the information added for a remote console destination is incorrect.
- When the remote console credential, port, protocol or proxy information has changed but the same information is not updated in the Arcserve UDP Console.
- When the network connection breaks between this console and the remote console.
- When the Arcserve UDP management service is stopped in the remote console.

You can use the test remote console connection option to verify if a remote console is connected and if the information added for the account is correct.

Follow these steps:

1. Log into Arcserve Unified Data Protection, and click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Remote Console**.

   The **Destinations: Remote Console** page displays added remote console in the center pane.
3. To manage, right click one of the added remote consoles or select a remote console and click **Actions**.

   Multiple options appear.
4. From the multiple options displayed, click **Test Remote Console Connection**.

   An **Information** message appears if the details of Remote Console are correct.

   An **Error** message appears if the details of Remote Console are incorrect. Provide the correct details to fix the connection and test again.
5. (Optional) If an error message appears, fix the broken connection by verifying one or all of the following reasons:
   - The account information is correct.
   - The network connection is established.

- Status of the Arcserve UDP Management service in the remote console is not stopped.

Now, the connection of remote console is successfully tested.

Chapter 10: Creating Plans to Protect Data

This section contains the following topics:

Understanding Plan and Tasks

To protect a node, you need to create a plan with a backup task. A plan is a group of tasks to manage backup, replication, and creation of virtual standby nodes. A plan consists of a single or multiple tasks. Tasks are a set of activities to define the source, destination, schedule, and advanced parameters.

You can create the following tasks:

Backup Tasks

Lets you create a backup task to protect Windows, Linux, and host-based virtual machine nodes. Based on the type of nodes you want to protect, use one of the following backup tasks:

Agent-Based Windows Backup

Defines a backup task to protect Windows nodes. In an agent-based backup method, an agent component is used to back up data. The agent is installed on the source node.

Host-Based Agentless Backup

Defines a backup task to protect host-based virtual machines in a VMware vCenter/ESX or Microsoft Hyper-V server. In an agentless backup method, you do not need to install an agent component on either the server or the virtual machine. However, you have to install the agent on a proxy server.

Agent-Based Linux

Defines a backup task to protect Linux nodes. The agent is installed on a Linux Backup Server and not on the source nodes that you want to protect.

Replicate from a remote RPS task

Lets you create a task to receive data from a remote recovery point server.

Replicate task

Lets you create a task to replicate backup data from a recovery point server to another recovery point server.

Virtual Standby task

Lets you create a task to create a virtual standby node.

File Copy task

Lets you copy selected files from the source node and store the copied files in a local or shared folder. You can also store the files in a cloud storage.

Copy Recovery Points task

Lets you copy the recovery points to a local or shared folder or Cloud.

Replicate to a remotely-managed RPS task

Lets you create a task to replicate or send data to a remote recovery point server.

File Archive task

Lets you copy the recovery points to a network share, cloud storage, or volume on a protected node. When the recovery points get copied to the destination, the source files are deleted.

Copy to Tape task

Lets you copy the recovery points to a tape with granular recovery capability directly from tape.

Assured Recovery task

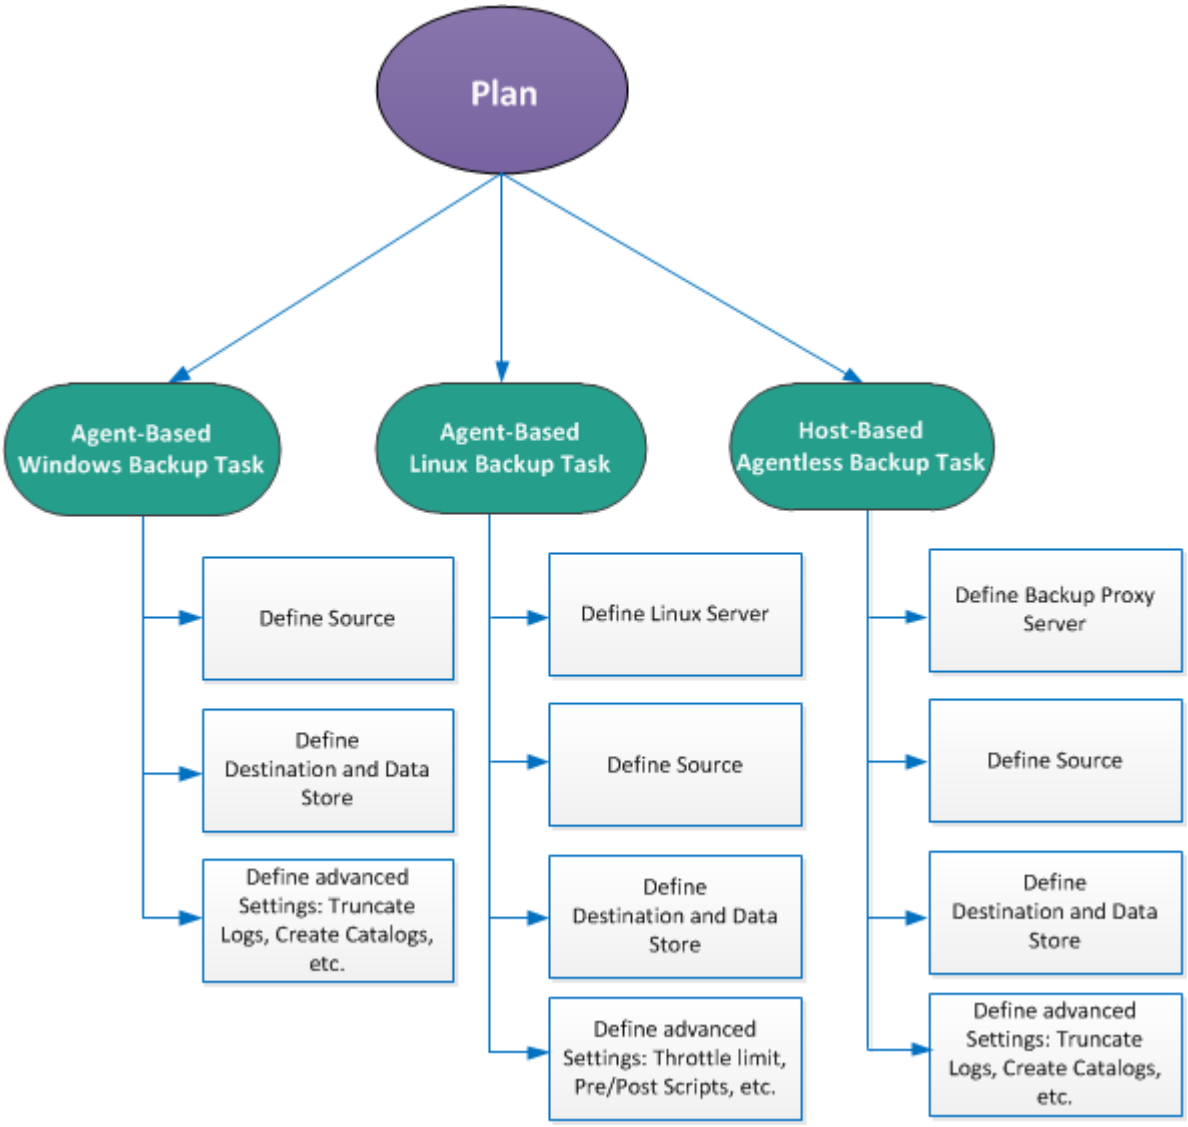
Lets you verify accessibility and assure recovery of the data.

The following table displays the list of follow-up tasks that you can add after Task 1:

Task 1	Follow-up Tasks
Backup: Agent-Based Windows	■ Replicate ■ Virtual Standby ■ Copy Recovery Points ■ File Copy ■ Replicate to a remotely-managed RPS ■ File Archive ■ Copy to Tape ■ Assured Recovery Test
Backup: Host-Based Agentless	■ Replicate ■ Virtual Standby ■ Copy Recovery Points ■ Replicate to a remotely-managed RPS ■ Copy to Tape ■ Assured Recovery Test
Backup: Agent-Based Linux	■ Replicate ■ Replicate to a remotely-managed RPS ■ Copy to Tape ■ Assured Recovery Test
Replicate data from a remote RPS	■ Virtual Standby ■ Replicate

	■ Assured Recovery Test
Backup: Office 365 Exchange Online	■ Replicate ■ Copy Recovery Points ■ Replicate to a remotely-managed RPS ■ Copy to Tape ■ Assured Recovery Test
Backup: Office 365 OneDrive	■ Replicate ■ Copy Recovery Points ■ Replicate to a remotely-managed RPS ■ Copy to Tape ■ Assured Recovery Test
Backup: Office 365 SharePoint Online	■ Replicate ■ Copy Recovery Points ■ Replicate to a remotely-managed RPS ■ Copy to Tape ■ Assured Recovery Test
Backup: Files on UNC or NFS Path	■ Replicate ■ Copy Recovery Points ■ Replicate to a remotely-managed RPS ■ Copy to Tape ■ Assured Recovery Test

The following diagram illustrates how different tasks form a backup plan. The diagram also shows parameters that you can define in each task.


Plan
Agent-Based Windows Backup Task
Agent-Based Linux Backup Task
Host-Based Agentless Backup Task
Define Source
Define Destination and Data Store
Define advanced Settings: Truncate Logs, Create Catalogs, etc.
Define Linux Server
Define Source
Define Destination and Data Store
Define advanced Settings: Throttle limit, Pre/Post Scripts, etc.
Define Backup Proxy Server
Define Source
Define Destination and Data Store
Define advanced Settings: Truncate Logs, Create Catalogs, etc.

How to Create a Windows Backup Plan

To protect your Windows nodes or clustered nodes, you need to create a plan. The plan for Windows nodes consists of a backup task. This backup task lets you specify the nodes you want to protect, the backup destination, and the backup schedule. The backup destination is a recovery point server where you want to store your backup data. The destination can also be a local destination or a remote share folder.

You can also back up an Oracle database. Before you create a plan to back up an Oracle database, review the following prerequisites:

- Prerequisite to back up an Oracle database

To backup Microsoft clustered nodes and shared disks, review the following pre-requisites:

- Review the Prerequisites to Back Up Microsoft Clustered Nodes and Shared Disks

What To Do Next?

1. Review the Prerequisites and Considerations
2. Create a Backup Plan
3. (Optional) Perform a Manual Backup
4. Verify the Backup

Review the Prerequisites and Considerations

Verify that you have completed the following prerequisite tasks:

- The latest version of Arcserve UDP Agent (Windows) is installed.
- Log into the Console.
- (Optional) Create data store to store the backup data.
- Review the prerequisites to back up an Oracle database.
- Review the Prerequisites to back up Microsoft clustered nodes and shared disk.
- (For backup of SQL when database is in full mode) Review How to enable log truncations when SQL Database is in full recovery mode.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

The following prerequisites are for hardware snapshots:

- Install a VSS hardware provider that supports hardware snapshot on the Arcserve UDP Agents. A typical configuration of a VSS hardware provider includes:
  - Specifying a server that controls the LUN.
  - Specifying the disk array credentials to access the disk array.

  Note: For more information on configuring the VSS hardware provider, contact your hardware provider vendor.

The following prerequisites are for specific network backup:

- Windows Agent and RPS must be in the same network.
- Windows protection node is scheduled.
- Local RPS with multiple NICs is connected.

Review the Prerequisites for Oracle Database

This section contains the following topics:

Review the Prerequisites for Oracle Database with UDP Windows Agent

To back up an Oracle database with consistent data, ensure that the ARCHIVELOG mode is enabled to archive the Redo logs.

Note: The data volume must include Oracle data files, control files, server parameter file, and online redo logs. The archived redo logs must be physically located on a separate volume.

Follow these steps to verify if the ARCHIVELOG mode is enabled:

a. Log into the Oracle server as an Oracle user with SYSDBA privileges.

b. Enter the following command at the SQL*Plus prompt:

ARCHIVE LOG LIST;

Archive log settings for the current instance is displayed.

c. Configure the following settings:

Database log mode: Archive Mode

Automatic archival: Enabled

d. Start the ARCHIVELOG mode.

Note: If the ARCHIVELOG mode is not enabled, start the ARCHIVELOG mode to back up the database.

Follow these steps to start the ARCHIVELOG mode:

a. Shut down the Oracle server.

b. Run the following statements in Oracle:

CONNECT SYS/SYS_PASSWORD AS SYSDBA

STARTUP MOUNT;

ALTER DATABASE ARCHIVELOG;

ALTER DATABASE OPEN;

By default, archive logs are written to the flash recovery area. If you do not want to write archive logs to the flash recovery area, set the LOG_ARCHIVE_DEST_n parameter to the location where you want to write archive logs.

SQL>ALTER SYSTEM SET LOG_ARCHIVE_DEST_
1='LOCATION=e:\app\administrator\oradata\<oracle_database_name>\arch'
SCOPE= BOTH;

System altered.

SQL> ARCHIVE LOG LIST;

Archive log settings for the current instance is displayed.

c. Configure the following settings:

Database log mode: Archive Mode

Automatic archival: Enabled

Archive destination: E:\app\oracle\oradata\<oracle_database_name>\arch

Oldest online log sequence: 21

Current log sequence: 23

The Oracle VSS writer service started and is functioning properly.

Note: If Oracle VSS Writer Service is not running, Arcserve UDP Agent (Windows) will automatically start it before taking the snapshot.

- Arcserve UDP Agent (Windows) is installed and a plan is scheduled.

  Ensure that you have selected the volumes that include all the Oracle data files, server parameter file, control files, archived redo logs, and online redo logs for the backup.

- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

If you want to perform a BMR for a disaster recovery, ensure that you have selected the system volumes and the volumes which includes all the oracle installation files.

Review the Prerequisites for Oracle Backups with RMAN (Windows and Linux)

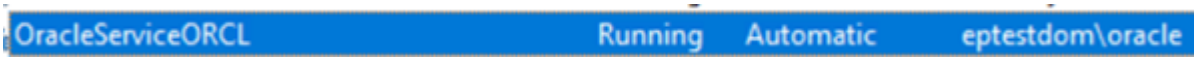
Review the following prerequisites before you restore the Oracle database items to alternate server:

Prerequisites for Configuring UDP and Oracle Windows

Add Plan

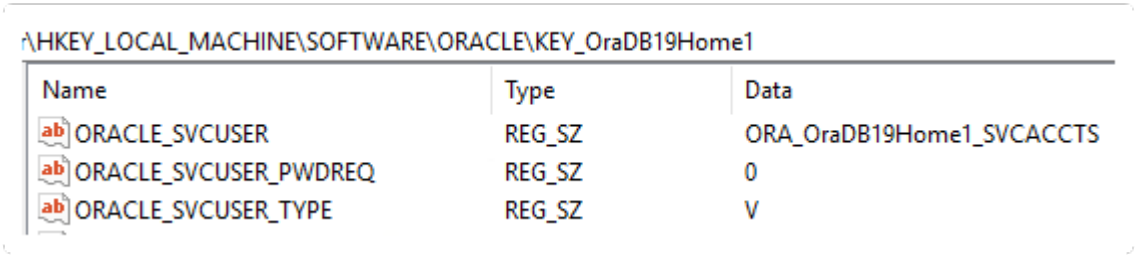
- The Oracle node(s), UDP Console and RPS must be in the same domain. When you log into the UDP console, do the following:
  1. Login as the "<domain>\Administrator" user, and then install the UDP and RPS systems.
  2. Log into the UDP Console and RPS using the "<domain>\Administrator" domain ID.
  3. Add the Oracle node into the UDP console using the "<domain>\Administrator" user ID.

     The "<domain>\Administrator" user is added into local "ORA_DBA" group in the Windows Oracle node(s).
- Supports below Oracle accounts used during Oracle installations (applicable for both Stand-alone and RAC).
  - Domain User: Make a note of domain user who owns the Oracle service (service account). For example, eptestdom\oracle.

  - Virtual Account: When Oracle is installed using a Virtual account, check for parameters in registry of the Oracle server before adding database to the UDP plan.

    Note: By default, the Virtual Accounts are enabled in Windows.

- Make sure to appropriately configure Listener, which helps to protect the Oracledatabase and connects to sqlplus and RMAN for that database.

Example: “rman target / “ must point to right database.

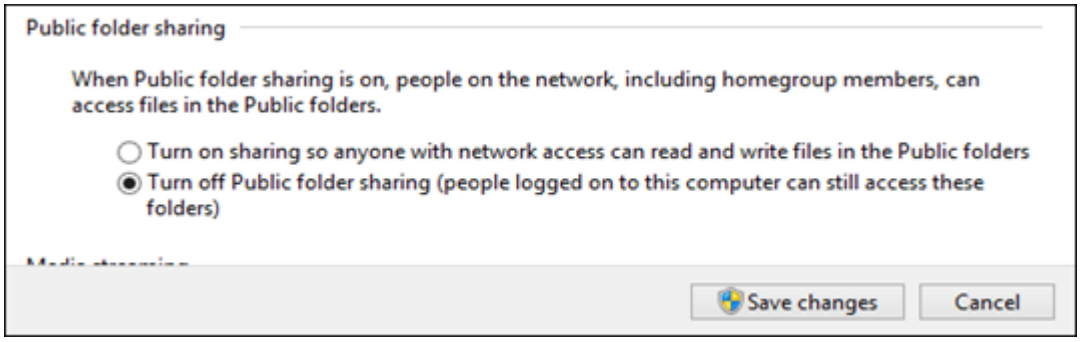
Backup

- When you backup the Oracle database(s), make sure that it is in OPEN state with the ARCHIVELOG mode enabled. For more information, see How to enable ARCHIVELOG mode.
- To successfully perform the Oracle RMAN backups with Arcserve UDP, you must enable Windows native authentication. To confirm whether the Windows authentication is enabled, make sure the below line is added in sqlnet.ora file in the ORACLE_HOME/network/admin path.

  SQLNET.AUTHENTICATION_SERVICES= (NTS)
- Public Folder Sharing Sharing is disabled to enable secure access to the database(s).

  To disable the use of Public folder for network sharing, follow these steps:

  1. Go to Network and Sharing Center, and then click **Change advanced sharing settings**.
  2. On the left pane, search for Public folder sharing setting for the active network location, and then click **Turn off Public folder sharing (people logged on to this computer can still access these folders)**.
  3. Click **Save changes**.



Restore Tables

Table level restores uses Auxiliary database, which requires the following permissions:

- Domain User for Oracle Installation: Add the service account Oracle user to the “Administrators" group in Oracle node(s). In our example, eptestdom\oracle is added to local administrator group.
- Virtual Account for Oracle Intallation: Open the registry key: HKLM\SYSTEM\CurrentControlSet\services\eventlog\Application, and then

right click and add permissions to the group such as 'ORA_OraDB18Home1_SVCACCTS'.

Restore to alternate Node

- The same version of Oracle with the same configuration is installed on the target system.
- The database is installed with the same SID as your source (original) machine.
- The Oracle database on the destination node must be in shutdown state.

Firewall settings to allow access to Windows Oracle node(s)

netsh advfirewall firewall add rule name="AS ora" dir=in action=allow program="C:**path\to\oracle\home**\bin\oracle.exe" enable=yes	
netsh advfirewall firewall add rule name="AS tns" dir=in action=allow program="C:**path\to\oracle\home**\tnslsnr.exe" enable=yes	
netsh advfirewall firewall add rule name="AS ora" dir=out action=allow program="C:**path\to\oracle\home**\bin\oracle.exe" enable=yes	
netsh advfirewall firewall add rule name="AS tns" dir=out action=allow program="C:**path\to\oracle\home**\bin\tnslsnr.exe" enable=yes	
netsh advfirewall firewall add rule name="ora_1521" dir=in action=allow protocol=TCP localport=**1521**	
netsh advfirewall firewall add rule name="ora_1521" dir=out action=allow protocol=TCP localport=**1521**	Port number may change based on the port configured.

Prerequisites for Oracle RMAN in Arcserve UDP Linux

Add Plan

- The UDP Console and RPS must be in the same domain. When you log into the UDP console, do the following:
  1. Login as the "<domain>\Administrator" user, and then install UDP and RPS systems.
  2. Log into the UDP Console and RPS using the "<domain>\Administrator" domain ID.
- The Oracle node /etc/hosts must have an entry for FQDN and “resolv.conf” with the appropriate DNS entry.
- Every database that needs to be protected must have an entry in the /etc/oratab file.

  Note: In RAC environment, make sure the content is updated for both the nodes.

  The syntax in /etc/oratab file is as follows:

  <DB_UNIQUE_NAME>:<ORACLE_HOME>:<ACTIVE/INACTIVE>

  Example: ORCL:/u01/app/oracle/ORCL/:Y
- Supported Oracle installations: All platforms are supported including the Virtual Account.

Backup

- When you backup the Oracle database(s), make sure that it is in OPEN state with the ARCHIVELOG mode enabled. For more information, see How to enable ARCHIVELOG mode.
- Install smbclient using the following command:

  #yum install smbclient
- Make sure to appropriately configure Listener, which helps to protect the Oracle database and connect to sqlplus and RMAN for that database.

  Example: “rman target / “ must account to right database.

Assured Recovery Jobs

- Make sure that the Oracle databases for backup is in MOUNT state or OPEN state.

Review the Prerequisites to Back Up Microsoft Clustered Nodes and Shared Disks

Review the following prerequisite steps when backing up Microsoft Clustered Nodes and Shared Disks:

- Install the Arcserve UDP Agent on all the clustered nodes.
- Add all agents or nodes into the same backup plan.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Note: The shared disks will be backed up along with the agent which owns the shares disks. If the shared disk is moved from Node A to Node B during a failover, then for the next backup job on Node B, the disk will be backed up as a full disk even though the job itself appears as an incremental. After another failover if the shared disk moves back to Node A, even then the disk will be backed up as a full disk even though the job itself appears as an incremental.

How to enable Log Truncations when SQL Database is in Full Recovery Mode

Symptom

When the database is in the Full mode and a full database backup is performed, the SQL truncation log cannot be truncated.

Solution

To resolve this problem, add two registry values to enable Arcserve UDP run the BACKUP LOG command to back up the transaction log. This command marks the space, which is already written to database file, as reusable.

Follow these steps to add the registry value:

1. Open the registry table editor on the agent machine using the following command:

   ```
   regedit
   ```

2. Navigate to the following keys depending on the agent-based or agentless backup:

   For agent-based backup for both 32-bit and 64-bit OS, navigate to the following key on the agent machine:

   HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll

   When using a version lower than Arcserve UDP v6.5 Update 2, then for agentless backup navigate to the below key. Create the registry table value inside the VM that you want to back up on the proxy server. If the registry table key is not available, then create the complete key path.

   - **32-bit OS:**

     HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll

   - **64-bit OS:**

     HKEY_LOCAL_MACHINE\SOFTWARE\WoW6432Node\Arcserve\Unified Data Protection\Engine\AFBackupDll

3. Create the following two registry values and for both set the value to 1:

   - dword value named BackupSQLLog4Purge
   - dword value named ForceShrinkSQLLog

     The registry value is added.

   The solution is in effect when the next purge job occurs.

Create a Backup Plan with a Backup Task

A backup plan includes a backup task that performs a backup of a physical node and stores data to a specified destination. Each task consists of parameters that define the source, destination, schedule, and other backup details.

Follow these steps:

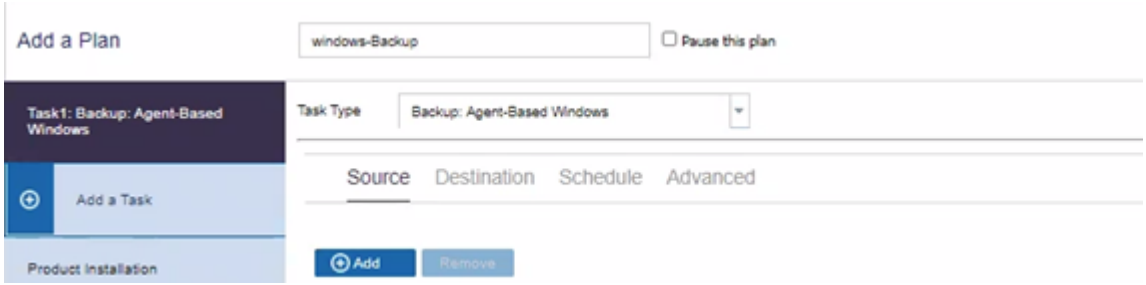
1. Log into Arcserve UDP Console.
2. Navigate to **resources** > **Plans** > **All Plans**.

   If you have created plans earlier, those plans are displayed on the center pane.
3. On the center pane, click **Add a Plan**.

   The Add a Plan page opens.
4. Enter a plan name.
5. (Optional) Select the **Pause this plan** check box.

   The plan will not run until you clear the check box to resume the plan.

   Note: If a plan is paused, then any in-progress job is not paused but all corresponding scheduled jobs associated with that plan are paused. However, you can manually run a job. For example, you can manually run backup job and replication job for a node even if the respective plan is paused. In such case, the following task to the on-demand (manual) job does not run. For example, there is a replication task after an on-demand backup job, the replication job does not run for the on-demand backup job. You need to manually run the replication job. When you resume the plan, the pending jobs do not resume immediately. After you resume the plan, the pending jobs run from the next scheduled time.
6. From the Task Type drop-down list, select **Backup: Agent-Based Windows**.



Now specify the Source, Destination, Schedule, and Advanced details.

Specify the Source

The Source page lets you specify the source nodes that you want to protect. You can select more than one node in a plan. If you have not added any nodes to the Console, you can add nodes when you create or modify a plan from the Source page. You can also save a plan without adding any source nodes. The plan gets deployed only after you add source nodes.

Follow these steps:

1. On the Source tab, click **Add**.
2. Select one of the following options:

   Select Nodes to Protect in Arcserve UDP

   Opens the Select Nodes to Protect in Arcserve UDP dialog. Select the nodes from the displayed list. Select this option if you have already added the nodes to the Console.

   Add Windows Node

   Opens the Add nodes to a plan dialog. Select this option if you have not added the nodes and you want to manually add the nodes to protect.

   Discover Nodes from Active Directory

   Opens the Add Nodes to Arcserve UDP Console dialog. Select this option if you want to discover and add nodes from the Active Directory.

3. If you selected the **Select Nodes to Protect in Arcserve UDP** option, on the Select Nodes to Protect dialog, do the following:

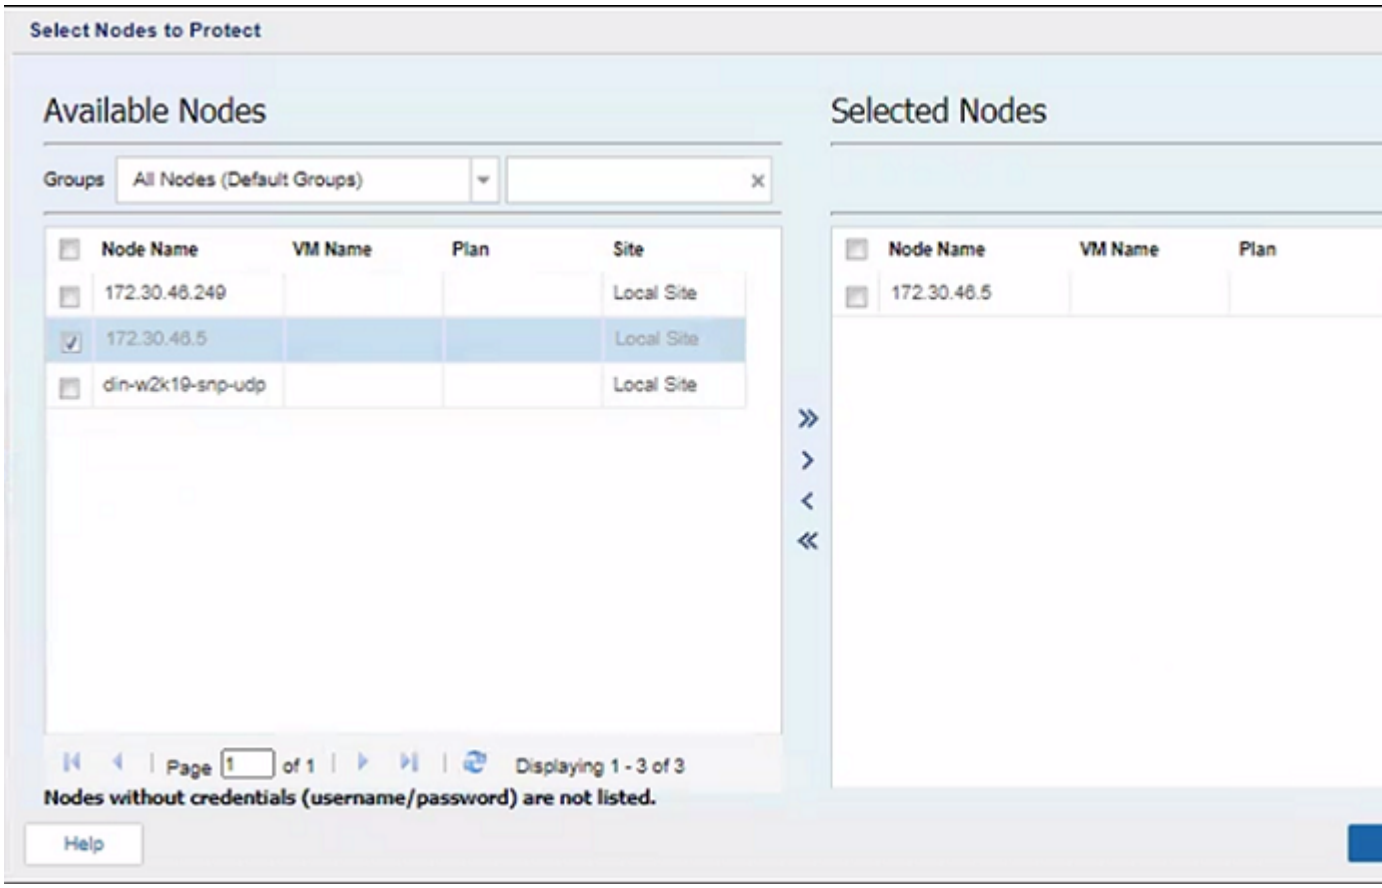
   a. (Optional) Select a filter from the Groups drop-down list to filter nodes. You can enter keywords to further filter your nodes.

      The nodes are displayed on the Available Nodes column.

   b. Select the nodes from the Available Nodes column, and then click the single right arrow (>) or double right arrow (>>) to move the selected nodes or all the nodes respectively to the Selected Nodes column.

      The nodes are displayed in the Selected Nodes column.

   c. Click **OK** to close the dialog.



4. To choose the Protection Type, select one of the following options:

   Back up all volumes

   Prepares a backup snapshot of all the volumes.

   Back up selected volumes

   Prepares a backup snapshot of the selected volume.

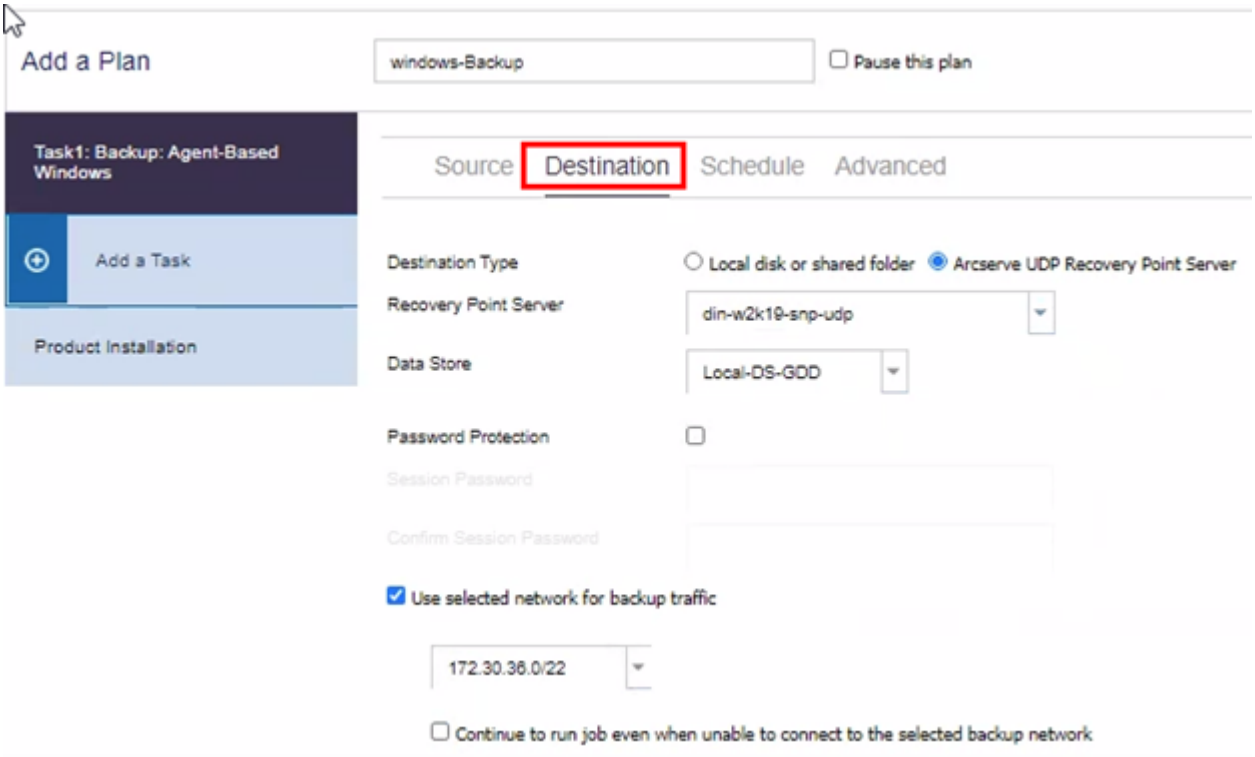
   The source is specified.

Specify the Destination

The destination is a location where you store the backup data. You must at least specify the destination to save the plan.

Follow these steps:

1. Click the **Destination** tab.



2. Select one of the following destination types:

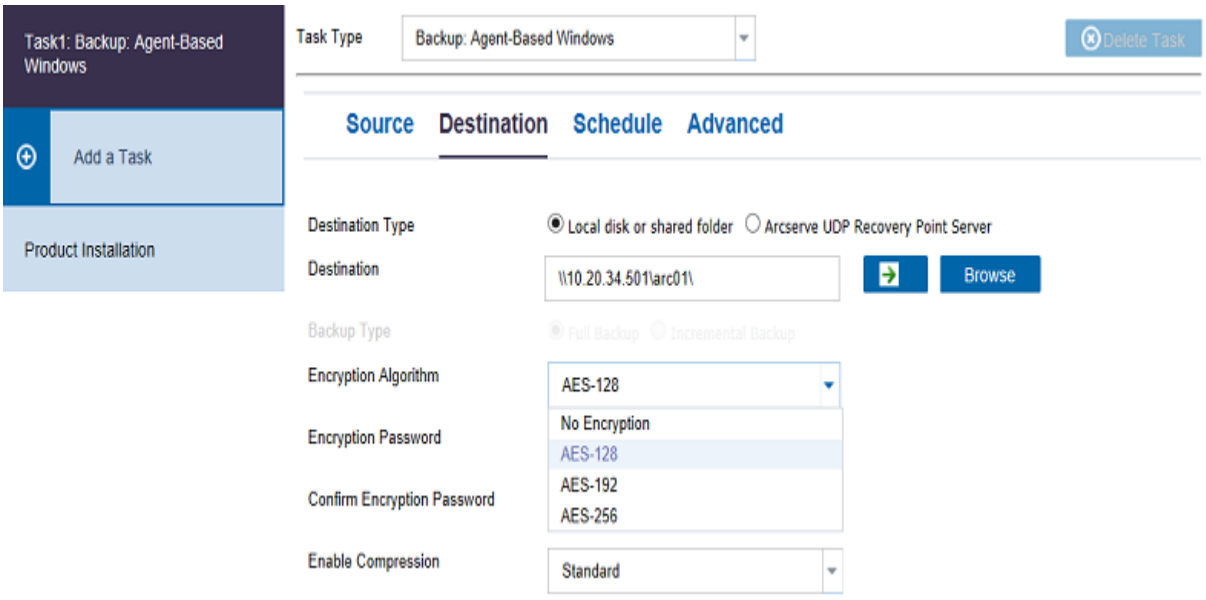
   Local disk or shared folder

   Specifies that the backup destination is a local destination or a shared folder. If you select this option, you can save data as either recovery points or recovery sets. The recovery points and recovery sets options are available on the Schedule tab.

   Arcserve UDP Recovery Point Server

   Specifies that the backup destination is a recovery point server. If you select this option, then data is stored as recovery points. You cannot store data as recovery sets.

3. If you select the **Local disk or shared folder** as a Destination Type, provide the following details:



- Destination: Provide the full path of the local or network destination. For the network destination, specify the credentials with the write access. Click **Browse** to locate the destination or click the forward arrow icon to test connection and provide the credentials for the folder destination provided.
- Encryption Algorithm: From the list of drop-down options, select the encryption algorithm. For more information, see Encryption Settings.
- (Optional) Encryption Password: Provide an encryption password.
- Confirm Encryption Password: Type the encryption password again to confirm.
- Enable Compression: Select a type of compression from the drop-down list. For more information, see Compression Type.

Note: If you store the data to a local disk or shared folder, you cannot replicate the data to another recovery point server. Replication is supported only if you store the data to a recovery point server.

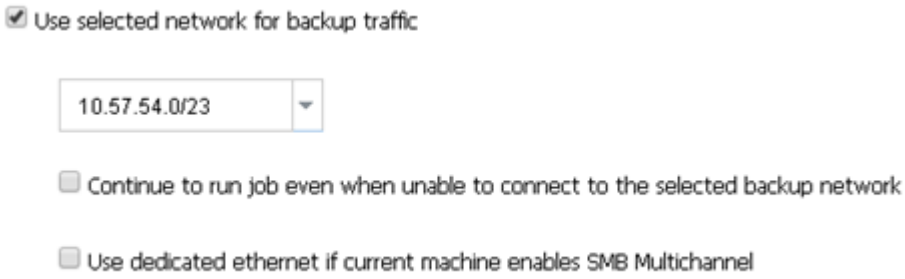
4. If you select the **Arcserve UDP Recovery Point Server** as a Destination Type, do the following:
   - Recovery Point Server: From the drop-down list, select the destination Recovery Point Server (RPS).
   - Data Store: From the drop-down list, select the destination data store.

- Password Protection: Select the **Password Protection** check box, provide the session password, and then retype the session password to confirm.

  Note: The session password is optional when the backup destination is an unencrypted RPS data store..

- (Optional) Select the **Use selected network for backup traffic** check box and follow these steps:

  a. To enable communication between the Windows Agent and Recovery Point Server, from the drop-down list, select the CIDR network.

  Note: If the selected backup network is not accessible and to continue the job with the available network or with the default network, click the **Continue to run job even when unable to connect to the selected backup network** check box.

  b. To disable the SMB Multichannel so that the data transfers only through the selected network, select the **Use dedicated ethernet if current machine enables SMB Multichannel** check box.

  Notes:

  - This option is not available by default. To enable this option, go to the following folder location: C:\Program Files\Arcserve\Unified Data Protection\Management\Configuration\ConsoleConfiguration.xml.Then, modify the value for *useDedicatedEthernet* as True.

```
- <SpecifyNetwork>
    <useDedicatedEthernet>false</useDedicatedEthernet>
  </SpecifyNetwork>
```

The Specify Network function becomes disabled in case of remote datastore, which used network shared folder as destination.

- When the Arcserve UDP management services restarts, the SMB Multichannel feature gets enabled.
- The SMB Multichannel feature is enabled in Windows, by default.

The destination is specified.

Specify the Schedule

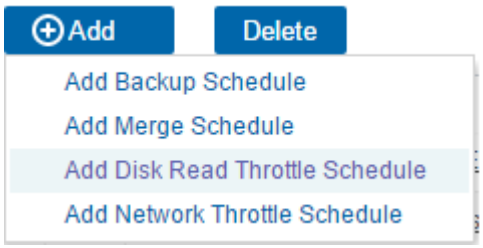
The Schedule page lets you define a schedule for Backup, Merge, and Throttle functions to repeat at specific intervals. After you define a schedule, the jobs run automatically as per the schedule. You can add multiple schedules and can provide retention settings.

A Backup Schedule refers to regular schedule that is repeated multiple times a day based on the number of hours or minutes you select. Besides the regular schedule, a backup schedule also provides options to add daily, weekly, and monthly schedules.

Note: For more information on scheduling and retention settings, see Understanding Advanced Scheduling and Retention.

Follow these steps:

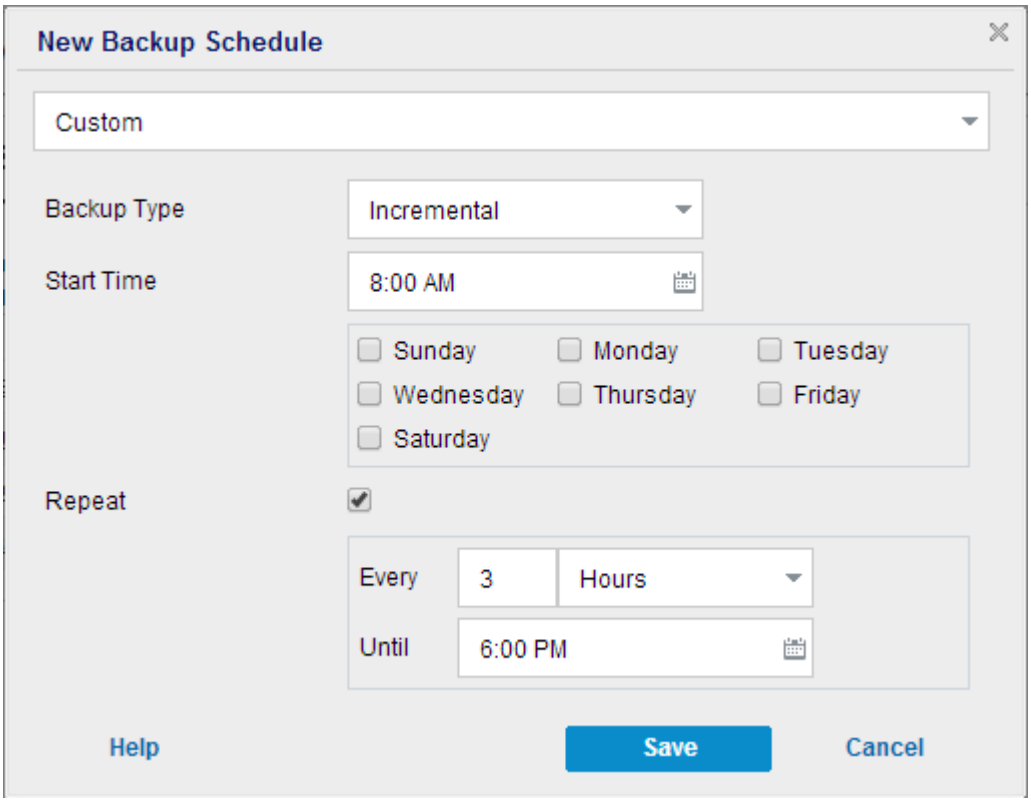
1. Add backup, merge, Disk Read throttle, and Network Throttle schedules.



Add Backup Schedule

a. Click **Add** and select **Add Backup Schedule**.

The **New Backup Schedule** dialog opens.



b. Select one of the following options:

Custom

Specifies the backup schedule that repeats multiple times a day.

Daily

Specifies the backup schedule that occurs once a day. By default, all the days of the week are selected for Daily backup. If you do not want to run the backup job on a specific day, clear the check box for that day of the week.

Weekly

Specifies the backup schedule that occurs once a week.

Monthly

Specifies the backup schedule that occurs once a month.

c. Select the backup type.

Full

Determines the backup schedule for Full Backups. As scheduled, Arcserve UDP performs a full backup of all used blocks from the source machine. A full backup typically consumes time depending on the backup size.

Verify

Determines the backup schedule for Verify Backups.

Arcserve UDP verifies that the protected data is valid and complete by performing a confidence check of the stored backup image to the backup source. If necessary, the image is resynchronized. A Verify Backup looks at the most recent backup of each individual block and compares the content and information to the source. This comparison verifies that the latest backed up blocks represent the corresponding information at the source. If the backup image for any block does not match the source (possibly because of changes in the system since the last backup), Arcserve UDP refreshes (resynchronizes) the backup of the block that does not match. You can also use a Verify Backup (infrequently) to get the guarantee of full backup without using the space required for a full backup.

Advantages: Produces a small backup image when compared to full backup because only the changed blocks (blocks that do not match the last backup) are backed up.

Disadvantages: Backup time is long because all source blocks are compared with the blocks of the last backup.

Incremental

Determines the backup schedule for Incremental Backups.

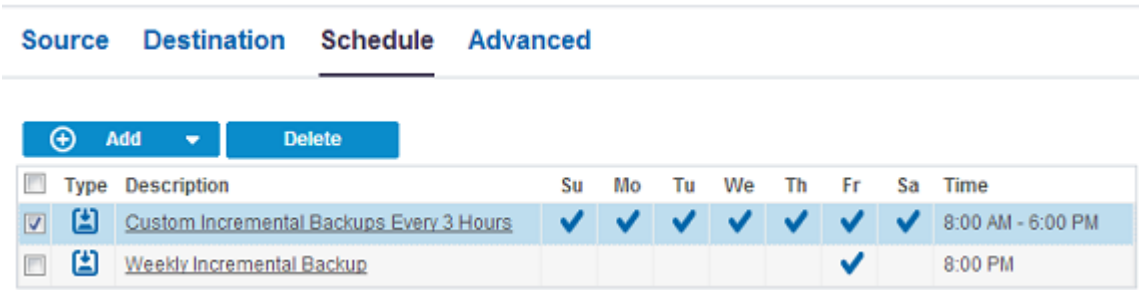
As scheduled, Arcserve UDP incrementally backs up only those blocks that have changed after the last successful backup. The advantages of Incremental Backups are that it is a fast backup and it produces a small backup image. This is the most optimal way to perform a backup.

d. Specify the backup start time.

e. (Optional) Select the **Repeat** check box and specify the repeat schedule.

f. Click **Save**.

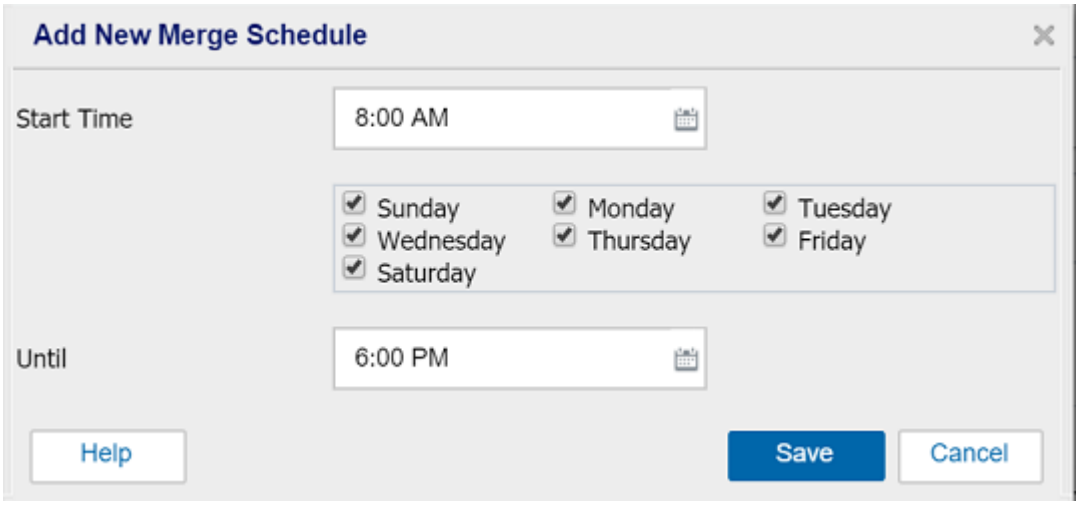
The Backup Schedule is specified and appears on the **Schedule** page.



Add Merge Schedule

a. Click **Add** and select **Add Merge Schedule**.

   The **Add New Merge Schedule** dialog opens.



b. Specify the start time to start the merge job.

c. Specify **Until** to specify an end time for the merge job.

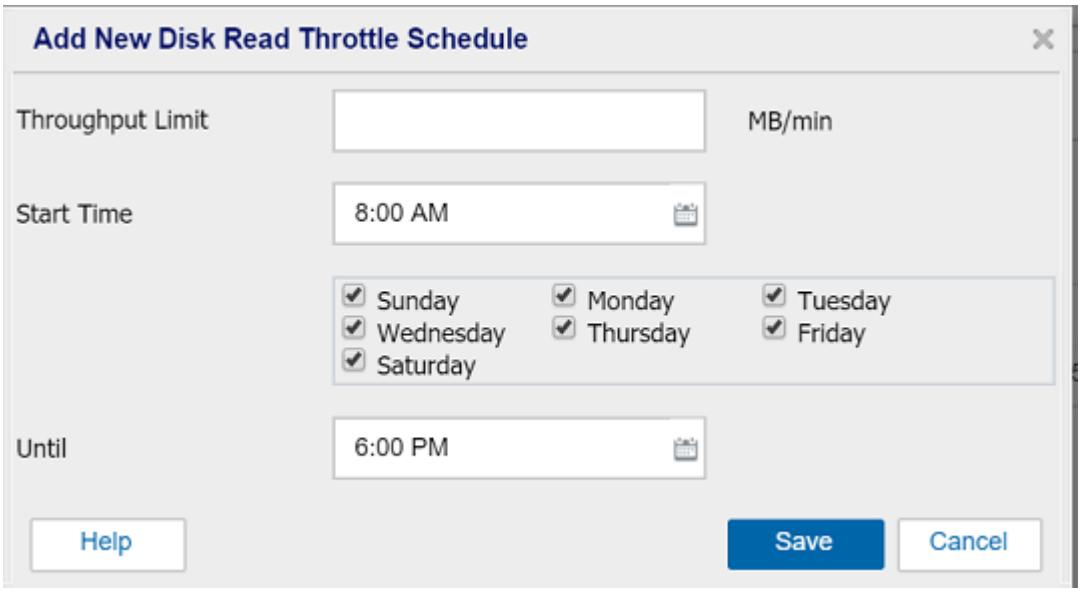
d. Click **Save**.

The Merge Schedule is specified and it is displayed on the **Schedule** page.

Add Disk Read Throttle Schedule

a. Click **Add** and select **Add Disk Read Throttle Schedule**.

   The **Add New Disk Read Throttle Schedule** dialog opens.



b. Specify the throughput limit in MB per minute unit.

c. Specify the start time to start the backup throughput job.

d. Specify **Until** to specify an end time for the throughput job.

e. Click **Save**.

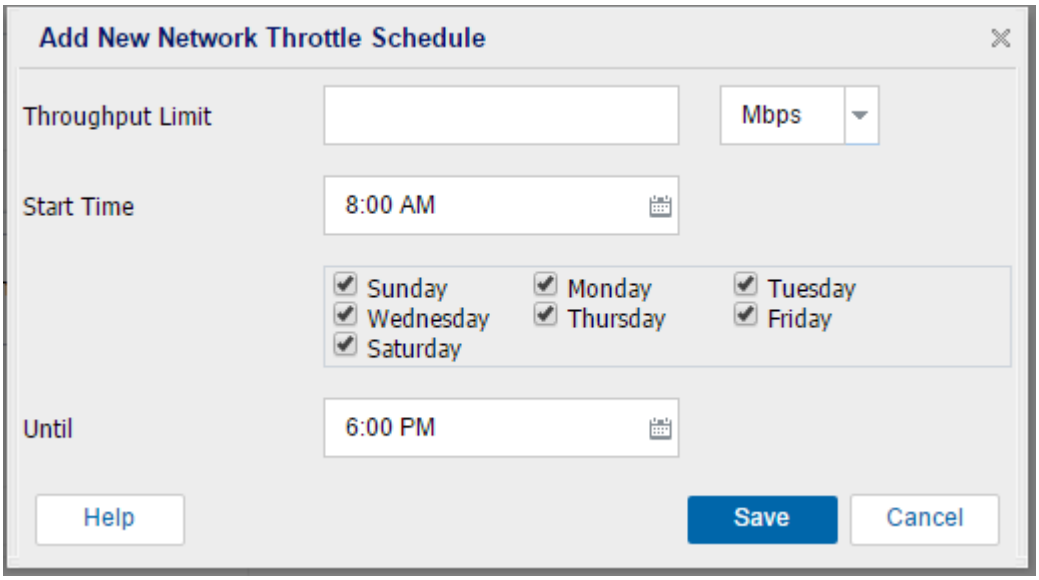
The Disk Read Throttle Schedule is specified and appears on the **Schedule** page.

Add Network Throttle Schedule

Note: Network Throttle Schedule appears only for Windows Agent based backup when you define a Deduplication enabled Data Store as the destination for the plan.

a. Click **Add** and select **Add Network Throttle Schedule**.

The **Add New Network Throttle Schedule** dialog opens.



b. Specify the throughput limit in Mbps or Kbps unit.

Note: Default minimum value: 500 kbps. To change the default value perform the following steps:

i. From the registry path SOFTWARE\Arcserve\Unified Data Protection\Management\Console, add a key MinNetworkThrottleValueInKpbs, type is REG_SZ, and set the value.

ii. Restart the Arcserve UDP Management service.

iii. Modify plan or create new plan.

The custom value takes effect.

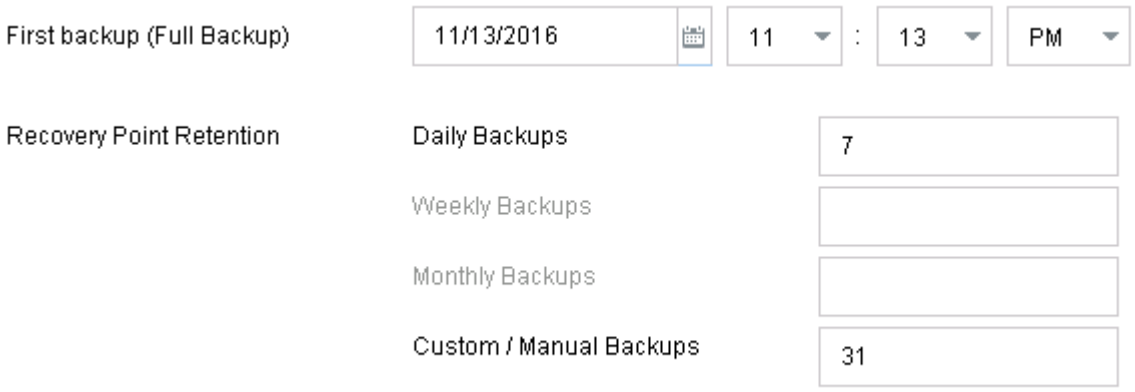
c. Specify the start time to start the backup throughput job.

d. Specify **Until** to specify an end time for the throughput job.

e. Click **Save**.

The Network Throttle schedule is specified and appears on the **Schedule** page.

2. Specify the start time for the scheduled backup.



3. Specify the recovery points retention settings for Custom, Daily, Weekly, and Monthly schedule.

   These options are enabled if you have added the corresponding backup schedule. If you modify the retention settings on this page, the changes are reflected on the Backup Schedule dialog.

   The schedule is specified.

Understanding Advanced Scheduling and Retention

The scheduling option lets you specify a Custom schedule or a Daily / Weekly / Monthly schedule, or both the schedules. In the Custom schedule, you can configure the backup schedule for each day of the week and you can add up to four backup schedules each day. You can select a specific day of a week and create a time window to define when to run backup and at what frequency.

Schedule	Supported Job	Comments
Backup	Backup job	Define time windows to run backup jobs.
Backup throttling	Backup job	Define time windows to control the backup speed.
Merge	Merge job	Define when to run merge jobs.
Daily schedule	Backup job	Define when to run daily backup jobs.
Weekly schedule	Backup job	Define when to run weekly backup jobs.
Monthly schedule	Backup job	Define when to run monthly backup jobs.

You can also specify the retention settings for the recovery points.

Note: Set the retention settings within each plan to control how data for the nodes assigned to that plan are retained at the target data store.

Schedules for Daily / Weekly / Monthly backups are independent to the Custom schedule, and each other. You can configure only to run Daily backup, Weekly backup or Monthly backup, without configuring the Custom schedule.

Backup Job Schedule

You can add four time windows per day in your backup schedule. A valid time window is from 12:00 AM until 11:59 PM. You cannot specify a time window such as 6:00 PM to 6:00 AM. In such cases, you must manually specify two different time windows.

For each time window, the start time is inclusive, and the end time is exclusive. For example, you have configured to run Incremental Backup every one hour between 6:00 AM and 9:00 AM and the backup will start at 6:00 AM. This means the backup runs at 6:00 AM, 7:00 AM, 8:00 AM, but NOT at 9:00 AM.

Note: If you want to run the backup job repeatedly until the end of day, set the schedule until 12:00 AM. For example, to run backup job every 15 minutes for the entire day, set the schedule from 12:00 AM to 12:00 AM, every 15 minutes.

Backup Throttle Schedule

Backup throttle schedule lets you control the backup throughput speed that in turn controls the resource usage (disk I/O, CPU, network bandwidth) of the server being backed up. This is useful if you do not want to affect the server

performance during business hours. You can add four time windows per day in your backup throttle schedule. For each time window, you can specify a value, in MB per minute. This value is used to control the backup throughput. Valid values are from 1 MB/minutes to 99999 MB/minutes.

If a backup job extends its specified time, then the throttle limit adjusts according to the specified time window. For example, you have defined the backup throttle limit as 500 MB/minute from 8:00 AM to 8:00 PM, and 2500 MB/minute from 8:00 PM to 10:00 PM. If a backup job starts at 7:00 PM and it runs for three hours, then, from 7:00 PM to 8:00 PM the throttle limit is 500 MB/minute and from 8:00 PM to 10:00 PM the throttle limit is 2500 MB/minute.

If you do not define any backup schedule and backup throughput schedule, the backup runs as fast as it can.

Merge Schedule

Lets you merge recovery points based on the provided schedule.

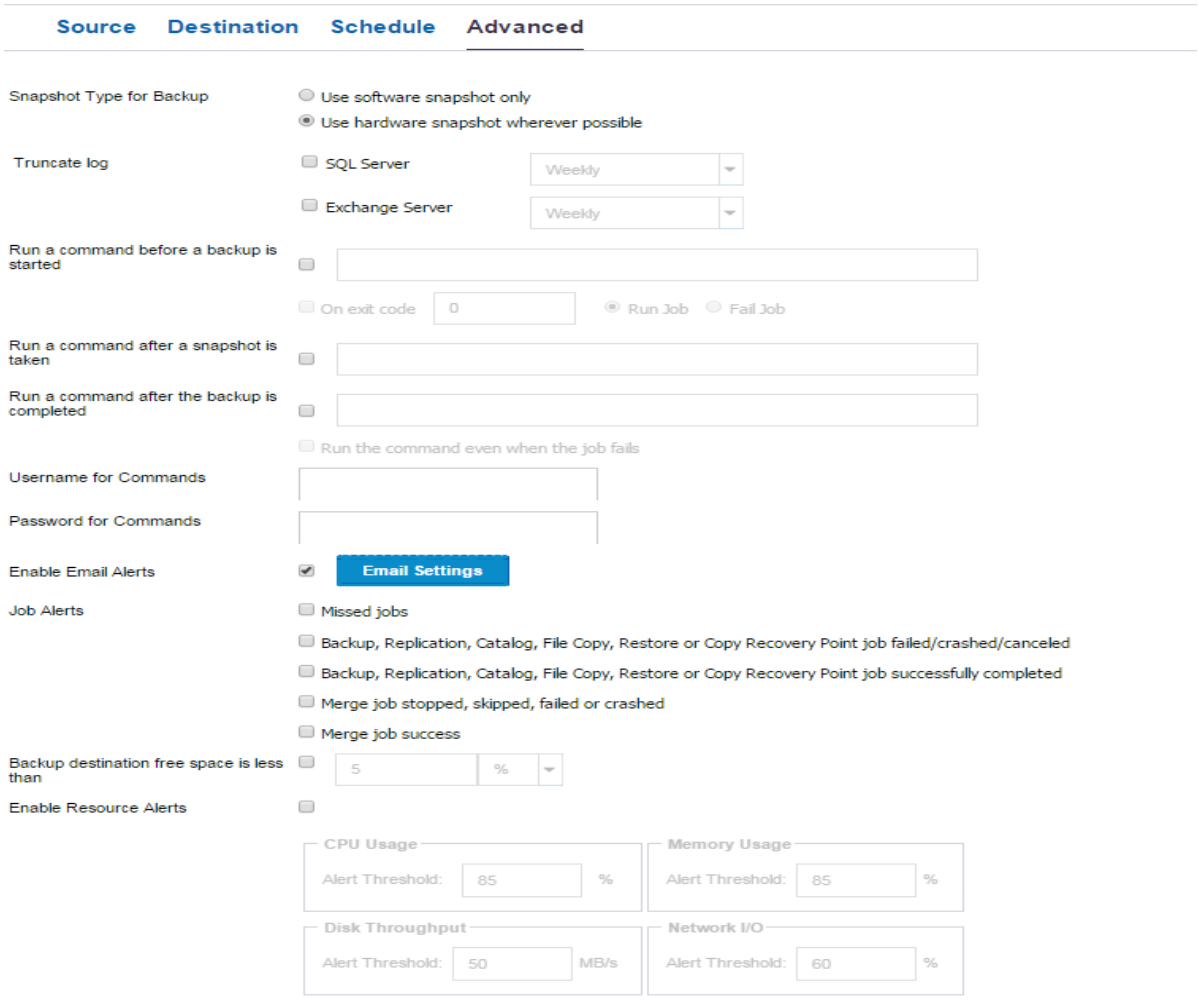
Consider the following points for the merge job:

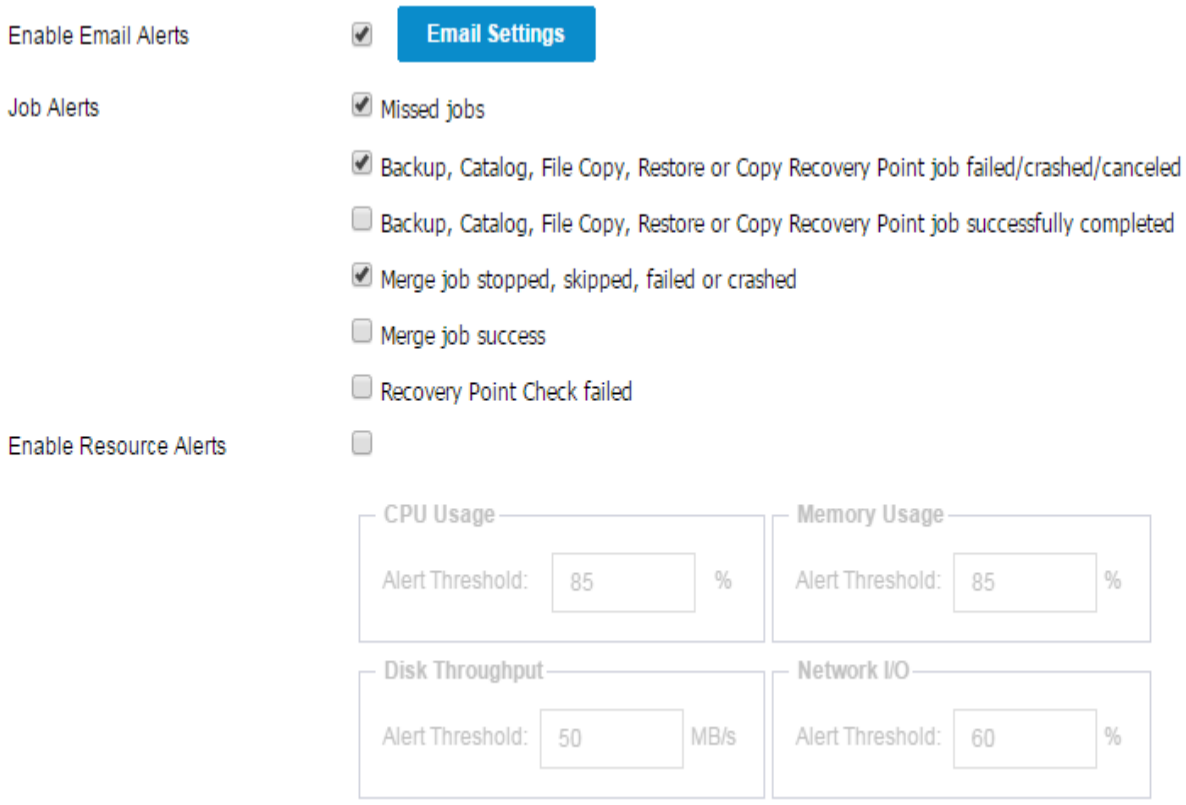
- At any given time only one merge job can run for a node.
- If a merge job starts, it has to complete before the next merge job can start. This means, if one or more sets of recovery points are merged, new recovery point cannot be added to this merge process, until the merge process of the current set of recovery point completes.
- If a merge job is processing more than one set of recovery points (for example set [1~4], set [5~11], and set [12~14]; they are three sets), recovery point server processes these sets one by one.
- If a merge job is resumed after a pause, the job detects at which point it is paused and resumes the merge from the break-point.

Specify the Advanced Settings

The **Advanced** tab lets you specify some advanced settings for the backup job. The advanced settings include providing truncate log settings, providing the location of any scripts, and email settings.

The following image displays the Advanced tab:


Source
Destination
Schedule
Advanced
Snapshot Type for Backup
Use software snapshot only
Use hardware snapshot wherever possible
Truncate log
SQL Server
Weekly
Exchange Server
Weekly
Run a command before a backup is started
On exit code
0
Run Job
Fail Job
Run a command after a snapshot is taken
Run a command after the backup is completed
Run the command even when the job fails
Username for Commands
Password for Commands
Enable Email Alerts
Email Settings
Job Alerts
Missed jobs
Backup, Replication, Catalog, File Copy, Restore or Copy Recovery Point job failed/crashed/canceled
Backup, Replication, Catalog, File Copy, Restore or Copy Recovery Point job successfully completed
Merge job stopped, skipped, failed or crashed
Merge job success
Backup destination free space is less than
5
%
Enable Resource Alerts
CPU Usage
Alert Threshold:
85
%
Memory Usage
Alert Threshold:
85
%
Disk Throughput
Alert Threshold:
50
MB/s
Network I/O
Alert Threshold:
60
%



Follow these steps:

1. Specify the following details.

Snapshot Type for Backup

Select one of the following options for the backup snapshot.

Use software snapshot only

Specifies that the backup type uses only the software snapshot. Arcserve UDP will not check for hardware snapshot. The software snapshot utilizes less resources on the virtual machines. You can use this option if the server has lower configurations and processing speed.

Use hardware snapshot wherever possible

Specifies that the backup type first checks for a hardware snapshot. If all the criteria are met, the backup type uses hardware snapshot.

Note: For more information on the hardware snapshot criteria, see the pre-requisite.

Truncate Log

Lets you specify the schedule to truncate logs for SQL Server and Exchange Server. You can specify the schedule as **Daily**, **Weekly**, or **Monthly**.

Enable purging logs at hourly interval for SQL via a registry switch

- Configure the Plan Settings. Check the "SQL Server" option in the "Truncate Log" Section under the "Advanced" Tab, and then select "Daily".
- Set the registry key on the SQL Server machine where the UDP Agent hosts. "PurgeSqlLogPerHour" is the interval in hours for purging SQL Log.

  Path: HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine

  Value Name: PurgeSqlLogPerHour (indicates the interval in hours for purging SQL Log.)

  Value type: REG_DWORD

User Name

Lets you specify the user who is authorized to run a script.

Password

Lets you specify the password of the user who is authorized to run the script.

Run a command before backup is started

Lets you run a script before the backup job starts. Specify the complete path where the script is stored. Click **On exit code** and specify the exit code for **Run Job** or **Fail Job**. **Run Job** indicates that the backup job will continue when the script returns the exit code. **Fail Job** indicates that the backup job will stop when the script returns the exit code.

Run a command after snapshot is taken

Lets you run a script after the backup snapshot is taken. Specify the complete path where the script is stored.

To run the post snapshot script even during a snapshot failure, select **Run the command even when the snapshot fails**.

Run a command after backup is over

Lets you run a script after the backup job is completed. Specify the complete path where the script is stored.

Enable Email Alerts

Lets you enable email alerts. You can configure email settings and can specify the types of alerts that you want to receive in an email. When you select this option, the following options are enabled for your selection.

Email Settings

Lets you configure the email settings. Click **Email Settings** and configure the email server and proxy server details. For more information about how to configure Email Settings, see Configure Email and Alert .

Job Alerts

Lets you select the types of job emails you want to receive.

Enable Resource Alerts

Lets you specify a threshold for CPU Usage, Memory Usage, Disk Throughput, Network I/O. You can provide the value in percentage. You will receive an email when the Alert Threshold value exceeds.

2. Click **Save**.

Note: When you select a node as a backup source or backup proxy, Arcserve UDP checks whether the agent is installed on the node and if it is the latest version. Arcserve UDP then displays a verification dialog that lists all the nodes that either have an outdated version of the agent or does not have the agent installed. To install/upgrade the agent on these nodes, select the installation method and click Save.

The changes are saved and a green check mark is displayed next to the task name. The plan page closes.

Note: If you have to add another task, you must select the plan from the **Resources** tab and modify the plan. To modify the plan, click the plan from the center pane. The plan opens and you can modify it.

The backup plan is created and automatically deployed to the source node. The backup runs per the schedule that you have configured in the **Schedule** tab. You can also perform a manual backup at any time.

(Optional) Perform a Manual Backup

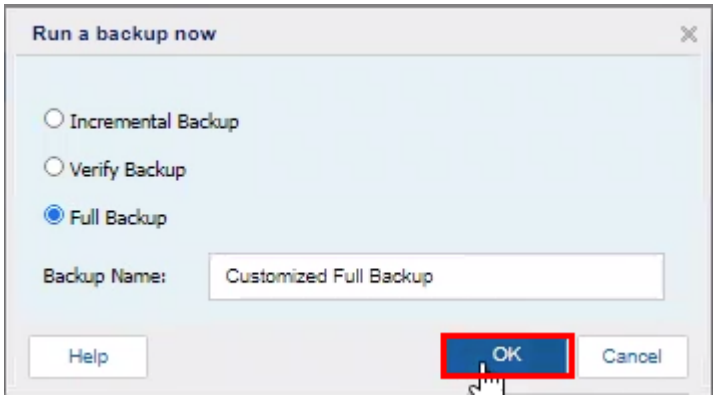
Typically, backups are performed automatically and are controlled by the schedule settings. In addition to the scheduled backup, a manual backup provides you the option to back up your nodes on a need basis. For example, if you have a repeat schedule for Full, Incremental, and Verify backups and you want to make major changes to your machine, you must perform an immediate manual backup without waiting for the next scheduled backup to occur.

Follow these steps:

1. Log into Arcserve UDP Console.
2. Navigate to **resources** > **Plan** > **All Plans**.

   The list of all plans is displayed on the center pane.
3. Right-click the plan, and then select **Backup Now**.

   The Run a backup now dialog opens.
4. On the Run a backup now dialog, select the type of backup you want to perform, and then specify the backup name if required.



The available options are:

Incremental Backup

Initiates an Incremental Backup of your machine. An Incremental Backup backs up only those blocks that have changed after the previous backup.

Note: If you are performing a backup for the first time, the full backup gets initiated automatically even though you selected the Incremental Backup option.

The advantages of Incremental Backups are that it is a fast backup and produces a small backup image. This is the most optimal way to perform backups.

Verify Backup

Initiates a Verify Backup of your machine by examining the most recent backup of each individual block and comparing the content and information to the original source. This comparison verifies that the latest backed up blocks represent the corresponding information at the source. If the backup image for any block does not match the source, Arcserve UDP Agent (Windows) refreshes (resynchronizes) the backup of the block that does not match.

Advantage: A very small backup image is produced when compared to a Full Backup because only the changed blocks (blocks that do not match the last backup) are backed up.

Disadvantage: The backup time is slow because all the source disk blocks are compared with the blocks of the last backup.

Full Backup

Initiates a Full Backup of your entire machine or the selected volumes.

Notes:

- If you add a new volume to the backup source, the newly added volume is fully backed up regardless of the overall backup method selected.
- If no backup name is specified, by default, it is automatically named as Customized Full/Incremental/Verify Backup.

5. Click **OK**.

   The Progress Information dialog appears.

6. Wait until the backup job completes, and then click **OK**.

7. To monitor the backup jobs, navigate to **jobs** > **Jobs in Progress**.

   Wait until the backup job is finished.

   Note: Only one job can be run at a time. If you manually attempt to launch a backup job when another job is currently running, an alert message informs you that another job is running and requests that you try again at a later time.

   The manual backup is successfully performed.

Verify the Backup

To verify your backup job, confirm that you have successfully created the backup plan. After you verify that the plan is created successfully, confirm whether the backup job is running as scheduled. You can verify the status of backup jobs from the **Jobs** tab.

Follow these steps to verify plans:

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and then click **All Nodes**.

   A list of all nodes is displayed on the center pane.
3. Verify that plans are mapped with nodes.

Follow one of these methods to verify backup jobs:

Method 1

1. Click the **jobs** tab.
2. From the left pane, click **All Jobs Completed**.

   The status of each job is listed on the center pane.
3. Verify that the backup job is successful.

   The backup job is verified.

Method 2

1. Click the **jobs** tab.
2. From the left pane, click **Jobs Successfully Completed**.

   The list of backup jobs completed appear on the center pane.
3. Click the job you want to verify.

   The right pane refreshes.
4. Under Job Details, click the **View Logs** hyperlink.
5. Verify that the backup job is successful.

   The backup job is verified.

How to Create a Nutanix Backup Plan

To protect your Nutanix nodes, you need to create a plan. The plan nodes consists of a backup task. This backup task lets you specify the nodes you want to protect, the backup destination, and the backup schedule. The backup destination is a recovery point server where you want to store your backup data. The destination can also be a local destination or a remote share folder.

To explore the functions available for Nutanix feature in Arcserve UDP, refer to How to Work With Nutanix Using Arcserve UDP.

How to Create a Virtual Standby to Nutanix AHV

To protect Virtual Standby on Nutanix AHV, refer to Protecting Virtual Standby on Nutanix AHV.

How to Create an Instant Virtual Machine on Nutanix AHV

To create an instant virtual machine on Nutanix AHV, refer to How to Create an Instant Virtual Machine on Nutanix AHV.

How to Create an Assured Recovery Plan to Protect Nutanix AHV for Linux Node

To create an assured recovery plan to protect Nutanix AHV for Linux Node, refer to How to Create an Assured Recovery Plan to Protect Nutanix AHV for Linux Node.

How to Perform Backup of Cluster Shared Volume

Arcserve UDP supports the backup of Cluster Shared Volume (CSV) from v6.5 Update 4. The CSV volumes are created upon Storage Space and are transparent to Arcserve UDPAgent (Windows). CSV is always considered for full backup at volume level, irrespective of the backup job type (full or incremental). The CSV backup protects the data of all popular file systems such as NTFS, NTFS Dedupe, ReFS and CSVFS.

You can create an agent-based plan or update an existing agent-based plan from Arcserve UDP Console to perform backup of specified CSV.

Follow these steps:

1. Select Cluster Shared Volume in Agent Machine. For more information, see Modify Agent Machine Settings to Backup CSV Volume.

   Note: You can also run the backup job from Windows Agent machine. For more information, see How to Perform Backup from Windows Agent Machine.

2. To perform backup from Arcserve UDP Console, navigate to the **Resources** tab.
3. From the left pane, navigate to **Plans**, and click **All Plans**.
4. On the center pane, click **Add a Plan** and specify a plan name.
5. From the **Task Type** drop-down list, select **Backup: Agent-Based Windows**.

6. From the **Source** tab of Plan, select the checkbox of **Back up all volumes** as **Protection Type**.
7. Specify the Destination, Schedule, and Advanced details. For more information, see How to Create a Windows Backup Plan.

   Now, after saving the plan when you run the backup job, the specified CSV volume is backed up for Windows agent.
8. Verify the status of backed up jobs. For more information, see Verify the Backup.

   Notes:

   - The activity log of the owner node displays the progress of the backup job and logs the success message upon completion of the backup. The activity log of the member node (which is not the owner) displays a warning message as below, but the backup status shows as successful.

     Warning: Failed to fetch physical location of Cluster Shared Volume [C:\ClusterStorage\Volume1]. This volume will be excluded from the backup.
   - If you want to exclude the CSV volume during the backup, perform the following steps:

     a. On the UDP Agent node, open the Registry to the following path:

        HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll

     b. Create a DWORD as 'BackupCSV'.

     c. Set the value to 0.

        The CSV volume is now excluded during the backup.

Modify Agent Machine Settings to Backup CSV Volume

Before you perform your first backup, you must configure the backup settings which are applied to each backup job. These settings can be retained for future backup or they can be modified at any time from the Arcserve UDP Agent (Windows) home page.

Follow these steps:

1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), click **Settings**.

The Settings window opens.

2. Click the **Backup Settings** tab.

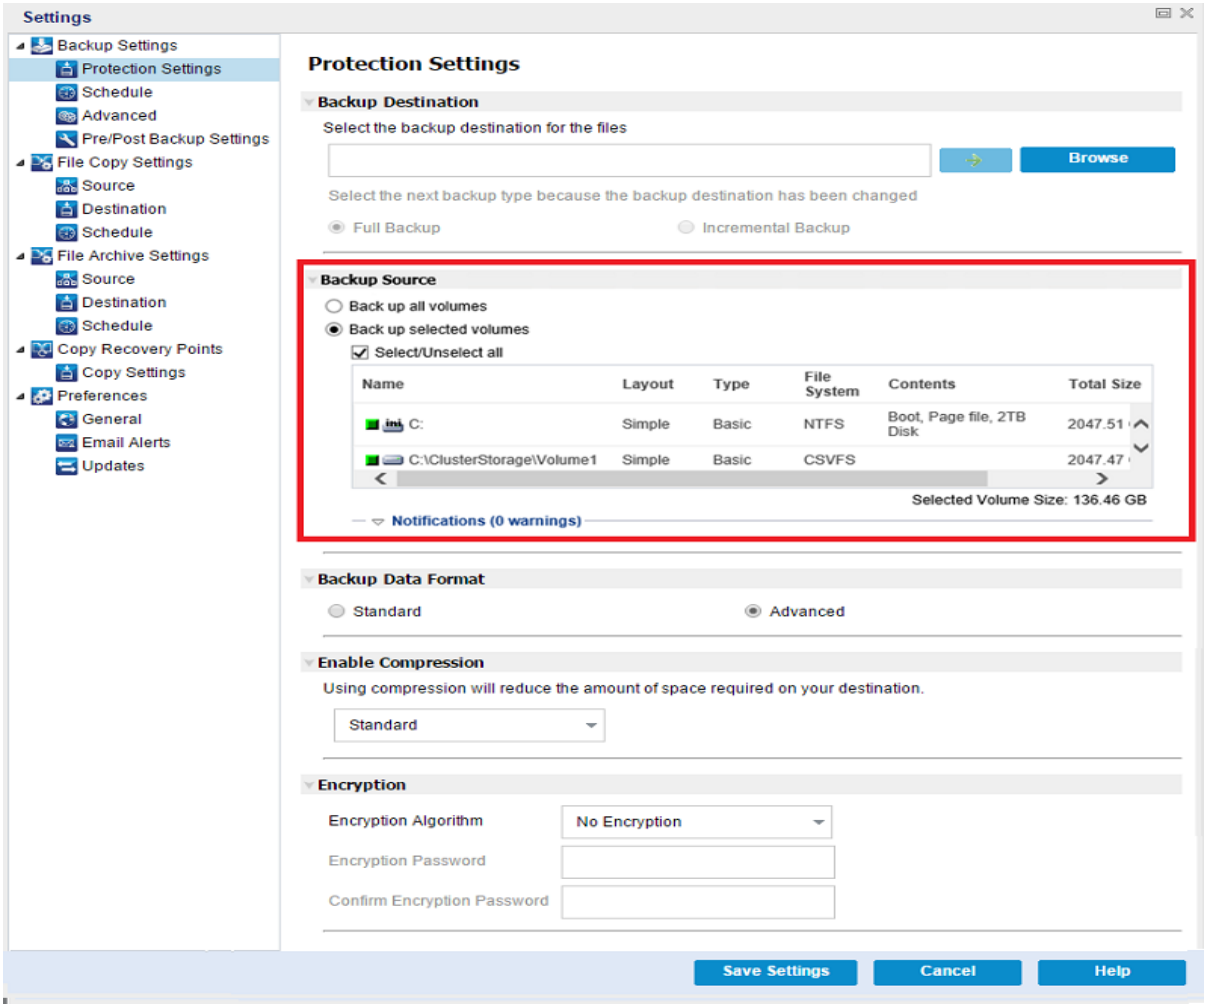
   Backup Settings window opens.

3. Click Protection Settings.

4. The Protection Settings window opens.

   Notes:

   - If the Arcserve UDP Agent (Windows) is being managed by Console, not all settings are available and will be displayed as read-only information.
   - When agent is managed by console and not protected in a plan, all the settings are still available except the Preference > Updates panel.



5. From the **Backup Source** group, select **Backup selected volumes**.

   List of available volumes in the Agent Machine appears.

6. Select the CSV Volume that you want to backup and click **Save Settings**.

Your backup protection settings are saved.

How to Perform Backup of Volumes without Driver Letter

From Arcserve UDP v6.5 Update 4, you can protect the volumes without drive letter using the option **Mounted in NTFS folder Volumes** at **Back up selected volumes**. Before selecting the option **Mounted in NTFS folder Volumes** from Arcserve UDP Console, you must mount NTFS folder volume. You can also customize further and back up only some volumes without drive letter instead of backing up all.

For more information view the following:

- Review Prerequisites
- Mount the Volume without Drive Letter in NTFS Folder
- Specify Volume without Drive Letter in Registry

Review Prerequisites

Review the following:

- For the datastore to be in running state across RPS reboots, enable automount on the RPS server. You can enable automount using the diskpart - automount command.

  Follow these steps:

  1. Open an elevated command prompt.
  2. In the elevated command prompt, do the following:
  3. Type *diskpart*, and then press **Enter**.
  4. Type *automount enable*, and press **Enter**.
  5. Close the elevated command prompt.
  6. To apply changes, restart the computer.

Mount the Volume without Drive Letter in NTFS Folder

You must mount the volume without drive letter in NTFS folder to back up the volume and contents of the mount point.

Follow these steps:

1. From **Windows Disk Manager**, select new disk that does not have a drive letter assigned.
2. From the right click options, click **Change Drive Letter and Paths**.

   Add Drive Letter and Paths dialog opens.
3. Click **Change**.

   Change Drive Letter or Path dailog opens.
4. Select **Mount in the following empty NTFS folder**, specify the following path and click **OK**.

   C:\MountPoint

   Mount Point appears in Windows Explorer. You have successfully mounted the volume without drive letter

Specify Volume without Drive Letter in Registry

You must specify the volume without drive letter in registry to protect the volume and contents in the mount point.

Follow these steps:

1. From **Windows Disk Manager**, select new disk that does not have a drive letter assigned.
2. From the right click options, click **Properties**.

   The disk properties window opens.
3. From the **Security** tab, copy **Object name**.
4. On the **Windows Agent** machine, navigate to the respective **Agent Node** and open **Registry**.

   The registry window opens.
5. Navigate to the below path and click **BackupVolumesWIthMountedVolumes**.

   HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll

   The Edit Multi String dialog appears.
6. In the **Value data** field, add the object name that you copied.
7. Click **OK**.

   The volume without drive letter is specified in registry.

   When you perform backup from Arcserve UDP Console, only the volumes specified in the registry are considered for backup.

How to Perform Backup of Selected Volume

This topic describes how to take backup of only the selected volumes instead of performing full backup. To perform backup of a specific volume, you must select the **Back up selected volumes** checkbox from the Source option while creating a plan.

Follow these steps:

1. While creating or modifying a plan, select the Source tab.
2. Select the check-box of Back up selected volume as Protected Type.

   Multiple fields appear to select drive and other options related to volume.
3. Select one or multiple drives from the list of drives.
4. Select one or multiple options from the following list of volumes:

   System Reserved Volume

   Refers to the volume required when booting operating system. Select the option if you want to perform BMR for the backup node or perform backup for a specific drive.

   Recovery Volume

   Refers to the volume required for Windows Recovery environment. Select the option if you want to back up the recovery volume.

   Boot Volume

   Refers to volume required when booting Windows Operating System. Select the option if you want to perform BMR for the backup node or perform backup for a specific drive.

   Mounted in NTFS Folder Volumes

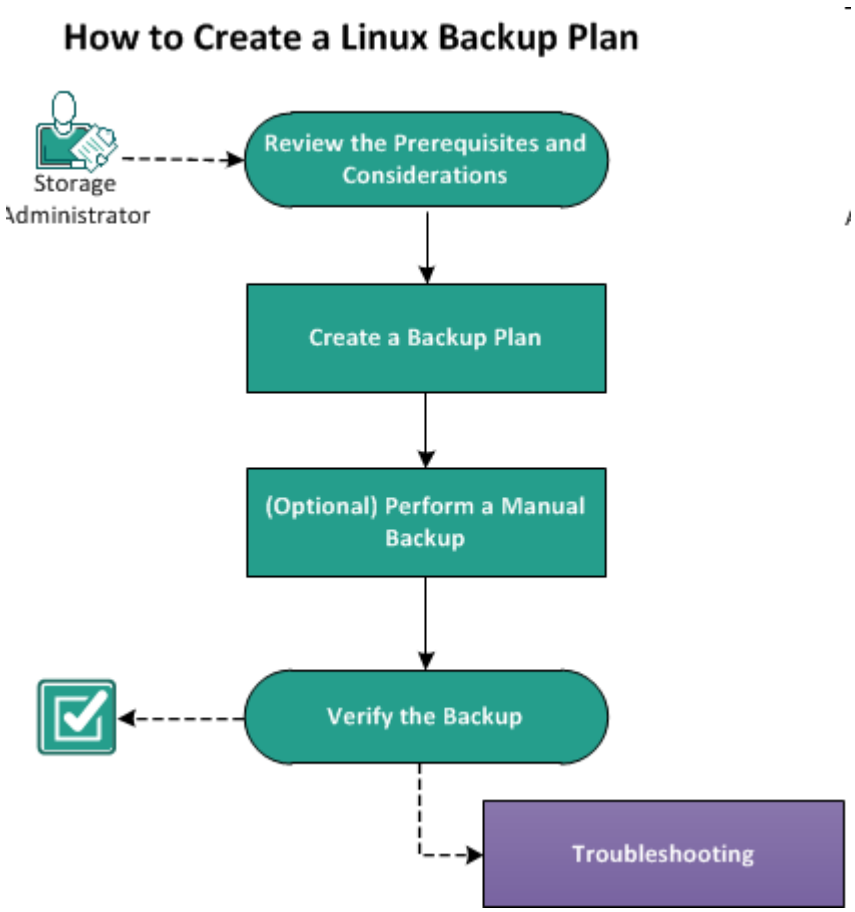
   Refers to Mounted in NTFS folder Volumes option. Select this option if you can perform back up of volumes without drive letter. Before selecting the option in the plan, you must mount NTFS folder volume. You can also customize further and back up only some volumes without drive letter instead of backing up all. For more information, view How to Backup Volumes without Driver Letter
5. Enter details in other tabs of plan and Save the plan.

   You have specified the drive or volume. Only the selected volume or drive is backed up when a backup job runs.

How to Create a Linux Backup Plan

To protect your Linux nodes, you need to create a plan. A backup plan for Linux nodes consists of a backup task. This backup task lets you specify the nodes you want to protect, the backup destination, and the backup schedule. The backup destination can be a local destination or remote share folder, or a data store in a recovery point server.

The following diagram illustrates the process to protect Linux nodes:



What To Do Next?

1. Review the Prerequisites and Considerations
2. Create a Backup Plan
3. (Optional) Perform a Manual Backup
4. Verify the Backup
5. Troubleshooting

Review the Prerequisites and Considerations

Complete the following prerequisites:

- Log into the Console.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Create a Backup Plan

A backup plan includes a backup task that performs a backup of the physical or virtual node and stores the data to the specified destination.

Follow these steps:

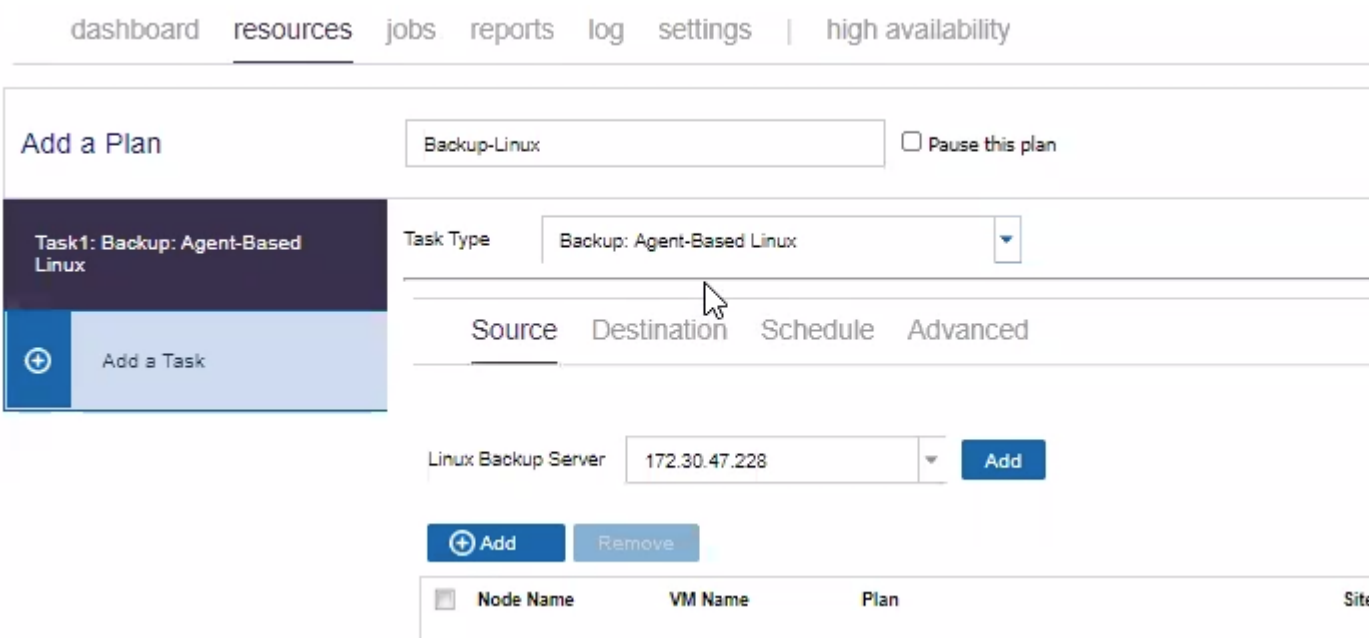
1. Log into Arcserve UDP Console.
2. Navigate to **resources** tab > **Plan** > **All Plans**.

   If you have created plans earlier, those plans are displayed on the center pane.
3. On the center pane, click **Add a Plan**.

   The Add a Plan page opens.
4. Enter a plan name.
5. (Optional) Select the **Pause this plan** check box.

   The plan will not run until you clear the check box to resume the plan.

   Note: If a plan is paused, then any in-progress job is not paused but all corresponding scheduled jobs associated with that plan are paused. However, you can manually run a job. For example, you can manually run backup job and replication job for a node even if the respective plan is paused. In such case, the following task to the on-demand (manual) job does not run. For example, there is a replication task after an on-demand backup job, the replication job does not run for the on-demand backup job. You need to manually run the replication job. When you resume the plan, the pending jobs do not resume immediately. After you resume the plan, the pending jobs run from the next scheduled time.
6. From the Task Type drop-down list, select **Backup: Agent-Based Linux**.



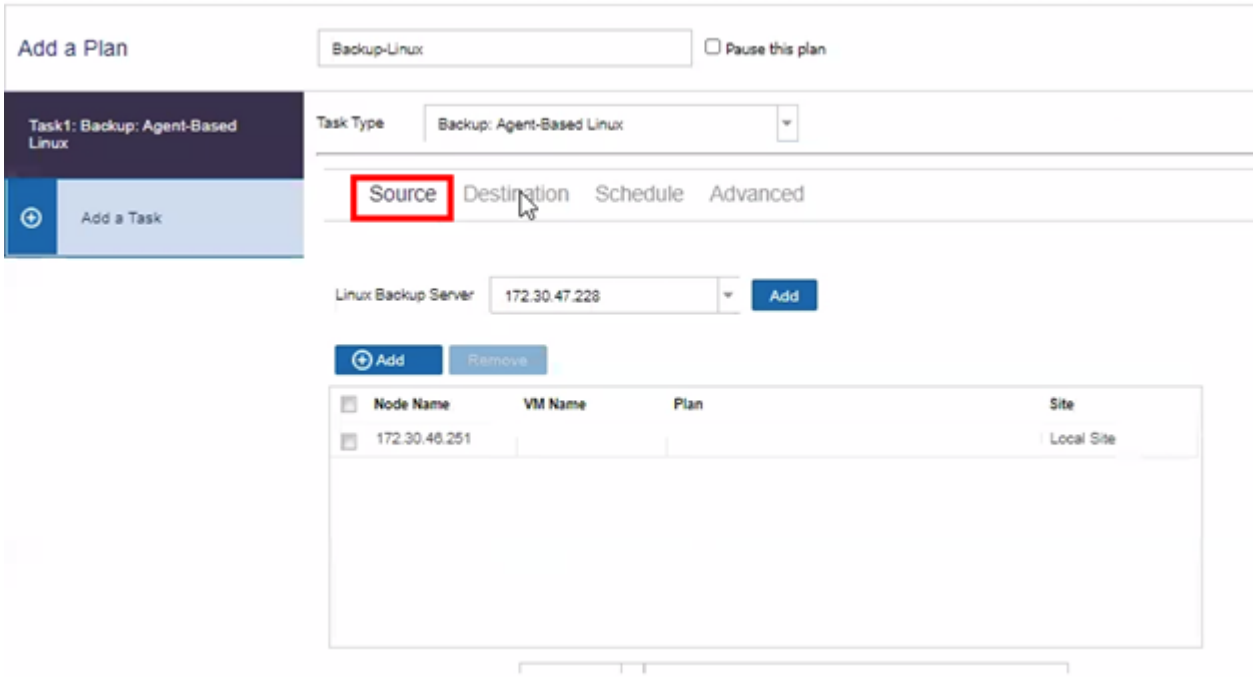
Now specify the Source, Destination, Schedule, and Advanced details.

Specify the Source

The Source page lets you specify the source nodes that you want to protect. You can select more than one node in a plan. If you have not added any nodes to the Console, you can add nodes from the Source page. You can save a plan without adding any source nodes, but the plan will not be deployed unless you add any nodes.

Follow these steps:

1. Click the Source tab.
2. From the Linux Backup Server drop-down list, select the Linux backup server.



3. (Optional) Click **Add** to add a new Linux Backup Server to the list.
4. Click **+Add**, and then select one of the following options:

Select Nodes to Protect in Arcserve UDP

The Select Nodes to Protect dialog opens. Select the nodes from the displayed list. Select this option if you have already added the nodes to the Console.

Add Linux Node

The Add Linux Nodes to Arcserve UDP Console dialog opens. Select this option if you have not added the nodes and you want to manually add the nodes to protect.

If selected, UDP Archiving Node is backed up by the Linux Backup Server installed on this node itself, even if you manually select another Linux Backup Server. To back up the node by another Linux Backup Server, you could add this node as a Linux Node instead of adding as UDP Archiving Node. For more information see, how to add nodes.

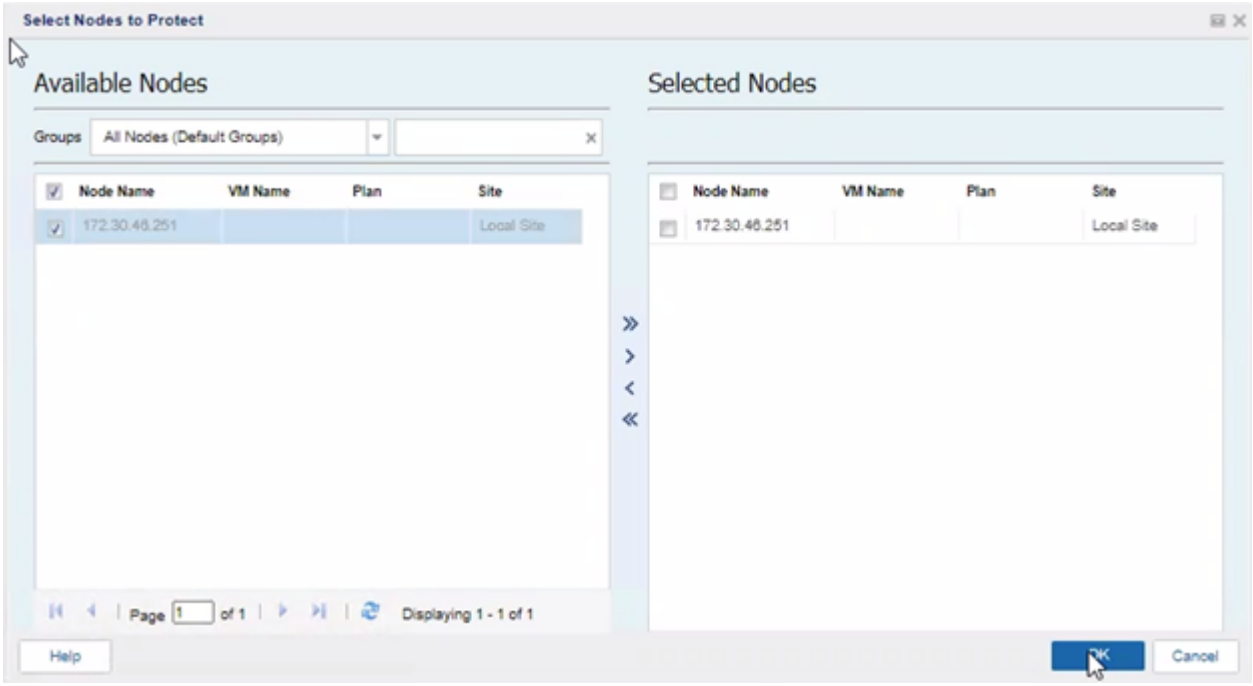
5. On the Select Nodes to Protect dialog, do the following:

   a. (Optional) Select a filter from the Groups drop-down list to filter nodes. You can enter keywords to further filter your nodes.

      The nodes are displayed on the Available Nodes column.

   b. Select the nodes from the Available Nodes column, click the right arrow, and then click **OK**.

      The selected nodes are displayed in the Selected Nodes column.



6. (Optional) Provide the details for the following options:

   Filter volumes for backup

   Select **Include** or **Exclude** from the drop-down list. Include specifies that only the specified volumes are included for backup. Any volume that is not specified will not be backed up. Exclude specifies that the volumes are excluded from the backup.

   Files/folders to be excluded

   Specify the files and folders that you do not want to backup for all the listed nodes. If you do not want to backup multiple files and folders, separate each

file and folder using a colon (:). Provide the full path of the file and folder that you want to exclude.

7. For Designated Backup Network, to select the network between Linux Backup Server (LBS) and Linux node, click the **Use selected network for backup traffic** check box, and then select the network from the drop-down list.

   Note: If the selected backup network is not accessible and to continue the job with the available network or the default network, click the **Continue to run job even when unable to connect to the selected backup network** check box.

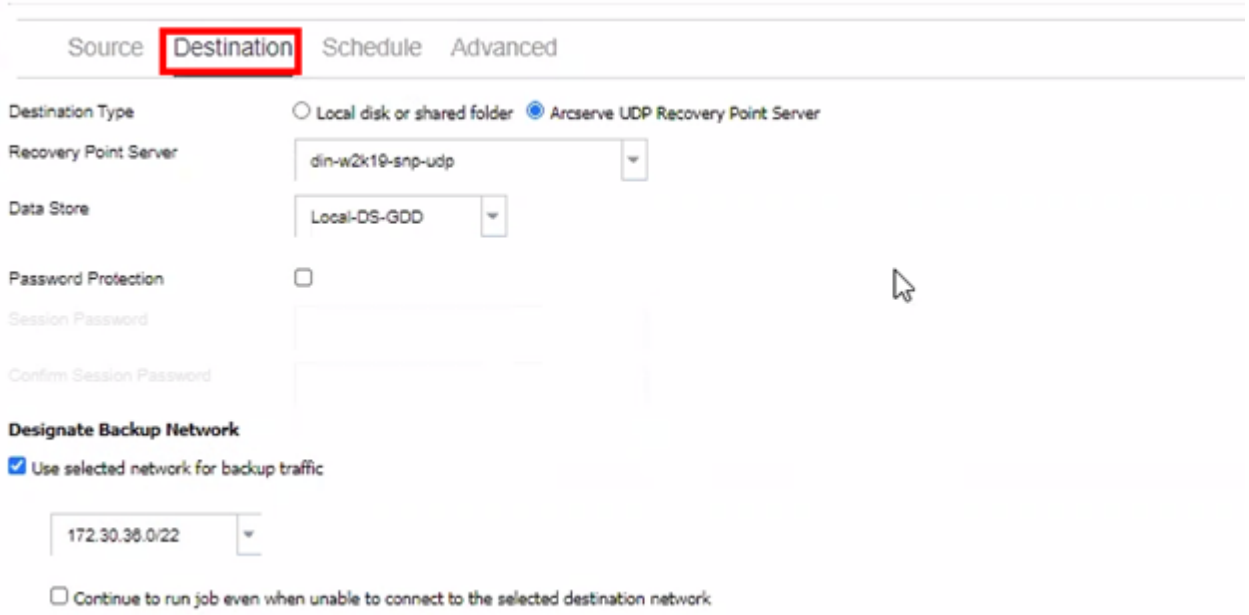
   The source is specified.

Specify the Destination

The destination is a location where you store the backup data. You must at least specify the destination to save the plan.

Follow these steps:

1. Click the **Destination** tab.



2. Select one of the following as Destination Type:

   Local disk or shared folder

   Specifies that the backup data is stored at a local disk or shared folder.

   Arcserve UDP Recovery Point Server

   Specifies that the backup destination is a recovery point server. If you select this option, then data is stored as recovery points. You cannot store data as recovery sets.

3. If you select the **Local disk or shared folder** as a Destination Type, then provide the following details:

   - If you have selected **NFS share**, then type the Backup Destination detail in the following format:

     IP address of the NFS Share:/full path of the storage location

     - Click the arrow button to validate the Backup Destination information.

       If the backup destination is invalid, an error message is displayed.

Note: Some versions of Data Domain NAS do not support the file locking mechanism of NFS. As a result, such NFS share cannot be used as a backup destination. For more information about this issue, see Compatibility Issues with Arcserve Unified Data Protection Agent for Linux.

- If you have selected **CIFS share**, then type the Backup Destination detail in the following format:

  //hostname/share_folder

  - Click the arrow button to validate the Backup Destination information.

    If the backup destination is invalid, an error message is displayed.

  Note: Make sure the shared folder name does not have any spaces.

- If you have selected **Source local**, follow these steps:

  a. Provide the path of the local destination.

     For example: /root

  b. Select a compression level from the Compression drop-down list to specify a type of compression that is used for backup.

     The available options for Compression are:

     Standard Compression

     Specifies a compression level from the Compression drop-down list to specify a type of compression that is used for backup.

     Maximum Compression

     Specifies that this option provides the highest CPU usage (lowest speed), but also has the lowest disk space usage for your backup image.

  c. Select an algorithm from the Encryption Algorithm drop-down list and type the encryption password, if necessary.

  d. Select the type of encryption algorithm that you want to use for backups.

     Data encryption is the translation of data into a form that is unintelligible without a deciphering mechanism. The Arcserve UDP Agent (Linux) data protection solution uses secure, AES (Advanced Encryption Standard) encryption algorithms to achieve the maximum security and privacy of your specified data.

     For available format options of Encryption, see Encryption Settings.

- A full backup and all its related incremental backups must use the same encryption algorithm.
- If the encryption algorithm for an incremental backup has changed, you must perform a full backup.

For example, if you change the algorithm format and then you run an incremental backup, then the backup type automatically converts to a full backup.

e. When an encryption algorithm is selected, you must provide (and confirm) an encryption password.

- The encryption password is limited to a maximum of 23 characters.
- A full backup and all its related incremental backups use the same password to encrypt data.

- If you want to backup to Amazon S3, do the following:

a. Select **CIFS share**, and then type the Amazon S3 storage in the following format

s3://S3 Region/S3 bucket name

b. Click the arrow, and provide Amazon S3 access information.

4. If you select the **Arcserve UDP Recovery Point Server** as a Destination Type, do the following:

- Recovery Point Server: From the drop-down list, select the destination recovery point server.
- Data Store: From the drop-down list, select the destination data store.
- Password Protection: Select the **Password Protection** check box, provide the session password, and then retype the password to confirm.

  Note: The session password is optional when the backup destination is an unencrypted RPS data store.
- For Designated Backup Network, to select the network between Linux Backup Server (LBS) and Recovery Point Server (RPS), click the **Use selected network for backup traffic** check box, and then select the network from the drop-down list.

  Note: If the selected backup network is not accessible and to continue the job with the available network or with the default network, click the **Continue to run job even when unable to connect to the selected backup network** check box.

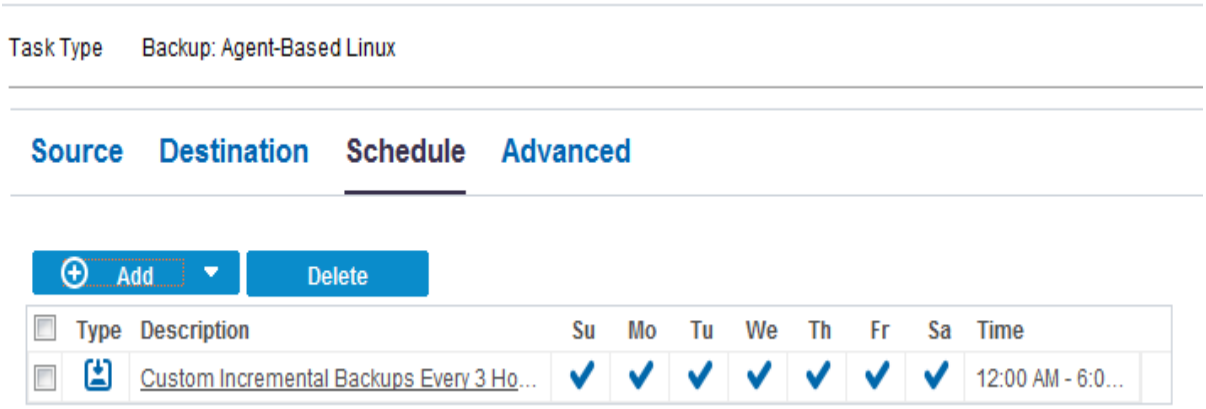
When the backup destination is a dedupe data store, the traffic is directly between Linux Source and Recovery Point Server (RPS) even though the Specify Network is configured between Linux Source and Linux Backup Server (LBS). However, when the backup destination is a non-dedupe datastore, the traffic goes through the Linux Backup Server (LBS).

The destination is specified.

Specify the Schedule

The Schedule page lets you define a backup schedule to repeat at specific intervals. After you define a schedule, the jobs run automatically as per the schedule. You can add multiple schedules and provide retention settings. If Local disk or shared folder is the backup destination, the default value is a custom incremental backup at every 10:00 PM. If RPS server is the backup destination, the default value is a daily Incremental backup at every 10:00 PM.

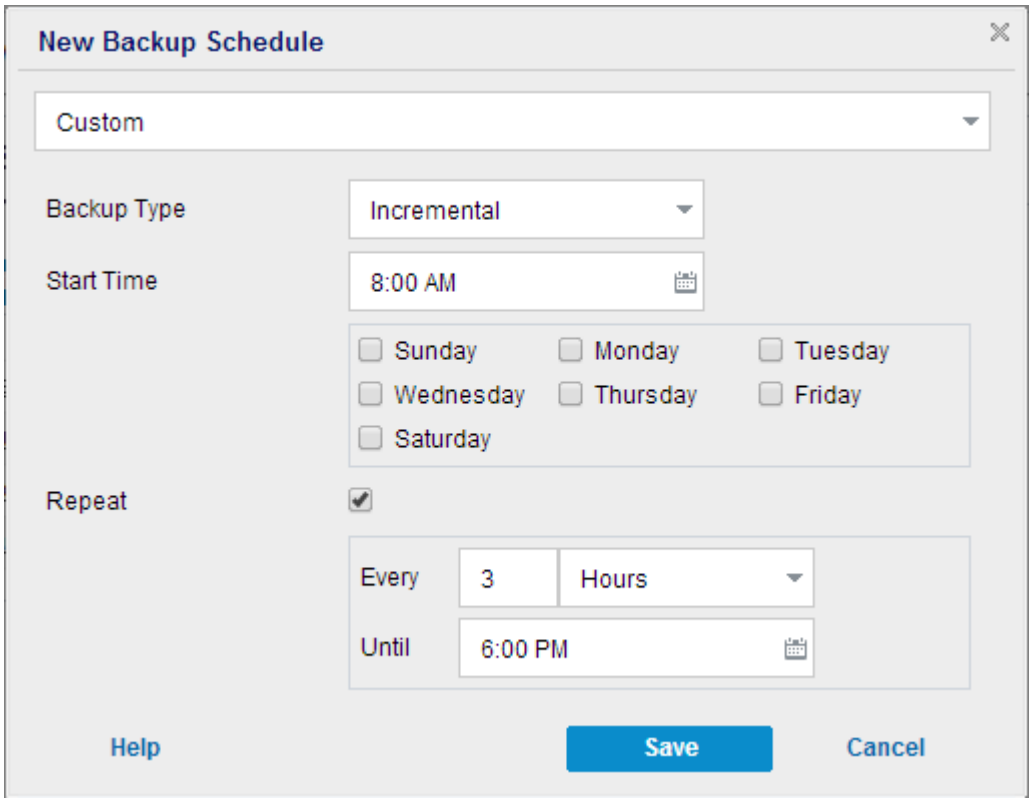
You can edit or delete a backup job schedule.



Follow these steps:

1. Click the **Schedule** tab.
2. Click **Add** and select **Add Backup Schedule**.

   The **New Backup Schedule** dialog opens.



a. Select one of the following options:

Custom

Specifies the backup schedule that repeats multiple times a day.

Daily

Specifies the backup schedule that occurs once a day. By default, all the days of the week are selected for Daily backup. If you do not want to run the backup job on a specific day, clear the check box for that day of the week.

Weekly

Specifies the backup schedule that occurs once a week.

Monthly

Specifies the backup schedule that occurs once a month.

b. Select the backup type.

Full

Determines the backup schedule for Full Backups. As scheduled, Arcserve UDP performs a full backup of all used blocks from the source machine. A full backup typically consumes time depending on the backup size.

Verify

Determines the backup schedule for Verify Backups.

Arcserve UDP verifies that the protected data is valid and complete by performing a confidence check of the stored backup image to the backup source. If necessary, the image is resynchronized. A Verify Backup looks at the most recent backup of each individual block and compares the content and information to the source. This comparison verifies that the latest backed up blocks represent the corresponding information at the source. If the backup image for any block does not match the source (possibly because of changes in the system since the last backup), Arcserve UDP refreshes (resynchronizes) the backup of the block that does not match. You can also use a Verify Backup (infrequently) to get the guarantee of full backup without using the space required for a full backup.

Advantages: Produces a small backup image when compared to full backup because only the changed blocks (blocks that do not match the last backup) are backed up.

Disadvantages: Backup time is long because all source blocks are compared with the blocks of the last backup.

Incremental

Determines the backup schedule for Incremental Backups.

As scheduled, Arcserve UDP incrementally backs up only those blocks that have changed since the last successful backup. The advantages of Incremental Backups are that it is a fast backup and it produces a small backup image. This is the most optimal way to perform a backup.

c. Specify the backup start time.

d. (Optional) Select the **Repeat** check box and specify the repeat schedule.

e. Click **Save**.

The Backup Schedule is specified and it is displayed on the **Schedule** page.

3. Specify the retention settings if the destination is a shared or a network folder.

Note: For more information about the recovery sets, see Understanding the Recovery Sets.

Specify the number of recovery sets to retain

Specifies the number of recovery sets retained.

Start a new recovery set on every:

Selected day of the week

Specifies the day of the week selected to start a new recovery set.

Selected day of the month

Specifies the day of the month selected to start a new recovery set. Specify 1 to 31, or the last day of the month.

Note: The Linux Backup Server checks for the number of recovery sets in the configured backup storage every 15 minutes and deletes any extra recovery set from the backup storage location.

4. If you have selected Arcserve Recovery Point Server as the destination, then follow these additional steps.

   Add Merge Schedule

   a. Click **Add** and select **Add Merge Schedule**.

      The **Add New Merge Schedule** dialog opens.

   b. Specify the start time to start the merge job.

   c. Specify **Until** to specify an end time for the merge job.

   d. Click **Save**.

      The Merge Schedule is specified and it is displayed on the **Schedule** page.

   Add Throttle Schedule

   a. Click **Add** and select **Add Throttle Schedule**.

      The Add New Throttle Schedule dialog opens.

   b. Specify the throughput limit in MB per minute unit.

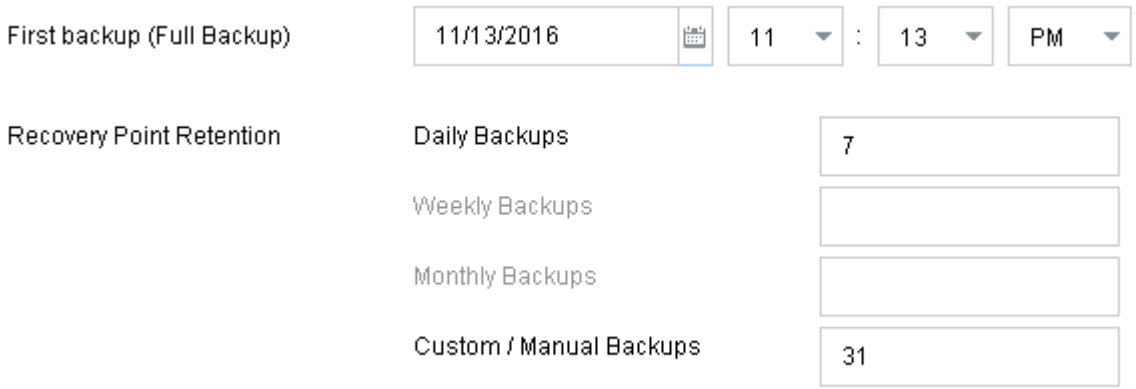
   c. Specify the start time to start the backup throughput job.

   d. Specify **Until** to specify an end time for the throughput job.

   e. Click **Save**.

      The Throttle Schedule is specified and it is displayed on the **Schedule** page.

5. Specify the start time for the scheduled backup.



6. Specify the recovery points retention settings for Custom, Daily, Weekly, and Monthly schedule.

   These options are enabled if you have added the corresponding backup schedule. If you modify the retention settings on this page, the changes are reflected on the Backup Schedule dialog.

   The backup schedule is specified.

Understanding the Recovery Sets

A recovery set is a storage setting where a group of recovery points backed-up over a specified period is stored as one set. A recovery set includes a series of backups, starting with a full backup, and then followed by a number of incremental, verify, or full backups. You can specify the number of recovery sets to retain.

The **Recovery Set Settings** ensures periodic maintenance of recovery sets. When the specified limit is exceeded, the oldest recovery set is deleted. The following values define the default, minimum, and maximum recovery sets in Arcserve Unified Data Protection Agent for Linux:

Default: 2

Minimum: 1

Maximum number of recovery sets: 100

Note: If you want to delete a recovery set to save backup storage space, reduce the number of retained sets and Backup Server automatically deletes the oldest recovery set. Do not attempt to delete the recovery set manually.

Example Set 1:

- Full
- Incremental
- Incremental
- Verify
- Incremental

Example Set 2:

- Full
- Incremental
- Full
- Incremental

A full backup is required to start a new recovery set. The backup that starts the set will be automatically converted to a full backup, even if there is no full backup configured or scheduled to be performed at that time. After the recovery set setting is changed (for example, changing the recovery set starting point from the first backup of Monday to the first backup of Thursday), the starting point of existing recovery sets will not be changed.

Note: An incomplete recovery set is not counted when calculating an existing recovery set. A recovery set is considered complete only when the starting backup of the next recovery set is created.

Example 1 - Retain 1 Recovery Set:

- Specify the number of recovery sets to retain as 1.

  Backup Server always keeps two sets to keep one complete set before starting the next recovery set.

Example 2 - Retain 2 Recovery Sets:

- Specify the number of recovery sets to retain as 2.

  Backup Server deletes the first recovery set when the fourth recovery set is about to start. This ensures that when the first backup is deleted and the fourth is starting, you still have two recovery sets (recovery set 2 and recovery set 3) available on disk.

  Note: Even if you choose to retain only one recovery set, you will need space for at least two full backups.

Example 3 - Retain 3 Recovery Sets:

- The backup start time is 6:00 AM, August 20, 2012.
- An incremental backup runs every 12 hours.
- A new recovery set starts on Friday. By default, the first backup job on Friday will be the start of the new recovery set.
- You want to retain 3 recovery sets.

  With the above configuration, an incremental backup will run at 6:00 AM and 6:00 PM every day. The first recovery set is created when the first backup (must be a full backup) is taken. Then the first full backup is marked as the starting backup of the recovery set. When the backup scheduled at 6:00 AM on Friday is run, it will be converted to a full backup and marked as the starting backup of the recovery set.

  Note: If you want to run the backup job repeatedly until the end of day, set the schedule until 12:00 AM. For example, to run backup job every 15 minutes for the entire day, set the schedule from 12:00 AM to 12:00 AM, every 15 minutes.

Specify the Advanced Settings

Using the Advanced tab, specify some additional settings for the backup job including the backup throughput and pre/post script settings.

Follow these steps:

1. Click the **Advanced** tab.
2. Specify the throttle backup value.

   Applicable only when the backup destination is a local or a shared folder.

   You can specify the maximum speed (MB/min) at which backups are written. You can throttle the backup speed to reduce CPU or network use. However, limiting the backup speed has an adverse effect on the backup window. As you lower the maximum backup speed, it increases the amount of time of perform the backup.

   Note: By default, the Throttle Backup option is not enabled and backup speed is not being controlled.
3. Specify your pre-backup settings and post-backup settings in **Pre/Post script Settings**.

   These scripts run script commands for actions to take before the start of the job and/or upon the completion of the job.

   Note: The Pre/Post Script Settings fields are populated only if you have already created a script file and placed it at the following location of Linux Backup Server:

   /opt/Arcserve/d2dserver/usr/prepost

   Note: For more information about creating the pre/post scripts, see *Manage Pre/Post Scripts for Automation*.
4. Click Enable Email Alerts to specify the Email Settings and select job alerts..

   Applicable only when the backup destination is Arcserve Recovery Point Server.
5. Click **Save**.

   The changes are saved.

   The backup plan is created and automatically deployed to the source node. The backup runs per the schedule that you have configured in the **Schedule** tab. You can also perform a manual backup at any time.

   Now, you can add the following tasks to the plan:

   - Replicate
   - Replicate to a remotely-managed RPS
   - Copy to Tape

(Optional) Manage Pre/Post Scripts for Automation

Pre/Post scripts let you run your own business logic at specific stages of a running job. You can specify when to run your scripts in **Pre/Post Script Settings** of the **Backup Wizard** and the **Restore Wizard** in the Console. The scripts can be run on the Backup Server depending on your setting.

Managing the pre/post script is a two part process, consisting of creating the pre/-post script and placing the script in the prepost folder.

Create Pre/Post Scripts

Follow these steps:

1. Log into the Backup Server as a root user.
2. Create a script file using the environment variables in your preferred scripting language.

Pre/Post Script Environment Variables

To create your script, use the following environment variables:

D2D_JOBNAME

Identifies the name of the job.

D2D_JOBID

Identifies the job ID. Job ID is a number provided to the job when you run the job. If you run the same job again, you get a new job number.

D2D_TARGETNODE

Identifies the node that is being backed up or restored.

D2D_JOBTYPE

Identifies the type of the running job. The following values identify the D2D_ JOBTYPE variable:

backup.full

Identifies the job as a full backup.

backup.incremental

Identifies the job as an incremental backup.

backup.verify

Identifies the job as a verify backup.

restore.bmr

Identifies the job as a bare-metal recovery (bmr). This is a restore job.

restore.file

Identifies the job as a file-level restore. This is a restore job.

D2D_SESSIONLOCATION

Identifies the location where the recovery points are stored.

D2D_PREPOST_OUTPUT

Identifies a temp file. The content of the first line of the temp file is displayed in the activity log.

D2D_JOBSTAGE

Identifies the stage of the job. The following values identify the D2D_JOBSTAGE variable:

pre-job-server

Identifies the script that runs on the Backup Server before the job starts.

post-job-server

Identifies the script that runs on the Backup Server after the job completes.

pre-job-target

Identifies the script that runs on the target machine before the job starts.

post-job-target

Identifies the script that runs on the target machine after the job completes.

pre-snapshot

Identifies the script that runs on the target machine before capturing the snapshot.

post-snapshot

Identifies the script that runs on the target machine after capturing the snapshot.

D2D_TARGETVOLUME

Identifies the volume that is backed up during a backup job. This variable is applicable for pre/post snapshot scripts for a backup job.

D2D_JOBRESULT

Identifies the result for a post job script. The following values identify the D2D_JOBRESULT variable:

success

Identifies the result as successful.

fail

Identifies the result as unsuccessful.

D2DSVR_HOME

Identifies the folder where Backup Server is installed. This variable is applicable for the scripts that run on the Backup Server.

The script is created.

Note: For all scripts, a return value of zero indicates success and a nonzero return value indicates failure.

Place the Script in the Prepost Folder and Verify

All the pre/post scripts for a Backup Server are centrally managed from the prepost folder at the following location:

/opt/Arcserve/d2dserver/usr/prepost

Follow these steps:

1. Place the file in the following location of the Backup Server:

   /opt/Arcserve/d2dserver/usr/prepost
2. Provide the execution permission to the script file.
3. Log into the Arcserve Unified Data Protection Agent for Linux web interface.
4. Open the **Backup Wizard** or the **Restore Wizard** and navigate to the **Advanced** tab.
5. Select the script file in the **Pre/Post Script Settings** drop-down list and then submit the job.
6. Click **Activity Log** and verify that the script is executed to the specified backup job.

   The script is executed.

   The pre/post scripts are successfully created and placed in the prepost folder.

(Optional) Perform a Manual Backup

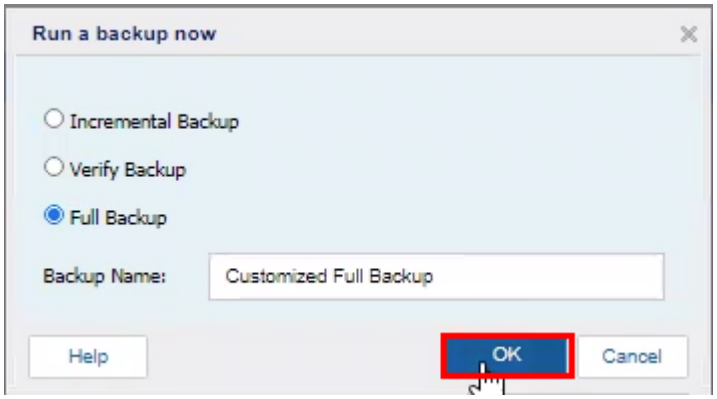
Typically, backups are performed automatically and are controlled by the schedule settings. In addition to the scheduled backup, a manual backup provides you the option to back up your nodes on a need basis. For example, if you have a repeat schedule for Full, Incremental, and Verify backups and you want to make major changes to your machine, you must perform an immediate manual backup without waiting for the next scheduled backup to occur.

Follow these steps:

1. Log into Arcserve UDP Console.
2. Navigate to **resources** > **Plan** > **All Plans**.

   The list of all plans is displayed on the center pane.
3. Right-click the plan, and then select **Backup Now**.

   The Run a backup now dialog opens.
4. On the Run a backup now dialog, select the type of backup you want to perform, and then specify the backup name if required.



The available options are:

Incremental Backup

Initiates an Incremental Backup of your machine. An Incremental Backup backs up only those blocks that have changed after the previous backup.

Note: If you are performing a backup for the first time, the full backup gets initiated automatically even though you selected the Incremental Backup option.

The advantages of Incremental Backups are that it is a fast backup and produces a small backup image. This is the most optimal way to perform backups.

Verify Backup

Initiates a Verify Backup of your machine by examining the most recent backup of each individual block and comparing the content and information to the original source. This comparison verifies that the latest backed up blocks represent the corresponding information at the source. If the backup image for any block does not match the source, Arcserve UDP Agent (Windows) refreshes (resynchronizes) the backup of the block that does not match.

Advantage: A very small backup image is produced when compared to a Full Backup because only the changed blocks (blocks that do not match the last backup) are backed up.

Disadvantage: The backup time is slow because all the source disk blocks are compared with the blocks of the last backup.

Full Backup

Initiates a Full Backup of your entire machine or the selected volumes.

Notes:

- If you add a new volume to the backup source, the newly added volume is fully backed up regardless of the overall backup method selected.
- If no backup name is specified, by default, it is automatically named as Customized Full/Incremental/Verify Backup.

5. Click **OK**.

   The Progress Information dialog appears.

6. Wait until the backup job completes, and then click **OK**.

7. To monitor the backup jobs, navigate to **jobs** > **Jobs in Progress**.

   Wait until the backup job is finished.

   Note: Only one job can be run at a time. If you manually attempt to launch a backup job when another job is currently running, an alert message informs you that another job is running and requests that you try again at a later time.

   The manual backup is successfully performed.

Verify the Backup

To verify the backup, confirm that you have successfully created the backup plan. After you verify that the plan is created successfully, confirm whether the backup job is running as scheduled. You can verify the status of backup jobs from the jobs tab.

Follow these steps: to verify plans

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and then click **All Nodes**.

   A list of all nodes is displayed on the center pane.
3. Verify that plans are mapped with nodes.

Follow one of these methods to verify backup jobs

Method 1

1. Click the **jobs** tab.
2. From the left pane, click **All Jobs**.

   The status of each job is listed on the center pane.
3. Verify that the backup job is successful.

   The backup job is verified.

Method 2

1. From the Arcserve UDP console, navigate to **jobs** tab > **Jobs Successfully Completed**.

   The list of backup jobs completed appear in the center pane.
2. Click the finished reverse replication job. The right pane refreshes.
3. Under Job Details, click the **View Logs** hyperlink.
4. Verify that the backup job is successful.

   The backup job is verified.

The backup job is verified.

Troubleshooting

Job Status, Job History, and Activity Log are Not Visible

Job Status, Job History, and Activity Log are Not Visible

Symptom

I cannot see the job status, job history, and activity log for Linux nodes in Arcserve UDP Console.

Solution

Linux Backup Server is unable to connect to Arcserve UDP using the hostname.

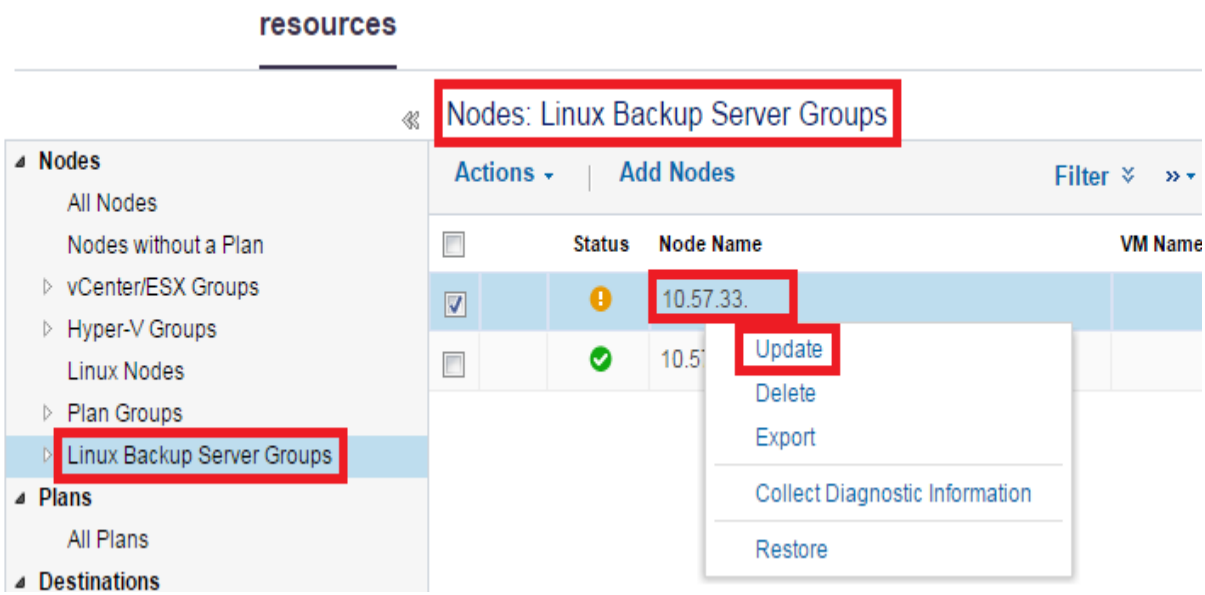
Follow these steps:

1. Create the server_ip.ini file at the following location of Arcserve UDP:

   "UDP installation path"\Management\Configuration\server_ip.ini

2. Enter the IP address of Arcserve UDP in this file.
3. Log into the Arcserve UDP Console and update Linux Backup Server and Linux nodes.

   Note: Linux Backup Server can be updated only from Linux Backup Server Groups, where all the Linux backup servers are listed.



The job status, job history, and activity log are visible.

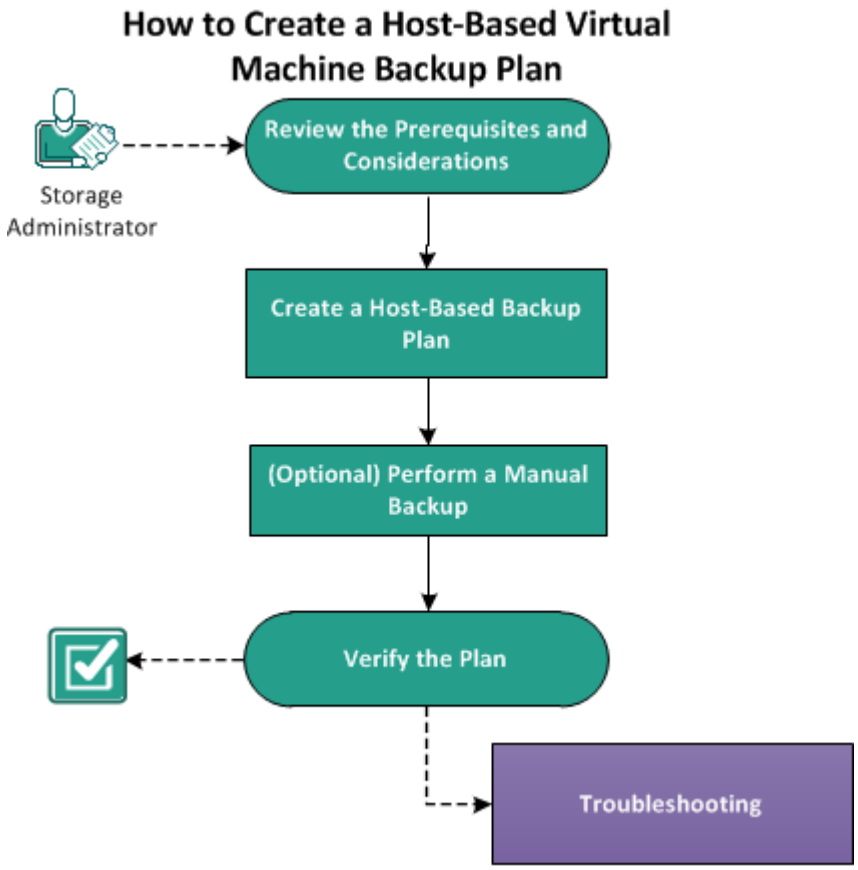
How to Create a Host-Based Virtual Machine Backup Plan

To protect your host-based virtual machine nodes, you need to create a host-based backup plan. A backup plan for host-based virtual machine nodes consists of a backup task. The backup task lets you specify the nodes that you want to protect, the backup destination, and the backup schedule. The backup destination can be a local destination or a remote share folder, or the recovery point server where you want to store your backup data.

You can also back up Oracle databases, SQL and Exchange Servers. To back up Oracle databases, ensure specific prerequisites. To back up SQL Server and Exchange Server no prerequisites are required. Review the following prerequisites to perform an application consistent backup of an Oracle database:

- Prerequisite to create an application consistent backup of an Oracle database

The following diagram illustrates the process to protect host-based virtual machine nodes.



What To Do Next?

1. Review the Prerequisites and Considerations
2. Create a Host-Based Backup Plan
3. (Optional) Perform a Manual Backup
4. Verify the Plan
5. Troubleshooting

Review the Prerequisites and Considerations

Verify if you have completed the following prerequisites:

- Log into the Console.
- Prepare a host-based backup proxy server where you have installed Arcserve Unified Data Protection Agent for Windows.
- To run functions such as Preflight Check, pre/post commands, or application log purge, update the virtual machine in the node list view of the Arcserve UDP Console with one of the following credentials for the guest virtual machine:
  - Built-in administrator user credentials.
  - Built-in domain administrator user credentials.
  - For other administrator credentials, disable the User Account Control (UAC) on the guest virtual machine.
- To be able to perform a database level restore (for Exchange and SQL Server) or granular level restore (Exchange) after a backup, the following prerequisites must be met:
  - The VM must support the application consistent backup. For more information on application consistent backup, see How to Create Application Consistent Snapshots for VMware or How to Create Application Consistent Snapshots for Hyper-V.
  - For VMware VM, the **VMware Tools** snapshot quiescing method must be used in the backup plan.
  - For Hyper-V VM, Arcserve UDP needs to automatically deploy a utility into the guest OS of VM to gather the application metadata during a backup. The guest OS of the VM needs to be accessed from either the backup proxy server or Hyper-V host using a network. At the same time, the VM node must be updated with proper administrative credentials in the node list view of Arcserve UDP Console. For some reasons, if the guest OS of the VM cannot be accessed from the backup proxy server and Hyper-V host both, follow these steps to manually install the utility in the guest OS of VM:
    a. Log into the backup proxy server and navigate to the following folder:

       <Arcserve UDP installation path>\Engine\BIN (Example, C:\Program Files\Arcserve\Unified Data Protection\Engine\BIN)

b. Locate the executable file **VMICService_32.exe** or **VMICService_64.exe**.

c. Copy the executable file to any a folder inside the guest OS of the VM. (For a 32-bit OS, copy **VMICService_32.exe**, otherwise copy **VMICService_64.exe**).

   For example, you can create an ISO image by including this executable file and mount it to the DVD device of the VM.

d. Log into the guest OS of the VM and run either **VMICService_32.exe -install** or **VMICService_64.exe -install**.

- Install the server component and create Data Stores if you want to store the backup data in the recovery point server.
- Review the prerequisites to back up an Oracle database.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

The following prerequisites are for hardware snapshots:

For Hyper-V

- Install a VSS hardware provider on the Hyper-V servers and to support the transportable snapshot install the VSS hardware provider on the backup proxy server. A typical configuration of a VSS hardware provider include:
  - Specifying a server that controls the LUN.
  - Specifying the disk array credentials to access the disk array.

  For more information on configuring the VSS hardware provider, contact your hardware provider vendor.
- The Hyper-V server and the proxy server must have a similar operating system version.
- If Hyper-V server belong to a cluster, the proxy server should not be part of the Hyper-V cluster.
- uninstall UDP CBT from "Programs and Features" in control panel in Hyper-V host.

For VMware

- Arcserve UDP supports NetApp iSCSI and NetApp NFS LUNs.
- To create a hardware snapshot for VMware, add the storage array to the Console. For more information on adding a storage array, see Add a Storage Array.
- To use the hardware snapshot, Flexclone license is mandatory for NetApp storage arrays running with data ONTAP operating in 7-Mode and Cluster mode.

Considerations for backing up the VM:

- **How a volume defragmentation can affect continued backups**

  The volume defragmentation by Windows native tool affects the size of the block-level backups because Arcserve UDP continues to incrementally back up all changed blocks. It means that blocks that shifted during the defragmentation are included in the backup, even if no data has changed in the files. As a result, the backup size increases. This behavior is expected.

- **How to protect Virtual Machines in Hyper-V 2016 using Windows Resilient Change Tracking**

  To protect VM in Hyper-V 2016, we recommend to leverage Windows Resilient Change Tracking (RCT) for incremental backup. Arcserve UDP automatically uses Windows Resilient Change Tracking (RCT) when backing up virtual machine with configuration version 8.0 or higher that runs in Hyper-V 2016 host/cluster. If your Hyper-V is upgraded from previous version, upgrade your VM configuration also using Microsoft documentation. When upgrade is complete for all the virtual machines protected by Arcserve UDP, we recommend to uninstall Arcserve UDP CBT Service from the Hyper-V host.

Considerations for specific network backup:

- Hypervisor, Windows Proxy and RPS must be in the same network.
- If ESXi is added with **Hostname** to the vCenter, the backup network selection must be default network of ESXi.

  Note: Default network is the network which is resolved by DNS.

- If ESXi is added with IP to the vCenter, the backup network selection should be the same network that ESXi uses to attach to the vCenter.

Scenario #	Has vCenter	What is used to add into vCenter	Could use backup network?	DNS / Hosts file work-around works?
1	No	N/A	Yes	N/A
2	Yes	Hostname of ESXi – The Production IP of the hostname is resolved	No	Yes. Update the DNS or the hosts file on the Proxy Server.
3	Yes	Hostname of ESXi – The Backup Network IP of the hostname is resolved	Yes	N/A
4	Yes	Production IP of ESXi	No	No
5	Yes	Backup Network IP of ESXi	Yes	N/A

Review Prerequisites to Perform Application Consistent Backup for Oracle Database

To back up an Oracle database with consistent data, ensure that the ARCHIVELOG mode is enabled to archive the Redo logs.

Note: The data volume must include Oracle data files, control files, server parameter file, and online redo logs. The archived redo logs must be physically located on a separate volume.

Follow these steps to verify if the ARCHIVELOG mode is enabled:

a. Log into the Oracle server as an Oracle user with SYSDBA privileges.

b. Enter the following command at the SQL*Plus prompt:

ARCHIVE LOG LIST;

Archive log settings for the current instance is displayed.

c. Configure the following settings:

Database log mode: Archive Mode

Automatic archival: Enabled

d. Start the ARCHIVELOG mode.

Note: If the ARCHIVELOG mode is not enabled, start the ARCHIVELOG mode to back up the database.

Follow these steps to start the ARCHIVELOG mode:

a. Shut down the Oracle server.

b. Run the following statements in Oracle:

```
CONNECT SYS/SYS_PASSWORD AS SYSDBA

STARTUP MOUNT;

ALTER DATABASE ARCHIVELOG;

ALTER DATABASE OPEN;
```

By default, archive logs are written to the flash recovery area. If you do not want to write archive logs to the flash recovery area, set the LOG_ARCHIVE_DEST_n parameter to the location where you want to write archive logs.

```
SQL>ALTER SYSTEM SET LOG_ARCHIVE_DEST_
1='LOCATION=e:\app\administrator\oradata\<oracle_database_name>\arch'
SCOPE= BOTH;

System altered.

SQL> ARCHIVE LOG LIST;
```

Archive log settings for the current instance is displayed.

c. Configure the following settings:

Database log mode: Archive Mode

Automatic archival: Enabled

Archive destination: E:\app\oracle\oradata\<oracle_database_name>\arch

Oldest online log sequence: 21

Current log sequence: 23

The Oracle VSS writer service started and is functioning properly.

Create a Host-Based Backup Plan

A backup plan includes a backup task that performs a backup of the virtual machine and stores the data to the specified destination. Each task consists of parameters that define the source, destination, schedule, and other backup details.

Follow these steps:

1. Click the **resources** tab on the Console.
2. From the left pane, navigate to **Plans**, and click **All Plans**.

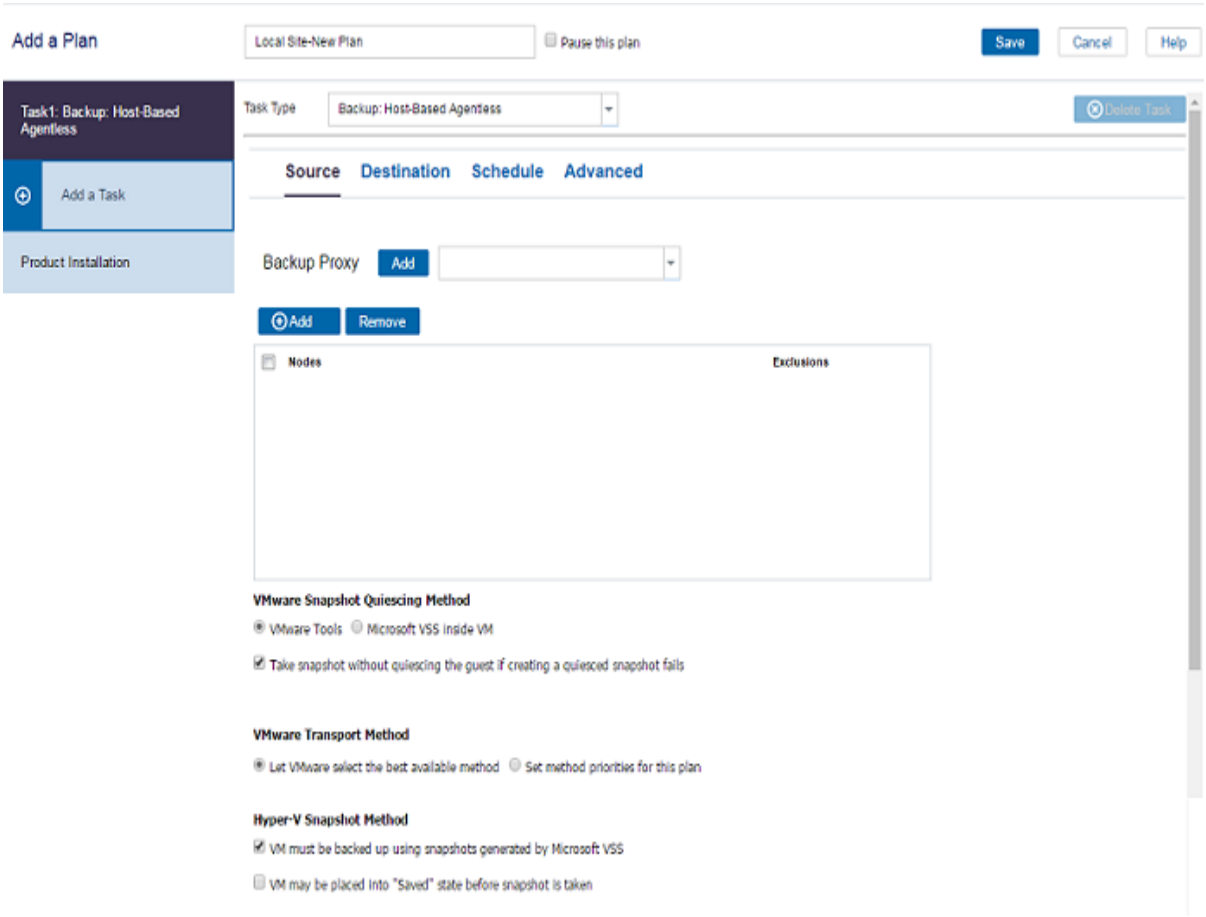
   If you have created plans earlier, those plans are displayed on the center pane.
3. On the center pane, click **Add a Plan**.

   Add a Plan opens.
4. Enter a plan name.
5. (Optional) Select **Pause this plan** check box.

   The plan will not run until you clear the check box to resume the plan.

   Note: If a plan is paused, then any in-progress job is not paused but all corresponding scheduled jobs associated with that plan are paused. However, you can manually run a job. For example, you can manually run backup job and replication job for a node even if the respective plan is paused. In such case, the following task to the on-demand (manual) job does not run. For example, there is a replication task after an on-demand backup job, the replication job does not run for the on-demand backup job. You need to manually run the replication job. When you resume the plan, the pending jobs do not resume immediately. After you resume the plan, the pending jobs run from the next scheduled time.

6. From the **Task Type** drop-down list, select **Backup, Host-Based Agentless**.



Now specify the Source, Destination, Schedule, and Advanced details.

Specify the Source

The Source page lets you specify the source nodes that you want to protect. You can select more than one node in a plan. If you have not added any nodes to the Console, you can add nodes when you create or modify a plan from the Source page. You can save a plan without adding any source nodes. The plan gets deployed only after you add source nodes.

Follow these steps:

1. Click the **Source** tab and add a backup proxy server.

   The proxy server is a node where you install the Arcserve Unified Data Protection Agent for Windows. If the Agent is not installed on this proxy server, when you save the plan, the agent is deployed to the proxy server. The agent deployment setting is in the Agent Installation task in the plan.

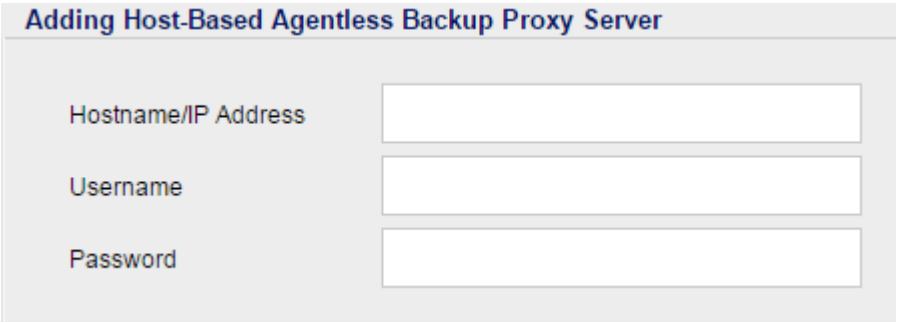
   - If Backup Proxy is already added, select the backup proxy from the drop-down list.

   - If the backup proxy is not added, then click **Add**.

     The **Adding Host-Based Agentless Backup Proxy Server** dialog opens.

     Specify the proxy server details and click **OK**.

   You do not have to add the port number and protocol. The port number and protocol are configured on the **settings** tab of the Console.

   Note: When you modify a plan by changing the Backup Proxy but the nodes included in the plan have jobs running, the plan deployment fail. Follow these steps to change proxy of a plan:

a. Pause the plan.

b. Wait until all the nodes in the plan complete the running backup jobs. (Or, you can cancel the running jobs).

c. Change the proxy of the plan and save it.

d. Resume the plan.

2. Click **Add Nodes** to select one of the following options to add nodes that you want to backup:

Select Nodes to Protect

Opens the **Select Nodes to Protect** dialog and you can select the nodes from the displayed list. Select this option if you have already added the nodes to the Console.

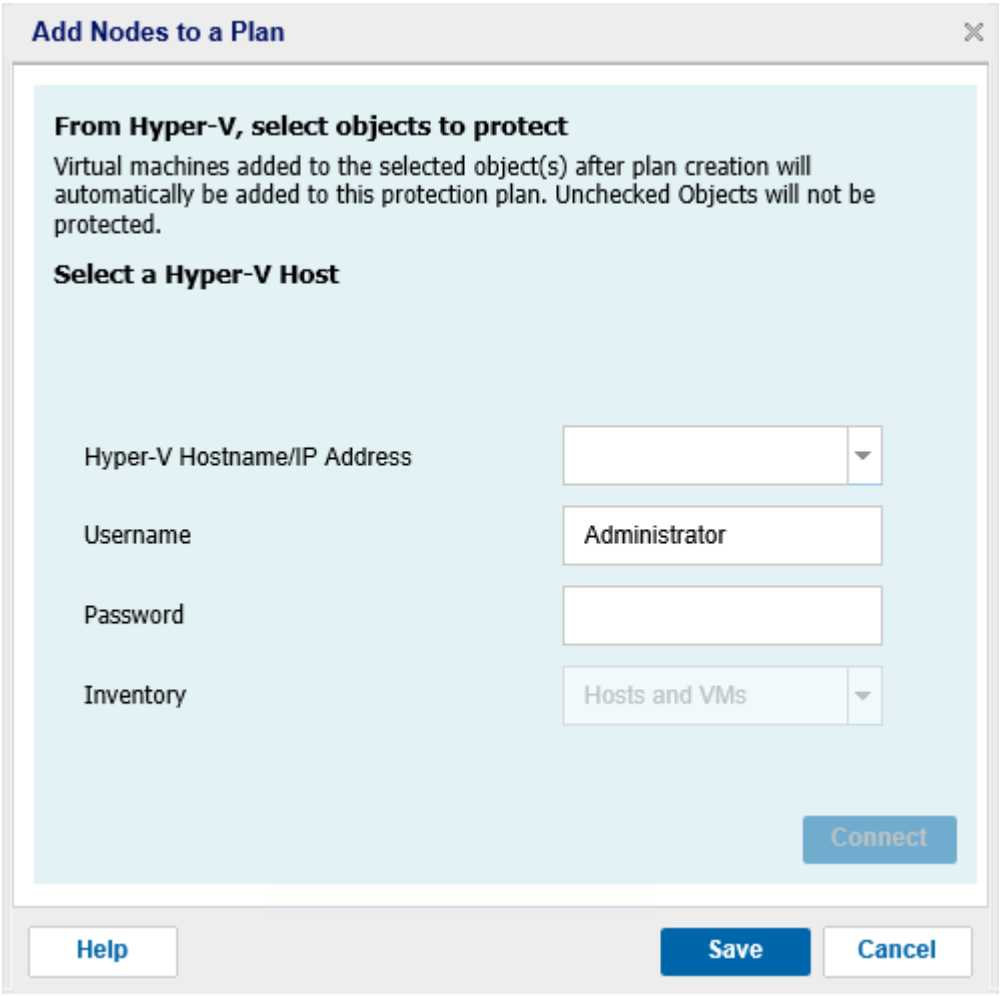
Add Nodes from a Hyper-V

Opens the **Add Nodes to Arcserve UDP Console** dialog. You can add individual VM nodes or VM container objects (including the Hyper-V cluster, the Hyper-V host, the storage location) to a plan. Once you add a VM container object to the plan, the plan protects all the assigned VMs automatically. When a new VM is created under the VM container object or moved to the VM container object from other place, Arcserve UDP protects the VM without manual intervention. If the VM is deleted from the VM container object (or moved out of the VM container object), Arcserve UDP stops protecting the VM.

To add a VM container object into a plan, follow these steps:

a. Specify the Hyper-V server details, select how you want to browse the VMs in the Inventory drop-down list box (Hosts and VMs, Storages and VMs), and click Connect.

A Hyper-V hierarchy tree is displayed.

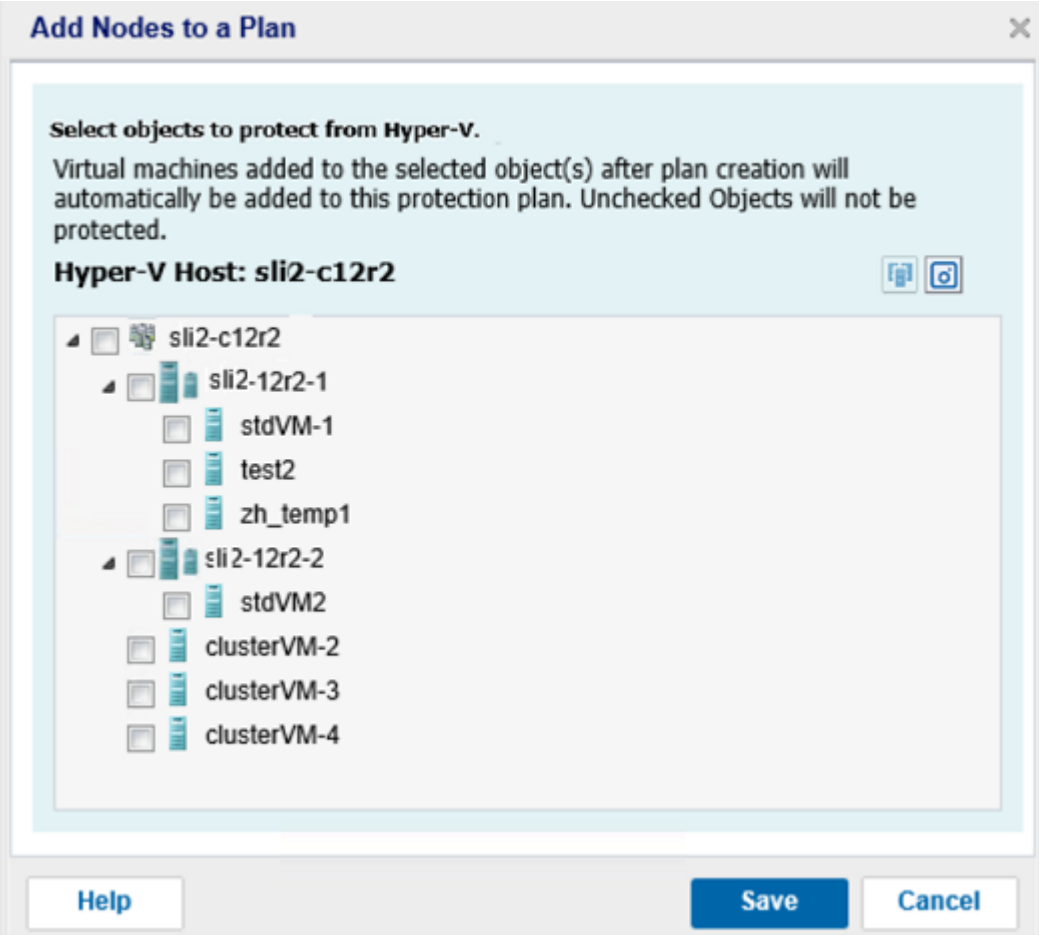


b. Expand the Hyper-V hierarchy tree to perform the following options:

Note: Depending upon the selection in the Inventory drop-down list, different hierarchy tree views are displayed - Hosts and VMs or Storages and VMs. To switch between different tree views, click the buttons on the top-right side of the tree.

i. **Add an Individual VM**

Select the check box of the VM that you want to add from the list, and click **Save**.



ii. **Add a VM container object**

Select the check box of the VM container object that you want to add from the list, and click Save. While you select the VM container object to protect, you can also clear selection of a child VM or child VM container object available below the object. As a result, the child VM or child VM container object is excluded from protection.

Note: Arcserve UDP uses Hyper-V cluster/host name, volume name or SMB share name to uniquely identify the container object in the Hyper-V list. As a result, you can add the same VM container object to plans more than once.

For example, add a Hyper-V host by the host name and then using IP address add the same Hyper-V host again. Another example is to add one SMB share, which uses the host name of the host machine, add the same SMB share, which uses IP address of the host machine.

Important! Arcserve UDP does not support adding the same container object to a plan twice to avoid unpredictable behaviors.

Add Nodes from vCenter/ESX

Opens the **Add Nodes to Arcserve UDP Console** dialog. You can add individual VM nodes, VM templates or VM container objects to a plan. Once you add a VM container object to the plan, the plan protects all the assigned VMs automatically. When a new VM is created under the VM container object or moved to the VM container object from other location, Arcserve UDP protects the VM without manual intervention. If the VM is deleted from the VM container object or moved out of the VM container object, Arcserve UDP stops protecting the VM.

Note: Apart from the VM container objects that exist in vSphere, Arcserve UDP also considers tag and tag category as VM container object, so that they are added into agentless backup plan. In this case, all VMs and template assigned with that tag are automatically protected. In addition, if an up level VM container object is assigned with a tag, all the VMs under that VM container object are considered to have the same tag virtually. For example, assigning a tag to a resource pool automatically provides the same tag to the VMs of that pool even when they actually do not have any tag assigned in vCenter. In addition, automatic protection by tag is supported only for vCenter 6.0 and 6.5.

To add a VM container object into a plan, follow these steps:

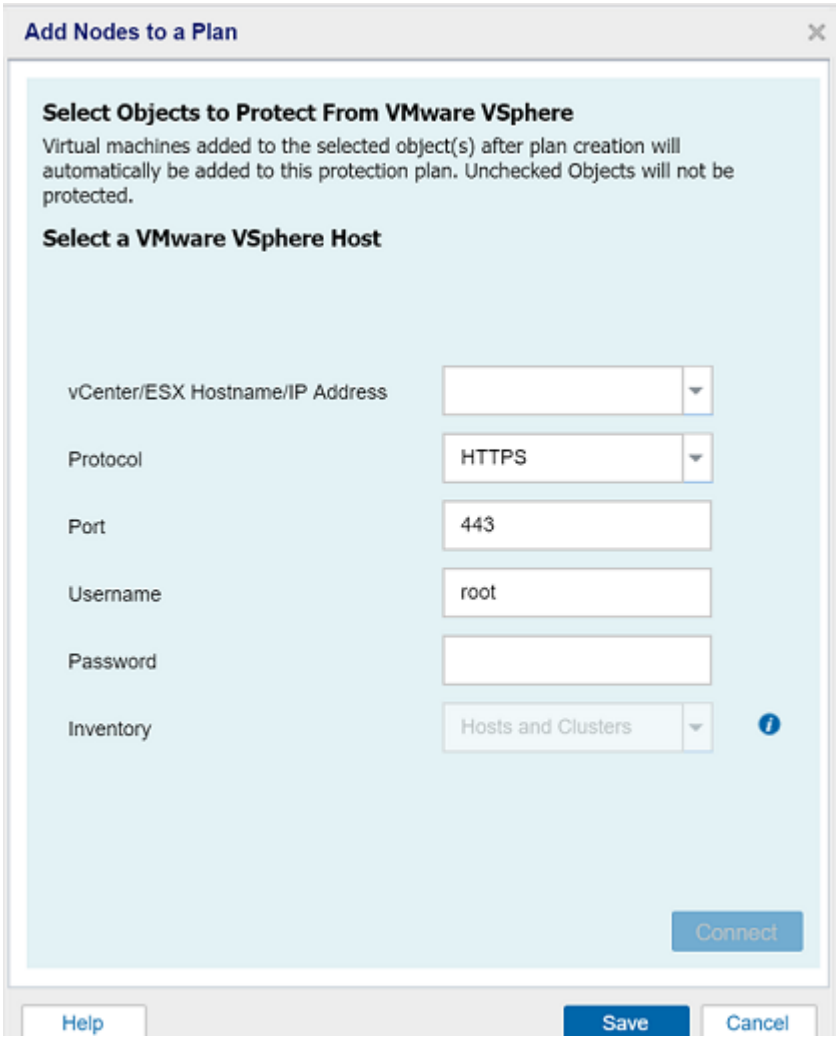
a. Specify the vCenter/ESX server details, select how you want to browse the VMs in Inventory drop down list box (Hosts and Clusters, VMs and Templates, VMs and Tags), and click **Connect**.

   A vSphere hierarchy tree is displayed.

   Note: Consider the following:

   - The VMware Virtual Disk Development Kit (VDDK) 6.x.x is bundled with Arcserve UDP 7.0. But, VDDK 6.x.x does not support HTTP. Also, vCenter and ESX support HTTPS connection only by default.
   - Select HTTPS **Protocol**, if you want to manually replace the built-in VDDK 6.x.x with another version of VDDK, and to configure

vCenter/ESX to allow HTTP connection manually.

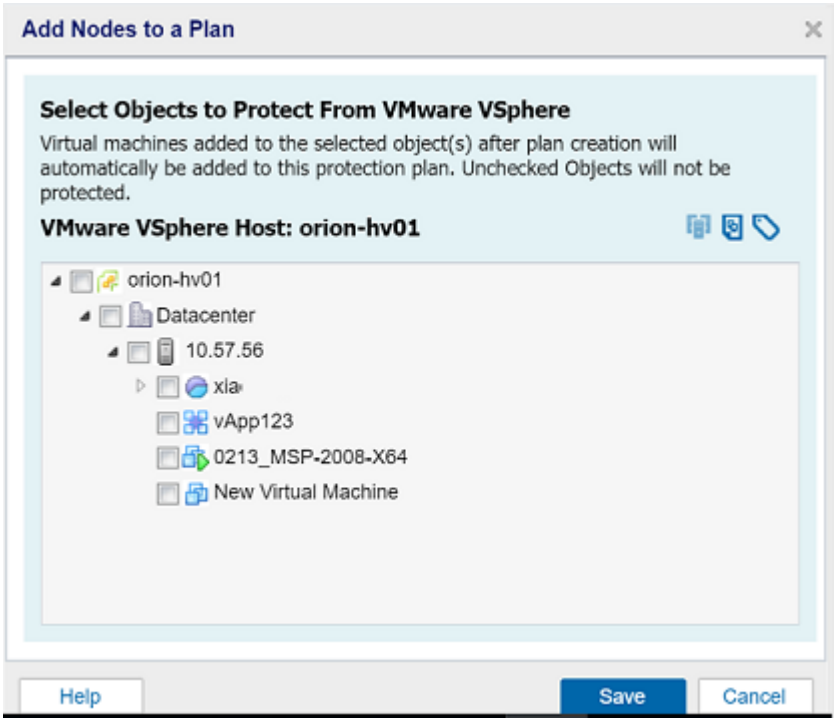


b. Expand the vSphere list to add the following:

Note: Depending upon what is selected in the Inventory drop-down list, different hierarchy tree views are displayed - Hosts and Clusters, VMs and Templates and VMs and Tags. You can switch among different tree views by clicking the buttons on the top-right side of the tree.

i. **Add an individual VM**

Select the check box of the VM that you want to add from the list, and click **Save**.



ii. **Add a VM container object**

Select the check box of the container that you want to add from the list, and click **Save**. While you select the container to protect, you can also clear selection of a child VM or child VM container object that is under it. As a result, the child VM or child VM container object is excluded from protection.

Note: Arcserve UDP uses vCenter/ESX name and vSphere MoRef ID (Managed Object Reference ID) to uniquely identify the VM container object in the **vSphere** list. This allows you to add the same VM container object to plan more than once.

For example:

a. Add a VM container object to a vCenter by connecting to vCenter using its host name, and then add the same VM container object again by connecting using the vCenter IP address.

b. Add a VM container object to a vCenter, and then add it again from the ESX host directly.

Important: Arcserve UDP does not support adding the same VM container object to a plan twice, as it may cause unpredictable behaviors.

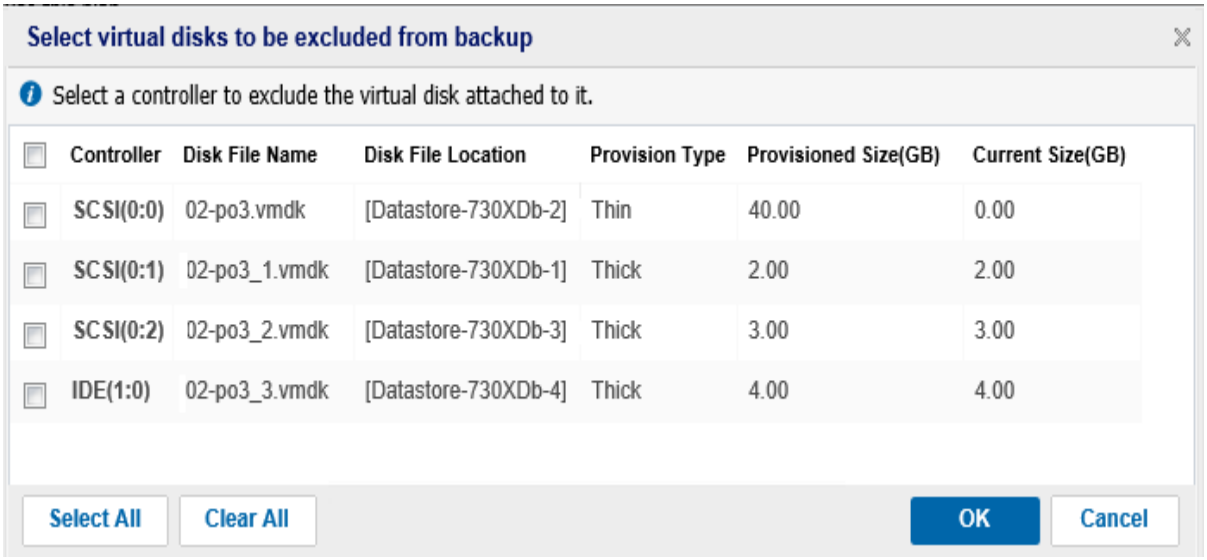
3. (Optional) Select virtual disk that you can exclude from the backup job.

By default the Agentless backup job backs up the whole VM including all its virtual disks. But, you can specify one or more virtual disks that are skipped during backup.

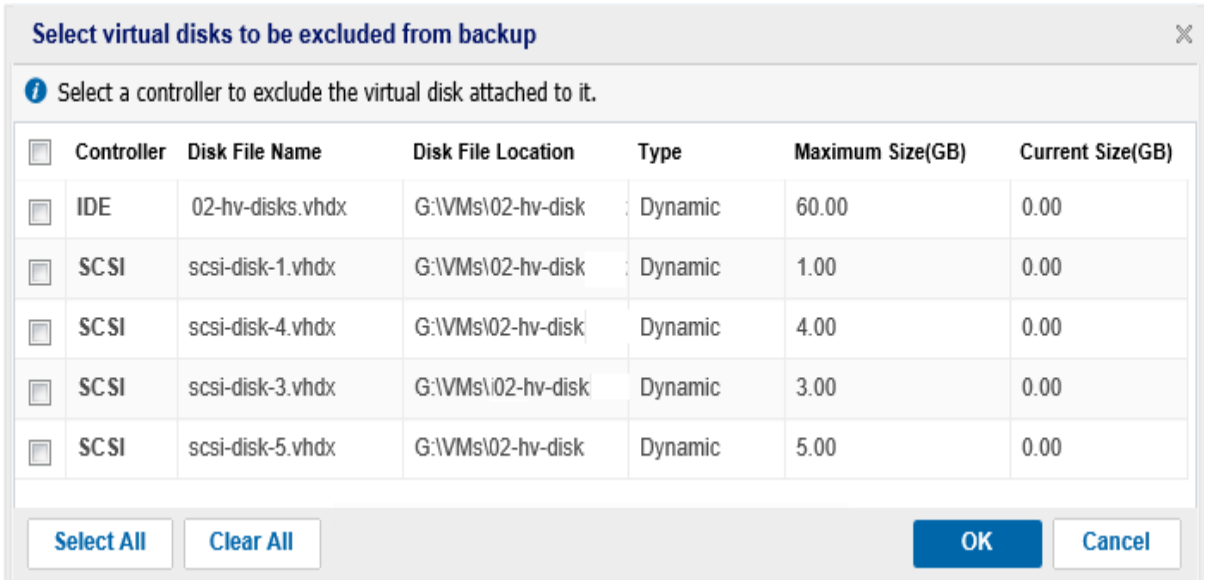
a. After adding a VM node into plan, click on the Configure button available at the right side of node.

   A dialog box pops up.

   Dialog box for VMware VM



   Dialog box for Hyper-V VM



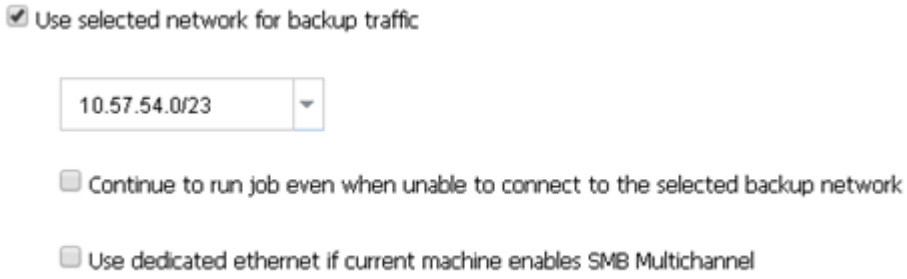
b. Select the check box for the virtual disk that you want to exclude from backup and click **OK** to save.

Notes:

- Virtual disk is excluded by the controller, instead of the name of virtual disk file.
- If the virtual disk containing the system volume of guest OS is excluded by backup, restored VM cannot boot up.
- If the VM has application (SQL Server or Exchange) installed and any virtual disk is excluded by backup, the DB level restore is not allowed.

4. (Optional) Select the **Use selected network for backup traffic** checkbox and follow these steps:

a. To enable communication between the Windows Proxy and Hypervisor Server, select the CIDR network from the drop-down menu.



b. If you want the backup task to continue even if the selected network is unavailable between Proxy Server and Hypervisor, select the **Continue to run job even when unable to connect to the selected backup network** checkbox.

c. To define the constraint on SMB Multichannel so that the data transfers only through selected network, select the **Use dedicated ethernet if current machine enables SMB Multichannel** check box.

Notes:

- This option is not available by default. To enable this option, go to the following folder location: C:\Program Files\Arcserve\Unified Data Protection\Management\Configuration\ConsoleConfiguration.xml. Then, modify the modify the value for useDedicatedEthernet as True.

```
- <SpecifyNetwork>
    <useDedicatedEthernet>false</useDedicatedEthernet>
  </SpecifyNetwork>
```

- The SMB Multichannel feature is enabled in Windows, by default.

5. Specify network settings for Host-based Backup Hyper-V VM Backup:

Registry setting: SMBSpecifiedIONetwork

Location: HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll or HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll\VMInstUUID

Type: Multi string value

Value Name: SMBSpecifiedIONetwork

Value Data: String in the form CIDR (ip/maskBitCount)

Example: 192.168.10.0/24

If the Hyper-V/Hyper-V Cluster VM disks are residing on the SMB share, then the user can also define another network to be used for data transfer between the SMB server and the proxy machine.

Note: The mapping between the specified IP and SMB Server hostname must be registered in the domain.

6. (Optional) Select one of the following quiescing methods for VMware. These options are applicable for VMware only.

VMware Tools

Indicates that Arcserve UDP uses the VMware tools for quiescing the virtual machine. If you have used the **Microsoft VSS inside VM** option in the previous backup job, the first consequent backup job with this option requires the credentials to access the virtual machine. This is because Arcserve UDP removes necessary tools from the VM. In addition, VMware Tools needs to be installed and update to date in the VM.

Microsoft VSS inside VM

Indicates that Arcserve UDP uses Microsoft VSS in the guest OS for quiescing the virtual machine. It is applicable only for virtual machines with Windows guest OS. VMware tools must be installed in the guest OS and the tools must be updated. When you use this option, the virtual machine must be powered on and it must be updated with the built-in administrator credentials. For more information on updating a node, see Update Nodes.

Note: The snapshot provided by VMware using this option may not be application consistent. In other words, the backup generated using this option may not be an application-consistent backup. The workaround is to use VMware Tools snapshot quiescing method, along with disabling the VSS writers *MSSearch Service Writer* and *Shadow Copy Optimization Writer* in guest OS of VM before this problem gets fixed.

Take snapshot without guest quiescence if quiescence snapshot fails

Indicates that, when backup job fails to take the snapshot with quiescence option, Arcserve UDP will continue the backup job by taking a snapshot without quiescing the virtual machine.

Notes:

- The **Microsoft VSS inside VM** option does not support the application database level and granular level of restore.
- Both the quiescing methods are not applicable when the virtual machine is powered off. If a backup job is initiated when the virtual machine is powered off, the backup job ignores both the quiescing methods.
- For both the quiescing methods, if the backup job cannot continue for any reason (for example, the credentials are incorrect), Arcserve UDP fails the backup job. For more information about the backup job failure, see the troubleshooting topic

7. (Optional) Select one of the transport methods for VMware. These options are applicable for VMware.

   Let VMware select the best available method

   Indicates that VMware selects the data transfer option. You do not have to manually set any data transfer option.

   Set method priorities for this plan

   Indicates that you can select the data transfer option and set the priority for each option. Use the arrow button to prioritize the transport mode.

   - HOTADD transport mode
   - NBD transport mode
   - NBDSSL transport mode
   - SAN transport mode

   Note: If you have specified the transport mode in both the Console and registry key, then the priority set from the Console overrides the priority set in the registry key. For more information on setting the priority using the registry key, see Define a Transport Mode for Host-Based Agentless Backup and Restore.

8. (Optional) Select Hyper-V snapshot method. These options are applicable for Hyper-V only.

   VM must be backed up using snapshots generated by Microsoft VSS method

   Indicates that Arcserve UDP uses the native snapshot methods of Microsoft - online and offline for the backup job. This is the default option. When this check box is not selected and when both Microsoft online and offline

methods are not available, the backup job uses the Arcserve UDP method to back up the virtual machine.

If the Microsoft offline method is used for backup and the virtual machine is required to be in a Saved state, select the **VM may be placed into "Saved" state before snapshot is taken** check box also. If you do not select this check box, the backup job fails.

Online backup is the recommended backup method because it supports the application consistent backup without the downtime of the virtual machine. The virtual machine is accessible during the backup. The online backup method must satisfy some prerequisites such as integration services must be installed and running. If any of the prerequisites are not satisfied, then only the offline backup method is used.

The Microsoft offline backup method has two approaches - save state approach and checkpoint approach. If the Hyper-V host has the Windows 2012R2 operating system with KB 2919355 or later, then the checkpoint approach is used; else the save state approach is used.

The major difference between these two approaches is that the save state approach requires the virtual machine to be inaccessible for a short time. The virtual machine must be placed into a saved state for a few minutes while taking the snapshot.

Apart from the Microsoft native snapshot methods, Arcserve UDP has its own snapshot method that can be used when the Microsoft native snapshot methods are not available.

Note: Both Microsoft offline method and Arcserve UDP method are crash-consistent backup methods. Both the methods cannot guarantee data integrity. The main difference between the methods is that the Microsoft offline method can be compared to the state that VM has been powered off abruptly whereas the Arcserve UDP method can be compared to the state that Hyper-V host has been powered off abruptly.

VM may be placed into "Saved" state before snapshot is taken

Indicates that the virtual machine is placed in the Saved state, if required, before taking the VSS snapshot. Select this option when the virtual machine does not support the online backup. If the virtual machine supports the online backup, then even on enabling this option the virtual machine will not be in the Saved state.

The source is specified.

Define a Transport Mode in the Registry for Host-Based Agentless Backup and Restore

You can define transport mode (transfer data) for UDP agent as proxy that executes host-based agentless backup or restore job for virtual machines residing on VMware ESX server. By default, host-based agentless backup and restore uses a mode that lets host-based agentless backup and restore to optimize the performance (increase the speed) of the data transfer. However, when you want to specify a particular transport mode for backup or restore, configure the registry key described in this topic.

Note: For backup, the transport mode defined in plan takes precedence over what is defined in registry.

Host-based VM backup can execute backups using the following transport modes:

- HOTADD transport mode
- NBD transport mode
- NBDSSL transport mode
- SAN transport mode

Be aware of the following considerations:

- This is an optional configuration task. By default, host-based VM backup executes backups using a transport mode that optimizes the performance of the backup operation.
- When you configure this registry key to use a specific transport mode and the mode is not available, the host-based VM backup uses an available default transport mode for the backup operation.
- You can define the transport mode for all VMs that are used for backup using the proxy server (proxy level) or define a specific VM (VM level). If you configure both the proxy server and the VM, the VM level registry takes precedence over the proxy level registry.

Follow these steps to define the transport mode at the proxy server level (applicable for both backup and restore):

1. Log in to the Arcserve Unified Data Protection Agent for Windows backup proxy server.
2. Open Windows Registry Editor and browse to the following key:

   [HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine]

3. Right-click VDDKEnforceTransport and click Modify on the pop-up menu to open the Edit String dialog.
4. In the Value Data field, specify the transport mode that you want to use during the backup job. Specify one or more of the following values separated by ":". (For example nbd or san:nbd:nbdssl:)

   hotadd

   HOTADD transport mode

   nbd

   NBD transport mode

   nbdssl

   NBDSSL transport mode

   san

   SAN transport mode
5. Click OK to apply the value and close the Edit String dialog.

   The transport mode is defined and is used the next time when the job runs.

   Note: To restore thin Virtual Machine Disks (VMDK), the non-advanced transport (LAN transport mode) mode is used by default. To enable the advanced transport mode for thin VMDK, update the registry key as shown in the following example:

   a. Open Windows Registry Editor and browse to the following key:

      ```
      [HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine]
      ```

   b. Create a key named AFRestoreDll.

   c. Create a string value named EnforceTransportForRecovery within the AFRestoreDll key.

   d. Specify the transport mode that you want to use during the recovery job. (For example: "san:nbd:nbdssl")

      Example

      [HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFRestoreDll]

      "EnforceTransportForRecovery"="san:hotadd:nbd:nbdssl"

   Follow these steps to define the transport mode at the VM level (applicable for backup only):

1. Log in to the Arcserve Unified Data Protection Agent for Windows backup proxy server for the virtual machines.

2. Open Windows Registry Editor and browse to the following key:

   [HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll\{VM-InstanceUUID}

3. Right-click VM-InstanceUUID and select New.
4. Click String Value on the pop-up menu.
5. Name the new string value as follows.

   EnforceTransport

6. Right-click EnforceTransport and click Modify on the pop-up menu to open the Edit String dialog.
7. In the Value Data field, specify the transport mode that you want to use during the backup job. Specify one of the following values:

   hotadd

   HOTADD transport mode

   nbd

   NBD transport mode

   nbdssl

   NBDSSL transport mode

   san

   SAN transport mode

8. Click OK to apply the value and close the Edit String dialog.

   The transport mode is defined and is used the next time when the job runs.

Specify the Destination

The destination is a location where you store the backup data. You must at least specify the destination to save the plan.

Follow these steps:

1. Select one of the following **Destination Type**:

   Local disk or shared folder

   Specifies that the backup destination is either a local destination or a shared folder. If you select this option, you can save data as either recovery points or recovery sets. The recovery points and recovery sets options are available on the **Schedule** tab.

   Arcserve UDP Recovery Point Server

   Specifies that the backup destination is a recovery point server. If you select this option, then data is stored as recovery points. You cannot store data as recovery sets.

2. If you have selected **Arcserve UDP Recovery Point Server**, then provide the following details:

   a. Select a recovery point server.

   b. Select a data store. The list displays all data stores that are created at the specified recovery point server.

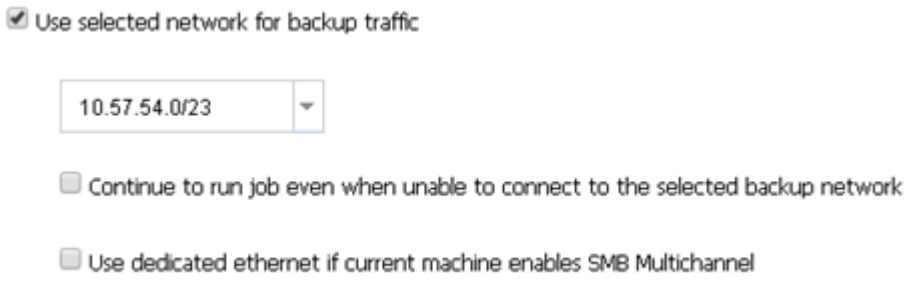
   c. Provide a session password.

   Note: The session password is optional when the backup destination is an unencrypted RPS data store.

   d. Confirm the session password.

   e. (Optional) Select the **Use selected network for backup traffic** checkbox and follow these steps:

      1. To enable communication between the Windows Proxy and Recovery Point Server, select the CIDR network from the drop-down menu.



2. If you want the backup task to continue even if the selected network is unavailable between Proxy Server and Recovery Point Server, select the **Continue to run job even when unable to connect to the selected backup network** checkbox.
3. To disable the SMB Multichannel so that the data transfers only through selected network, select the **Use dedicated ethernet if current machine enables SMB Multichannel** check box.

Notes:

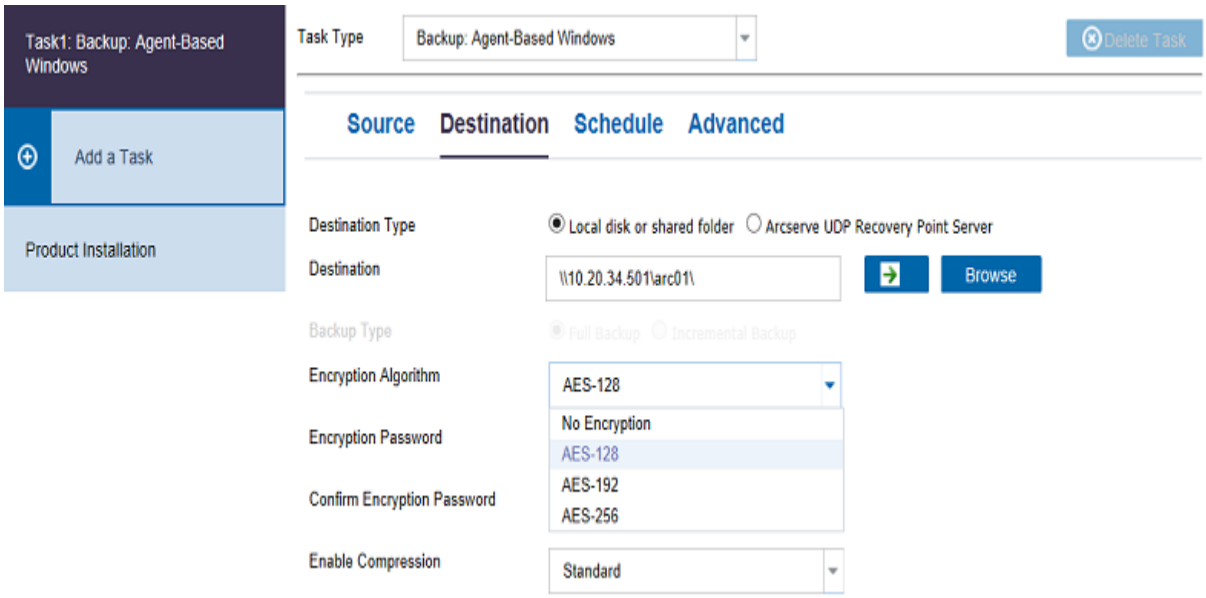
- This option is not available by default. To enable this option, go to the following folder location: C:\Program Files\Arcserve\Unified Data Protection\Management\Configuration\ConsoleConfiguration.xml.Then, modify the value for *useDedicatedEthernet* as True.

```
- <SpecifyNetwork>
      <useDedicatedEthernet>false</useDedicatedEthernet>
  </SpecifyNetwork>
```

  The Specify Network function becomes disabled in case of remote datastore, which used network shared folder as destination.

- The SMB Multichannel feature is enabled in Windows, by default.

3. If you have selected **Local disk or shared folder**, then provide the following details:



a. Provide the full path of the local or network destination. For the network destination, specify the credentials with the write access. You can click Browse to locate the destination or click the forward arrow icon to test connection and provide the credentials for the folder destination provided.

b. From the list of drop-down options, select the encryption algorithm. For more information, see Encryption Settings.

c. Optionally, provide an encryption password.

d. Enter the encryption password again to confirm.

e. Select a type of compression. For more information, see Compression Type.

Note: If you store the data to a local disk or shared folder, you cannot replicate the data to another recovery point server. Replication is supported only if you store the data to a recovery point server.

The destination is specified.

Specify the Schedule

The Schedule page lets you define a schedule for Backup, Merge, and Throttle functions to repeat at specific intervals. After you define a schedule, the jobs run automatically per the schedule. You can add multiple schedules and provide retention settings.

A Backup Schedule refers to regular schedule that is repeated multiple times daily based on the number of hours or minutes you select. Besides the regular schedule, a backup schedule also provides options to add daily, weekly, and monthly schedules.

Note: For more information on scheduling and retention settings, see Understanding Advanced Scheduling and Retention.

Follow these steps:

1. (Optional) Select the option to manage recovery points. This option is visible only if you have selected Local or shared folder as your backup destination.

   Retain by Recovery Points

   The backup data is stored as recovery points.

   Retain by Recovery Sets

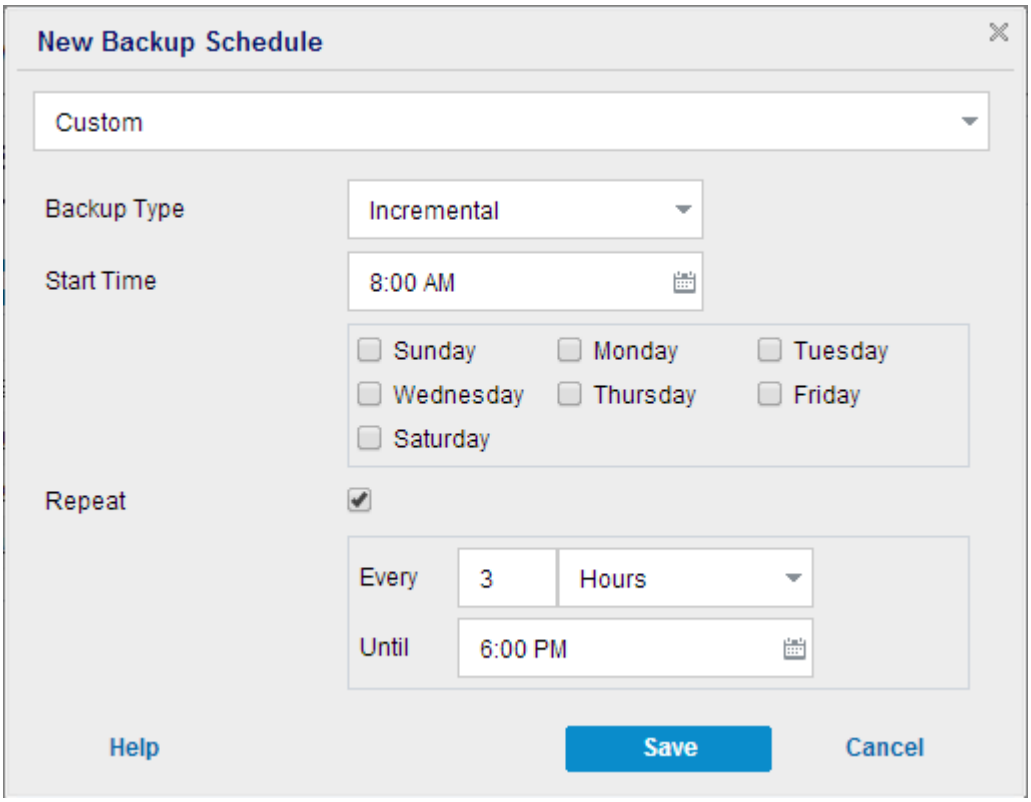
   The backup data is stored as recovery sets.

2. Add backup, merge, and throttle schedules.

   Add Backup Schedule

   a. Click **Add** and select **Add Backup Schedule**.

      The **New Backup Schedule** dialog opens.



b. Select one of the following options:

Custom

Specifies the backup schedule that repeats multiple times a day.

Daily

Specifies the backup schedule that occurs once a day. By default, all the days of the week are selected for Daily backup. If you do not want to run the backup job on a specific day, clear the check box for that day of the week.

Weekly

Specifies the backup schedule that occurs once a week.

Monthly

Specifies the backup schedule that occurs once a month.

c. Select the backup type.

Full

Determines the backup schedule for Full Backups. As scheduled, Arcserve UDP performs a full backup of all used blocks from the source machine. A full backup typically consumes time depending on the backup size.

Verify

Determines the backup schedule for Verify Backups.

As scheduled, Arcserve UDP verifies that the protected data is valid and complete by performing a confidence check of the stored backup image to the original backup source. If necessary, the image is resynchronized. A Verify Backup looks at the most recent backup of each individual block and compares the content and information to the source. This comparison verifies that the latest backed up blocks represent the corresponding information at the source. If the backup image for any block does not match the source (possibly because of changes in the system since the last backup), Arcserve UDP refreshes (resynchronizes) the backup of the block that does not match. You can also use a Verify Backup (very infrequently) to get the guarantee of full backup without using the space required for a full backup.

Advantages: Produces a small backup image when compared to full backup because only the changed blocks (blocks that do not match the last backup) are backed up.

Disadvantages: Backup time is long because all source blocks are compared with the blocks of the last backup.

Incremental

Determines the backup schedule for Incremental Backups.

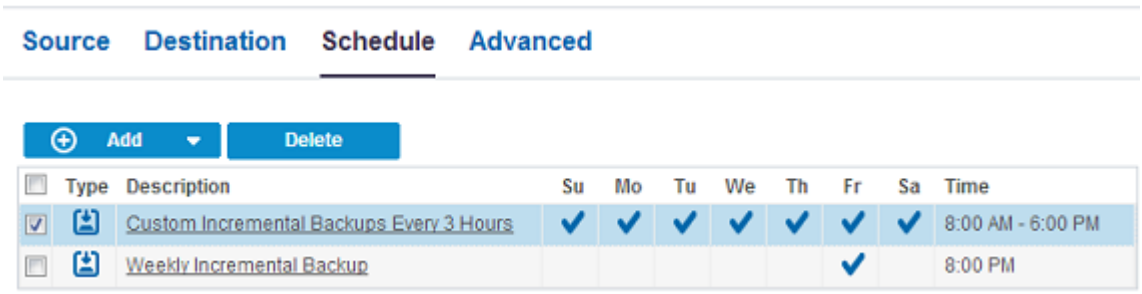
As scheduled, Arcserve UDP incrementally backs up only those blocks that have changed since the last successful backup. The advantages of Incremental Backups are that it is a fast backup and it produces a small backup image. This is the most optimal way to perform backups and you should use this by default.

d. Specify the backup start time.

e. (Optional) Select the **Repeat** check box and specify the repeat schedule.

f. Click **Save**.

The Backup Schedule is specified and it is displayed on the **Schedule** page.



Add Merge Schedule

a. Click **Add** and select **Add Merge Schedule**.

b. The **Add New Merge Schedule** dialog opens.

c. Specify the start time to start the merge job.

d. Specify **Until** to specify an end time for the merge job.

e. Click **Save**.

The Merge Schedule is specified and it is displayed on the **Schedule** page.

Add Throttle Schedule

a. Click **Add** and select **Add Throttle Schedule**.

b. The Add New Throttle Schedule dialog opens.

c. Specify the throughput limit in MB per minute unit.

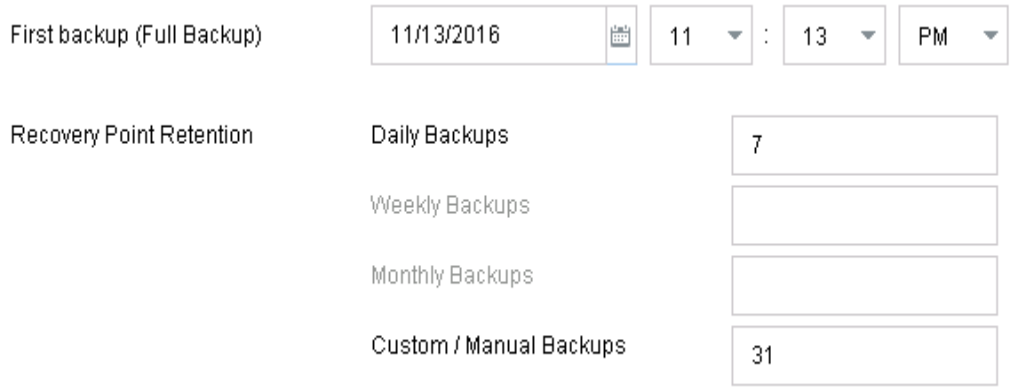
d. Specify the start time to start the backup throughput job.

e. Specify **Until** to specify an end time for the throughput job.

f. Click **Save**.

The Throughput Schedule is specified and it is displayed on the **Schedule** page.

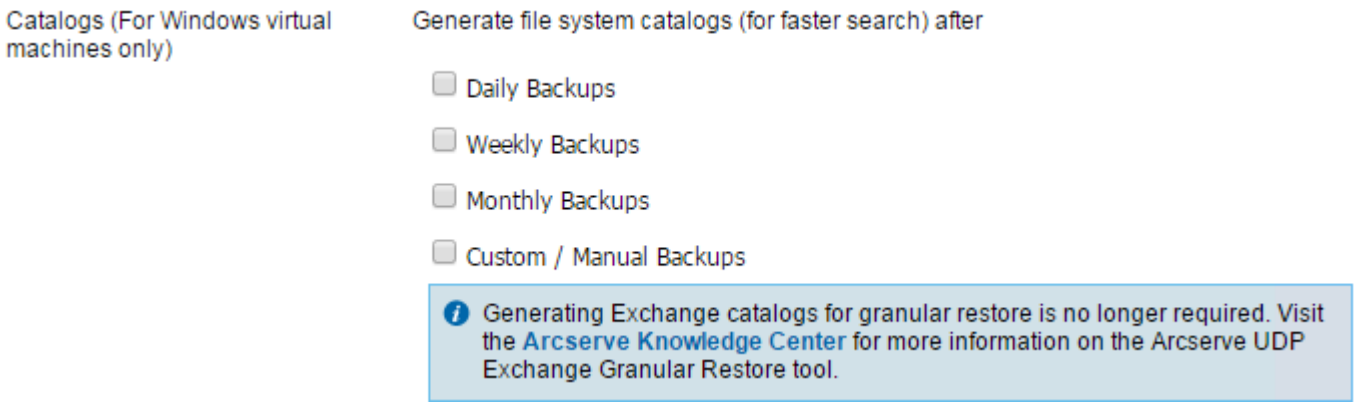
3. Specify the start time for the scheduled backup.



4. Specify the recovery points retention settings for Custom, Daily, Weekly, and Monthly schedule.

   These options are enabled if you have added the corresponding backup schedule. If you modify the retention settings on this page, the changes are reflected on the Backup Schedule dialog.

5. Specify the catalog details.



Catalogs let you generate the file system catalog. The File System catalog is required to perform faster and easier search. The catalogs are enabled depending on the type of backup that you have specified.

The schedule is specified.

Understanding Advanced Scheduling and Retention

The scheduling option lets you specify a Custom schedule or a Daily / Weekly / Monthly schedule, or both the schedules. In the Custom schedule, you can configure the backup schedule for each day of the week and you can add up to four backup schedules each day. You can select a specific day of a week and create a time window to define when to run backup and at what frequency.

Schedule	Supported Job	Comments
Backup	Backup job	Define time windows to run backup jobs.
Backup throttling	Backup job	Define time windows to control the backup speed.
Merge	Merge job	Define when to run merge jobs.
Daily schedule	Backup job	Define when to run daily backup jobs.
Weekly schedule	Backup job	Define when to run weekly backup jobs.
Monthly schedule	Backup job	Define when to run monthly backup jobs.

You can also specify the retention settings for the recovery points.

Note: Set the retention settings within each plan to control how data for the nodes assigned to that plan are retained at the target data store.

Schedules for Daily / Weekly / Monthly backups are independent to the Custom schedule, and each other. You can configure only to run Daily backup, Weekly backup or Monthly backup, without configuring the Custom schedule.

Backup Job Schedule

You can add four time windows per day in your backup schedule. A valid time window is from 12:00 AM until 11:59 PM. You cannot specify a time window such as 6:00 PM to 6:00 AM. In such cases, you must manually specify two different time windows.

For each time window, the start time is inclusive, and the end time is exclusive. For example, you have configured to run Incremental Backup every one hour between 6:00 AM and 9:00 AM and the backup will start at 6:00 AM. This means the backup runs at 6:00 AM, 7:00 AM, 8:00 AM, but NOT at 9:00 AM.

Note: If you want to run the backup job repeatedly until the end of day, set the schedule until 12:00 AM. For example, to run backup job every 15 minutes for the entire day, set the schedule from 12:00 AM to 12:00 AM, every 15 minutes.

Backup Throttle Schedule

Backup throttle schedule lets you control the backup throughput speed that in turn controls the resource usage (disk I/O, CPU, network bandwidth) of the server being backed up. This is useful if you do not want to affect the server

performance during business hours. You can add four time windows per day in your backup throttle schedule. For each time window, you can specify a value, in MB per minute. This value is used to control the backup throughput. Valid values are from 1 MB/minutes to 99999 MB/minutes.

If a backup job extends its specified time, then the throttle limit adjusts according to the specified time window. For example, you have defined the backup throttle limit as 500 MB/minute from 8:00 AM to 8:00 PM, and 2500 MB/minute from 8:00 PM to 10:00 PM. If a backup job starts at 7:00 PM and it runs for three hours, then, from 7:00 PM to 8:00 PM the throttle limit is 500 MB/minute and from 8:00 PM to 10:00 PM the throttle limit is 2500 MB/minute.

If you do not define any backup schedule and backup throughput schedule, the backup runs as fast as it can.

Merge Schedule

Lets you merge recovery points based on the provided schedule.

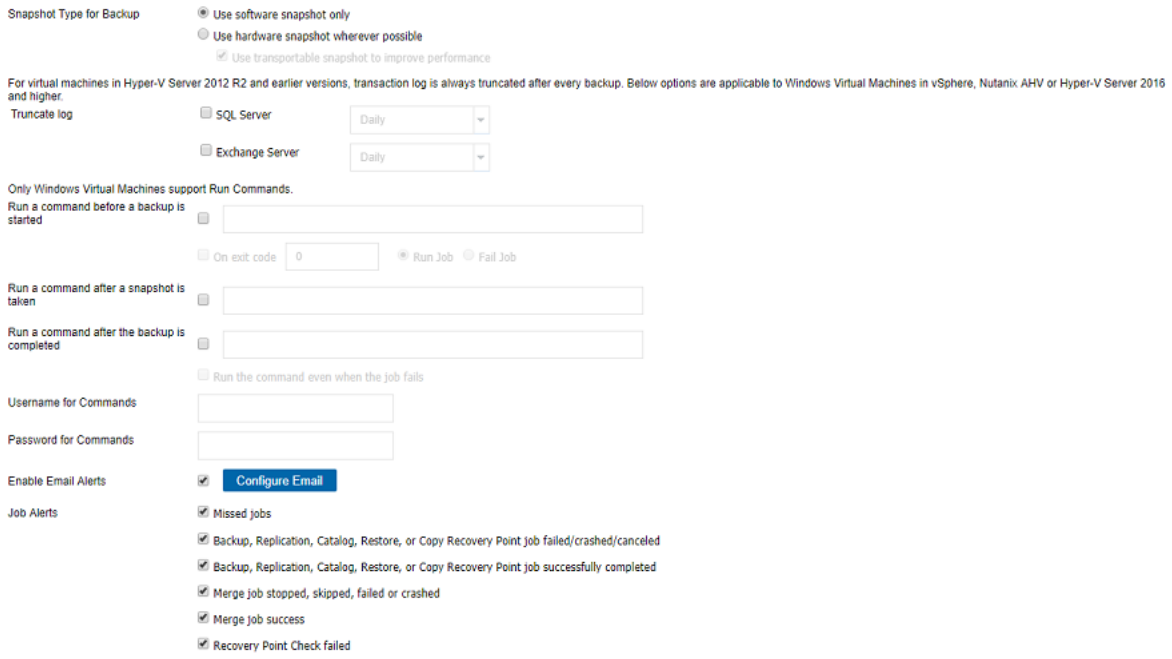
Consider the following points for the merge job:

- At any given time only one merge job can run for a node.
- If a merge job starts, it has to complete before the next merge job can start. This means, if one or more sets of recovery points are merged, new recovery point cannot be added to this merge process, until the merge process of the current set of recovery point completes.
- If a merge job is processing more than one set of recovery points (for example set [1~4], set [5~11], and set [12~14]; they are three sets), recovery point server processes these sets one by one.
- If a merge job is resumed after a pause, the job detects at which point it is paused and resumes the merge from the break-point.

Specify the Advanced Settings

The **Advanced** tab lets you specify some advanced settings for the backup job. The advanced settings include providing snapshot type for backup, truncate log settings, providing the location of any scripts, and email settings. Review the prerequisites before you select the hardware snapshot type.

The following image displays the Advanced tab:



Follow these steps:

1. Specify the following details.

   Snapshot Type for Backup

   Select one of the following options for the backup snapshot.

   Use software snapshot only

   Specifies that the backup type uses only the software snapshot. Arcserve UDP will not check for hardware snapshot. The software snapshot utilizes less resources on the virtual machines. You can use this option if the server has lower configurations and processing speed.

   Use hardware snapshot wherever possible

   Specifies that the backup type first checks for a hardware snapshot. If all the criteria are met, the backup type uses hardware snapshot.

Note: For more information on the hardware snapshot criteria, see the pre-requisite.

For Hyper-V, Arcserve UDP uses the hardware provider to take a VSS snapshot of the volumes on the Hyper-V host and import the hardware snapshot to the proxy server. The proxy server must have a suitable hardware provider installed. For VMware, Arcserve UDP creates vSphere software snapshot for a brief period and then creates a hardware snapshot. This hardware snapshot is mounted on the VMware ESX server and the software snapshot is deleted. Arcserve UDP then uses the contents in the hardware snapshot to back up the VM related files.

Use transportable snapshots to improve performance

Specifies that the hardware snapshot uses a transportable snapshot. A transportable snapshot increases the backup throughput. This option is applicable for Hyper-V servers only.

Truncate Log

Let you specify the schedule to truncate logs for SQL Server and Exchange Server. You can specify the schedule as **Daily**, **Weekly**, or **Monthly**. This is applicable only for VMware.

Run a command before a backup is started

Let you run a script before the backup job starts. Specify the path where the script is stored inside the guest OS of the virtual machine. Click **On exit code** and specify the exit code for **Run Job** or **Fail Job**. **Run Job** indicates that the backup job will continue when the script returns the exit code. **Fail Job** indicates that the backup job will stop when the script returns the exit code. This is applicable only for Windows VM.

Notes: (also applicable for the After snapshot is taken and After backup is completed commands)

- We suggest to specify the full path of the command/script. For example, use C:\Windows\System32\Ping.exe, instead of just Ping.exe.
- To avoid the situation that backup job gets stuck because command/script hangs, by default the command/script will be terminated if it cannot finish in 3 minutes. If you want to change the default timeout setting, follow these steps at proxy server or VM level:

  At Proxy Server level (applicable for all backup jobs running in this proxy server)

a. Open the registry key from the following location:

   [HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll]

b. Add a DWORD value with name PrePostCMDTimeoutInMinute and specify its value with timeout time in minute.

At VM level

a. Open the registry key from the following location:

   [HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll\<vm instance uuid>]

b. Add a DWORD value with name PrePostCMDTimeoutInMinute and specify its value with timeout time in minute.

   Note: If you add the registry value in both the VM and proxy level registry, then the setting in the VM level registry will have the priority over the setting in the Proxy level registry.

Run a command after a snapshot is taken

Lets you run a script after the backup snapshot is taken. Specify the path where the script is stored inside the guest OS of the virtual machine. This is applicable only for Windows VM.

Run a command after a backup is completed

Lets you run a script after the backup job is completed. Specify the path where the script is stored inside the guest OS of the virtual machine. This is applicable only for Windows VM.

Run a command even when the job fails

If this check box is selected, the script specified in "Run a command after a backup is completed" is executed even when the backup job fails. Otherwise, that script is executed only when backup job completes successfully.

Username for Commands

Lets you specify the username to run the commands.

Password for Commands

Let you specify the password to run the commands.

Enable Email Alerts

Lets you enable email alerts. You can configure email settings and specify the types of alerts that you want to receive in an email. When you select this option, the following options are enabled for your selection.

Email Settings

Lets you configure the email settings. Click **Email Settings** and configure the email server and proxy server details.

Job Alerts

Lets you select the types of job alert emails that you want to receive.

2. Click **Save**.

The updates are saved and a green check mark is displayed next to the task name. The plan page closes.

Note:

- If you have to add another task, select the plan from the resources tab and modify the plan. To modify the plan, click the plan from the center pane. The plan opens and you can modify it.
- When you select a node that does not have Arcserve UDP agent installed, as the backup proxy, UDP console automatically deploys an agent into that node after the plan is saved.

The host-based agentless backup plan for the virtual machine is created. The backup runs per the schedule that you have configured in the **Schedule** tab. You can also perform a manual backup at any time.

Run Script Command and Log Truncation with Additional Administrator Account

Additional administrator account refers to those accounts that are not default administrators. The following two accounts are involved when you run the commands or scripts:

1. Account set by Update Node
2. Account set on the Advanced tab of a Plan

VMware and Hyper-V virtual machines have separate conditions to use the additional administrator accounts.

VMware Virtual Machines

If both accounts are set, use the first account to log in to the virtual machine (vSphere SDK is used to communicate with virtual machine, so that the network access is not required between proxy server and virtual machine). Then use the second account to run the command or script in the virtual machine.

If either of the account is not set, use the available account to log in to the virtual machine and run the command or script.

It is recommended to use the built-in administrator account or built-in domain administrator account for both accounts.

If you use any additional administrator account (non-built-in administrator account), the procedure is different.

Follow these steps:

1. To log in to the virtual machine using the added administrator account, follow the step in the Update Node topic to ensure that the account has the required permissions.
2. To run the command or script using the additional administrator account, ensure that this account has the required permission. Log in to the guest virtual machine using the additional administrator account, run the command or script, and confirm that the command or script can complete successfully.

Hyper-V Virtual Machines

You need only one account for Hyper-V virtual machines. If both accounts are set, use the second account (set on the Advanced tab of a plan) to connect to the virtual machine and launch the command or script. Windows Management Instrumentation (WMI) is used to communicate with the virtual machine, so that the network access is required between proxy server and virtual machine.

If either of the account is not set, use the available administrator account to connect to the virtual machine and launch the command or script.

Follow these steps:

1. Access the virtual machine with remote WMI. Ensure that you have the required permissions with the additional administrator account. See the Update Node topic for the requirements of the account.
2. To run the command or script using the additional administrator account, ensure that this account has the required permission. Log in to the guest virtual machine using the additional administrator account, run the command or script, and confirm that the command or script can complete successfully.

   Also, verify if WMI is allowed by firewall on the guest VM. If not enabled, follow these steps:

1. Log into the guest VM.
2. Open the Control panel.
3. Open the Windows Firewall.
4. Click Allow an app or feature through Windows Firewall.
5. Enable Windows Management Instrumentation(WMI).
6. Click OK.

Define a Limit to the Quantity of Concurrent Backups

You can define a limit to the quantity of backup jobs that run concurrently. This capability lets you optimize the performance of the host-based VM backup proxy server in your backup environment. By default, Host-Based VM Backup can run up to four VMware VM backup jobs, ten Hyper-V VM backup jobs, and four Nutanix AHV VM backup jobs concurrently. In environments that contain many virtual machines that are associated with a proxy server, a high quantity of concurrent backups can have an adverse effect on network and backup performance.

Note: When the quantity of concurrent jobs exceeds the defined limit, the jobs that exceed the limit enter a job queue.

Note: If the maximum number of concurrent VMware backup jobs exceeds the ESX server connection limit, communication failure can occur between the ESX server and the backup proxy, and the file system of the ESX server data store can remain locked. In such cases, restart the ESX server or migrate the locked virtual machine to another data store to unlock the VM. For more details, refer to the VMware document (VMware KB:1022543).

Follow these steps:

1. Log in to the Arcserve UDP virtual machine proxy system.
2. Open Windows Registry Editor and browse to the following key:

   ```
   HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine
   ```

3. Locate the following keys:

   ```
   VMwareMaxJobNum
   ```

   ```
   HyperVMaxJobNum
   ```

   ```
   AHVMaxJobNum
   ```

   Note: Both the keys are already created and the default value is 4 and 10, respectively.

4. Right-click VMMaxJobNum, HyperVMaxJobNum or AHVMaxJobNum, and then click Modify on the pop-up menu.

   The Edit String dialog opens.

5. In the Value Data field, specify the quantity of backup jobs that you want to allow to run concurrently.

   - **Minimum limit**--1
   - **Maximum limit**--none

- **Default**--10 for Hyper-V and 4 for VMware and Nutanix AHV

6. Click OK.

   The limit is defined.

   The limit of concurrent backup jobs is defined.

(Optional) Perform a Manual Backup

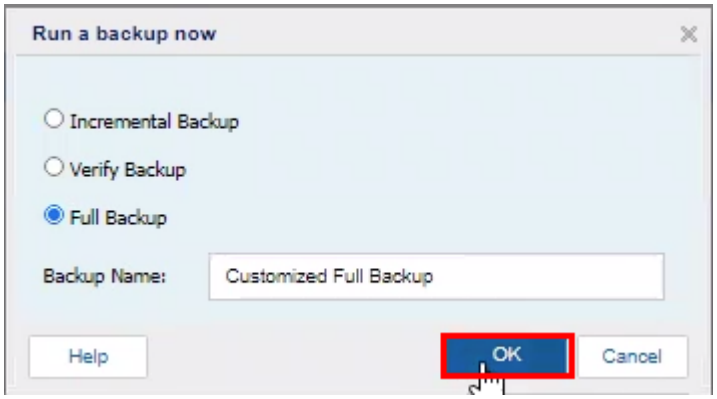
Typically, backups are performed automatically and are controlled by the schedule settings. In addition to the scheduled backup, a manual backup provides you the option to back up your nodes on a need basis. For example, if you have a repeat schedule for Full, Incremental, and Verify backups and you want to make major changes to your machine, you must perform an immediate manual backup without waiting for the next scheduled backup to occur.

Follow these steps:

1. Log into Arcserve UDP Console.
2. Navigate to **resources** > **Plan** > **All Plans**.

   The list of all plans is displayed on the center pane.
3. Right-click the plan, and then select **Backup Now**.

   The Run a backup now dialog opens.
4. On the Run a backup now dialog, select the type of backup you want to perform, and then specify the backup name if required.



The available options are:

Incremental Backup

Initiates an Incremental Backup of your machine. An Incremental Backup backs up only those blocks that have changed after the previous backup.

Note: If you are performing a backup for the first time, the full backup gets initiated automatically even though you selected the Incremental Backup option.

The advantages of Incremental Backups are that it is a fast backup and produces a small backup image. This is the most optimal way to perform backups.

Verify Backup

Initiates a Verify Backup of your machine by examining the most recent backup of each individual block and comparing the content and information to the original source. This comparison verifies that the latest backed up blocks represent the corresponding information at the source. If the backup image for any block does not match the source, Arcserve UDP Agent (Windows) refreshes (resynchronizes) the backup of the block that does not match.

Advantage: A very small backup image is produced when compared to a Full Backup because only the changed blocks (blocks that do not match the last backup) are backed up.

Disadvantage: The backup time is slow because all the source disk blocks are compared with the blocks of the last backup.

Full Backup

Initiates a Full Backup of your entire machine or the selected volumes.

Notes:

- If you add a new volume to the backup source, the newly added volume is fully backed up regardless of the overall backup method selected.
- If no backup name is specified, by default, it is automatically named as Customized Full/Incremental/Verify Backup.

5. Click **OK**.

   The Progress Information dialog appears.

6. Wait until the backup job completes, and then click **OK**.

7. To monitor the backup jobs, navigate to **jobs** > **Jobs in Progress**.

   Wait until the backup job is finished.

   Note: Only one job can be run at a time. If you manually attempt to launch a backup job when another job is currently running, an alert message informs you that another job is running and requests that you try again at a later time.

   The manual backup is successfully performed.

   Note: You can also trigger manual backup. Right-click on a plan and select **Backup Now** from the options. In this case, if the plan protects container objects of vSphere, UDP will trigger backup for all VMs available under that container.

Verify the Plan

To verify your backup, confirm that you have successfully created the backup plan. After you verify that the plan is created successfully, confirm whether the backup job is running as scheduled. You can verify the status of backup jobs from the jobs tab.

Follow these steps: to verify plans

1. Click the Resources tab.
2. From the left pane, navigate to Nodes, and click All Nodes.

   A list of all nodes is displayed on the center pane.
3. Verify that plans are mapped with nodes.

Follow these steps: to verify backup jobs

1. Click the jobs tab.
2. From the left pane, click All Jobs.

   The status of each job is listed on the center pane.
3. Verify that the backup job is successful.

   The backup job is verified.

How to Create a Virtual Standby to AWS EC2 Plan

The virtual standby converts the recovery points to virtual machine formats on specified cloud and prepares a snapshot to easily recover your data when needed. This feature provides the high availability capability also and ensures that the virtual machine can take over immediately when the source machine fails. The standby virtual machine is created by converting the recovery points to Amazon AWS EC2 virtual machine format.

Note: The virtual standby task runs only if the backup task creates a valid recovery point snapshot. If the backup task fails, then the virtual standby task is skipped.

What To Do Next?

1. Review the Prerequisites and Considerations
2. Create a Plan with a Backup Task
3. Add a Virtual Standby to EC2 Task to the Plan
4. (Optional) Run the Virtual Standby to EC2 Task Manually
5. Pause and Resume the Virtual Standby Job
6. Verify the Plan
7. Terminate EC2 Resources

Review the Prerequisites and Considerations

Verify if you have completed the following prerequisites:

- Log into the Console.
- Install the server component and create Data Stores if you want to store the backup data to recovery point servers.
- You have a valid recovery point to create a virtual standby machine. The recovery points can be from one of the following tasks:
  - Backup, Agent-based Windows
  - Backup, Host-Based Agentless
  - Replicate
  - Replicate from a remote Recovery Point Server
- Back up the full machine to enable the Virtual Standby task. You cannot create a Virtual Standby task if the backup is not a full backup.
- Install the system volume on the first disk. Verify that the system volume and boot volume of the source machine are on the same disk.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.
- Configure the Security group setting on EC2 to open the related ports for inbound access, including TCP 8014 and TCP 4091.
- Verify if the Amazon AWS account can access AWS S3 and AWS EC2. Arcserve UDP does not provide the account.
- You must install the system volume on the first disk. Verify that the system volume and boot volume of the source machine are on the same disk.
- Specific AWS API permissions are required for Amazon IAM users to achieve control and interaction with AWS APIs for VSB to EC2. For details, refer to How to configure IAM granular permissions for IAM users with VSB to EC2.

Considerations:

- Virtual Standby to EC2 task does NOT support the source machine boot from UEFI firmware due to Amazon AWS EC2 limitation.
- .NET framework 4.5 is needed on the source machine to install Amazon PV driver for Virtual Standby to EC2 task.
- The system and boot volume on the source node in VSB to EC2 plan must be at the first disk due to AWS limitation.

- The customization to TCP/IP setting cannot apply to the first network interface on EC2 instance due to AWS limitation.
- The elastic IP address assignment is only available for the first network interface on EC2 instance.
- The customization to TCP/IP setting is only applicable to the scenario when Direct Access or VPN is configured between primary site and the AWS network. Otherwise, the customization makes the EC2 instance inaccessible.
- You cannot modify Enable/Disable auto assign public IP property for the existed nodes in the plan and the update only affects the new nodes added into the plan.
- You must install PowerShell version 3.0 or above on the source machine before configuring the VSB to EC2 for the following instance types:

  C5, C5d, C5n, F1, G3, G4, H1, I3, I3en, Inf1, m4.16xlarge, M5, M5a, M5ad, M5d, M5dn, M5n, P2, P3, R4, R5, R5a, R5ad, R5d, R5dn, R5n, T3, T3a, X1, X1e, and z1d

Configure IAM granular permissions for IAM users with VSB to EC2

This section explains the steps and API permission policy required for the Arcserve UDP agent installed on the VSB Cloud proxy within Amazon EC2 Web Services. The permissions help perform the actions required for data transfer and Virtual Standby to the AWS EC2 Cloud.

Using the procedure, you can help an Amazon IAM user get control and interaction with the AWS API. The permissions policy is not only applied to the user directly, but also to a Role and Group within the IAM Security interface within Amazon Web Services.

Follow these steps:

1. Log into Amazon Web Services as an administrator.
2. Select My Security Credentials, click **Users** on the left side, and then click the **Create New Users** button.
3. Enter desired user name.

   Note: Verify if the option for **Generate an access key for each User** is selected.
4. Click the **Create** button.
5. Click **Download Credentials**.

   The credentials contain your Access key and Secret that you need later within the UDP console.
6. From the Users view, select the user from the list of users and then click the Permissions tab available at the bottom.
7. From the Custom Policy option, create an inline custom policy for the user.
8. Enter a name for the policy and paste the following content in Policy Document.

```
{
"Version": "2012-10-17",
"Statement": [
{
    "Sid": "Stmt1477881304097",
    "Action": [
      "ec2:AssignPrivateIpAddresses",
      "ec2:AssociateAddress",
```

"ec2:AttachNetworkInterface",

"ec2:AttachVolume",

"ec2:AuthorizeSecurityGroupEgress",

"ec2:AuthorizeSecurityGroupIngress",

"ec2:CreateNetworkInterface",

"ec2:CreateSnapshot",

"ec2:CreateTags",

"ec2:CreateVolume",

"ec2:DeleteNetworkInterface",

"ec2:DeleteSnapshot",

"ec2:DeleteTags",

"ec2:DeleteVolume",

"ec2:DescribeAccountAttributes",

"ec2:DescribeAddresses",

"ec2:DescribeAvailabilityZones",

"ec2:DescribeBundleTasks",

"ec2:DescribeClassicLinkInstances",

"ec2:DescribeConversionTasks",

"ec2:DescribeCustomerGateways",

"ec2:DescribeDhcpOptions",

"ec2:DescribeExportTasks",

"ec2:DescribeFlowLogs",

"ec2:DescribeHosts",

"ec2:DescribeHostReservations",

"ec2:DescribeHostReservationOfferings",

"ec2:DescribeIdentityIdFormat",

"ec2:DescribeIdFormat",

"ec2:DescribeImageAttribute",

"ec2:DescribeImages",

"ec2:DescribeImportImageTasks",

"ec2:DescribeImportSnapshotTasks",

"ec2:DescribeInstanceAttribute",

"ec2:DescribeInstanceStatus",

"ec2:DescribeInstances",

"ec2:DescribeInternetGateways",

"ec2:DescribeKeyPairs",

"ec2:DescribeMovingAddresses",

"ec2:DescribeNatGateways",

"ec2:DescribeNetworkAcls",

"ec2:DescribeNetworkInterfaceAttribute",

"ec2:DescribeNetworkInterfaces",

"ec2:DescribePlacementGroups",

"ec2:DescribePrefixLists",

"ec2:DescribeRegions",

"ec2:DescribeReservedInstances",

"ec2:DescribeReservedInstancesListings",

"ec2:DescribeReservedInstancesModifications",

"ec2:DescribeReservedInstancesOfferings",

"ec2:DescribeRouteTables",

"ec2:DescribeSecurityGroups",

"ec2:DescribeSnapshotAttribute",

"ec2:DescribeSnapshots",

"ec2:DescribeSpotDatafeedSubscription",

"ec2:DescribeSpotFleetInstances",

"ec2:DescribeSpotFleetRequestHistory",

"ec2:DescribeSpotFleetRequests",

"ec2:DescribeSpotInstanceRequests",

"ec2:DescribeSpotPriceHistory",

"ec2:DescribeStaleSecurityGroups",

"ec2:DescribeSubnets",

"ec2:DescribeTags",

"ec2:DescribeVolumeAttribute",

"ec2:DescribeVolumeStatus",

"ec2:DescribeVolumes",

"ec2:DescribeVpcAttribute",

"ec2:DescribeVpcClassicLink",

"ec2:DescribeVpcEndpointServices",

"ec2:DescribeVpcEndpoints",

"ec2:DescribeVpcPeeringConnections",

"ec2:DescribeVpcs",

"ec2:DescribeVpnConnections",

"ec2:DescribeVpnGateways",

"ec2:DetachClassicLinkVpc",

"ec2:DetachInternetGateway",

"ec2:DetachNetworkInterface",

"ec2:DetachVolume",

"ec2:DetachVpnGateway",

"ec2:DisableVgwRoutePropagation",

"ec2:DisableVpcClassicLink",

"ec2:DisableVpcClassicLinkDnsSupport",

"ec2:DescribeVpcClassicLinkDnsSupport",

"ec2:DetachNetworkInterface",

"ec2:DetachVolume",

"ec2:DisassociateAddress",

"ec2:ModifyInstanceAttribute",

"ec2:ModifyNetworkInterfaceAttribute",

"ec2:ModifySnapshotAttribute",

"ec2:ModifySubnetAttribute",

"ec2:ModifyVolumeAttribute",

"ec2:RevokeSecurityGroupEgress",

"ec2:RevokeSecurityGroupIngress",

"ec2:RunInstances",

"ec2:StartInstances",

```
            "ec2:StopInstances",
            "ec2:TerminateInstances"
        ],
        "Effect": "Allow",
        "Resource": [
        "*"
        ]
    },
    {
        "Sid": "Stmt1477880716900",
        "Action": [
            "s3:CreateBucket",
            "s3:DeleteBucket",
            "s3:DeleteObject",
            "s3:GetObject",
            "s3:ListBucket",
            "s3:PutObject"
        ],
        "Effect": "Allow",
        "Resource": [
            "*"
        ]
    },
    {
        "Sid": "Stmt1477883239716",
        "Action": [
            "iam:GetUser"
        ],
        "Effect": "Allow",
        "Resource": [
            "*"
```

```
        ]
    }
    ]
    }
```

9. Click **Apply Policy**.
10. In UDP console, use this IAM user's access key and security access key to create VSB to EC2 plan.

Create a Plan with a Backup Task

A plan includes different types of tasks that you want to perform. To create a virtual standby machine, you create a plan that includes a backup task and a virtual standby task. A backup task performs a backup of the source nodes and stores the data to the specified destination. This backup data is then used by the virtual standby feature and converts it to a virtual machine format.

You can create a virtual standby machine from an agent-based Windows backup, host-based agentless backup. You can also create a virtual standby machine from data that are replicated using the **Replicate** task. The following procedure is an example to create agent-based Windows backup.

Notes:

- For more information on host-based agentless backup, see How to Create a Host-Based Virtual Machine Backup Plan.
- For more information on replicating a backup data, see How to Replicate Data Between Data Stores Managed from a UDP Console and Managed From Different UDP Consoles.
- Virtual standby to EC2 does NOT support the source machine boot from UEFI firmware due to Amazon AWS EC2 limitations.

Follow these steps:

1. Click the **resources** tab on the Console.
2. From the left pane, navigate to **Plans**, and click **All Plans**.

   If you have created plans earlier, those plans are displayed on the center pane.
3. On the center pane, click **Add a Plan**.

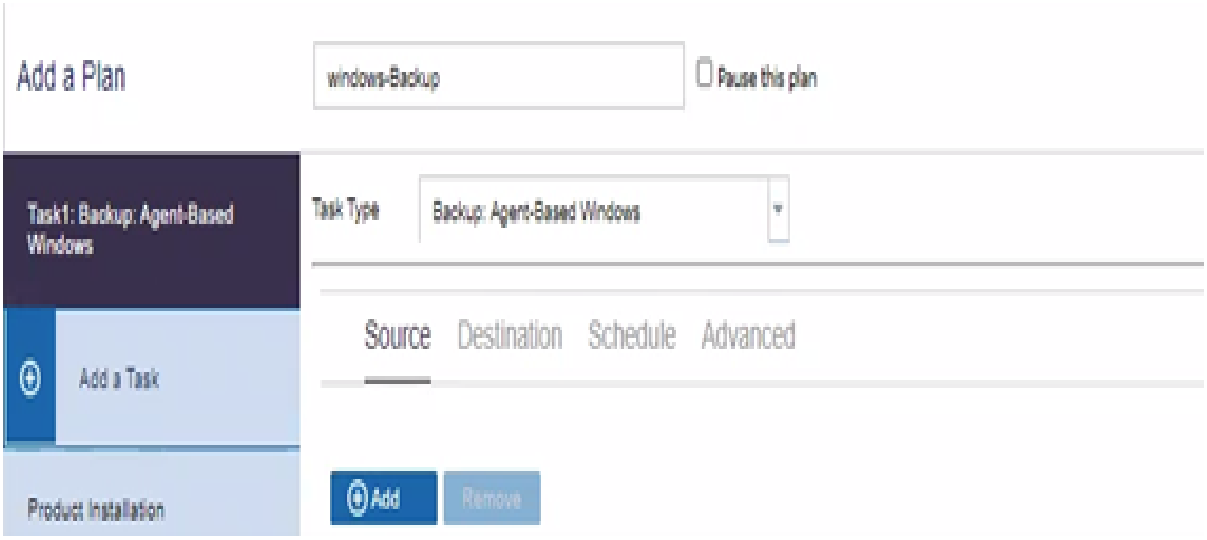
   Add a Plan opens.
4. Enter a plan name.
5. (Optional) Select **Pause this plan** check box.

   The plan will not run until you clear the check box to resume the plan.

   Note: If a plan is paused, then any in-progress job is not paused but all corresponding scheduled jobs associated with that plan are paused. However, you can manually run a job. For example, you can manually run backup job and replication job for a node even if the respective plan is paused. In such case, the following task to the on-demand (manual) job does not run. For example, there is a replication task after an on-demand backup job, the replication job does not run for the on-

demand backup job. You need to manually run the replication job. When you resume the plan, the pending jobs do not resume immediately. After you resume the plan, the pending jobs run from the next scheduled time.

6. From the **Task Type** drop-down list, select **Backup, Agent-Based Windows**.



Now, specify the Source, Destination, Schedule, and Advanced details.

Specify the Source

The Source page lets you specify the source nodes that you want to protect. You can select more than one node in a plan. If you have not added any nodes to the Console, you can add nodes when you create or modify a plan from the Source page. You can also save a plan without adding any source nodes. The plan gets deployed only after you add source nodes.

Follow these steps:

1. On the Source tab, click **Add**.
2. Select one of the following options:

   Select Nodes to Protect in Arcserve UDP

   Opens the Select Nodes to Protect in Arcserve UDP dialog. Select the nodes from the displayed list. Select this option if you have already added the nodes to the Console.

   Add Windows Node

   Opens the Add nodes to a plan dialog. Select this option if you have not added the nodes and you want to manually add the nodes to protect.

   Discover Nodes from Active Directory

   Opens the Add Nodes to Arcserve UDP Console dialog. Select this option if you want to discover and add nodes from the Active Directory.

3. If you selected the **Select Nodes to Protect in Arcserve UDP** option, on the Select Nodes to Protect dialog, do the following:

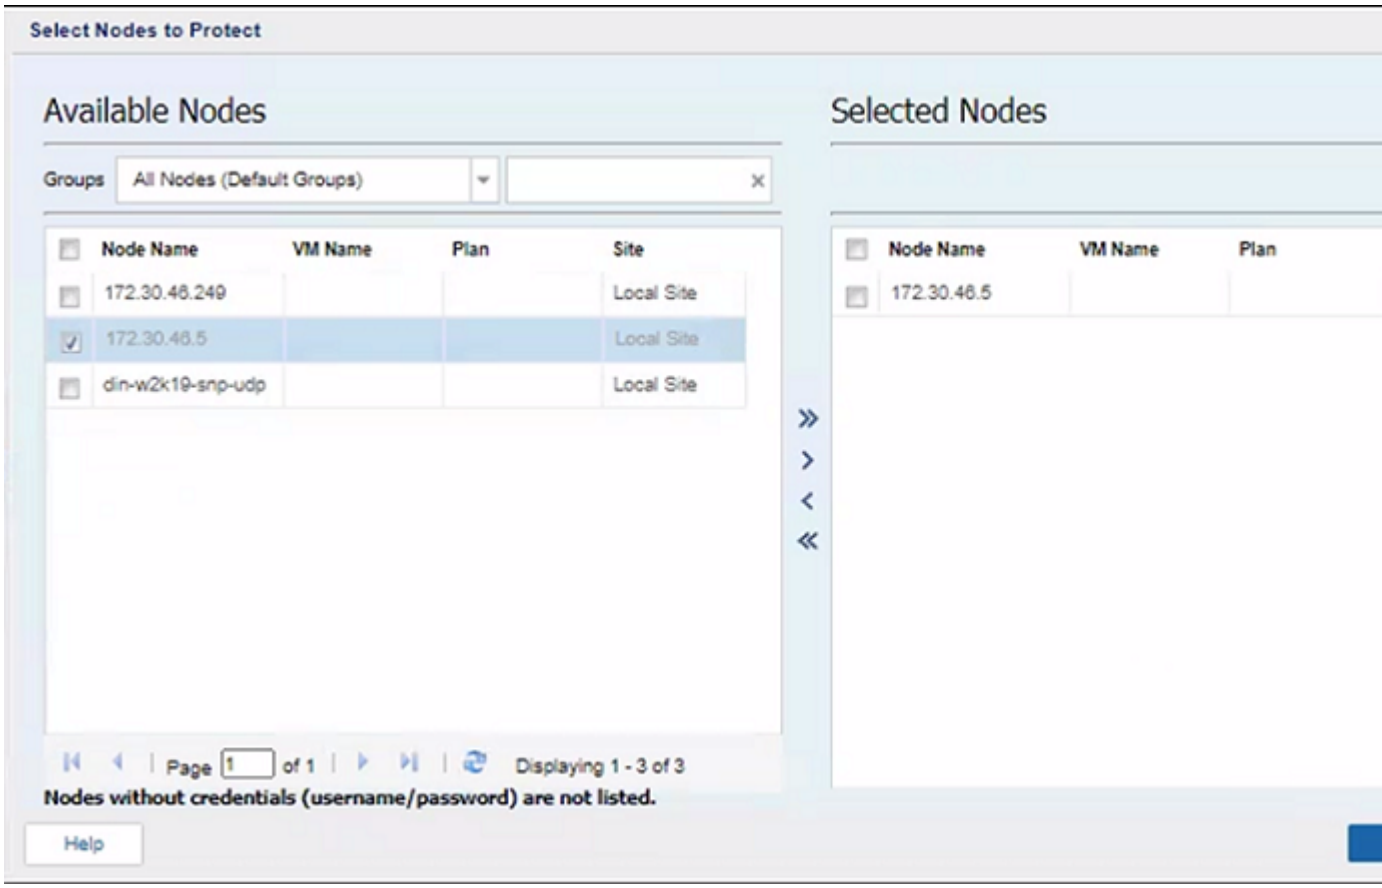
   a. (Optional) Select a filter from the Groups drop-down list to filter nodes. You can enter keywords to further filter your nodes.

      The nodes are displayed on the Available Nodes column.

   b. Select the nodes from the Available Nodes column, and then click the single right arrow (>) or double right arrow (>>) to move the selected nodes or all the nodes respectively to the Selected Nodes column.

      The nodes are displayed in the Selected Nodes column.

   c. Click **OK** to close the dialog.



4. To choose the Protection Type, select one of the following options:

   Back up all volumes

   Prepares a backup snapshot of all the volumes.

   Back up selected volumes

   Prepares a backup snapshot of the selected volume.

   The source is specified.

Specify the Destination

The destination is a location where you store the backup data. You must at least specify the destination to save the plan.

Follow these steps:

1. Select one of the following **Destination Type**:

   Local disk or shared folder

   Specifies that the backup destination is either a local destination or a shared folder. If you select this option, you can save data as either recovery points or recovery sets. The recovery points and recovery sets options are available on the **Schedule** tab.

   Arcserve UDP Recovery Point Server

   Specifies that the backup destination is a recovery point server. If you select this option, then data is stored as recovery points. You cannot store data as recovery sets.

2. If you have selected **Arcserve UDP Recovery Point Server**, then provide the following details:
   a. Select a recovery point server
   b. Select a data store. The list displays all data stores that are created at the specified recovery point server.
   c. Provide a session password.
   d. Confirm the session password.
3. If you have selected **Local disk or shared folder**, then provide the following details:
   a. Provide the full path of the local or network destination. For the network destination, specify the credentials with the write access.
   b. Select the encryption algorithm. For more information, see Encryption Settings.
   c. Optionally, provide an encryption password.
   d. Confirm the encryption password.
   e. Select a type of compression. For more information, see Compression Type.

Note: If you store the data to a local disk or shared folder, you cannot replicate the data to another recovery point server. Replication is supported only if you store the data to a recovery point server.

The destination is specified.

Specify the Schedule

The Schedule page lets you define a schedule for Backup, Merge, and Throttle functions to repeat at specific intervals. After you define a schedule, the jobs run automatically per the schedule. You can add multiple schedules and can provide retention settings.

A Backup Schedule refers to regular schedule that is repeated multiple times a day based on the number of hours or minutes you select. Besides the regular schedule, a backup schedule also provides options to add daily, weekly, and monthly schedules.

Note: For more information on scheduling and retention settings, see Understanding Advanced Scheduling and Retention.

Follow these steps:

1. (Optional) Select the option to manage recovery points. This option is visible only if you have selected Local or shared folder as your backup destination.

   Retain by Recovery Points

   The backup data is stored as recovery points.

   Retain by Recovery Sets

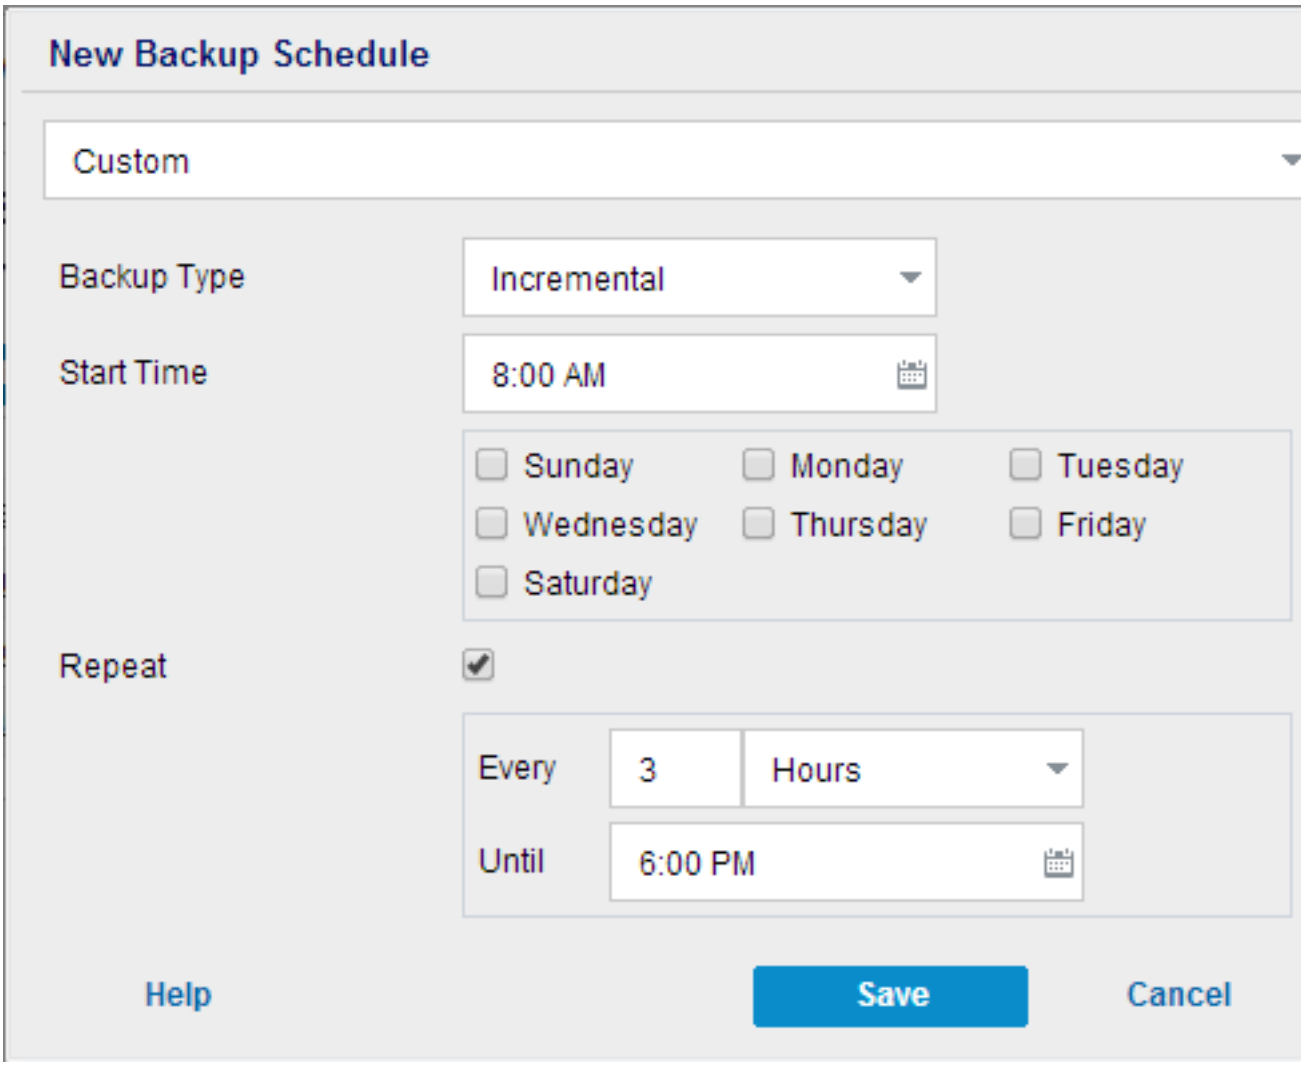
   The backup data is stored as recovery sets.

2. Add backup, merge, and throttle schedules.

   Add Backup Schedule

   a. Click **Add** and select **Add Backup Schedule**.

      The **New Backup Schedule** dialog opens.



b. Select one of the following options:

Custom

Specifies the backup schedule that repeats multiple times a day.

Daily

Specifies the backup schedule that occurs once a day. By default, all the days of the week are selected for Daily backup. If you do not want to run the backup job on a specific day, clear the check box for that day of the week.

Weekly

Specifies the backup schedule that occurs once a week.

Monthly

Specifies the backup schedule that occurs once a month.

c. Select the backup type.

Full

Determines the backup schedule for Full Backups. As scheduled, Arcserve UDP performs a full backup of all used blocks from the source machine. A full backup typically consumes time depending on the backup size.

Verify

Determines the backup schedule for Verify Backups.

Arcserve UDP verifies that the protected data is valid and complete by performing a confidence check of the stored backup image to the backup source. If necessary, the image is resynchronized. A Verify Backup looks at the most recent backup of each individual block and compares the content and information to the source. This comparison verifies that the latest backed up blocks represent the corresponding information at the source. If the backup image for any block does not match the source (possibly because of changes in the system since the last backup), Arcserve UDP refreshes (resynchronizes) the backup of the block that does not match. You can also use a Verify Backup (infrequently) to get the guarantee of full backup without using the space required for a full backup.

Advantages: Produces a small backup image when compared to full backup because only the changed blocks (blocks that do not match the last backup) are backed up.

Disadvantages: Backup time is long because all source blocks are compared with the blocks of the last backup.

Incremental

Determines the backup schedule for Incremental Backups.

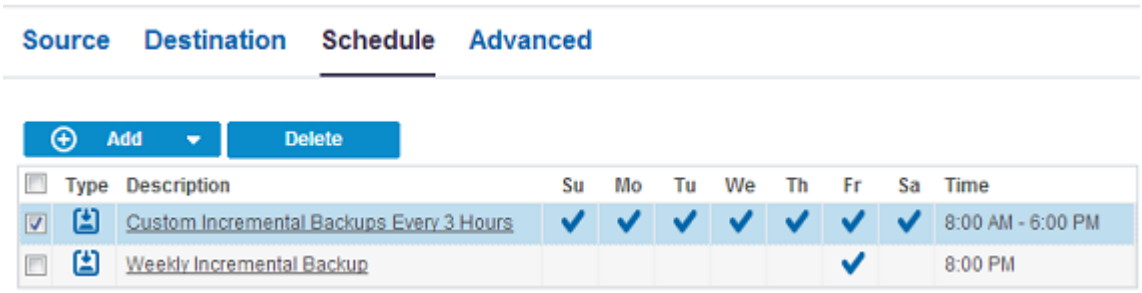
As scheduled, Arcserve UDP incrementally backs up only those blocks that have changed since the last successful backup. The advantages of Incremental Backups are that it is a fast backup and it produces a small backup image. This is the most optimal way to perform a backup.

d. Specify the backup start time.

e. (Optional) Select the **Repeat** check box and specify the repeat schedule.

f. Click **Save**.

The Backup Schedule is specified and displayed on the **Schedule** page.



Add Merge Schedule

a. Click Add and select Add Merge Schedule.

   The **Add New Merge Schedule** dialog opens.

b. Specify the start time to start the merge job.

c. Specify **Until** to specify an end time for the merge job.

d. Click **Save**.

   The Merge Schedule is specified and displayed on the **Schedule** page.

Add Throttle Schedule

a. Click **Add** and select **Add Throttle Schedule**.

   The Add New Throttle Schedule dialog opens.

b. Specify the throughput limit in MB per minute unit.

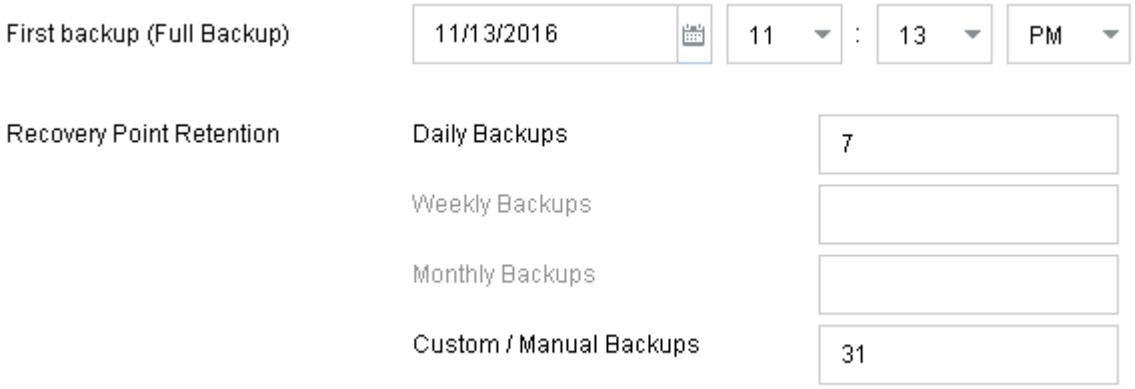
c. Specify the start time to start the backup throughput job.

d. Specify **Until** to specify an end time for the throughput job.

e. Click **Save**.

   The Throttle Schedule is specified and displayed on the **Schedule** page.

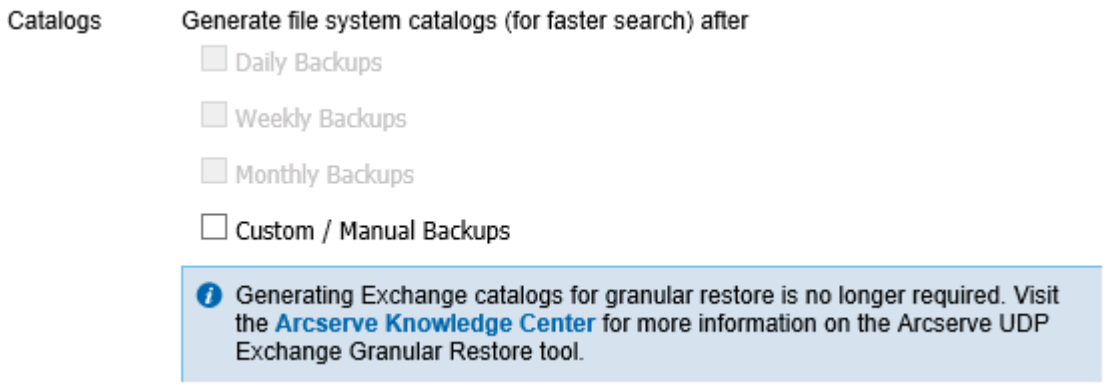
3. Specify the start time for the scheduled backup.



4. Specify the recovery points retention settings for Custom, Daily, Weekly, and Monthly schedule.

   These options are enabled if you have added the corresponding backup schedule. If you modify the retention settings on this page, the changes are reflected on the Backup Schedule dialog.

5. Specify the catalog details.



   Catalogs let you generate the File System catalog. The File System catalog is required to perform faster and better search. If you select the catalog check boxes, the catalogs are enabled depending on the type of backup that you have specified. Clear the check box to disable generating the catalog.

   The schedule is specified.

Understanding Advanced Scheduling and Retention

The scheduling option lets you specify a Custom schedule or a Daily / Weekly / Monthly schedule, or both the schedules. In the Custom schedule, you can configure the backup schedule for each day of the week and you can add up to four backup schedules each day. You can select a specific day of a week and create a time window to define when to run backup and at what frequency.

Schedule	Supported Job	Comments
Backup	Backup job	Define time windows to run backup jobs.
Backup throttling	Backup job	Define time windows to control the backup speed.
Merge	Merge job	Define when to run merge jobs.
Daily schedule	Backup job	Define when to run daily backup jobs.
Weekly schedule	Backup job	Define when to run weekly backup jobs.
Monthly schedule	Backup job	Define when to run monthly backup jobs.

You can also specify the retention settings for the recovery points.

Note: Set the retention settings within each plan to control how data for the nodes assigned to that plan are retained at the target data store.

Schedules for Daily / Weekly / Monthly backups are independent to the Custom schedule, and each other. You can configure only to run Daily backup, Weekly backup or Monthly backup, without configuring the Custom schedule.

Backup Job Schedule

You can add four time windows per day in your backup schedule. A valid time window is from 12:00 AM until 11:59 PM. You cannot specify a time window such as 6:00 PM to 6:00 AM. In such cases, you must manually specify two different time windows.

For each time window, the start time is inclusive, and the end time is exclusive. For example, you have configured to run Incremental Backup every one hour between 6:00 AM and 9:00 AM and the backup will start at 6:00 AM. This means the backup runs at 6:00 AM, 7:00 AM, 8:00 AM, but NOT at 9:00 AM.

Note: If you want to run the backup job repeatedly until the end of day, set the schedule until 12:00 AM. For example, to run backup job every 15 minutes for the entire day, set the schedule from 12:00 AM to 12:00 AM, every 15 minutes.

Backup Throttle Schedule

Backup throttle schedule lets you control the backup throughput speed that in turn controls the resource usage (disk I/O, CPU, network bandwidth) of the server being backed up. This is useful if you do not want to affect the server

performance during business hours. You can add four time windows per day in your backup throttle schedule. For each time window, you can specify a value, in MB per minute. This value is used to control the backup throughput. Valid values are from 1 MB/minutes to 99999 MB/minutes.

If a backup job extends its specified time, then the throttle limit adjusts according to the specified time window. For example, you have defined the backup throttle limit as 500 MB/minute from 8:00 AM to 8:00 PM, and 2500 MB/minute from 8:00 PM to 10:00 PM. If a backup job starts at 7:00 PM and it runs for three hours, then, from 7:00 PM to 8:00 PM the throttle limit is 500 MB/minute and from 8:00 PM to 10:00 PM the throttle limit is 2500 MB/minute.

If you do not define any backup schedule and backup throughput schedule, the backup runs as fast as it can.

Merge Schedule

Lets you merge recovery points based on the provided schedule.

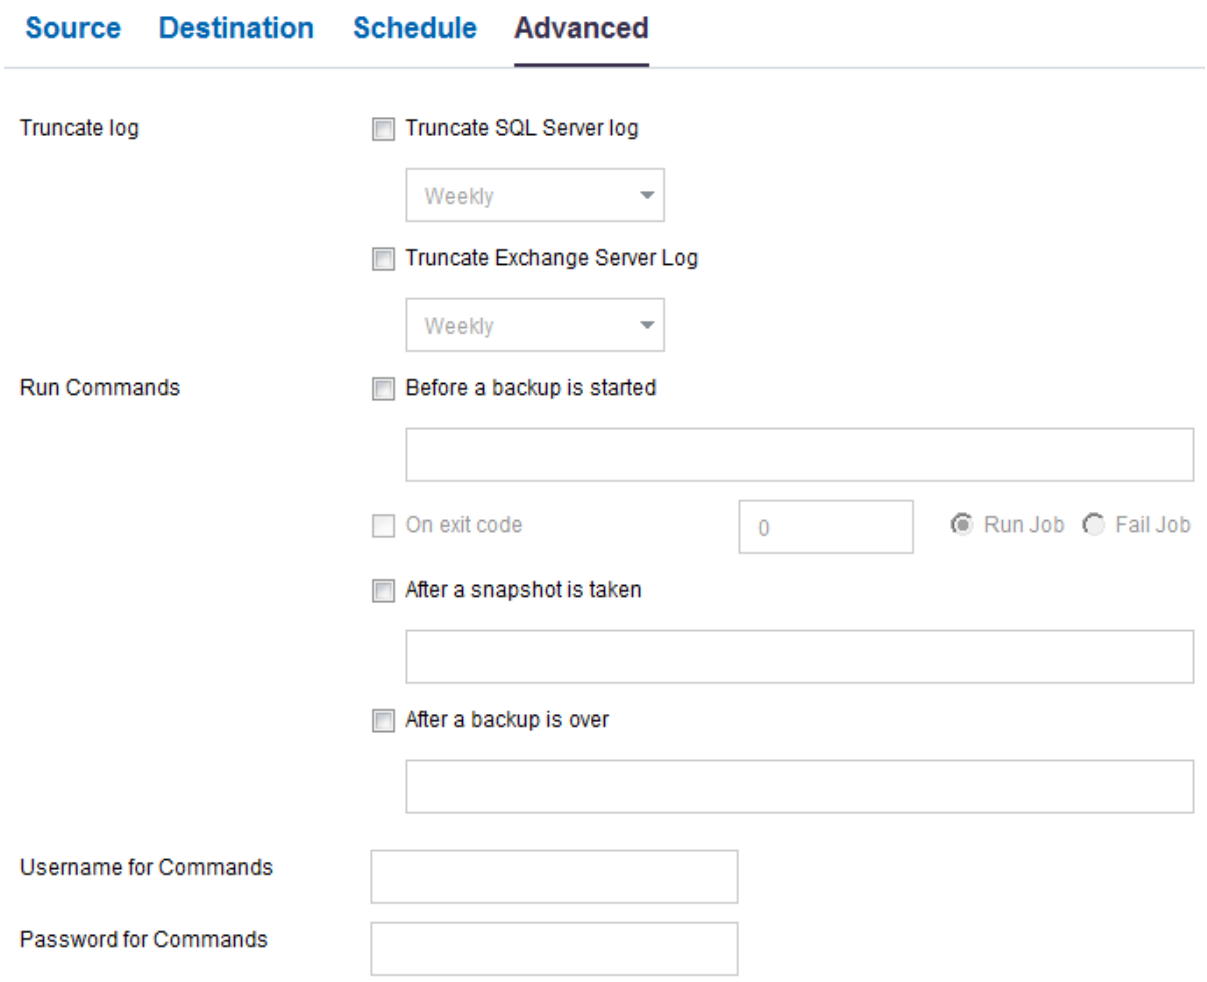
Consider the following points for the merge job:

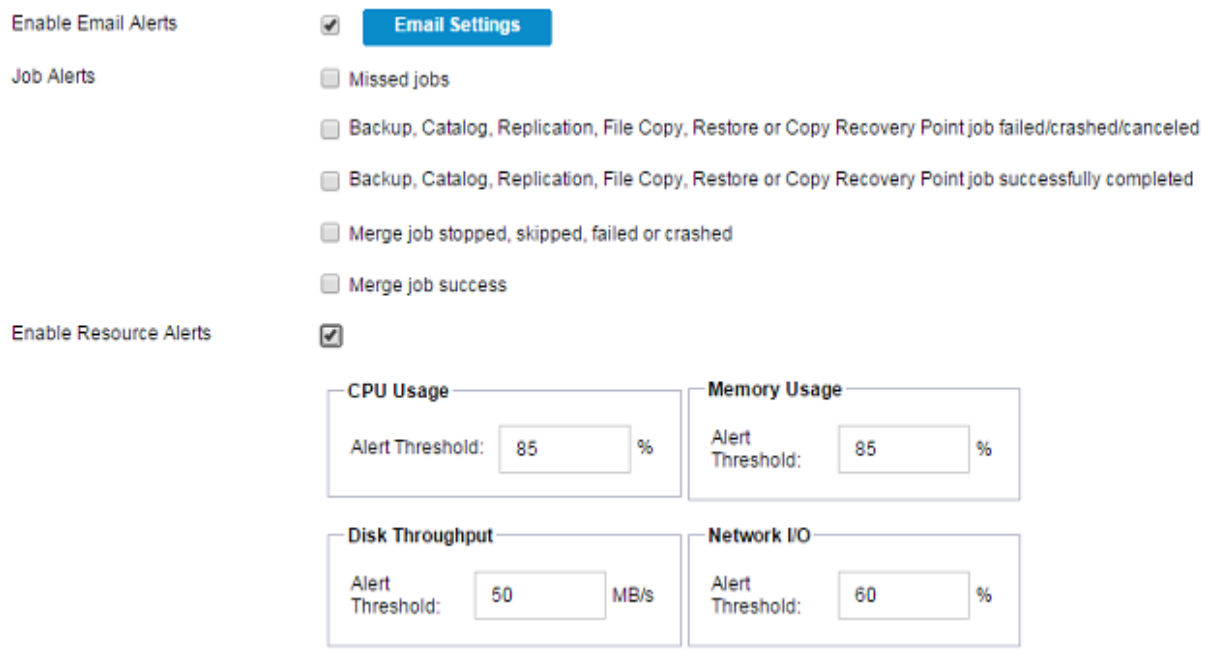
- At any given time only one merge job can run for a node.
- If a merge job starts, it has to complete before the next merge job can start. This means, if one or more sets of recovery points are merged, new recovery point cannot be added to this merge process, until the merge process of the current set of recovery point completes.
- If a merge job is processing more than one set of recovery points (for example set [1~4], set [5~11], and set [12~14]; they are three sets), recovery point server processes these sets one by one.
- If a merge job is resumed after a pause, the job detects at which point it is paused and resumes the merge from the break-point.

Specify the Advanced Settings

The **Advanced** tab lets you specify some advanced settings for the backup job. The advanced settings include providing truncate log settings, providing the location of any scripts, and email settings.

The following image displays the Advanced tab:





Follow these steps:

1. Specify the following details.

Truncate Log

Lets you specify the schedule to truncate logs for SQL Server and Exchange Server. You can specify the schedule as **Daily**, **Weekly**, or **Monthly**.

User Name

Lets you specify the user who is authorized to run a script.

Password

Lets you specify the password of the user who is authorized to run the script.

Run a command before backup is started

Lets you run a script before the backup job starts. Specify the path where the script is stored. Click **On exit code** and specify the exit code for **Run Job** or **Fail Job**. **Run Job** indicates that the backup job will continue when the script returns the exit code. **Fail Job** indicates that the backup job will stop when the script returns the exit code.

Run a command after snapshot is taken

Lets you run a script after the backup snapshot is taken. Specify the path where the script is stored.

Run a command after backup is over

Lets you run a script after the backup job is completed. Specify the path where the script is stored.

Enable Email Alerts

Lets you enable email alerts. You can configure email settings and can specify the types of alerts that you want to receive in an email. When you select this option, the following options are enabled for your selection.

Email Settings

Lets you configure the email settings. Click **Email Settings** and configure the email server and proxy server details.

Job Alerts

Lets you select the types of job emails you want to receive.

Enable Resource Alerts

Lets you specify a threshold for CPU Usage, Memory Usage, Disk Throughput, Network I/O. You can provide the value in percentage. You will receive an email when the Alert Threshold value exceeds.

2. Click **Save**.

Note: When you select a node as a backup source or backup proxy, Arcserve UDP checks whether the agent is installed on the node and if it is the latest version. Arcserve UDP then displays a verification dialog that lists all the nodes that either have an outdated version of the agent or does not have the agent installed. To install/upgrade the agent on these nodes, select the installation method and click Save.

The changes are saved and a green checkmark is displayed next to the task name. The plan page closes.

Note: If you have to add another task, you must select the plan from the **Resources** tab and modify the plan. To modify the plan, click the plan from the center pane. The plan opens and you can modify it.

The backup plan is created and automatically deployed to the source node. The backup runs per the schedule that you have configured in the **Schedule** tab. You can also perform a manual backup at any time.

Add a Virtual Standby to EC2 Task to the Plan

Create a virtual standby to EC2 task so that the backup data is converted to a virtual machine format and a virtual machine is created.

Notes:

- Virtual Standby to EC2 does not support automatic starting of the Virtual Machine.
- If you pause the plan, the Virtual Standby job will not start. When you resume the plan again, the Virtual Standby job does not resume automatically. You have to manually run another backup job to start the Virtual Standby job. Also, if the plan is paused the Pause/Resume Virtual Standby option will not be available. If you do not want the virtual machine to start automatically after the plan is paused, then manually pause the heartbeat for the nodes.

Follow these steps:

1. Click **Add a Task** from the left pane.

   A new task is added to the left pane.

2. From the **Task Type** drop-down menu, select **Virtual Standby**.

   The Virtual Standby task is added.

3. From the **Source** tab select one source for the virtual standby task.
4. Click the **Virtualization Server** tab.
5. Select EC2 as the virtualization type and enter details.

   Virtualization Type - EC2

   Account Name

   Select the existing Amazon AWS account to access AWS EC2. You can also add a new account by clicking **Add**. For more information, see Add a Cloud Account for Amazon EC2.

   EC2 Region

   Select the EC2 region where your cloud proxy is located. Arcserve UDP supports all EC2 global regions and EC2 China region.

   Note: The account that you specify must be an administrative account or an account with administrative privileges on EC2.

   VSB Cloud Proxy

   Specify one EC2 instance in the selected region as the cloud proxy.

Note: The EC2 instance must have Arcserve UDP agent installed.

Username and Password

Specify the credential to login at the VSB cloud proxy.

Protocol

Specify HTTP or HTTPS as the protocol that you want to use for communication between the source Arcserve Unified Data Protection agent on VSB cloud proxy.

Port

Specify the port that you want to use for data transfer between the source server and the VSB cloud proxy.

Note: As the cloud proxy is used for data transfer, the related ports must be enabled to access inbound in the AWS EC2 security group, including TCP 8014 and 4091.

6. Click the **Virtual Machine** tab and enter the details for Basic setting, Cloud Storage setting and Networks setting.

   Amazon AWS EC2

   Apply the following Virtual Machine options to Amazon AWS EC2:

   VM Name Prefix

   Specify the prefix that you want to add to the display name for the virtual machine on the AWS EC2.

   Default value: UDPVM_

   Recovery Point Snapshots

   Specify the number of recovery point snapshots (recovery points) for the Virtual Standby machine. The maximum number of recovery point snapshots count is 29 for AWS EC2.

   Combine all unconverted sessions into a single recovery point snapshot

   Specify whether to combine all unconverted sessions into a single recovery point snapshot when next scheduled VSB job takes place.

   Default: Selected

   Instance Type

   Amazon EC2 provides a wide selection of instance types optimized to suit different use cases. Instances are virtual servers that can run applications. They have varying combinations of CPU, memory, storage, and networking capacity, and provide the flexibility to choose the appropriate mix of resources for your

applications. For more information about instance types and how they meet your computing needs, view the link.

EBS Volume Types

General Purpose (SSD) volumes can burst to 3000 IOPS, and deliver a consistent baseline of 3 IOPS/GiB. Provisioned IOPs (SSD) volumes can deliver up to 20000 IOPS, and are the best for EBS-optimized instances. Magnetic volumes, previously known as standard volumes, deliver 100 IOPS on an average, and can burst to hundreds of IOPS. For more information about EBS volume types, view the link.

Network

Lets you define the VPC, subnets, the NICs and security group for the virtual standby virtual machine on AWS EC2.

Note: When the auto assign public IP is enabled, due to the limitation of AWS EC2, only one NIC is mapped to AWS EC2 and others are discarded..

Same number of network adapters as source at last backup

Select this option to define how to map the virtual NIC to the network on EC2. Specify this option when the virtual machine contains virtual NICs and a virtual network.

Note: Those settings are available to configure only when the auto assign public IP is disabled.

7. Click **Save**.

The changes are saved and the virtual standby task is automatically deployed to the virtual standby server.

You have successfully created and deployed the virtual standby plan.

Add Data Transfer to Cloud Throttling Schedule

Throttle schedule lets you control the data transfer to cloud throughput speed that in turn controls the resource usage (network bandwidth) of the server being transfer to cloud. This is useful if you do not want to affect the server performance during business hours. You can add four time windows per day in your data transfer to cloud throttling schedule. For each time window, you can specify a value, in Mbps/Kbps per minute. This value is used to control the transfer to cloud throughput. Valid values are from 1 Mbps/Kbps to 99999 Mbps/Kbps.

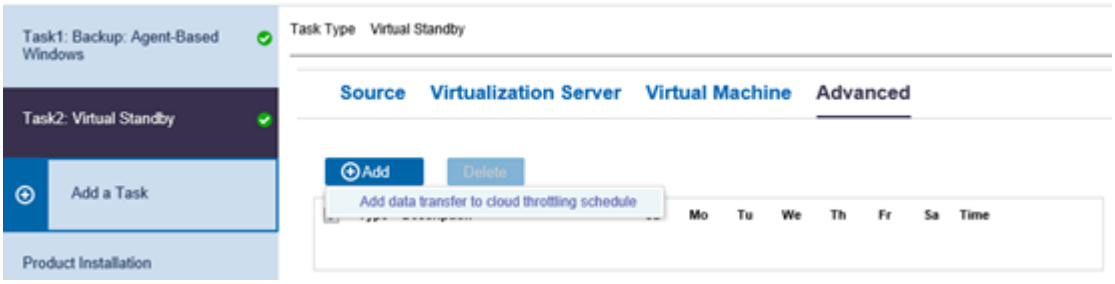
If the data transfer job extends its specified time, then the throttle limit adjusts according to the specified time window. For example, you have defined the data transfer throttle limit as 500 Mbps from 8:00 AM to 8:00 PM, and 2500 Mbps from 8:00 PM to 10:00 PM. If a data transfer job starts at 7:00 PM and it runs for three hours, then, from 7:00 PM to 8:00 PM the throttle limit is 500 Mbps and from 8:00 PM to 10:00 PM the throttle limit is 2500 Mbps.

If there are more than one source node in the backup task with virtual standby to the cloud, they will divide throttle limit equally. For example, you have defined the data transfer throttle limit as 500 Mbps and there are to source nodes in the plan. When their transfer data to cloud simultaneously, the throttle limit is 250 Mbps for every node. After one node transfer task finished, the other running node's throttle limit changes to 500 Mbps.

If you do not define any throttling schedule, the data transfer to cloud job runs as fast as it can.

Follow these steps:

1. Click Add and select Add data transfer to cloud throttling schedule.



The Add New Throttle Schedule dialog appears.

2. Specify the throughput limit in Mbps/Kbps unit.

   The Standby VM Network Configuration - <node name> page opens.

3. Specify the start time to start the backup throughput job.
4. Specify Until to specify an end time for the throughput job.

5. Click Save.

   The Throttle Schedule is specified and displayed on the Schedule page.

Configure the Standby VM Network

You can power on the Standby VM on AWS EC2 with customized network settings. You can configure the following network settings on the standby VM:

- Specify the virtual network and NIC (Network Interface Card), and TCP/IP settings for each network adapter from the **Network Adapter Settings** tab.
- Update the DNS servers to redirect clients from the source computer to the virtual standby virtual machines based on the TCP/IP settings from the **DNS Update Settings** tab.

Follow these steps:

1. From the **resources** tab, navigate to the **Virtual Standby** node group.

   The Virtual Standby nodes are displayed on the center pane.

2. On the center pane, select the node and click **Standby VM Network Configuration**.

   The Standby VM Network Configuration - <node name> page opens.

3. On the **Network Adapter Settings** tab, select the virtual network from the **Standby VM - Virtual Network** list.
4. Select the subnet from the subnet list.
5. Select the elastic IP address from the **Elastic IP** list.
6. Select **Customize the TCP/IP settings**.
7. Click the **Add address** button and add **IP Addresses**, **Gateway Addresses**, **DNS Addresses**, and **WINS Addresses**.

   Note: If you add **DNS Addresses**, then configure the DNS servers in the **DNS Update Settings** tab.

8. Click **Save**.

   The Standby VM Network Configuration - <node name> page closes.

   The Standby VM network is configured.

(Optional) Run the Virtual Standby to EC2 Task Manually

To manually run a virtual standby job, you have to first perform a manual backup. The virtual standby to EC2 task is associated with a backup task. If a plan includes a backup task and a virtual standby to EC2 task, then when you manually run the backup job, the virtual standby job runs automatically after the completion of the backup job.

Follow these steps:

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and click **All Nodes**.

   If you have added any plans, these plans will be displayed in the center pane.
3. Select the nodes that you want to backup and that has a plan assigned to it.
4. On the center pane, click **Actions**, **Backup Now**.

   The **Run a backup now** dialog opens.
5. Select the backup type and provide a name for the backup job.
6. Click **OK**.

   The backup job runs.

   The virtual standby to EC2 job runs immediately after the backup job is over.

   The virtual standby to EC2 job is manually run.

Pause and Resume Virtual Standby Job

Virtual conversion is the process where virtual standby converts the Arcserve Unified Data Protection recovery points from source nodes to virtual machine formats named recovery point snapshots. In the event a source node fails, the virtual standby feature uses the recovery point snapshots to power on a virtual machine for the source node.

As a best practice, allow the virtual conversion process to operate continuously. However, if you want to pause the virtual conversion process on local and remote virtual standby servers temporarily, you can do so from the Console. After you correct the problems on the source node, you can resume the virtual conversion process.

When you pause virtual standby jobs (conversion jobs), the pause operation does not pause the conversion job that is currently in progress. The pause operation applies to only the job that is expected to run at the end of the next backup job. As a result, the next conversion job does not start until you explicitly resume the (paused) conversion job.

If you resume virtual standby for nodes and if there are multiple backup sessions without recovery point snapshot, you will get a dialog to select the smart copy option. If you click Yes, virtual standby will convert the combined session into a single recovery point snapshot. If you click No, virtual standby will convert each session individually

Note: Optionally, you can pause and resume virtual standby jobs directly from the nodes. For more information, see Pause and Resume Virtual Standby Jobs from the Nodes.

Follow these steps:

1. Log in to Arcserve Unified Data Protection.
2. Click the **resources** tab.
3. From the left pane, navigate to **Virtual Standby** and click **All Nodes**.

   If you have added any nodes, then the nodes will be displayed in the center pane.
4. Select the node that you want to pause or resume.
5. On the center pane, click **Actions**, **Virtual Standby**, **Pause** or **Resume**.

   The virtual standby function for the selected node is paused or resumed.

Verify the Plan

To verify your Virtual Standby feature, confirm that you have successfully created the Virtual Standby plan. After you verify that the plan is created successfully, check whether the backup job is running as scheduled. After the backup job successfully completes, the Virtual Standby job runs. You can check the status of the backup job and virtual standby job from the **jobs** tab.

Follow these steps to verify plans:

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and click **All Nodes**.

   A list of all nodes is displayed on the center pane.
3. Verify that plans are mapped with nodes.

Follow these to verify Virtual Standby jobs:

1. Click the **jobs** tab.
2. From the left pane, click **All Jobs**.

   The status of each job is listed on the center pane.
3. Verify that the backup job and Virtual Standby job is successful.

   The plan for virtual standby is successfully verified.

   The virtual standby machine is created.

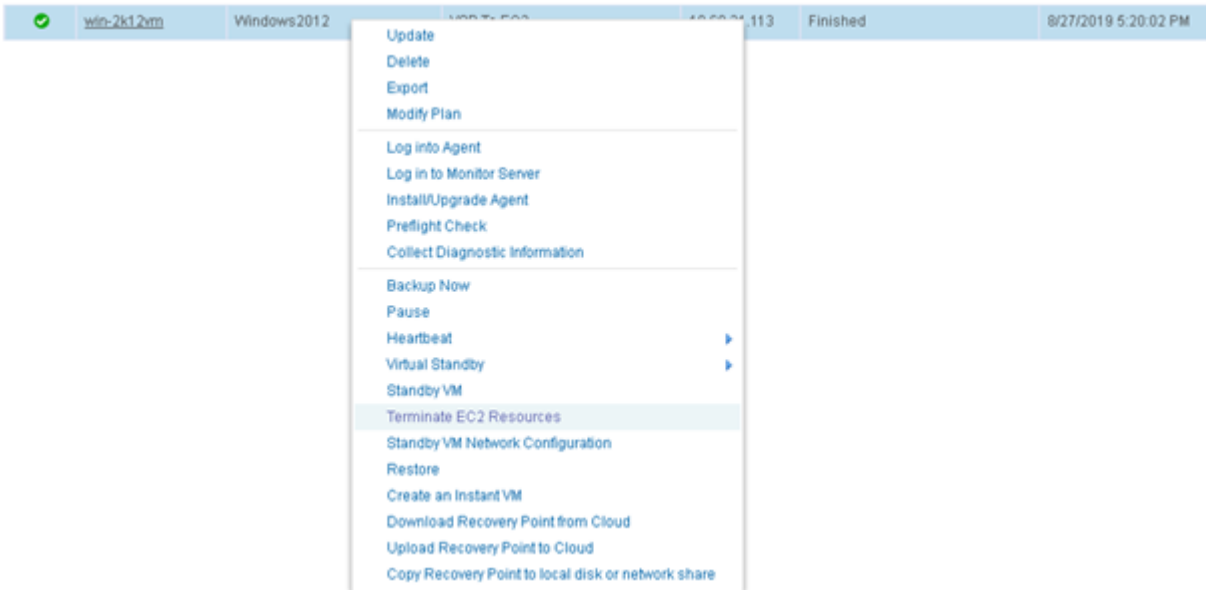
Terminate EC2 Resources

Follow these steps:

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes** and click **All Nodes**.

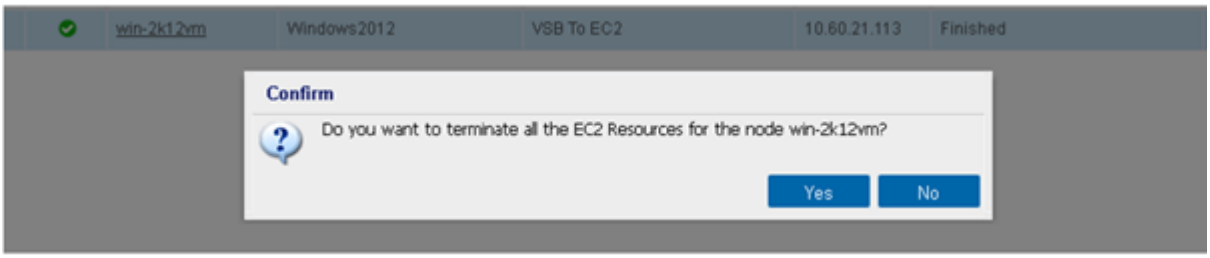
   All the nodes are displayed on the center pane.
3. Right-click a node and select Terminate EC2 Resources.

   A Confirm dialog opens.



4. Click **Yes** to clean the AWS Cloud resources that are generated from Virtual Standby to EC2 task. Click **No** to stop the Termination process.

   Note: If a node is not having any successful VSB to EC2 jobs, such node does not have the Terminate EC2 Resources option.



How to Create a Virtual Standby to Microsoft Azure Plan

The virtual standby converts the recovery points to virtual machine formats and prepares a snapshot to easily recover your data when needed. This feature provides the high availability capability also and ensures that the virtual machine can take over immediately when the source machine fails. The standby virtual machine is created by converting the recovery points to a Microsoft virtual machine format.

Note: The virtual standby task runs only if the backup task creates a valid recovery point snapshot. If the backup task fails, then the virtual standby task is skipped.

What To Do Next?

1. Review the Prerequisites and Considerations
2. Create a Plan with a Backup Task
3. Add a Virtual Standby to Azure Task
4. (Optional) Run the Virtual Standby Job Manually
5. Pause and Resume the Heartbeat
6. Pause and Resume the Virtual Standby Job
7. Verify the Plan

Review the Prerequisites and Considerations

Verify if you have completed the following prerequisites:

- Log into the Console.
- Install the server component and create Data Stores if you want to store the backup data to recovery point servers.
- You have a valid recovery point to create a virtual standby machine. The recovery points can be from one of the following tasks:
  - Backup, Agent-based Windows
  - Backup, Host-Based Agentless
  - Replicate
  - Replicate from a remote Recovery Point Server
- Verify if the prerequisites to add a cloud account on Microsoft Azure are met. For more information, see Prerequisites for adding a Cloud Account on Microsoft Azure.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Create a Plan with a Backup Task

A plan includes different types of tasks that you want to perform. To create a virtual standby machine, you create a plan that includes a backup task and a virtual standby task. A backup task performs a backup of the source nodes and stores the data to the specified destination. This backup data is then used by the virtual standby feature and converts it to a virtual machine format.

For more information, refer to create a plan with a backup task.

Add a Virtual Standby Task to the Plan

Create a Virtual Standby to Azure task so that the backup data is converted to a virtual machine format and a virtual machine is created. The virtual standby feature also monitors the heartbeat of the source node so that when the source node is down, the virtual machine immediately takes over as the source node.

Notes:

- Virtual standby cannot automatically power on recovery point snapshots taken from host-based virtual machine nodes, nodes replicated from a remote recovery point server, and the Source of the Virtual Standby task is the one replicated to a different Site. You have to manually power on recovery point snapshots for such nodes.
- If you pause the plan, the Virtual Standby job will not start. When you resume the plan again, the Virtual Standby job is not resumed automatically. You have to manually run another backup job to start the Virtual Standby job. Also, if the plan is paused the Pause/Resume Virtual Standby option will not be available. If you do not want the virtual machine to start automatically after the plan is paused, then you have to manually pause the heartbeat for the nodes.

Follow these steps:

1. Click **Add a Task** from the left pane.

   A new task is added to the left pane.

2. From the **Task Type** drop-down menu, select **Virtual Standby**.

   The Virtual Standby task is added.

3. From the **Source** tab select one source for the virtual standby task.

4. Click the **Virtualization Server** tab and enter the virtualization server and monitoring server details.

   Virtualization Type - Azure

   Virtualization Type

   Specify Azure as the virtualization type.

   Account Name

   Select an existing Azure account. You can also add a new account by clicking **Add**.

   For more information, see how to add a cloud account.

Resource Group

Specify a resource group. You should have a resource group in Azure.

For instructions, see Resource group in Azure in the Microsoft documentation.

Region

Select the Azure region where your standby VM want to be located. For more information about Region, see Regions in Azure.

Monitor

Specify the host name of the server that monitors the status of the source server.

Notes:

- You can use any physical computer or virtual machine as the monitor server .
- You cannot use the backup source server as the monitor server.
- Monitor server configuration is not required if the nodes are replicated from a remote recovery point server or the Source of the Virtual Standby task is the one replicated to a different Site.
- Monitor server configuration is not required if the Virtual Standby Source is the replicate task and the replication target RPS server is inside Azure.

User Name

Specify the user name to log into the monitoring system.

Password

Specify the password for the user name to log into the monitoring system.

Protocol

Specify HTTP or HTTPS as the protocol that you want to use for communication between the Arcserve Unified Data Protection and the monitoring server.

Port

Specify the port that you want to use for data transfer between the Arcserve Unified Data Protection and the monitoring server.

5. Click the **Virtual Machine** tab and enter the details for the VM Basic Settings, VM DataStore for VMware, VM path for Hyper-V, and VM Network.

VM Name Prefix

Specify the prefix that you want to add to the display name for the virtual machine on the Azure.

Default value: UDPVM_

Recovery Point Snapshots

Specify the number of recovery point snapshots (recovery points) for the standby virtual machine. The maximum number of recovery point snapshots count is 29 for Azure.

Default value: 5

Combine all unconverted sessions into a single recovery point snapshot

Specify whether to combine all unconverted sessions into a single recovery point snapshot when next scheduled VSB job takes place.

Default: Selected

Virtual Machine Size

Microsoft Azure provides a wide selection of Virtual Machine Size optimized to suit different use cases. They have varying combinations of CPU, memory, storage, and networking capacity. For more information about Virtual Machine Size and how they meet your computing needs, view Sizes for Windows virtual machine in Azure.

Storage Account Name

Select a Storage Account Name. You should have a Storage Account Name in Azure. For Storage Account kind, either select Storage (general purpose v1) or StorageV2 (general purpose v2). For more information, see Storage account in Azure in the Microsoft documentation.

Virtual Network

Select a Virtual Network. You should have a Virtual Network in Azure. For more information, see Virtual Network in Azure in the Microsoft documentation.

Subnet

Select a Subnet according to selected Virtual Network. You should have a Subnet in Azure. For more information, see Subnet in Azure in the Microsoft documentation.

Network Security Group

Select a Network Security Group. You should have a Network Security Group in Azure. Configure the security group rules to open the related ports, including 3389 for remote desktop, 8014, 8015 for Arcserve UDP communication. For instructions, see Network Security Group in the Microsoft documentation.

Enable auto assign Public IP

When the auto assign public IP is enabled, the public IP will be assigned to Standby VM automatically when it is started in Azure.

6. Click the **Advanced** tab and provide the following details:

Automatically start the Virtual Machine

Specify if you want to start the virtual machine automatically.

Note: This option is unavailable for host-based virtual machine nodes and nodes replicated from a remote recovery point server and the Source of the Virtual Standby task is the one replicated to a different Site. The Virtual Standby Source is the replicate task and the replication target RPS server inside Azure.

Timeout

Specify the time that the monitor server must wait for a heartbeat before it powers on a recovery point snapshot.

Frequency

Specify the frequency that the source server communicates heartbeats to the monitor server.

Example: The Timeout value specified is 60. The Frequency value specified is 10. The source server will communicate heartbeats in 10-second intervals. If the monitoring server does not detect a heartbeat within 60 seconds of the last heartbeat that was detected, the monitor server powers on a virtual machine using the latest recovery point snapshot.

Customize job parameters

You can customize job parameters for the following options:

- *Number of threads uploading for each job*: Default Value: 4
- *Buffer size for each thread*: Default Value: 4096 KB

Enable Email Alerts

Lets you receive email alerts depending on the settings that you provide. When you select this option, further categories of email alerts are enabled for your selection.

- **Missing heartbeat for source machine**--Virtual standby sends alert notifications when the monitor server does not detect a heartbeat from the source server.

  Note: For nodes from Replicate from a remote Recovery Point Server or if the source of the Virtual Standby task is the one that is replicated to a different site, this option is not available.

- **VM powered on for source machine configured with auto power ON**--Virtual Standby sends alert notifications when it powers on a virtual machine

that was configured to power on automatically when a heartbeat is not detected.

Note: For nodes from Replicate from a remote Recovery Point Server or if the source of the Virtual Standby task is the one that is replicated to a different site, this option is not available. This option is unavailable for host-based virtual machine nodes also.

- **VM powered on for source machine configured with manual power ON**--Virtual Standby sends alert notifications when it manually powers on a virtual machine.
- **Virtual Standby errors/failure/crash**--Virtual Standby sends alert notifications when it detects an error that occurred during the conversion process.
- **Virtual Standby success**--Virtual Standby sends alert notifications when it detects that a virtual machine powered on successfully.
- **The Virtual Standby did not start successfully from the Recovery Point Snapshot**--Virtual Standby sends alert notifications when it detects that a virtual machine was not powered automatically and the Automatically start the Virtual Machine Stand-in Recovery option is specified.

7. Click **Save**.

   The changes are saved and the virtual standby task is automatically deployed to the virtual standby server.

   Note: When the virtual standby task is complete, the virtual machine standby volume is created. The standby virtual machine is created only after the virtual machine is powered on from Arcserve UDP.

   You have successfully created and deployed the Virtual Standby to Azure plan.

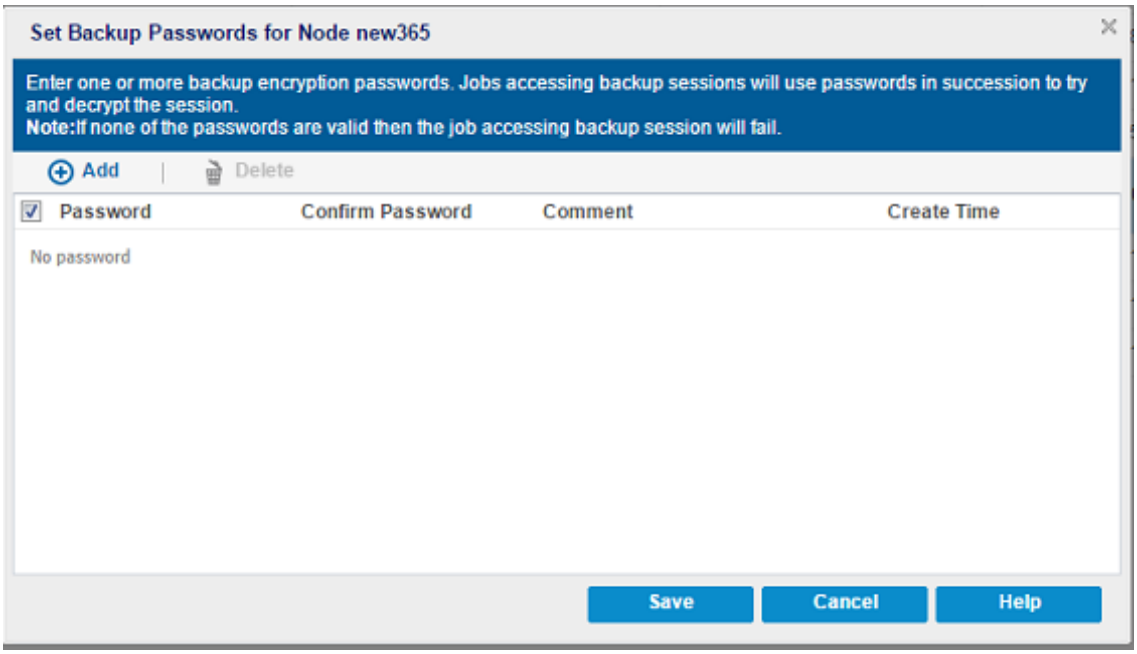
Set Backup Passwords for One or More Nodes

To ensure that the converter can convert the replicated recovery points, virtual standby lets you specify backup passwords for the data that the converter can use to convert the data.

Follow these steps:

1. On the Console, click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and click **All Nodes**.
3. From the center pane, right-click the node and click **Set Backup Passwords**.

   The **Set Backup Passwords for Node** dialog opens.

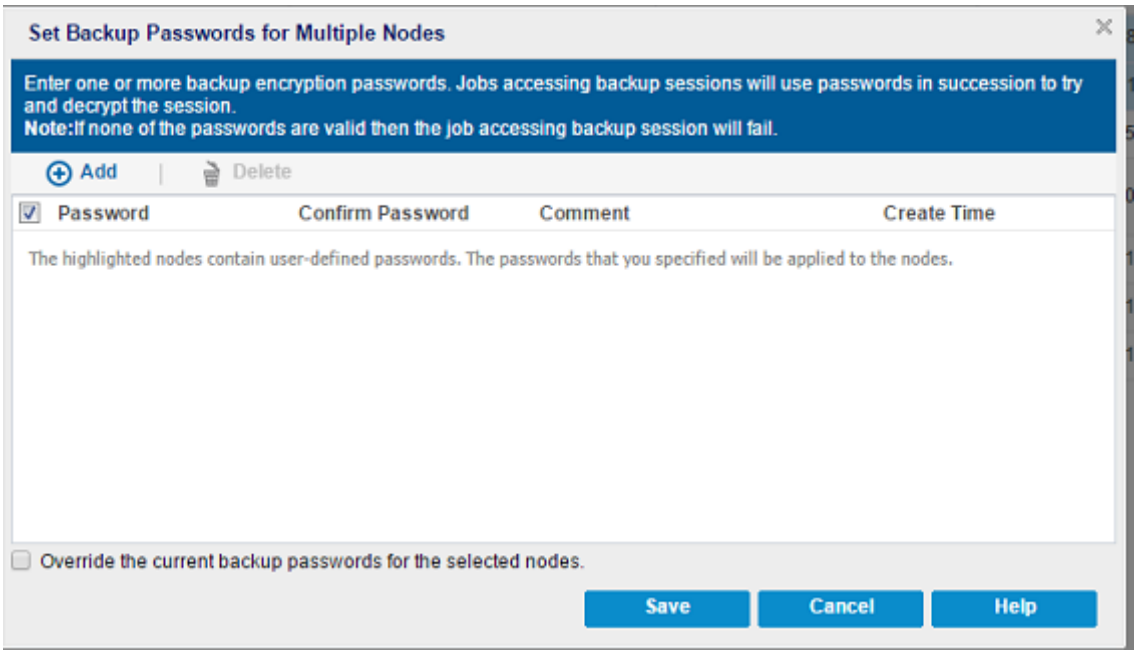


4. Perform the following tasks in the **Set Backup Passwords** dialog for one or more nodes:

   Add--Click **Add** to add one or more backup passwords to the selected nodes.

   Delete--Click **Delete** to delete one or more backup passwords from the selected nodes.

   Note: For multiple nodes, you can override the current backup passwords for multiple nodes by selecting the **Override the current backup passwords** for the selected nodes check box.



5. Click **Save**.

   The dialog closes and the backup passwords are set for the selected remote nodes.

(Optional) Run the Virtual Standby Job Manually

To manually run a virtual standby job, you have to first perform a manual backup. The virtual standby task is associated with a backup task. If a plan includes a backup task and a virtual standby task and you manually run the backup job, the virtual standby job runs automatically after the completion of the backup job.

Follow these steps:

1. Click the **Resources** tab.
2. From the left pane, navigate to **Nodes**, and click **All Nodes**.

   Displays the plans that you added.
3. Select the nodes that you want to backup. The selected node must have assigned a plan.
4. On the center pane, click **Actions**, **Backup Now**.

   The **Run a backup now** dialog opens.
5. Select the backup type and provide a name for the backup job.
6. Click **OK**.

   The backup job runs.

   The virtual standby job runs immediately after the backup job is over.

   The virtual standby job is manually run.

Pause and Resume Heartbeat

The Arcserve Unified Data Protection solution lets you pause and resume the heartbeats that are detected by the monitoring server. The heartbeat is the process where the source server and monitoring server communicate about the health of the source server. If the monitoring server does not detect a heartbeat after a specified length of time, the virtual standby feature provisions the virtual machine to function as the source node.

Examples: When to Pause or Resume Heartbeats

The following examples describe when to pause and resume heartbeats:

- Pause the heartbeat when you want to offline a node (source server) for maintenance.
- Resume the heartbeat after the maintenance tasks are complete and the node (source server) is online.

Be aware of the following behavior:

- You can pause and resume heartbeats at the group level or at the individual node level.
- You can pause and resume heartbeats for one or more nodes in one step.
- The Arcserve Unified Data Protection solution does not power on recover point snapshots while the heartbeat is in a paused state.
- When you upgrade the agent installations on source nodes, Arcserve Unified Data Protection pauses the heartbeat for the nodes. To help ensure that monitor servers monitor the upgraded nodes, resume the heartbeat for the nodes after you complete the upgrades on the nodes.

Follow these steps:

1. Log in to Arcserve Unified Data Protection.
2. Click the **resources** tab.
3. From the left pane, navigate to **Virtual Standby** and click **All Nodes**.

   If you have added any nodes, then the nodes will be displayed in the center pane.
4. Select the node that you want to pause or resume.
5. On the center pane, click **Actions**, **Heartbeat**, **Pause** or **Resume**.

   The heartbeat of the selected node is paused or resumed.

Pause and Resume Virtual Standby Job

Virtual conversion is the process where virtual standby converts the Arcserve Unified Data Protection recovery points from source nodes to virtual machine formats named recovery point snapshots. In the event a source node fails, the virtual standby feature uses the recovery point snapshots to power on a virtual machine for the source node.

As a best practice, allow the virtual conversion process to operate continuously. However, if you want to pause the virtual conversion process on local and remote virtual standby servers temporarily, you can do so from the Console. After you correct the problems on the source node, you can resume the virtual conversion process.

When you pause virtual standby jobs (conversion jobs), the pause operation does not pause the conversion job that is currently in progress. The pause operation applies to only the job that is expected to run at the end of the next backup job. As a result, the next conversion job does not start until you explicitly resume the (paused) conversion job.

If you resume virtual standby for nodes and if there are multiple backup sessions without recovery point snapshot, you will get a dialog to select the smart copy option. If you click Yes, virtual standby will convert the combined session into a single recovery point snapshot. If you click No, virtual standby will convert each session individually

Note: Optionally, you can pause and resume virtual standby jobs directly from the nodes. For more information, see Pause and Resume Virtual Standby Jobs from the Nodes.

Follow these steps:

1. Log in to Arcserve Unified Data Protection.
2. Click the **resources** tab.
3. From the left pane, navigate to **Virtual Standby** and click **All Nodes**.

   If you have added any nodes, then the nodes will be displayed in the center pane.
4. Select the node that you want to pause or resume.
5. On the center pane, click **Actions**, **Virtual Standby**, **Pause** or **Resume**.

   The virtual standby function for the selected node is paused or resumed.

Verify the Plan

To verify your Virtual Standby feature, confirm that you have successfully created the Virtual Standby plan. After you verify that the plan is created successfully, check whether the backup job is running as scheduled. After the backup job successfully completes, the Virtual Standby job runs. You can check the status of the backup job and virtual standby job from the **jobs** tab.

Follow these steps to verify plans:

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and click **All Nodes**.

   A list of all nodes is displayed on the center pane.
3. Verify that plans are mapped with nodes.

Follow these to verify Virtual Standby jobs:

1. Click the **jobs** tab.
2. From the left pane, click **All Jobs**.

   The status of each job is listed on the center pane.
3. Verify that the backup job and Virtual Standby job is successful.

   The plan for virtual standby is successfully verified.

   The virtual standby machine is created.

How to Create a Virtual Standby Plan

The virtual standby converts the recovery points to virtual machine formats and prepares a snapshot to easily recover your data when needed. This feature provides the high availability capability also and ensures that the virtual machine can take over immediately when the source machine fails. The standby virtual machine is created by converting the recovery points to an VMware or a Hyper-V virtual machine format.

Note: The virtual standby task runs only if the backup task creates a valid recovery point snapshot. If the backup task fails, then the virtual standby task is skipped.

What To Do Next?

1. Review the Prerequisites and Considerations
2. Create a Plan with a Backup Task
3. Add a Virtual Standby Task to the Plan
4. (Optional) Run the Virtual Standby Job Manually
5. Pause and Resume the Heartbeat
6. Pause and Resume the Virtual Standby Job
7. Verify the Plan
8. Applying Best Practices

Review the Prerequisites and Considerations

Verify if you have completed the following prerequisites:

- Log into the Console.
- Install the server component and create Data Stores if you want to store the backup data to recovery point servers.
- You have a valid recovery point to create a virtual standby machine. The recovery points can be from one of the following tasks:
  - Backup, Agent-based Windows
  - Backup, Host-Based Agentless
  - Replicate
  - Replicate from a remote Recovery Point Server
- Back up the full machine to enable the Virtual Standby task. You cannot create a Virtual Standby task if the backup is not a full backup.
- Verify if you have the minimum permission to perform required VSB tasks. For more information, see Minimum Permission Required for VSB tasks.
- Virtual Standby is not supported for Linux sources. To view a list of supported operating systems, databases, and browsers, see Compatibility Matrix.

Considerations:

- Virtual Standby uses thin-provisioned disks for vSphere.
- Virtual Standby uses dynamically expanding virtual hard disks for Hyper-V.
- VMware uses separate machine for monitor server with agent installed.
- Hyper-v itself acts as monitor server with agent installed.

Minimum Permission Required for VSB tasks

The table displays list of minimum permission required to perform all the VSB tasks.

Note: Global permissions are set at vCenter level.

Tasks	Permission
Datastore	Allocate space
	Browse datastore
	Low level file operations
Global	Disable methods
	Enable methods
	Licenses
Host>Configuration	Storage partition configuration
Network	Assign network
Resource	Assign virtual machine to resource pool
Virtual machine > Configuration	Add existing disk
	Add new disk
	Add or remove device
	Advanced
	Change CPU count
	Disk change tracking
	Memory
Virtual machine > Interaction	Power off
	Power on
	Console Interaction
Virtual machine > Inventory	Create from existing
	Create new
	Remove
Virtual machine > Provisioning	Allow disk access
	Allow read-only disk access
	Allow virtual machine download
Virtual machine > Snapshot management	Create snapshot
	Remove snapshot
	Revert to snapshot

Create a Plan with a Backup Task

A plan includes different types of tasks that you want to perform. To create a virtual standby machine, you create a plan that includes a backup task and a virtual standby task. A backup task performs a backup of the source nodes and stores the data to the specified destination. The virtual standby feature uses the backup data and converts to a virtual machine format.

You can create a virtual standby machine from an agent-based Windows backup, host-based agentless backup. You can also create a virtual standby machine from data that are replicated using the **Replicate** task. The following procedure is an example to create agent-based Windows backup.

Notes:

- For more information on host-based agentless backup, see How to Create a Host-Based Virtual Machine Backup Plan.
- For more information on replicating a backup data, see How to Replicate Data Between Data Stores Managed from a UDP Console and Managed From Different UDP Consoles.

Follow these steps:

1. Click the **resources** tab on the Console.
2. From the left pane, navigate to **Plans**, and click **All Plans**.

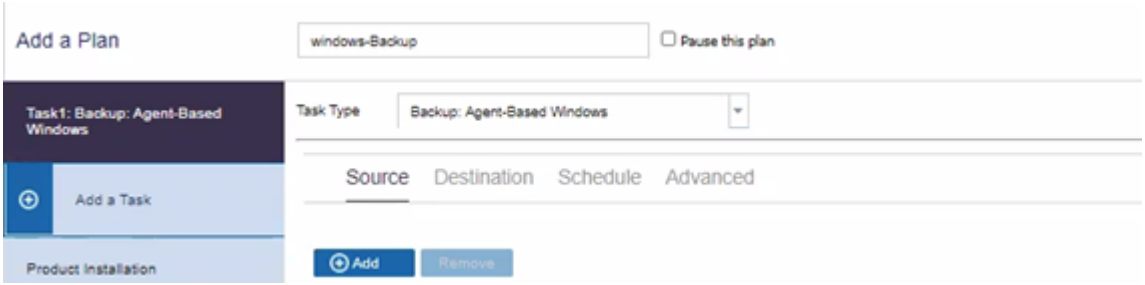
   If you have created plans earlier, those plans are displayed on the center pane.
3. On the center pane, click **Add a Plan**.

   Add a Plan opens.
4. Enter a plan name.
5. (Optional) Select **Pause this plan** check box.

   The plan will not run until you clear the check box to resume the plan.

   Note: If a plan is paused, then any in-progress job is not paused but all corresponding scheduled jobs associated with that plan are paused. However, you can manually run a job. For example, you can manually run backup job and replication job for a node even if the respective plan is paused. In such case, the following task to the on-demand (manual) job does not run. For example, there is a replication task after an on-demand backup job, the replication job does not run for the on-demand backup job. You need to manually run the replication job. When you resume the plan, the pending jobs do not resume immediately. After you resume the plan, the pending jobs run from the next scheduled time.

6. From the **Task Type** drop-down list, select **Backup, Agent-Based Windows**.



Now, specify the Source, Destination, Schedule, and Advanced details.

Specify the Source

The Source page lets you specify the source nodes that you want to protect. You can select more than one node in a plan. If you have not added any nodes to the Console, you can add nodes when you create or modify a plan from the Source page. You can also save a plan without adding any source nodes. The plan gets deployed only after you add source nodes.

Follow these steps:

1. On the Source tab, click **Add**.
2. Select one of the following options:

   Select Nodes to Protect in Arcserve UDP

   Opens the Select Nodes to Protect in Arcserve UDP dialog. Select the nodes from the displayed list. Select this option if you have already added the nodes to the Console.

   Add Windows Node

   Opens the Add nodes to a plan dialog. Select this option if you have not added the nodes and you want to manually add the nodes to protect.

   Discover Nodes from Active Directory

   Opens the Add Nodes to Arcserve UDP Console dialog. Select this option if you want to discover and add nodes from the Active Directory.

3. If you selected the **Select Nodes to Protect in Arcserve UDP** option, on the Select Nodes to Protect dialog, do the following:

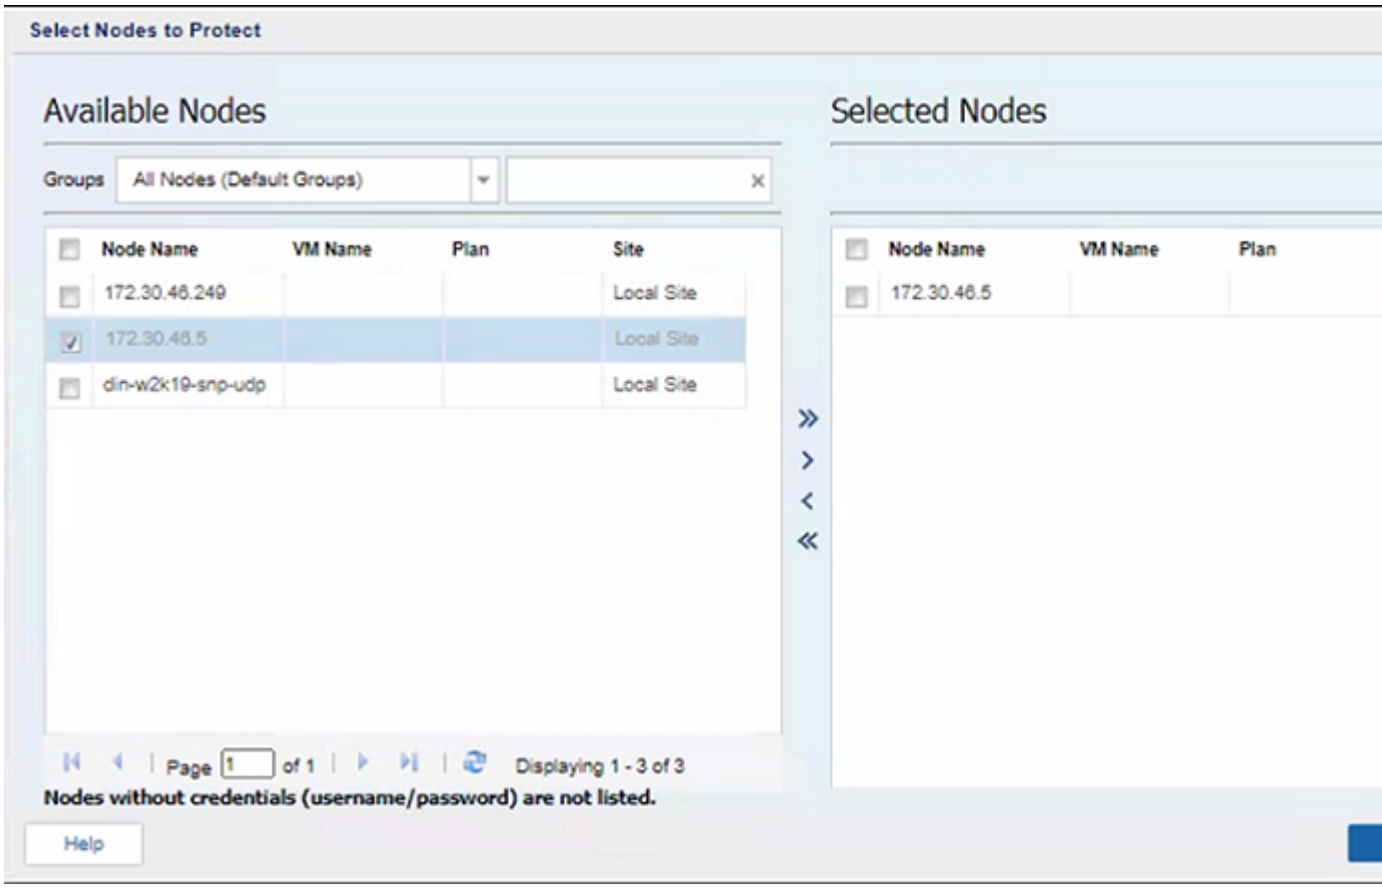
   a. (Optional) Select a filter from the Groups drop-down list to filter nodes. You can enter keywords to further filter your nodes.

      The nodes are displayed on the Available Nodes column.

   b. Select the nodes from the Available Nodes column, and then click the single right arrow (>) or double right arrow (>>) to move the selected nodes or all the nodes respectively to the Selected Nodes column.

      The nodes are displayed in the Selected Nodes column.

   c. Click **OK** to close the dialog.



4. To choose the Protection Type, select one of the following options:

   Back up all volumes

   Prepares a backup snapshot of all the volumes.

   Back up selected volumes

   Prepares a backup snapshot of the selected volume.

   The source is specified.

Specify the Destination

The destination is a location where you store the backup data. You must at least specify the destination to save the plan.

Follow these steps:

1. Select one of the following **Destination Type**:

   Local disk or shared folder

   Specifies that the backup destination is either a local destination or a shared folder. If you select this option, you can save data as either recovery points or recovery sets. The recovery points and recovery sets options are available on the **Schedule** tab.

   Arcserve UDP Recovery Point Server

   Specifies that the backup destination is a recovery point server. If you select this option, then data is stored as recovery points. You cannot store data as recovery sets.

2. If you have selected **Arcserve UDP Recovery Point Server**, then provide the following details:
   a. Select a recovery point server
   b. Select a data store. The list displays all data stores that are created at the specified recovery point server.
   c. Provide a session password.
   d. Confirm the session password.
3. If you have selected **Local disk or shared folder**, then provide the following details:
   a. Provide the full path of the local or network destination. For the network destination, specify the credentials with the write access.
   b. Select the encryption algorithm. For more information, see Encryption Settings.
   c. Optionally, provide an encryption password.
   d. Confirm the encryption password.
   e. Select a type of compression. For more information, see Compression Type.

Note: If you store the data to a local disk or shared folder, you cannot replicate the data to another recovery point server. Replication is supported only if you store the data to a recovery point server.

The destination is specified.

Specify the Schedule

The Schedule page lets you define a schedule for Backup, Merge, and Throttle functions to repeat at specific intervals. After you define a schedule, the jobs run automatically per the schedule. You can add multiple schedules and can provide retention settings.

A Backup Schedule refers to regular schedule that is repeated multiple times a day based on the number of hours or minutes you select. Besides the regular schedule, a backup schedule also provides options to add daily, weekly, and monthly schedules.

Note: For more information on scheduling and retention settings, see Understanding Advanced Scheduling and Retention.

Follow these steps:

1. (Optional) Select the option to manage recovery points. This option is visible only if you have selected Local or shared folder as your backup destination.

   Retain by Recovery Points

   The backup data is stored as recovery points.

   Retain by Recovery Sets

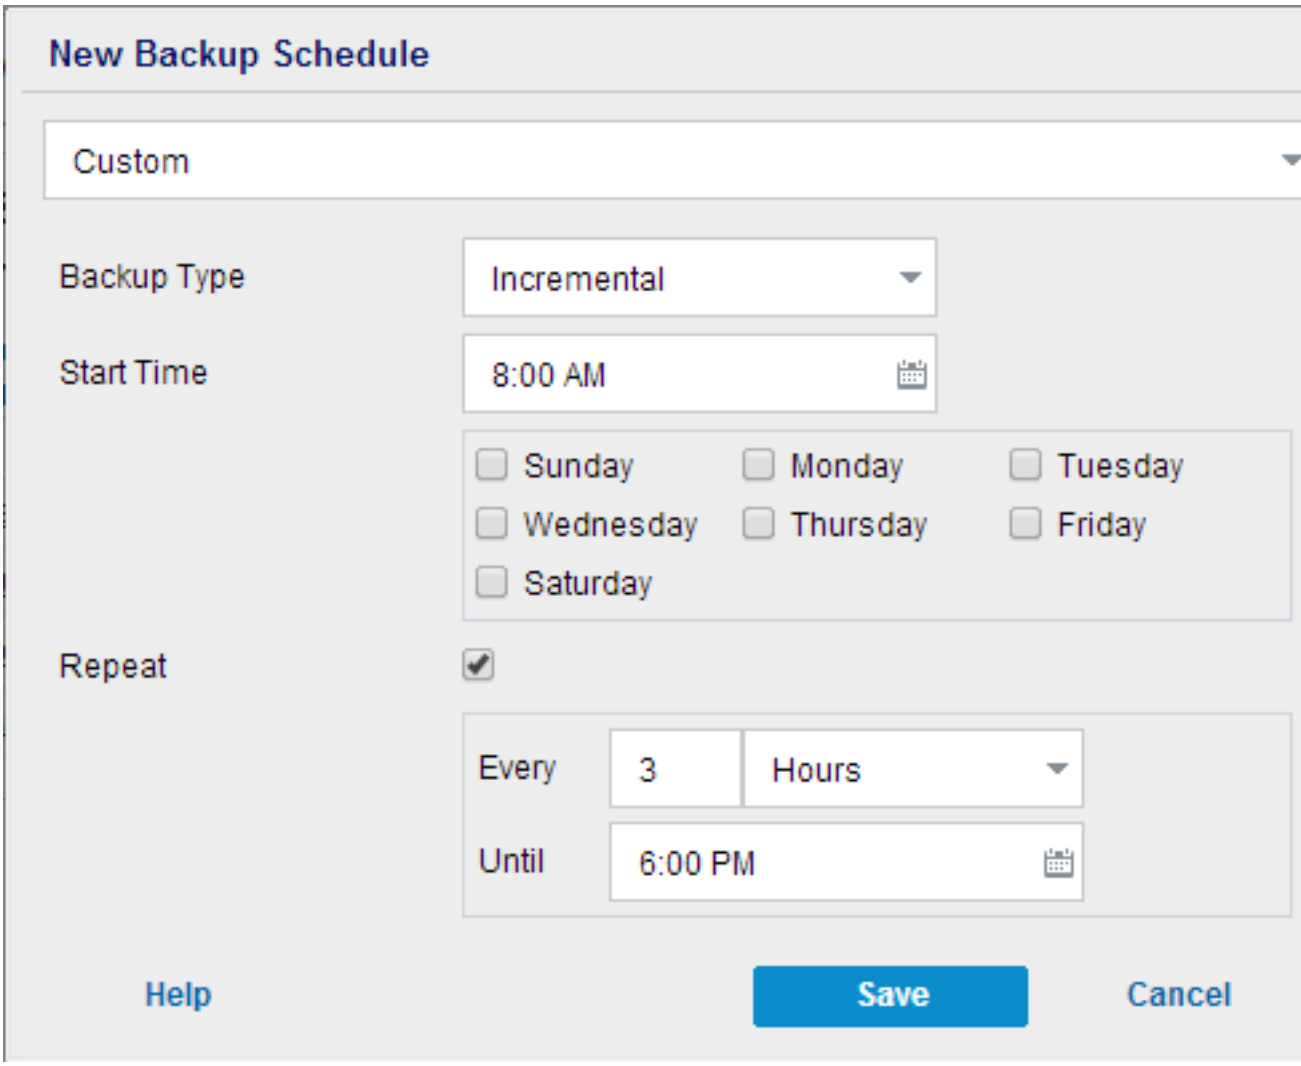
   The backup data is stored as recovery sets.

2. Add backup, merge, and throttle schedules.

   Add Backup Schedule

   a. Click **Add** and select **Add Backup Schedule**.

      The **New Backup Schedule** dialog opens.



b. Select one of the following options:

Custom

Specifies the backup schedule that repeats multiple times a day.

Daily

Specifies the backup schedule that occurs once a day. By default, all the days of the week are selected for Daily backup. If you do not want to run the backup job on a specific day, clear the check box for that day of the week.

Weekly

Specifies the backup schedule that occurs once a week.

Monthly

Specifies the backup schedule that occurs once a month.

c. Select the backup type.

Full

Determines the backup schedule for Full Backups. As scheduled, Arcserve UDP performs a full backup of all used blocks from the source machine. A full backup typically consumes time depending on the backup size.

Verify

Determines the backup schedule for Verify Backups.

Arcserve UDP verifies that the protected data is valid and complete by performing a confidence check of the stored backup image to the backup source. If necessary, the image is resynchronized. A Verify Backup looks at the most recent backup of each individual block and compares the content and information to the source. This comparison verifies that the latest backed up blocks represent the corresponding information at the source. If the backup image for any block does not match the source (possibly because of changes in the system since the last backup), Arcserve UDP refreshes (resynchronizes) the backup of the block that does not match. You can also use a Verify Backup (infrequently) to get the guarantee of full backup without using the space required for a full backup.

Advantages: Produces a small backup image when compared to full backup because only the changed blocks (blocks that do not match the last backup) are backed up.

Disadvantages: Backup time is long because all source blocks are compared with the blocks of the last backup.

Incremental

Determines the backup schedule for Incremental Backups.

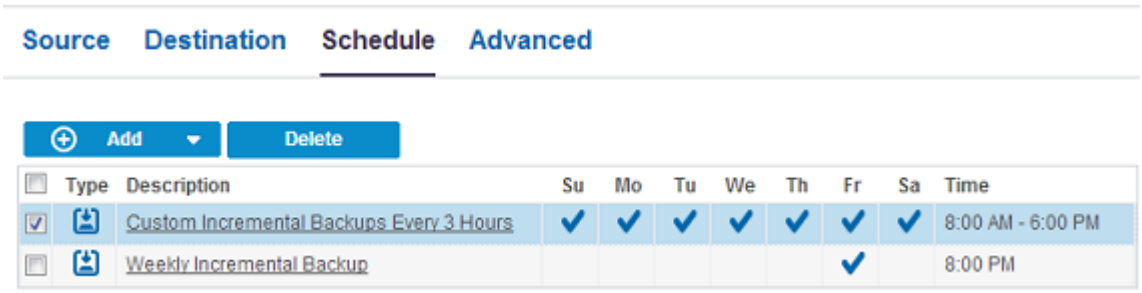
As scheduled, Arcserve UDP incrementally backs up only those blocks that have changed since the last successful backup. The advantages of Incremental Backups are that it is a fast backup and it produces a small backup image. This is the most optimal way to perform a backup.

d. Specify the backup start time.

e. (Optional) Select the **Repeat** check box and specify the repeat schedule.

f. Click **Save**.

The Backup Schedule is specified and displayed on the **Schedule** page.



Add Merge Schedule

a. Click Add and select Add Merge Schedule.

   The **Add New Merge Schedule** dialog opens.

b. Specify the start time to start the merge job.

c. Specify **Until** to specify an end time for the merge job.

d. Click **Save**.

   The Merge Schedule is specified and displayed on the **Schedule** page.

Add Throttle Schedule

a. Click **Add** and select **Add Throttle Schedule**.

   The Add New Throttle Schedule dialog opens.

b. Specify the throughput limit in MB per minute unit.

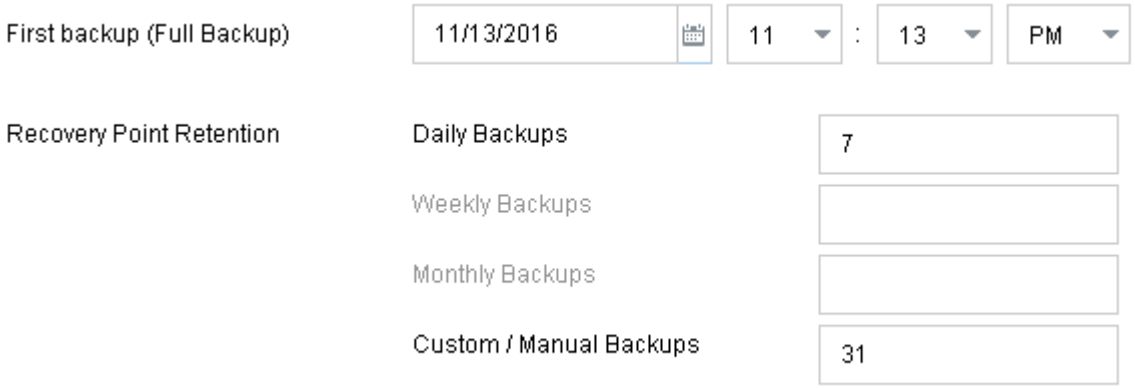
c. Specify the start time to start the backup throughput job.

d. Specify **Until** to specify an end time for the throughput job.

e. Click **Save**.

   The Throttle Schedule is specified and displayed on the **Schedule** page.

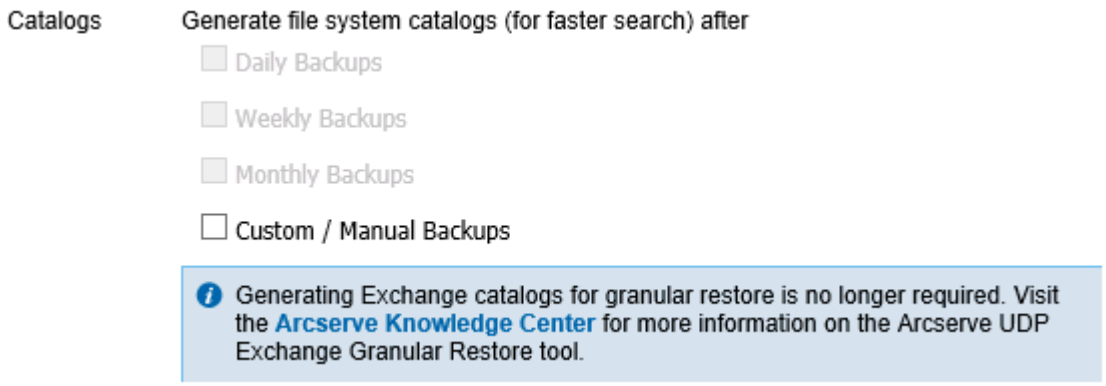
3. Specify the start time for the scheduled backup.



4. Specify the recovery points retention settings for Custom, Daily, Weekly, and Monthly schedule.

   These options are enabled if you have added the corresponding backup schedule. If you modify the retention settings on this page, the changes are reflected on the Backup Schedule dialog.

5. Specify the catalog details.



   Catalogs let you generate the File System catalog. The File System catalog is required to perform faster and better search. If you select the catalog check boxes, the catalogs are enabled depending on the type of backup that you have specified. Clear the check box to disable generating the catalog.

   The schedule is specified.

Understanding Advanced Scheduling and Retention

The scheduling option lets you specify a Custom schedule or a Daily / Weekly / Monthly schedule, or both the schedules. In the Custom schedule, you can configure the backup schedule for each day of the week and you can add up to four backup schedules each day. You can select a specific day of a week and create a time window to define when to run backup and at what frequency.

Schedule	Supported Job	Comments
Backup	Backup job	Define time windows to run backup jobs.
Backup throttling	Backup job	Define time windows to control the backup speed.
Merge	Merge job	Define when to run merge jobs.
Daily schedule	Backup job	Define when to run daily backup jobs.
Weekly schedule	Backup job	Define when to run weekly backup jobs.
Monthly schedule	Backup job	Define when to run monthly backup jobs.

You can also specify the retention settings for the recovery points.

Note: Set the retention settings within each plan to control how data for the nodes assigned to that plan are retained at the target data store.

Schedules for Daily / Weekly / Monthly backups are independent to the Custom schedule, and each other. You can configure only to run Daily backup, Weekly backup or Monthly backup, without configuring the Custom schedule.

Backup Job Schedule

You can add four time windows per day in your backup schedule. A valid time window is from 12:00 AM until 11:59 PM. You cannot specify a time window such as 6:00 PM to 6:00 AM. In such cases, you must manually specify two different time windows.

For each time window, the start time is inclusive, and the end time is exclusive. For example, you have configured to run Incremental Backup every one hour between 6:00 AM and 9:00 AM and the backup will start at 6:00 AM. This means the backup runs at 6:00 AM, 7:00 AM, 8:00 AM, but NOT at 9:00 AM.

Note: If you want to run the backup job repeatedly until the end of day, set the schedule until 12:00 AM. For example, to run backup job every 15 minutes for the entire day, set the schedule from 12:00 AM to 12:00 AM, every 15 minutes.

Backup Throttle Schedule

Backup throttle schedule lets you control the backup throughput speed that in turn controls the resource usage (disk I/O, CPU, network bandwidth) of the server being backed up. This is useful if you do not want to affect the server

performance during business hours. You can add four time windows per day in your backup throttle schedule. For each time window, you can specify a value, in MB per minute. This value is used to control the backup throughput. Valid values are from 1 MB/minutes to 99999 MB/minutes.

If a backup job extends its specified time, then the throttle limit adjusts according to the specified time window. For example, you have defined the backup throttle limit as 500 MB/minute from 8:00 AM to 8:00 PM, and 2500 MB/minute from 8:00 PM to 10:00 PM. If a backup job starts at 7:00 PM and it runs for three hours, then, from 7:00 PM to 8:00 PM the throttle limit is 500 MB/minute and from 8:00 PM to 10:00 PM the throttle limit is 2500 MB/minute.

If you do not define any backup schedule and backup throughput schedule, the backup runs as fast as it can.

Merge Schedule

Lets you merge recovery points based on the provided schedule.

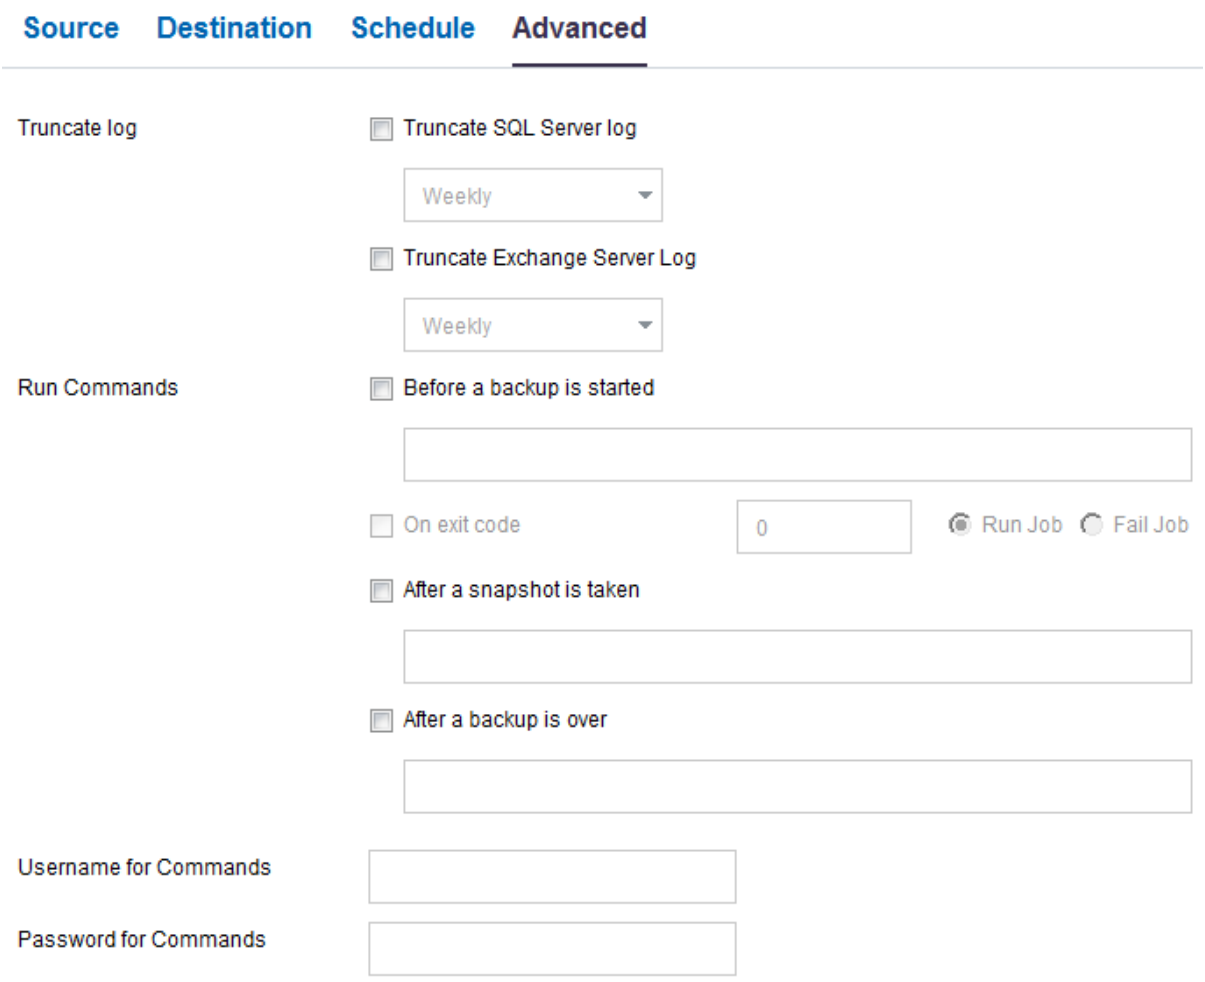
Consider the following points for the merge job:

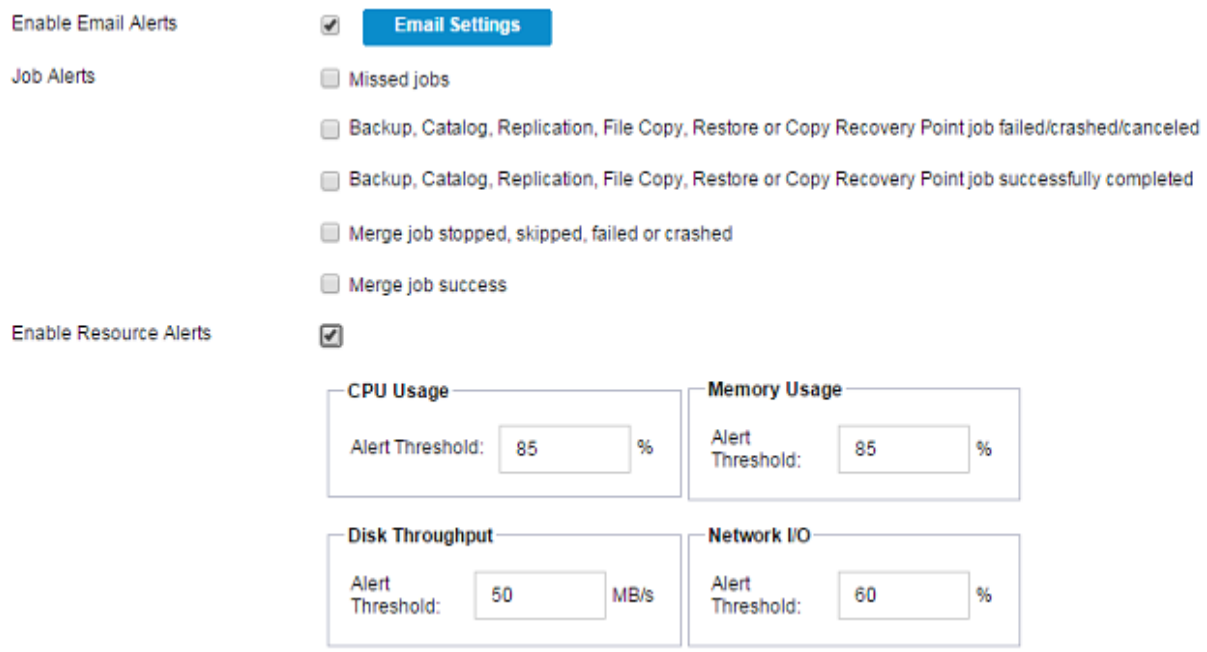
- At any given time only one merge job can run for a node.
- If a merge job starts, it has to complete before the next merge job can start. This means, if one or more sets of recovery points are merged, new recovery point cannot be added to this merge process, until the merge process of the current set of recovery point completes.
- If a merge job is processing more than one set of recovery points (for example set [1~4], set [5~11], and set [12~14]; they are three sets), recovery point server processes these sets one by one.
- If a merge job is resumed after a pause, the job detects at which point it is paused and resumes the merge from the break-point.

Specify the Advanced Settings

The **Advanced** tab lets you specify some advanced settings for the backup job. The advanced settings include providing truncate log settings, providing the location of any scripts, and email settings.

The following image displays the Advanced tab:





Follow these steps:

1. Specify the following details.

Truncate Log

Lets you specify the schedule to truncate logs for SQL Server and Exchange Server. You can specify the schedule as **Daily**, **Weekly**, or **Monthly**.

User Name

Lets you specify the user who is authorized to run a script.

Password

Lets you specify the password of the user who is authorized to run the script.

Run a command before backup is started

Lets you run a script before the backup job starts. Specify the path where the script is stored. Click **On exit code** and specify the exit code for **Run Job** or **Fail Job**. **Run Job** indicates that the backup job will continue when the script returns the exit code. **Fail Job** indicates that the backup job will stop when the script returns the exit code.

Run a command after snapshot is taken

Lets you run a script after the backup snapshot is taken. Specify the path where the script is stored.

Run a command after backup is over

Lets you run a script after the backup job is completed. Specify the path where the script is stored.

Enable Email Alerts

Lets you enable email alerts. You can configure email settings and can specify the types of alerts that you want to receive in an email. When you select this option, the following options are enabled for your selection.

Email Settings

Lets you configure the email settings. Click **Email Settings** and configure the email server and proxy server details.

Job Alerts

Lets you select the types of job emails you want to receive.

Enable Resource Alerts

Lets you specify a threshold for CPU Usage, Memory Usage, Disk Throughput, Network I/O. You can provide the value in percentage. You will receive an email when the Alert Threshold value exceeds.

2. Click **Save**.

Note: When you select a node as a backup source or backup proxy, Arcserve UDP checks whether the agent is installed on the node and if it is the latest version. Arcserve UDP then displays a verification dialog that lists all the nodes that either have an outdated version of the agent or does not have the agent installed. To install/upgrade the agent on these nodes, select the installation method and click Save.

The changes are saved and a green checkmark is displayed next to the task name. The plan page closes.

Note: If you have to add another task, you must select the plan from the **Resources** tab and modify the plan. To modify the plan, click the plan from the center pane. The plan opens and you can modify it.

The backup plan is created and automatically deployed to the source node. The backup runs per the schedule that you have configured in the **Schedule** tab. You can also perform a manual backup at any time.

Add a Virtual Standby Task to the Plan

Create a virtual standby task so that the backup data is converted to a virtual machine format and a virtual machine is created. The virtual standby feature also monitors the heartbeat of the source node so that when the source node is down, the virtual machine immediately takes over as the source node.

Notes:

- Virtual standby cannot automatically power on recovery point snapshots taken from host-based virtual machine nodes, nodes replicated from a remote recovery point server, and the Source of the Virtual Standby task is the one replicated to a different Site. You have to manually power on recovery point snapshots for such nodes.
- If you pause the plan, the Virtual Standby job will not start. When you resume the plan again, the Virtual Standby job is not resumed automatically. You have to manually run another backup job to start the Virtual Standby job. Also, if the plan is paused the Pause/Resume Virtual Standby option will not be available. If you do not want the virtual machine to start automatically after the plan is paused, then you have to manually pause the heartbeat for the nodes.

Follow these steps:

1. Click **Add a Task** from the left pane.

   A new task is added to the left pane.

2. From the **Task Type** drop-down menu, select **Virtual Standby**.

   The Virtual Standby task is added.

3. From the **Source** tab select one source for the virtual standby task.

4. Click the **Virtualization Server** tab and enter the virtualization server and monitoring server details.

   Virtualization Type - VMware

   ESX Host/vCenter

   Specify the host name of the ESX or vCenter Server system.

   User Name

   Specify the user name that is required to log in to the VMware system.

   Note: The account that you specify must be an administrative account or an account with administrative privileges on the ESX or vCenter Server system.

Password

Specify the password for the user name that is required to log in to the VMware system.

Protocol

Specify HTTP or HTTPS as the protocol that you want to use for communication between the source Arcserve Unified Data Protection agent and the monitoring server.

Port

Specify the port that you want to use for data transfer between the source server and the monitoring server.

ESX Node

The values in this field vary based on the value specified in the ESX Host/vCenter field:

ESX Server systems

When you specify an ESX Server system in the ESX Host/vCenter field, this field displays the host name of the ESX Server system.

vCenter Server systems

When you specify a vCenter Server system the ESX Host/vCenter field, this field lets you specify (from a drop-down list) the ESX Server system that you want to associate with this plan.

Monitor

Specify the host name of the server that monitors the status of the source server.

Notes:

- You can use any physical computer or virtual machine as the monitor server .
- You cannot use the backup source server as the monitor server.
- Monitor server configuration is not required if the nodes are replicated from a remote recovery point server or the Source of the Virtual Standby task is the one replicated to a different Site.
- Monitor server configuration is not required if the Virtual Standby Source is the replicate task and the replication target RPS server is inside Azure.

User Name

Specify the user name to log into the monitoring system.

Password

Specify the password for the user name to log into the monitoring system.

Protocol

Specify HTTP or HTTPS as the protocol that you want to use for communication between the Arcserve Unified Data Protection and the ESX Server system (monitoring server).

Port

Specify the port that you want to use for data transfer between the Arcserve Unified Data Protection and the ESX Server system (monitoring server).

Use monitor server as proxy for data transfer

Specify this option to let the monitor server copy the conversion data from the Arcserve Unified Data Protection agent node to the ESX Server data store. With this option enabled, the virtual standby feature transfers the conversion data from the agent node to the ESX Server data store using the fibre channel communication, which is faster than using the LAN communication to transfer data. Only the write operation for the conversion happens over the fibre channel. The read operation happens over the LAN.

Note: The Use monitor server as proxy for data transfer option is enabled by default. You can disable this option to allow the Arcserve Unified Data Protection agent node to copy the conversion data directly to the data store on the ESX Server system.

Virtualization Type - Hyper-V

HyperV Host Name

Specify the host name of the Hyper-V system.

User Name

Specify the user name that is required to log in to the Hyper-V system.

Note: The account that you specify must be an administrative account or an account with administrative privileges on the Hyper-V system.

Password

Specify the password for the User Name that is required to log in to the Hyper-V system.

Protocol

Specify HTTP or HTTPS as the protocol that you want to use for communication between the Arcserve Unified Data Protection server and the Hyper-V Server system (monitoring server).

Port

Specify the port that you want to use for data transfer between the Arcserve Unified Data Protection server and the Hyper-V Server system (monitoring server).

Note: To validate the Hyper-V system, click **Connect**.

5. Click the **Virtual Machine** tab and enter the details for the VM Basic Settings, VM DataStore for VMware, VM path for Hyper-V, and VM Network.

   VMware Systems:

   Apply the following Virtual Machine options to VMware systems:

   VM Name Prefix

   Specify the prefix that you want to add to the display name for the virtual machine on the ESX Server system.

   Default value: UDPVM_

   Resource Pool

   Specify the name of resource pool where standby virtual machine is to be grouped.

   CPU Count

   Specify the minimum and maximum CPU count supported by the standby virtual machine.

   Memory

   Specify the total amount of RAM in MB to be allocated for the standby virtual machine.

   Note: The amount of RAM specified must be a multiple of two.

   Recovery Point Snapshots

   Specify the number of recovery point snapshots (recovery points) for the standby virtual machine. The maximum number of recovery point snapshots count is 29 for VMware virtualization servers.

   All virtual disks share the same datastore

   Select this option to copy all of the disks related to the virtual machine to one data store.

   Clear the check box to copy the disk-related information for the virtual machine to the corresponding data store. Specify the location where you want to store the conversion data.

   Network

Lets you define the NICs, virtual networks, and paths that the ESX Server system uses to communicate with the virtual machines.

Note: The VMware SR-IOV passthrough and Flexible network adapter is not supported.

Same number of network adapters as source at last backup

Select this option to define how to map the virtual NIC to the virtual network. Specify this option when the virtual machine contains virtual NICs and a virtual network.

Clear the check box to define the name of the virtual network that you want the NIC to use to communicate.

Hyper-V Systems:

Apply the following Virtual Machine options to Hyper-V systems:

Basic Settings

Complete the following Basic settings:

VM Name Prefix

Specify the prefix that you want to add to the display name for the virtual machine on the Hyper-V system.

Default value: UDPVM_

CPU Count

Specify the minimum and maximum CPU count supported by the standby virtual system.

Memory

Specify the total amount of RAM in MB to be allocated to the standby virtual machine.

Note: The amount of RAM specified must be a multiple of four.

Recovery Point Snapshots

Specify the number of recovery point snapshots for the standby virtual machine. The maximum number of recovery point snapshots is 24 for Hyper-V virtualization servers.

All virtual disks share the same path

Select this option to specify the location on the Hyper-V server where you want to store the conversion data.

Clear the check box to specify the location on the Hyper-V server where you want to store the conversion data for each virtual disk.

Note: The Arcserve UDP solution does not support creating virtual disk images (VHD/VHDX files) on compressed volumes and volumes that are encrypted by the file system. If the path specified resides on compressed or encrypted Hyper-V volumes, Arcserve UDP prevents you from creating the virtual standby task.

VM Network

Lets you define the NICs, virtual networks, and paths that the Hyper-V server uses to communicate with the virtual machines. Specify one of the following options and complete the required fields.

Same number of network adapters as source at last backup

Select this option to define how to map the virtual NIC to the virtual network. Specify this option when the virtual machine contains virtual NICs and a virtual network.

Clear the check box to define the name of the virtual network that you want the NIC to use to communicate.

6. Click the **Advanced** tab and provide the following details:

Automatically start the Virtual Machine

Specify if you want to start the virtual machine automatically.

Note: This option is unavailable for host-based virtual machine nodes and nodes replicated from a remote recovery point server and the Source of the Virtual Standby task is the one replicated to a different Site.

Timeout

Specify the time that the monitor server must wait for a heartbeat before it powers on a recovery point snapshot.

Frequency

Specify the frequency that the source server communicates heartbeats to the monitor server.

Example: The Timeout value specified is 60. The Frequency value specified is 10. The source server will communicate heartbeats in 10-second intervals. If the monitoring server does not detect a heartbeat within 60 seconds of the last heartbeat that was detected, the monitor server powers on a virtual machine using the latest recovery point snapshot.

Enable Email Alerts

Lets you receive email alerts depending on the settings that you provide. When you select this option, further categories of email alerts are enabled for your selection.

- **Missing heartbeat for source machine**--Virtual standby sends alert notifications when the monitor server does not detect a heartbeat from the source server.

  Note: For nodes from Replicate from a remote Recovery Point Server or if the source of the Virtual Standby task is the one that is replicated to a different site, this option is not available.

- **VM powered on for source machine configured with auto power ON**--Virtual Standby sends alert notifications when it powers on a virtual machine that was configured to power on automatically when a heartbeat is not detected.

  Note: For nodes from Replicate from a remote Recovery Point Server or if the source of the Virtual Standby task is the one that is replicated to a different site, this option is not available. This option is unavailable for host-based virtual machine nodes also.

- **VM powered on for source machine configured with manual power ON**--Virtual Standby sends alert notifications when it manually powers on a virtual machine.

- **Virtual Standby errors/failure/crash**--Virtual Standby sends alert notifications when it detects an error that occurred during the conversion process.

- **Virtual Standby success**--Virtual Standby sends alert notifications when it detects that a virtual machine powered on successfully.

- **The Virtual Standby did not start successfully from the Recovery Point Snapshot**--Virtual Standby sends alert notifications when it detects that a virtual machine was not powered automatically and the Automatically start the Virtual Machine Stand-in Recovery option is specified.

- **Hypervisor is not reachable**--Virtual Standby sends alert notifications when it detects that it cannot communicate with the ESX Server system or the Hyper-V system.

- **VM storage free space less than**--Virtual Standby sends alert notifications when it detects insufficient free disk space on the defined hypervisor path. The detection occurs when the amount of free disk space is less than the user-defined threshold. The threshold can be defined either an absolute value (MB) or as a percentage of the capacity of the volume.

7. Click **Save**.

   The changes are saved and the virtual standby task is automatically deployed to the virtual standby server.

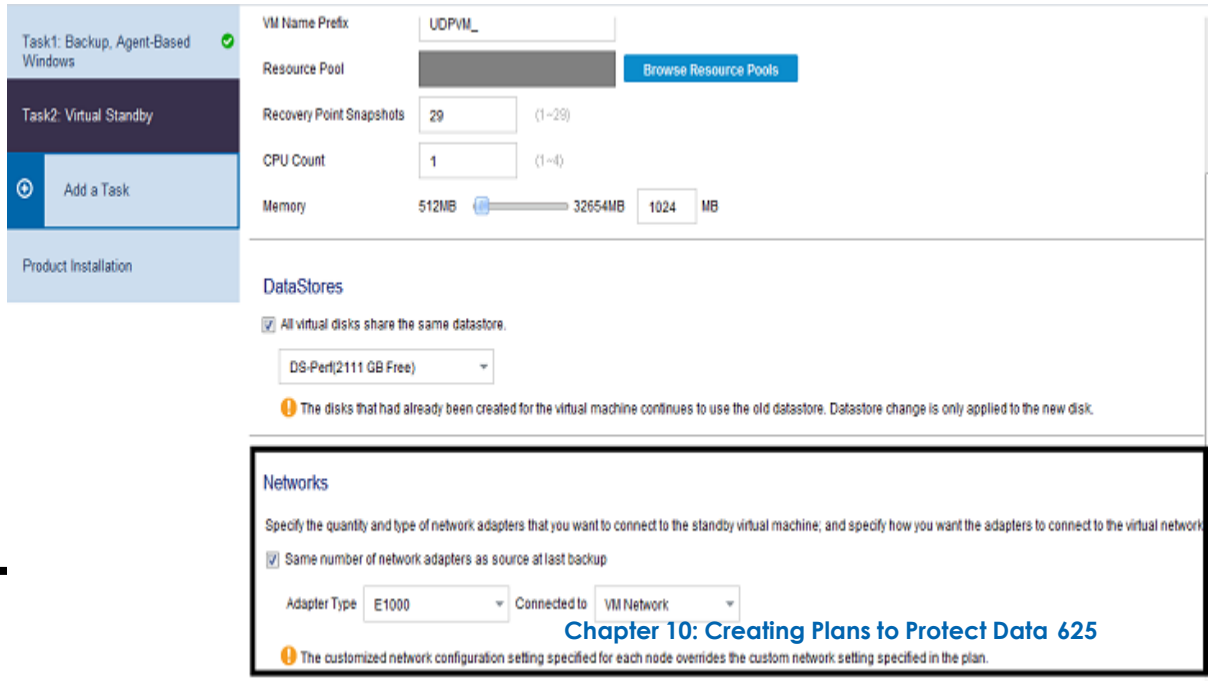
   You have successfully created and deployed the virtual standby plan.

How the Application Determines the Quantity of NICs to Power ON

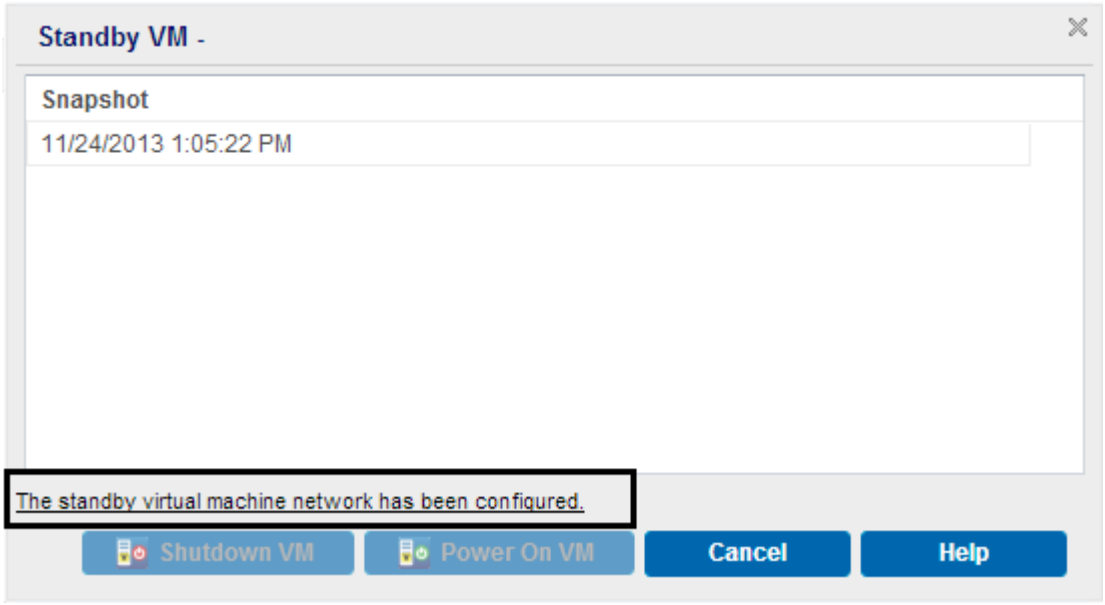
While powering on virtual machines, virtual standby determines the quantity of NICs (network interface cards) to power on based on whether the standby virtual machine network is configured. The following table illustrates how virtual standby determines the quantity of NICs that are required to power on standby virtual machines:

Values Defined in the Plan for VM Network	The Power on the standby virtual machine with customized network configurations option is not specified	The Power on the standby virtual machine with customized network configurations option is specified
The values defined are the same as the source machine.	Virtual standby powers on the quantity on NICs defined for the source machine as of the last backup job.	Virtual standby powers on the quantity NICs based on the larger of the following values: ■ The quantity defined under custom network configuration. ■ The quantity of NICs defined for the source machine as of the last backup job.
The values defined are custom values.	Virtual standby powers on the quantity of custom networks that are defined in the plan.	Virtual standby powers on the quantity NICs based on the larger of the following values: ■ The quantity defined under custom network configuration. ■ The quantity of NICs defined for the custom policy.

The following dialog (Edit Virtual Standby task of Modify a Plan) in the Virtual Standby task consist of custom configurations for NICs to power on:



The following dialog (Standby VM - <host_name>) illustrates the location where you specify the Power on the standby virtual machine with customized network configurations option:

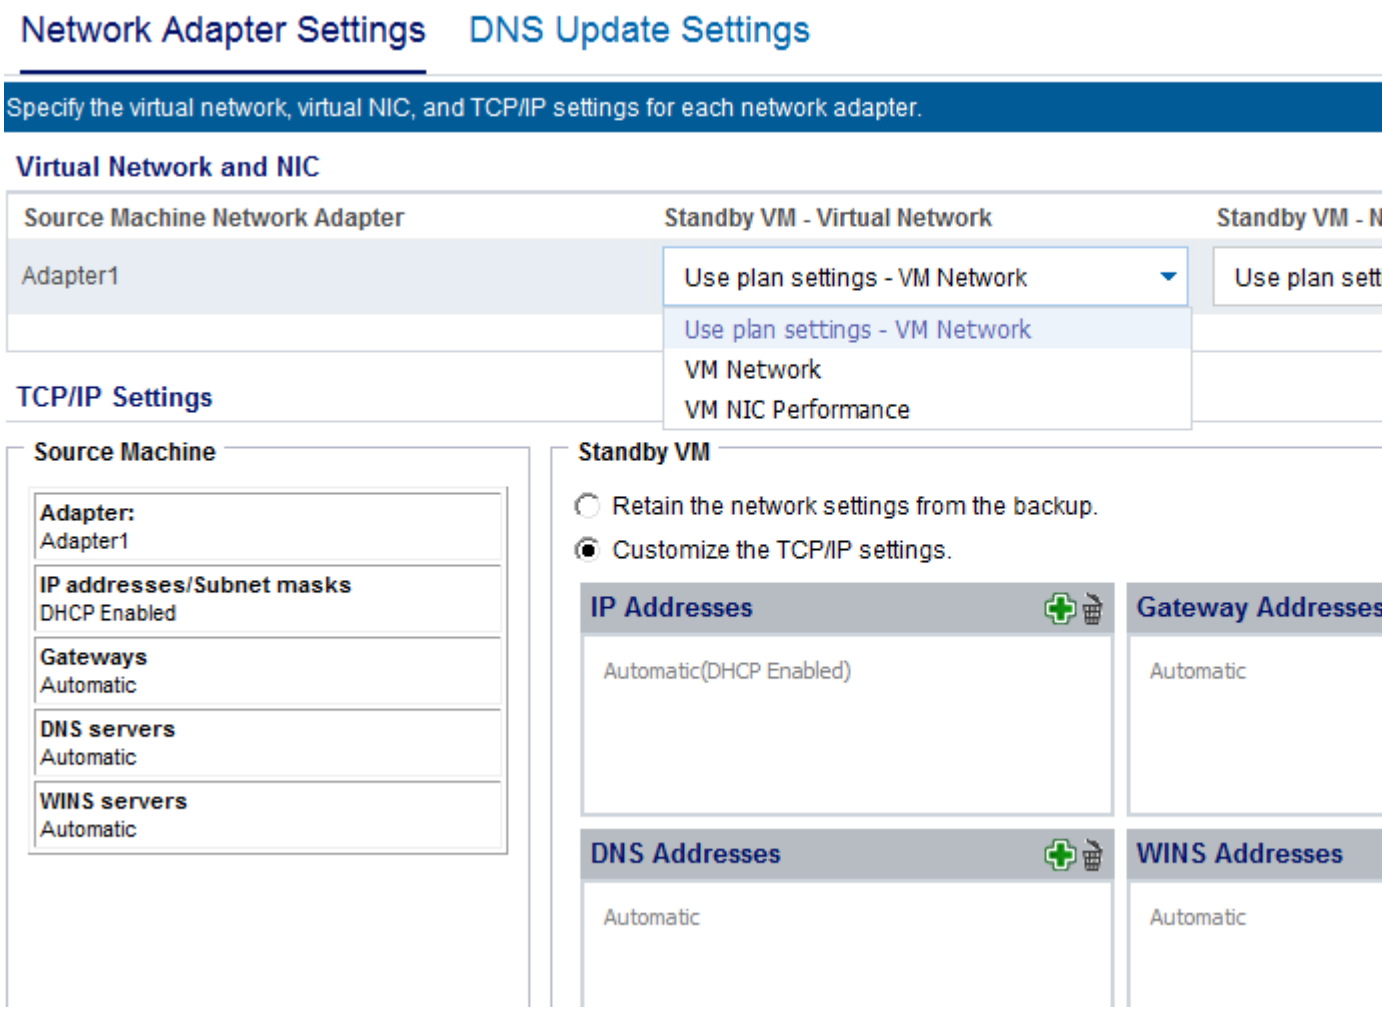


Configure the Standby VM Network

You can power on the Standby VM with customized network settings. You can configure the following network settings on the standby VM:

- Specify the virtual network and NIC (Network Interface Card), and TCP/IP settings for each network adapter from the **Network Adapter Settings** tab.
- Update the DNS servers to redirect clients from the source computer to the virtual standby virtual machines based on the TCP/IP settings from the **DNS Update Settings** tab.

The following diagram displays the **Network Adapter Settings** tab of **Standby VM Network Configuration**:



Follow these steps:

1. From the **resources** tab, navigate to the **Virtual Standby** node group.

   The Virtual Standby nodes are displayed on the center pane.

2. On the center pane, select the node and click **Standby VM Network Configuration**.

   The Standby VM Network Configuration - <node name> page opens.

3. On the **Network Adapter Settings** tab, select the virtual network from the **Standby VM - Virtual Network** list.
4. Select the NIC type from the **Standby VM - NIC Type** list.
5. Select **Customize the TCP/IP settings**.
6. Click the **Add address** button and add **IP Addresses**, **Gateway Addresses**, **DNS Addresses**, and **WINS Addresses**.

   Note: If you add **DNS Addresses**, then configure the DNS servers in the **DNS Update Settings** tab.

7. Click **Save**.

   The Standby VM Network Configuration - <node name> page closes.

   The Standby VM network is configured.

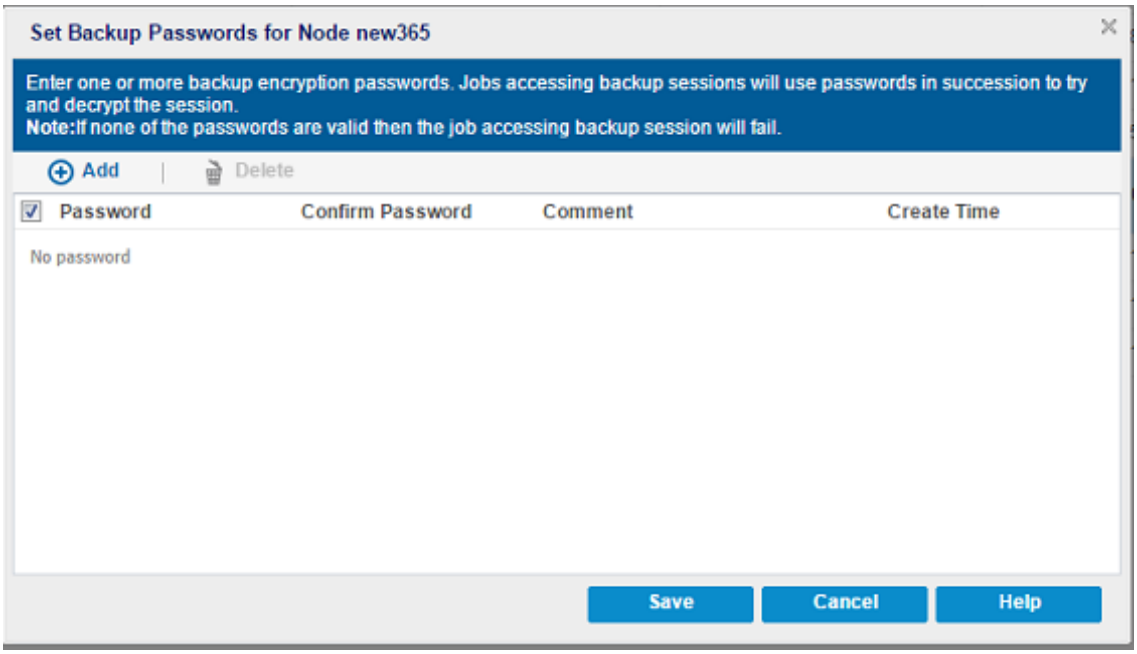
Set Backup Passwords for One or More Nodes

To ensure that the converter can convert the replicated recovery points, virtual standby lets you specify backup passwords for the data that the converter can use to convert the data.

Follow these steps:

1. On the Console, click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and click **All Nodes**.
3. From the center pane, right-click the node and click **Set Backup Passwords**.

   The **Set Backup Passwords for Node** dialog opens.

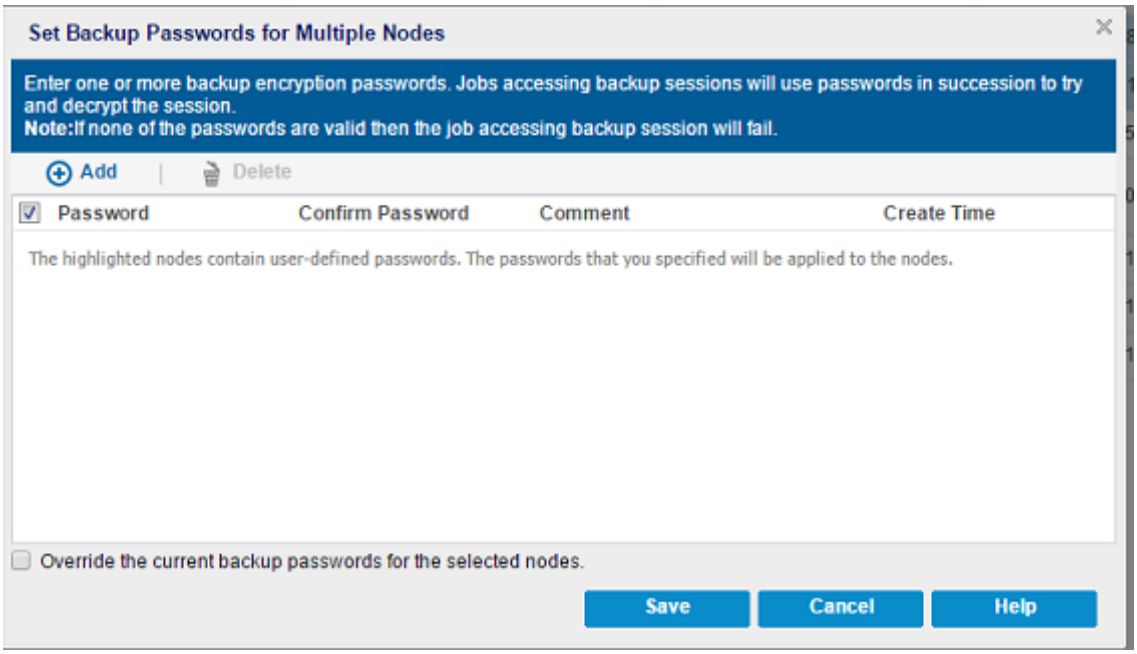


4. Perform the following tasks in the **Set Backup Passwords** dialog for one or more nodes:

   Add--Click **Add** to add one or more backup passwords to the selected nodes.

   Delete--Click **Delete** to delete one or more backup passwords from the selected nodes.

   Note: For multiple nodes, you can override the current backup passwords for multiple nodes by selecting the **Override the current backup passwords** for the selected nodes check box.



5. Click **Save**.

   The dialog closes and the backup passwords are set for the selected remote nodes.

(Optional) Run the Virtual Standby Job Manually

To manually run a virtual standby job, you have to first perform a manual backup. The virtual standby task is associated with a backup task. If a plan includes a backup task and a virtual standby task and you manually run the backup job, the virtual standby job runs automatically after the completion of the backup job.

Follow these steps:

1. Click the **Resources** tab.
2. From the left pane, navigate to **Nodes**, and click **All Nodes**.

   Displays the plans that you added.
3. Select the nodes that you want to backup. The selected node must have assigned a plan.
4. On the center pane, click **Actions**, **Backup Now**.

   The **Run a backup now** dialog opens.
5. Select the backup type and provide a name for the backup job.
6. Click **OK**.

   The backup job runs.

   The virtual standby job runs immediately after the backup job is over.

   The virtual standby job is manually run.

Pause and Resume Heartbeat

The Arcserve Unified Data Protection solution lets you pause and resume the heartbeats that are detected by the monitoring server. The heartbeat is the process where the source server and monitoring server communicate about the health of the source server. If the monitoring server does not detect a heartbeat after a specified length of time, the virtual standby feature provisions the virtual machine to function as the source node.

Examples: When to Pause or Resume Heartbeats

The following examples describe when to pause and resume heartbeats:

- Pause the heartbeat when you want to offline a node (source server) for maintenance.
- Resume the heartbeat after the maintenance tasks are complete and the node (source server) is online.

Be aware of the following behavior:

- You can pause and resume heartbeats at the group level or at the individual node level.
- You can pause and resume heartbeats for one or more nodes in one step.
- The Arcserve Unified Data Protection solution does not power on recover point snapshots while the heartbeat is in a paused state.
- When you upgrade the agent installations on source nodes, Arcserve Unified Data Protection pauses the heartbeat for the nodes. To help ensure that monitor servers monitor the upgraded nodes, resume the heartbeat for the nodes after you complete the upgrades on the nodes.

Follow these steps:

1. Log in to Arcserve Unified Data Protection.
2. Click the **resources** tab.
3. From the left pane, navigate to **Virtual Standby** and click **All Nodes**.

   If you have added any nodes, then the nodes will be displayed in the center pane.
4. Select the node that you want to pause or resume.
5. On the center pane, click **Actions**, **Heartbeat**, **Pause** or **Resume**.

   The heartbeat of the selected node is paused or resumed.

Pause and Resume Virtual Standby Job

Virtual conversion is the process where virtual standby converts the Arcserve Unified Data Protection recovery points from source nodes to virtual machine formats named recovery point snapshots. In the event a source node fails, the virtual standby feature uses the recovery point snapshots to power on a virtual machine for the source node.

As a best practice, allow the virtual conversion process to operate continuously. However, if you want to pause the virtual conversion process on local and remote virtual standby servers temporarily, you can do so from the Console. After you correct the problems on the source node, you can resume the virtual conversion process.

When you pause virtual standby jobs (conversion jobs), the pause operation does not pause the conversion job that is currently in progress. The pause operation applies to only the job that is expected to run at the end of the next backup job. As a result, the next conversion job does not start until you explicitly resume the (paused) conversion job.

If you resume virtual standby for nodes and if there are multiple backup sessions without recovery point snapshot, you will get a dialog to select the smart copy option. If you click Yes, virtual standby will convert the combined session into a single recovery point snapshot. If you click No, virtual standby will convert each session individually

Note: Optionally, you can pause and resume virtual standby jobs directly from the nodes. For more information, see Pause and Resume Virtual Standby Jobs from the Nodes.

Follow these steps:

1. Log in to Arcserve Unified Data Protection.
2. Click the **resources** tab.
3. From the left pane, navigate to **Virtual Standby** and click **All Nodes**.

   If you have added any nodes, then the nodes will be displayed in the center pane.
4. Select the node that you want to pause or resume.
5. On the center pane, click **Actions**, **Virtual Standby**, **Pause** or **Resume**.

   The virtual standby function for the selected node is paused or resumed.

Verify the Plan

To verify your Virtual Standby feature, confirm that you have successfully created the Virtual Standby plan. After you verify that the plan is created successfully, check whether the backup job is running as scheduled. After the backup job successfully completes, the Virtual Standby job runs. You can check the status of the backup job and virtual standby job from the **jobs** tab.

Follow these steps to verify plans:

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and click **All Nodes**.

   A list of all nodes is displayed on the center pane.
3. Verify that plans are mapped with nodes.

Follow these to verify Virtual Standby jobs:

1. Click the **jobs** tab.
2. From the left pane, click **All Jobs**.

   The status of each job is listed on the center pane.
3. Verify that the backup job and Virtual Standby job is successful.

   The plan for virtual standby is successfully verified.

   The virtual standby machine is created.

Applying Best Practices

Exclude Files from Antivirus Scanning

Exclude Files from Antivirus Scanning

Antivirus software can interfere with the smooth running of virtual standby process by either temporarily blocking access to files or by quarantining or deleting files that are incorrectly classified as suspicious or dangerous. You can configure most antivirus software to exclude specific files, or folders so that you can skip scanning certain data. It is important to configure your antivirus software properly so that it does not interfere with backup and restore operations, or any other types of processes.

In a Hyper-V server, the antivirus software corrupts the VM configuration file. The Hyper-V server changes the VM state to 'save' mode and the VM becomes corrupted and useless. In such cases, you have to delete the VM and perform a full conversion to create a new VM.

To ensure that the local and remote virtual standby works properly and to avoid the VM from entering the save mode, exclude the following files that target Hyper-V virtual machines:

- Virtual machine configuration files directory:

  (Default) C:\ProgramData\Microsoft\Windows\Hyper-V

  Arcserve UDP Virtual Standby virtual machine configuration files directory

- Virtual machine virtual hard disk files directory:

  (Default) C:\Users\Public\Documents\Hyper-V\Virtual Hard Disks

  Arcserve UDP Virtual Standby virtual machine virtual hard disk files directory

- Snapshot files directory:

  (Default) %systemdrive%\ProgramData\Microsoft\Windows\Hyper-V\Snapshots

  Arcserve UDP Virtual Standby virtual machine snapshot files directory

How to Create an Assured Recovery Plan

To verify accessibility and assure recovery of the data, you need to create an assured recovery plan. The plan for assured recovery is based on a backup/replication plan. This recovery task lets you add an assured recovery task to an existing backup or replication plan. The assured recovery task comprises of a source, test settings, schedule, and advanced settings. You can also run an Assured Recovery job manually.

What To Do Next?

1. Review the Prerequisites and Considerations
2. Add an Assured Recovery Task to a Plan
3. (Optional) Run the Assured Recovery Test Job Manually

Review the Prerequisites and Considerations

Verify that you have completed the following prerequisite tasks:

- Log into the Console.
- Installed the server component and created data stores, if you want to store the backup data to recovery point servers.
- You have a valid recovery point to create an Instant Virtual Machine or Instant Virtual Disk. You can select the recovery points from one of the following tasks:
  - Backup, Agent-based Windows
  - Backup, Host-Based Agentless
  - Replicate
  - Backup, Office 365 Exchange Online
  - Backup, Office 365 OneDrive
  - Backup, Office 365 SharePoint Online
  - Backup, Files on UNC or NFS Path
  - Backup, Agent-Based Linux
- Verify that the Arcserve UDP Agents are already installed on the Proxy Server.
- Verify that the operating system of the Proxy Server is 64-bit Windows Server 2008 R2 or higher version.

  Note: For UNC path backup plan and Office 365 (Exchange Online, OneDrive, SharePoint Online) backup plan, the Assured Recovery proxy server should be Windows 2012 or higher version.
- Verify that the Proxy Server has enough space for the Instant Virtual Machine or Instant Virtual Disk.

  Note: The necessary space highly depends on the RAM size that you configured in AR task for AR IVM test type. For each IVM, the hypervisor needs maximum to same size of the RAM size to hold the temporary data in memory swap files. For example, if you have 5 nodes in the same plan with AR-IVM task, and the VM memory size is set to 4 GB, then you need at least 5*4 = 20 GB free size to hold the memory swap files. In addition, you may need free size of 10 MB to hold the VM configuration files.
- Verify that the Windows operating system version on the Assured Recovery proxy server is the same or above the Windows operating system version on the source VM.

Add an Assured Recovery Task to the Plan

An assured recovery task includes an assured recovery task to an existing backup/replicate plan. Each task consists of parameters that define the source, task settings, schedule, and advanced settings. Create an Assured Recovery task based on backup/replication task to verify accessibility and assured recovery of the data and provide the data integrity check.

Follow these steps:

1. Click the **resources** tab on the Console.
2. From the left pane, navigate to **Plans**, and click **All Plans**.

   If you have created plans earlier, those plans are displayed on the center pane.
3. From the center pane, open an existing backup or replication plan.
4. In the selected plan, click **Add a Task** from the left pane.
5. (Optional) Select the **Pause this plan** check box.

   The plan will not run until you clear the check box to resume the plan.

   Note: If you pause the plan, the Assured Recovery job will not start. When you resume the plan again, the Assured Recovery job is not resumed automatically. You have to manually run another backup job to start the Assured Recovery job.
6. From the **Task Type** drop-down list, select **Assured Recovery Test**.
7. Now specify the Source, Task settings (IVM or IVD), Schedule, and Advanced details.
8. Click **Save**

   The changes are saved and the Assured Recovery task is automatically deployed.

Specify the Source

The Source tab lets you specify the source recovery point that you want to protect.

Follow these steps:

1. Click the **Source** tab.
2. Click the drop-down list of Recovery Point Source for Assured Recovery Test, and select the desired option.
3. Select Types of Recovery Points for Assured Recovery test. You can select either of the following two options:
   - Assured Recovery Test from selected backup types(s)

     You can select either Daily Backups, Weekly Backups, Monthly Backups. The daily, weekly, or monthly options are enabled depending on the schedule of the source backup. For example, if the source backup has only monthly backup scheduled, then the Assured Recovery test from selected backup types(s) option has only Monthly Backups enabled.
   - Assured Recovery Test from the Latest Recovery Points

   The source is specified

Specify the Task Setting for Test Type – Instant Virtual Machine

Selecting Instant Virtual Machine (IVM) as Test Type lets you start the virtual machine from the recovery point and optionally executes the custom script. The recovery point is considered in good state if the VM boots successfully and the verification script returns successful result within the user-specified time interval.

Note: This step is not required if you are selecting **Instant Virtual Disk** as Test Type.

You can select a Hypervisor Type from VMware vSphere and Microsoft Hyper-V, and enter details of Proxy Settings, and VM Settings to specify the task setting for IVM.

Follow these steps:

1. Select **Instant Virtual Machine** as Test Type.

   Note: When the source task is Agent-Based Linux backup/replication, the Instant Virtual Machine option is the only supported Test Type.

2. Select one of the Hypervisor options, and enter related details:

   VMware vSphere

   Note: If AR task setting is IVM to vSphere, the needed VMware license is similar to the one that Instant VM feature needed.

   a. Perform one of the following options to provide a vCenter/ESX(i) Server:

      - If you have already added a VMware node to the selected Site in the Console, select the node from the vCenter/ESX(i) Server drop-down list.
      - If you have not added any VMware nodes, then click **Add**.

        The **Specify the VM Destination** dialog opens.
      - Specify the virtual machine details and click OK.

        All the ESX(i) or resource pools are displayed on the central pane of the VM Location page

   b. Select one of the options from the ESX(i), cluster, resource pool, or virtual App as the location.

      The VMware vSphere machine is specified.

      Note: Specify the data store of ESX(i) server if the source task is Agent-Based Linux backup/replication or Host-Based Agentless backup/replication.

   Specify Microsoft Hyper-V

Perform one of the following options to provide Microsoft Hyper-V:

- If you have already added a Hyper-V node to the selected Site in the Console, select the node from the Hyper-V Server/Cluster drop-down list.
- If you have not added any Hyper-V nodes, then click **Add**

  The Specify the VM Destination dialog opens.

  Note: When you connect to the Hyper-V Instant VM using a local non-built-in administrator account, the remote UAC needs to be disabled. For more information on how to disable the remote UAC for non-built-in administrator, see How to disable a remote UAC for a non-built-in administrator.

- Specify the Hyper-V server details and click **OK**

The Hyper-V virtual machine is specified.

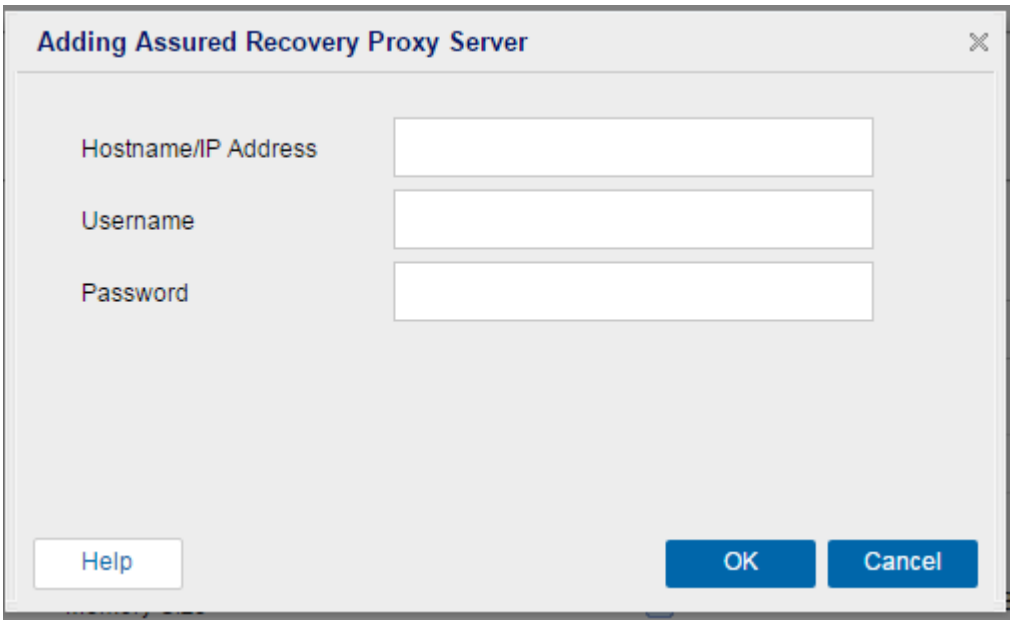
Specify Nutanix AHV

This is only for Linux VM support. For more information, see Protecting an Instant Virtual Machine on Nutanix AHV for Linux Node.

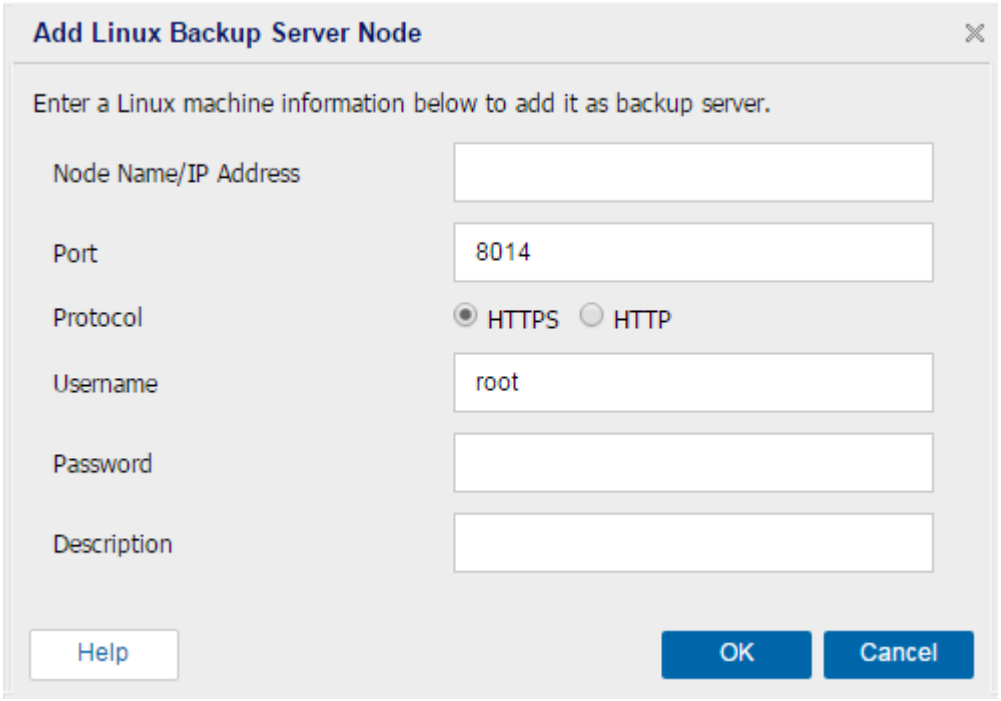
3. Perform one of the following options to provide **Proxy Server**:
   - If you have already added a proxy node to the selected Site in the Console, select the node from the proxy server drop-down list.
   - If you have not added any proxy nodes, then click **Add**.

     The **Adding Assured Recovery Proxy Server** dialog opens.

     For Windows Proxy Server



For Linux Proxy Server



- Specify the Proxy server details
- Click **OK**.

The nodes are displayed on the **Available Nodes** area.

Notes:

- You have to select a Proxy Server only when the hypervisor is VMware vSphere.

- When the hypervisor is VMware vSphere, you must get the Windows Network File System (NFS) role installed on the Recovery Point Server. Instant VM process automatically installs the NFS. To manually install the Network File System, see How to manually install Network File System on a Windows Server.
- Select Proxy server(s) based on selected node type(s) included in the plan. For example, Windows proxy server for Windows node and Linux Proxy Server for Linux node.

4. (Optional) Perform the following steps to specify Gateway Settings for Assured Recovery.

   Note: Applicable only to Linux.

   a. Using DHCP settings or specify IP Address, Mask and Default Gateway manually.

   b. Select a Virtual Network from drop-down list.

   Note: Gateway Settings is required only when the Assured Recovery Test source is Agent-Based Linux backup/replication or Host-Based Agentless backup/replication.

5. Specify the details for the Assured Recovery **VM Settings**.

   VM Name Prefix

   Specifies the VM Name Prefix. The name of the source node with a prefix is the default name of the Instant VM. Some special characters are not allowed in the name, such as '@', '\' and so on.

   Default value: UDPARVM_

   VM Files Folder

   Specify the folder location of the Assured Recovery VM on the proxy server. You can browse the volume information of the proxy server.

   CPU Count

   Specifies the number of CPU that you would require in the Assured Recovery VM.

   Memory Size

   Specifies the size of memory that you would require in the Assured Recovery VM.

   Network Settings

   Specifies the Network Settings in the Assured Recovery VM. You can select "Connect Assured Recovery VM to Network" or not. You can also use DHCP

settings for Assured Recovery VM or use TCP/IP settings from the backup session.

Adapter Type

Specifies the Adapter Type in the Assured Recovery VM. The available Adapter Type may vary depending on the hypervisor.

Specify the Task Settings For Test Type – Instant Virtual Disk

Selecting Instant Virtual Disk as Test Type for Test Settings lets you mount the recovery point as local disk, verifies integrity of volume/file system, and optionally executes the custom script. The recovery point is considered as in good state if you successfully mount the recovery point and all the tests return successful result.

Note:

- This step is not required if you are selecting **Instant Virtual Machine** as Test Type.
- The **Instant Virtual Disk** option does not support non-windows node.

Follow these steps:

1. Perform one of the following steps to provide server for **Windows Proxy Server:**
   a. Select one of the available nodes from the drop-down list.

      Note: If you have already added a proxy node to the selected Site in the Console, you can view list of nodes from the **Windows Proxy Server** drop-down list.

   b. If you have not added any proxy nodes, perform the following steps:
      i. Click **Add**.
      ii. The **Adding Assured Recovery Proxy Server** dialog opens.
      iii. Specify the proxy server details and click **OK**.
2. **Browse** the folder location of the Virtual Hard Disk on the proxy server.

Specify Assured Recovery Test Job Schedule

The schedule tab lets you specify the assured recovery test job schedule. If you do not specify a schedule, then the task starts immediately when you complete the primary task.

Follow these steps:

1. Click the **Schedule** tab.
2. Click **Add**, and then **Add Assured Recovery Test Job Schedule**.

   Add Assured Recovery Test Job Schedule dialog is displayed.
3. Specify the schedule and click **Save**.

   The dialog closes and schedule tab displays the specified.

   Notes:

   - If Assured Recovery source is Latest Recovery Points + Enable Assured Recovery schedule, the default value is 1. As a result, by default the latest one is verified.

     For value <=0, all unverified recovery points are tested.
   - If Assured Recovery source is Daily/Weekly/Monthly + Enable Assured Recovery schedule, the default value is 9 (7 daily + 1 weekly + 1 monthly).

     For value <=0, all unverified recovery points are tested.
   - If you do not want to use the default value, you can add a DWORD 32bit registry key at the following location:

     HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AssuredRecovery\MaxNumberOfRecoveryPoint

Specify the Advanced Settings

The Advanced tab helps you apply other settings required to complete the plan. This is optional, if you do not need to add further settings. Some of the options that you view on the **Advanced** tab vary upon your selection in the previous tab. For example, the Heartbeat Properties option is visible only if you have selected Instant Virtual Machine as your Test Type in **Test Settings**.

Heartbeat Properties

Note: The option is visible only if you have selected Instant Virtual Machine as your Test Type.

Timeout

Specifies the longest time that the proxy server can wait for the heartbeat of the Assured Recovery VM. If Assured Recovery VM is not available longer than the defined timeout, the Assured Recovery test job fails.

Frequency

Specifies the frequency with which the proxy server checks the heartbeat status of Assured Recovery VM.

Note: This options is not visible if you have selected Agent-Based Linux backup/replication as Recovery Point Source of Assured Recovery Test.

Check Points

Note: The option is visible only if you have selected Instant Virtual Disk as your Test Type.

Verify file system

Verifies if the file system and size of volume are similar to what the source machine has.

Run check disk command

Checks the data integrity by running check disk command on the volumes exposed by instant virtual disk.

Custom Command Location on Proxy

Note: This options is not visible if you have selected Agent-Based Linux backup/replication as Recovery Point Source of Assured Recovery Test.

Specifies the custom script that is stored on the proxy server. You can browse the volume information of the proxy server. If IVM assured recovery type is selected, the script is copied into VM and executed when VM boots. If IVHD assured recovery type is selected, the script is executed on the proxy server.

Note: The Assured Recovery job can only support the executable Windows bat script or executable binary. But, this limitation does not restrict your usage of Assured Recovery job. You can use a Windows batch script written to start your powershell script, such as the name test.bat with the following content:

Powershell.exe -NoProfile -ExecutionPolicy ByPass -Command "& '%ScriptPath%\AR-check.ps1'"

Additionally, the following arguments are passed to the script as per your requirement.

- set PlanName=%1%

- set NodeName=%2%

- set ProxyServer=%3%

- set RecoveryPointName=%3%

- set MountPointRootPathName=%5% (Applicable only to Instant Virtual Disk test type)

As a result, the bat is invoked and run as below:

For Instant Virtual Machine test type:

"%ScriptPath%\test.bat PlanName NodeName ProxyServer RecoveryPointName

For Instant Virtual Disk test type:

%ScriptPath%\test.bat PlanName NodeName ProxyServer RecoveryPointName MountPointRootPathName

Exit code

Specifies the exit code for Succeed Job or Fail Job.

Succeed Job

Specifies that the assured recovery job is set to succeed when the script returns the exit code.

Fail Job

Specifies that the assured recovery job is set to fail when the script returns the exit code.

Fail Job if it runs longer

Specifies if the script runs longer than the specified times.

Default value: 15 Minutes.

Linux Pre/Post Script Setting On Proxy

Note: This option is visible only if you have selected Agent-Based Linux backup/replication or Host-Based Agentless backup/replication as Recovery Point Source of Assured Recovery Test.

Run on Linux Backup Server after Assured Recovery job is over

Specifies the script located on Linux Proxy that runs after completion of the Assured Recovery job on Linux Backup Server.

Run on Assured Recovery VM after Assured Recovery VM is booted

Specifies the script located on Linux Proxy that runs after the Assured Recovery VM is booted on Assured Recovery VM.

Fail job if it runs longer

Specifies if the script runs longer than the specified times.

Email Alerts

Lets you enable email alerts. You can configure email settings and specify the type of alerts that you want to receive in an email. When you select this option, the following options are enabled for your selection.

Notes: Unavailable for jobs running with Linux Backup Server.

Email Settings

Lets you configure the email settings. Click **Email Settings** and in the **Email Settings** dialog configure the email server and proxy server details.

Set Backup Passwords for One or More Nodes

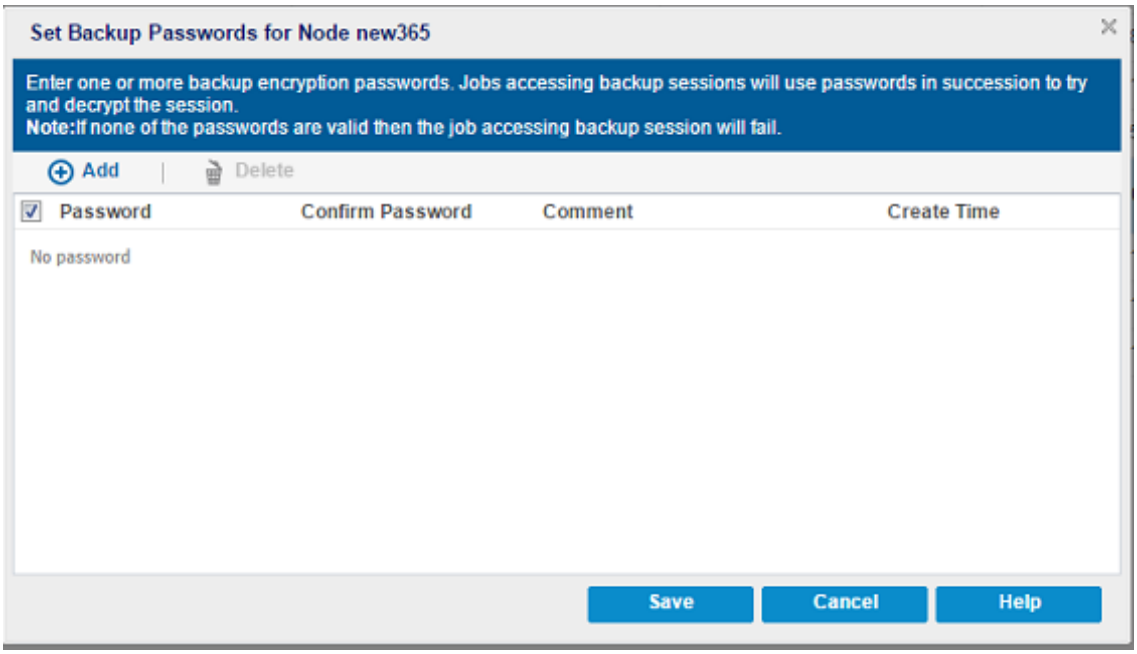
To ensure that the Assured Recovery job can test the replicated recovery points, Assured Recovery lets you specify backup passwords for the data. Jobs accessing backup sessions use passwords continuously to decrypt the session.

Note: If none of the passwords are valid then the job accessing the backup session fails.

Follow these steps:

1. On the Console, click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and click **All Nodes**.
3. From the center pane, right-click the node and click **Set Backup Passwords**.

   The **Set Backup Passwords for Node** dialog opens.

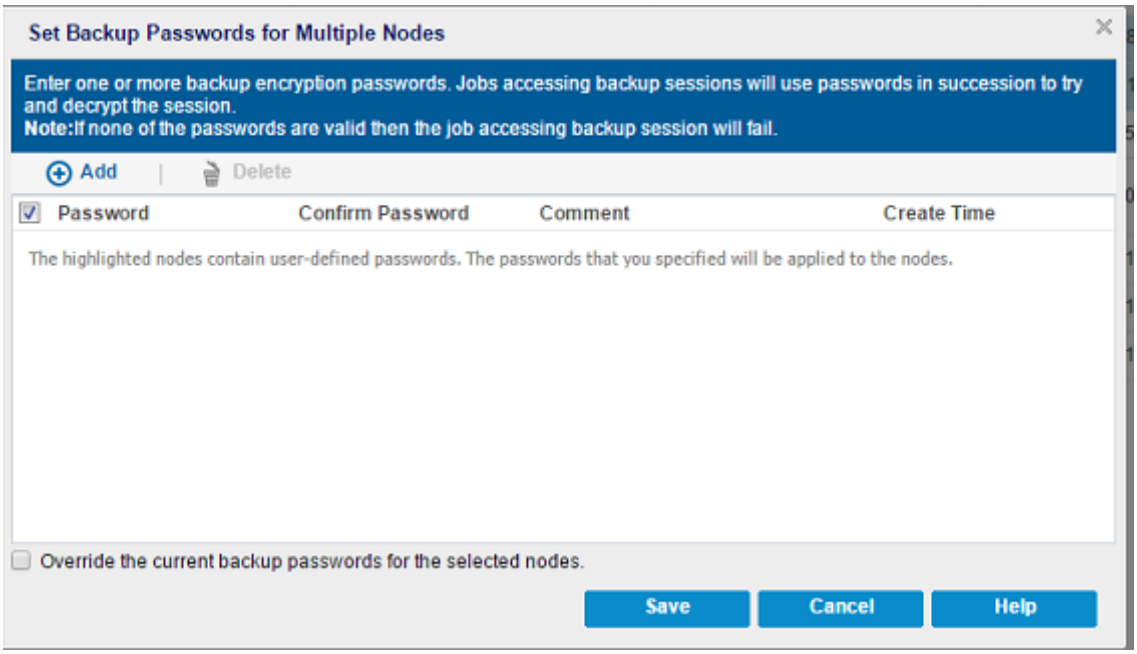


4. Perform the following tasks in the **Set Backup Passwords** dialog for one or more nodes:

   Add--Click **Add** to add one or more backup passwords to the selected nodes.

   Delete--Click **Delete** to delete one or more backup passwords from the selected nodes.

   Note: For multiple nodes, you can override the current backup passwords for multiple nodes by selecting the **Override the current backup passwords** for the selected nodes check box.



5. Click **Save**.

   The dialog closes and the backup passwords are set for the selected remote nodes.

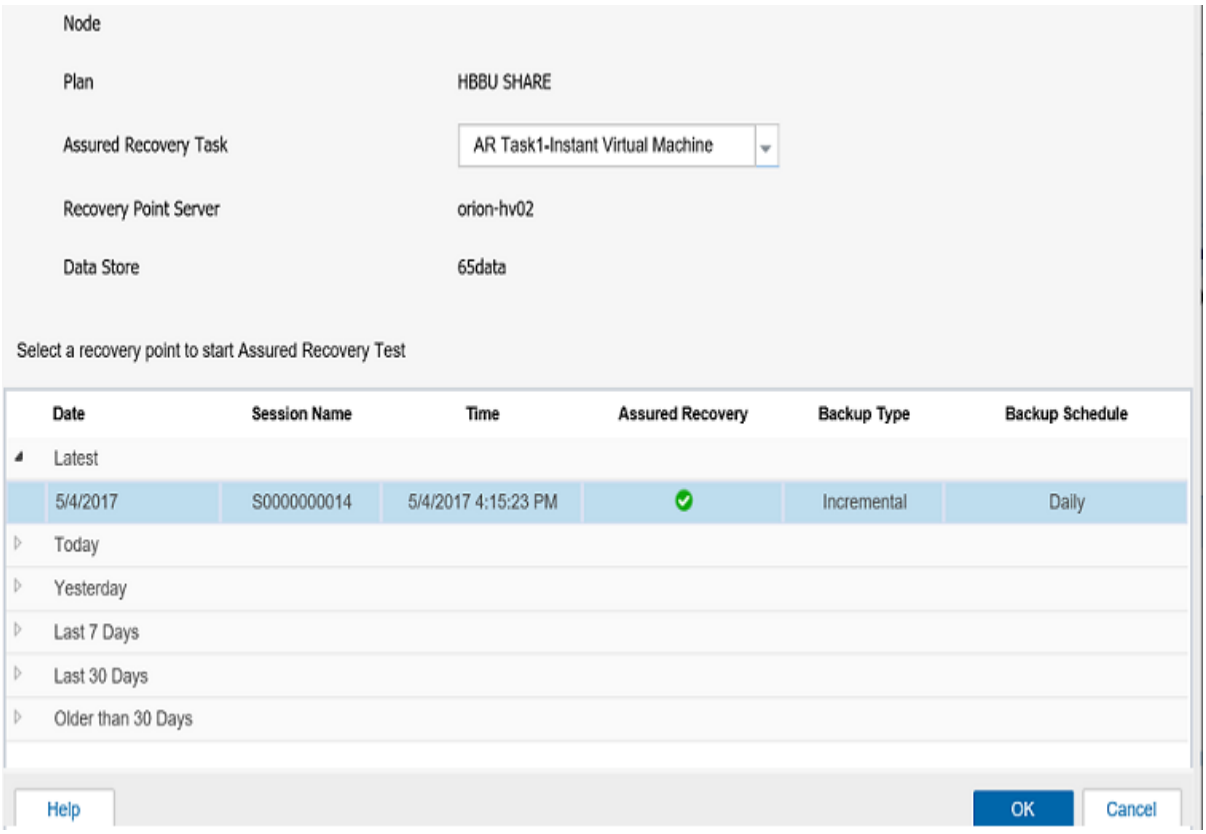
(Optional) Run the Assured Recovery Test Job Manually

In Arcserve UDP all plans, including Assured Recovery jobs, are performed automatically and are controlled by the schedule settings. For Assured Recovery jobs, besides the scheduled run, Arcserve UDP lets you perform manual test of nodes and plans for Assured Recovery Test. This topic provides separate procedures to perform a manual Assured Recovery Test of node and plan.

Follow these steps to perform a manual Assured Recovery Test of node:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to Nodes, and click **All Nodes**.
3. From the center pane, select the nodes that has an Assured Recovery plan assigned and click **Actions**.
4. From the options displayed for Actions, click **Run Assured Recovery Test Now**.

   The Assured Recovery dialog opens.



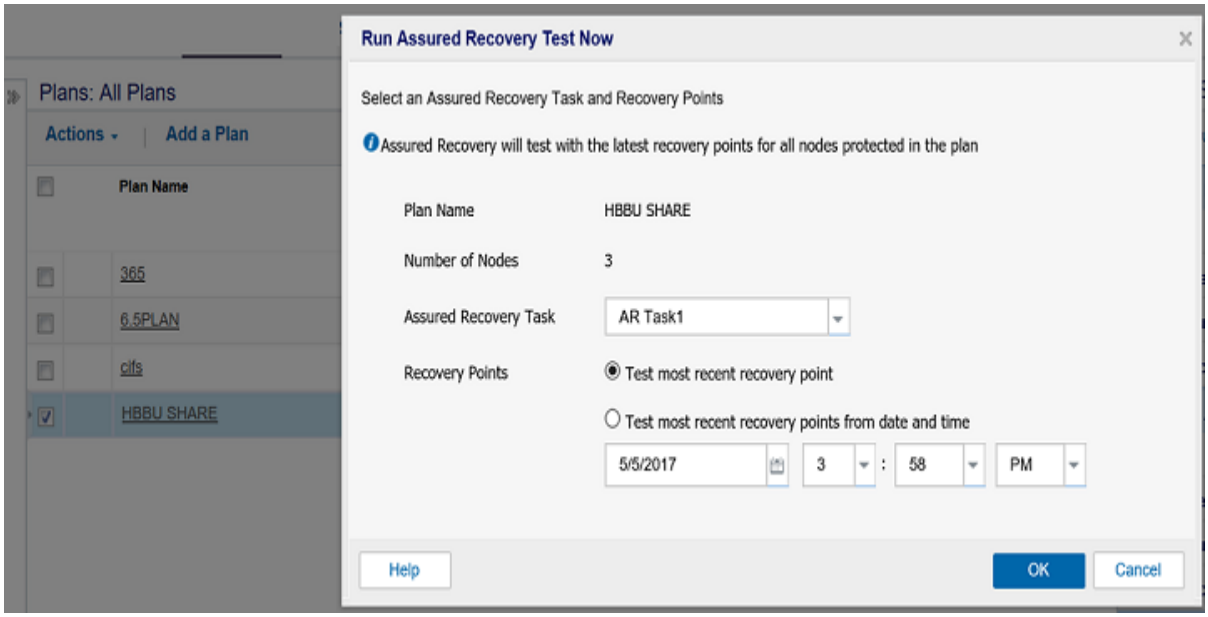
5. Select an Assured Recovery task and a recovery point, and click **OK**.

   The Assured recovery test of node runs.

   Follow these steps to perform a manual Assured Recovery Test of plan:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to Plans, and click **All Plans**.
3. From the center pane, select an Assured Recovery plan and click **Actions**.
4. From the options displayed for Actions, click **Run Assured Recovery Test Now**.

   The Assured Recovery dialog opens.



5. Select an Assured Recovery task and one of the options for recovery point, and click **OK**.

   The Assured recovery test of plan runs.

How to Create a UNC/NFS Path Backup Plan

To protect your directories and files located on the UNC/NFS path, you need to create a plan. The plan for UNC/NFS path consists of a backup task. This backup task lets you specify the nodes you want to protect, the backup destination, and the backup schedule. The backup destination is a recovery point server where you want to store your backup data. The destination is a remote share folder. For UNC/NFS path, the maximum backup size is 128 TB.

What To Do Next?

1. Review the Prerequisites and Considerations
2. Create a File Share
3. Create a Backup Plan
4. (Optional) Perform a Manual Backup
5. Create Other tasks on UNC/NFS Plan
6. Verify the Backup

Review the Prerequisites and Considerations

Verify that you have completed the following prerequisite tasks:

1. Log into the Console.
2. Prepare a UNC/NFS path backup proxy server where you have installed Arcserve Unified Data Protection Agent for Windows.
3. Have User credential with at least Read permission for the UNC/NFS path that you plan to protect.

   Note: To add, update, and delete a UNC/NFS Path node, see How to Add and Manage UNC/NFS Path.
4. Have a recovery point server with non-dedupe datastore or dedupe datastore.
5. Create datastore to store the backup data.
6. Create a file share on Nutanix Prism Web Portal.

Consideration:

- Extra license consumption for UNC paths. To resolve, view troubleshooting.
- Backup proxy is always a Windows machine.
- For the Nutanix NFS plan to work, NfsUtils needs to be pre-installed on the backup proxy machine (as is the case for the built-in NFS plan).
- For Nutanix File Server, if any of the FSVMs from the cluster fails, the other FSVM takes the primary ownership of File Server, and HA gets enabled for the FSVM cluster.

Create a File Share

The Share/Export page lets you create distributed or standard file shares. A distributed share or home share is a repository where the personal files of a user reside, whereas the standard share or general-purpose share is a repository shared by the group. A distributed share is distributed at the top-level directories whereas the standard share is on a single file server VM (FSVM). This section provides information about how to create a file share to perform the UNC/NFS backup for Nutanix Files.

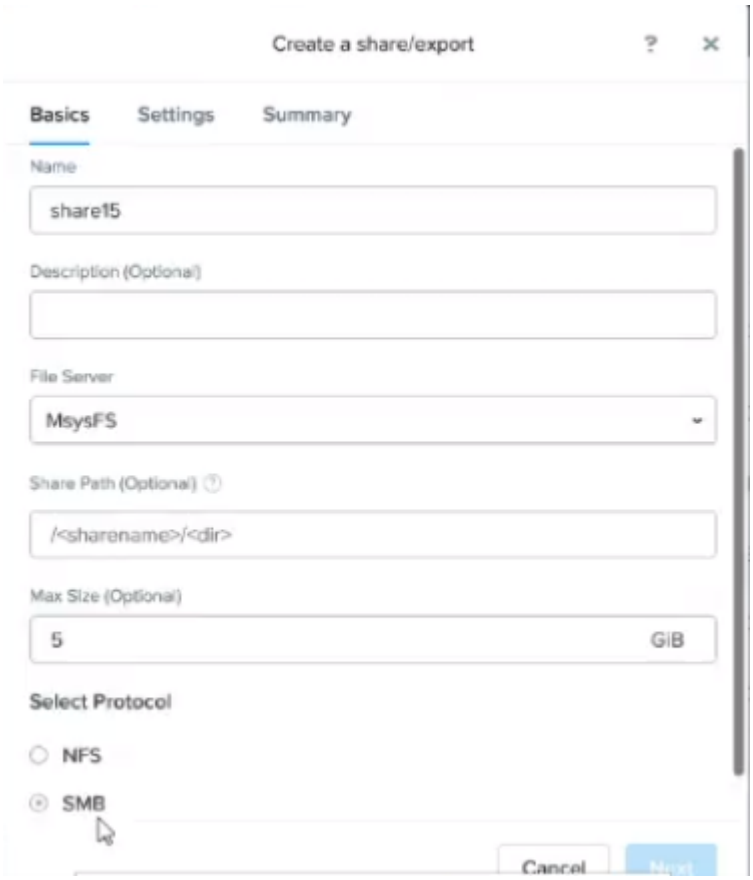
Follow these steps:

1. Log into Nutanix Prism web portal.
2. Navigate to **File Server** on the main menu, and then click **Share/Export** on the File Server dashboard.

   The Create a share/export window appears displaying the Basics tab.
3. On the Basics tab, do the following, and then click **Next**:
   - Name: Enter a name for the share.

     Note: The naming conventions for the share name are as follows:
     - Each name must be unique.
     - Names are not case-sensitive.
     - Files allow unicode characters.
     - Maximum name length is 80 characters.
     - A blank space or space character cannot appear as the first or last character in the name.
     - Names that end with the $ are hidden shares.
   - Description (Optional): Enter a description for the share if required.
   - File Server: From the drop-down list, select a file server in which the share resides.
   - Share Path (Optional): Provide the file path to the share.
   - Max Size (Optional): Enter the maximum share size. By default, the size is in GiB.
   - Select Protocol: Select **SMB** as a protocol.

     Note: To enable the multiprotocol feature, select the **Enable multiprotocol access for NFS** check box.



4. On the Settings tab, retain defaults or make changes as needed, and then click **Next**.

   Use "Distributed" share type instead of "Standard"

   Lets you create a distributed share.

   Note: By default, the shares on single-FSVM deployments are standard.

   Enable Self Service Restore

   Lets you restore files from previous snapshots.

   Enable File System Compression

   Lets you save space and reduce data on the share through in-line compression of written data.

   Note: Make sure to disable compression at the container level. If you enabled the compression, the Enable File System Compression option gets disabled.

   Enable Access Based Enumeration (ABE)

   Lets you enable Access-Based Enumeration (ABE).

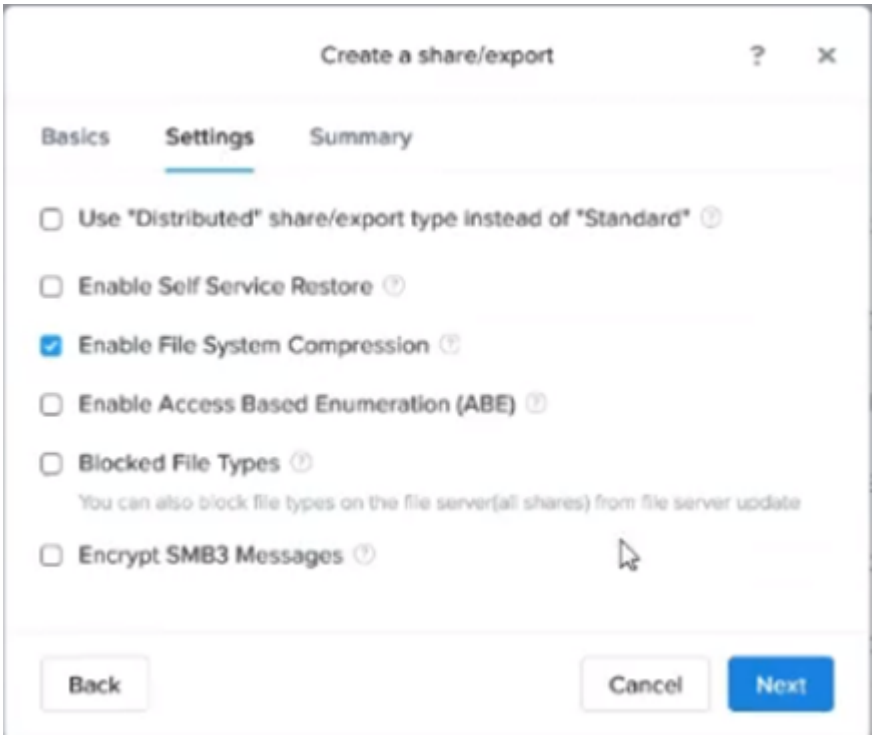
   Blocked File Types

Lets you specify the blocked file types.

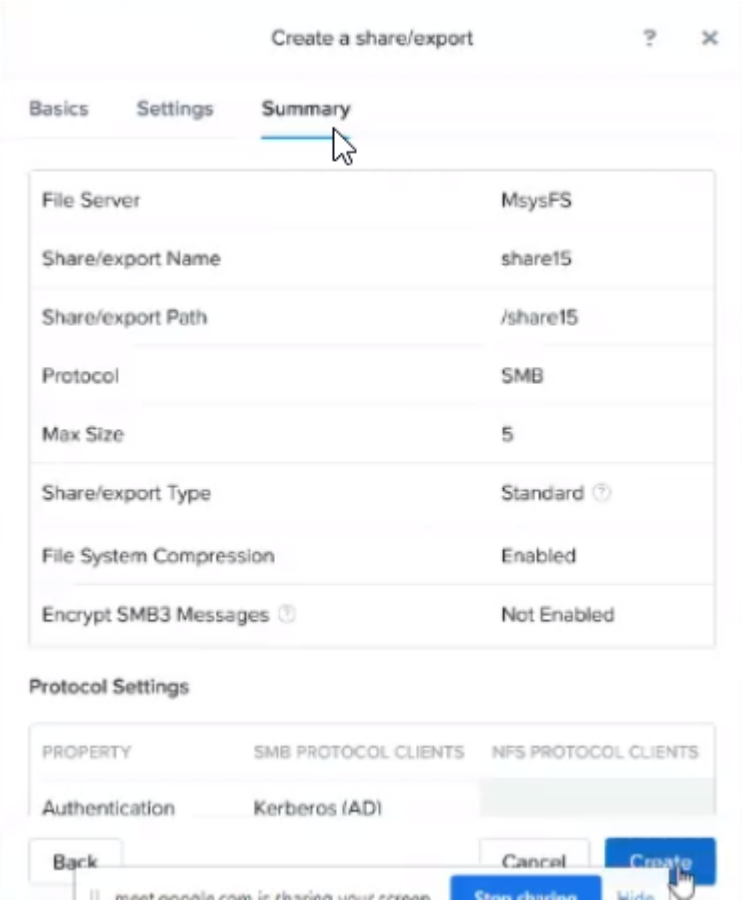
Enter a comma-separated list of file types that you want to block from the share.

Encrypt SMB3 Messages

Lets you enable message encryption between the file server and the client.



5. On the Summary page, verify the details, and then click **Create**.



Note: To make changes to the settings of the previous pages, click **Back**.

The file share is created successfully.

Note: Next, copy the mount path and then map the newly created share in the local UDP. For more information, see Step 6 (a1) of Specify the Source.

Create a Backup Plan with a UNC/NFS Path Task

A backup plan includes a backup task that performs a backup of a physical node and stores data to a specified destination. Each task consists of parameters that define the source, destination, schedule, and other backup details.

Follow these steps:

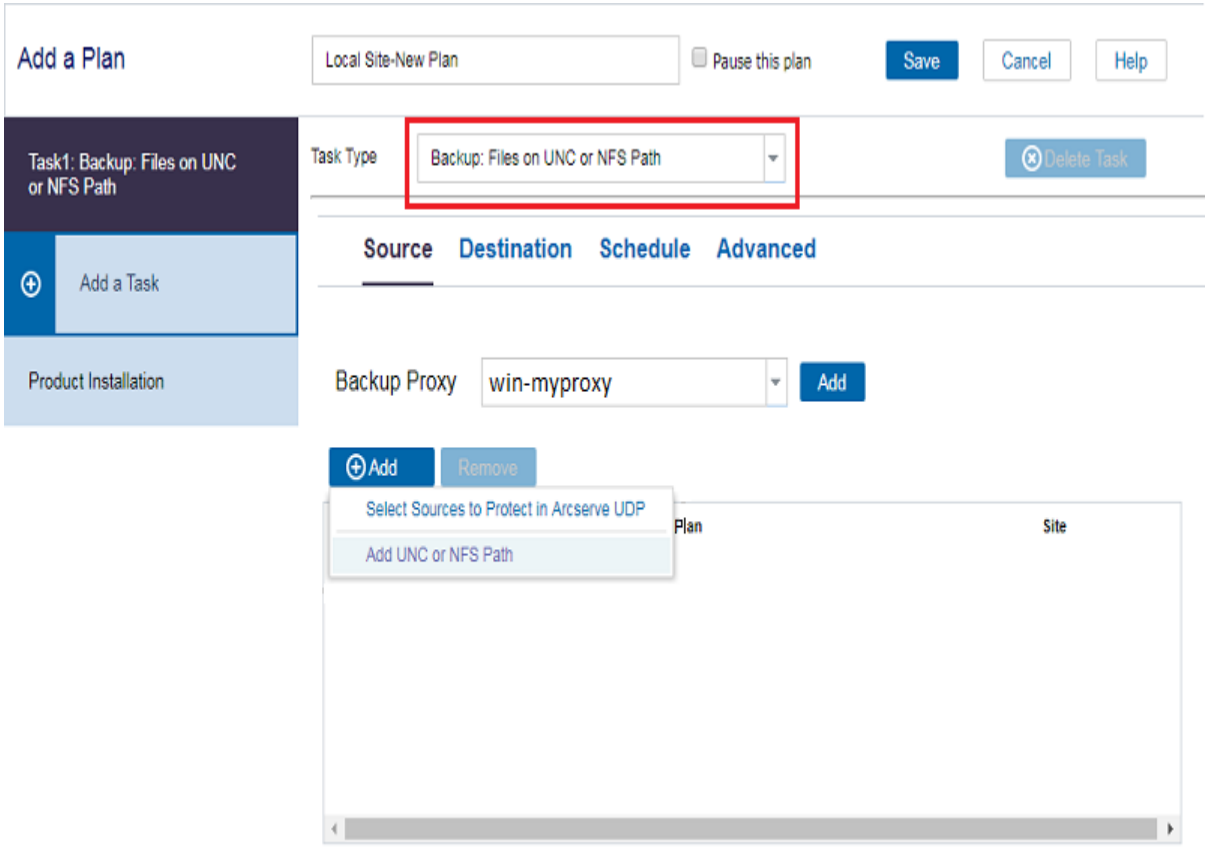
1. Log into Arcserve UDP Console.
2. Click the **resources** tab.
3. From the left pane, navigate to **Plans**, and then click **All Plans**.

   If you have created plans earlier, those plans are displayed on the center pane.
4. On the center pane, click **Add a Plan**.

   The Add a Plan page opens.
5. Enter a plan name.
6. (Optional) Select the **Pause this plan** check box.

   The plan will not run until you clear the check box to resume the plan.

   Note: If a plan is paused, then any in-progress job is not paused but all corresponding scheduled jobs associated with that plan are paused. However, you can manually run a job. For example, you can manually run backup job and replication job for a node even if the respective plan is paused. In such case, the following task to the on-demand (manual) job does not run. For example, there is a replication task after an on-demand backup job, the replication job does not run for the on-demand backup job. You need to manually run the replication job. When you resume the plan, the pending jobs do not resume immediately. After you resume the plan, the pending jobs run from the next scheduled time.
7. From the Task Type drop-down list, select **Backup: Files on UNC/NFS Path**.



Now specify the Source, Destination, Schedule, and Advanced details.

Specify the Source

The Source page lets you specify the source nodes that you want to protect. You can select more than one node in a plan. If you have not added any nodes to the Console, you can add nodes when you create or modify a plan from the Source page. You can also save a plan without adding any source nodes. The plan gets deployed only after you add source nodes.

Follow these steps:

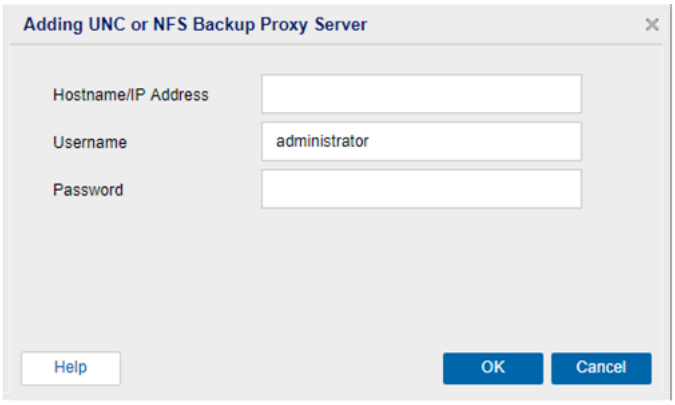
1. Click the **Source** tab.
2. If Backup Proxy is already added, select the backup proxy from the Backup Proxy drop-down list.

   The proxy server is a node where you install Arcserve Unified Data Protection Agent for Windows.

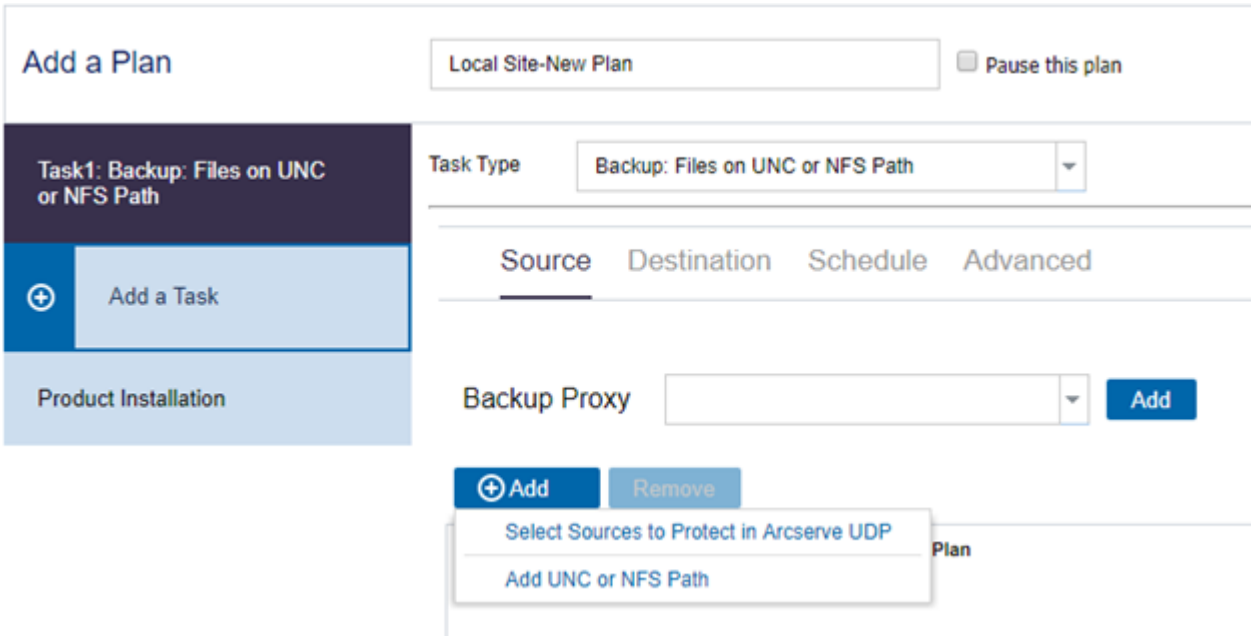


3. If the backup proxy is not added, perform the following steps:
   a. Click **Add**.

      The Adding UNC or NFS Backup Proxy Server dialog opens.



   b. Specify the proxy server details, and then click **OK**.
4. To add the UNC or NFS Path node, click **+ Add**, and then select one of the following:



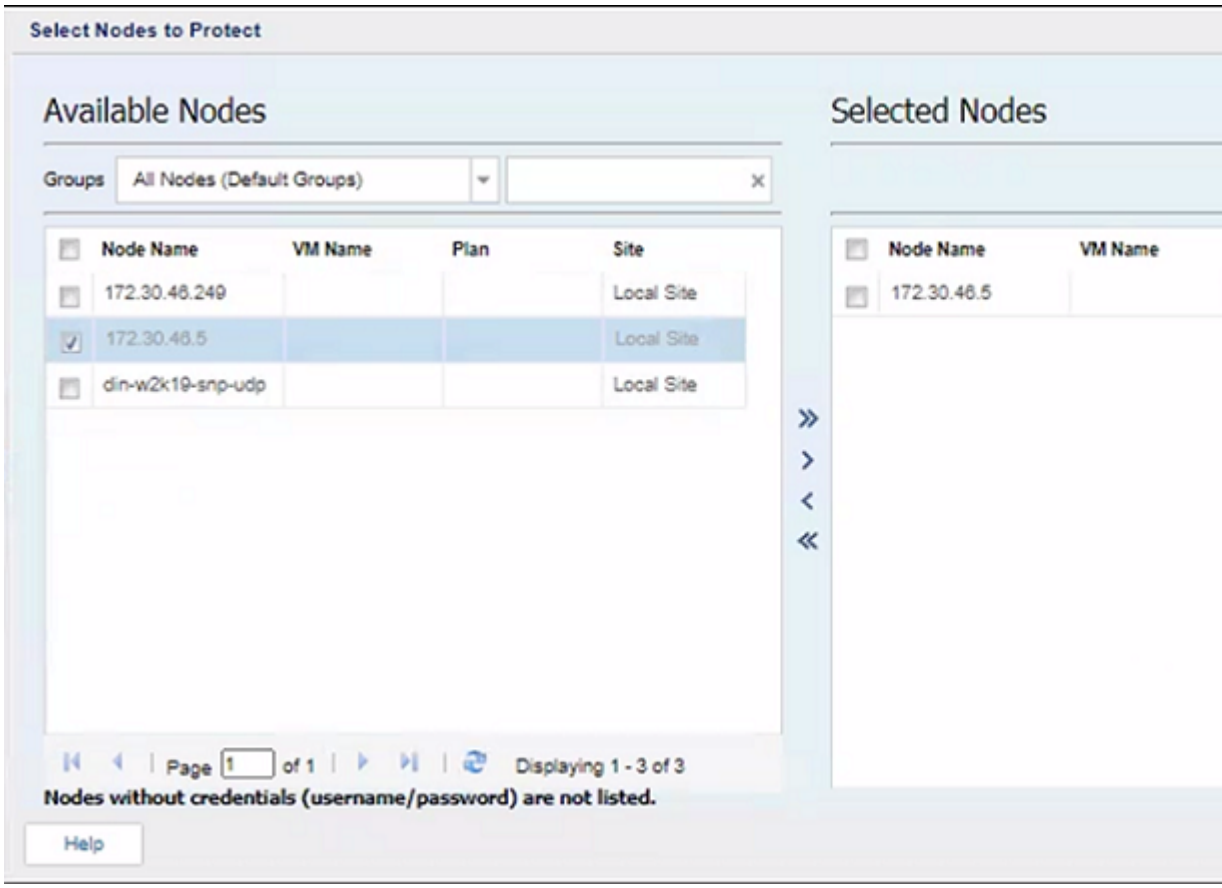
Select Sources to Protect in Arcserve UDP

Opens the Select Nodes to Protect dialog and you can select the nodes from the displayed list. Select this option if you have already added the nodes to the Console.

Add UNC or NFS Path

Opens the Add Nodes to Arcserve UDP Console dialog. Select this option if you have not added the nodes and want to manually add the nodes to protect.

5. If you select the **Select Sources to Protect in Arcserve UDP** option, perform the following steps:
   a. (Optional) Select a filter from the Groups drop-down list to filter nodes. You can enter keywords to further filter your nodes.



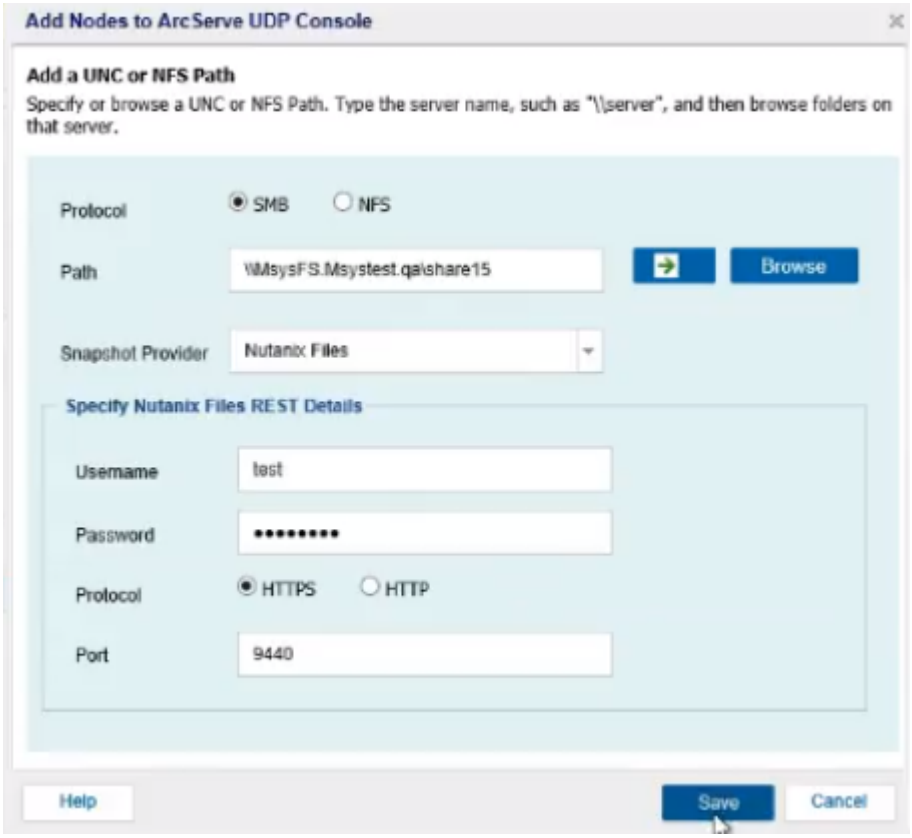
The nodes are displayed on the Available Nodes side.

b. Select the nodes from the Available Nodes side, and then click the **Add all nodes** (>>) or **Add selected nodes** (>) icon.

   The selected nodes are displayed on the Selected Nodes side.

c. Click **OK** to close the dialog.

6. If you select the **Add UNC or NFS Path** option, perform the following steps on the Add a UNC or NFS Path page:



a. Select **SMB** or **NFS** as a protocol.

- If you select **SMB**, do the following:

1. Enter the UNC Path in the following format:

   *Hostname**share*

   Note: To get the UNC path, do the following:

   a. Select the newly created share in the center pane.

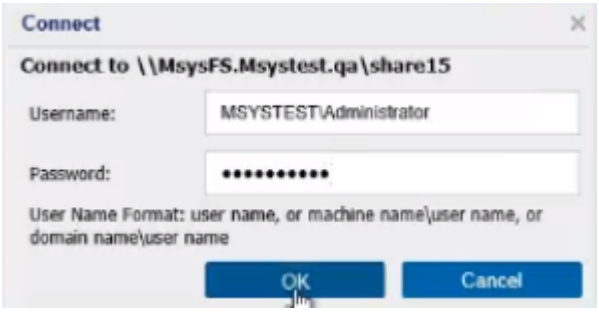
   b. Under the SHARE/EXPORT DETAILS column, from the Summary pane, navigate to **Mount Path** and then copy the mount path.

   If the UNC path is valid, a right arrow and the Browse option gets enabled.

2. Click the right arrow (>) to validate the UNC Path.

   The Connect dialog appears.

3. Enter the username and password as needed, and then click **OK**.



The Add Nodes to Arcserve UDP Console dialog displays the verified UNC Path.

- If you select NFS, do the following:

  1. Enter the UNC Path in the following format:

     *Hostname**share*

  2. Select Encoding from the available options - ANSI, GB2312-80, KSC5601, BIG5, SHIFT-JIS, EUC-KR, EUC-TW, EUC-JP.

     The default value is ANSI.

     Note: During the backup of an NFS Shared Folder, all such files/folders whose name has an unsupported language encoding are skipped or their name appears as junk.

b. From the Snapshot Provider drop-down list, select **Nutanix Files**.

c. For the Specify Nutanix Files REST Details box, specify the following credentials:

Username

Refers to the user name created by the REST API user using the Nutanix prism.

Password

Refers to password for the username.

Protocol

Specifies the protocol that you want to use to communicate with the destination server. The available options are HTTP and HTTPS.

Note: For a more secure communication, select the HTTPS protocol.

Port

Refers to the port number. The default port value is 9440 to access the REST API calls and Prism.

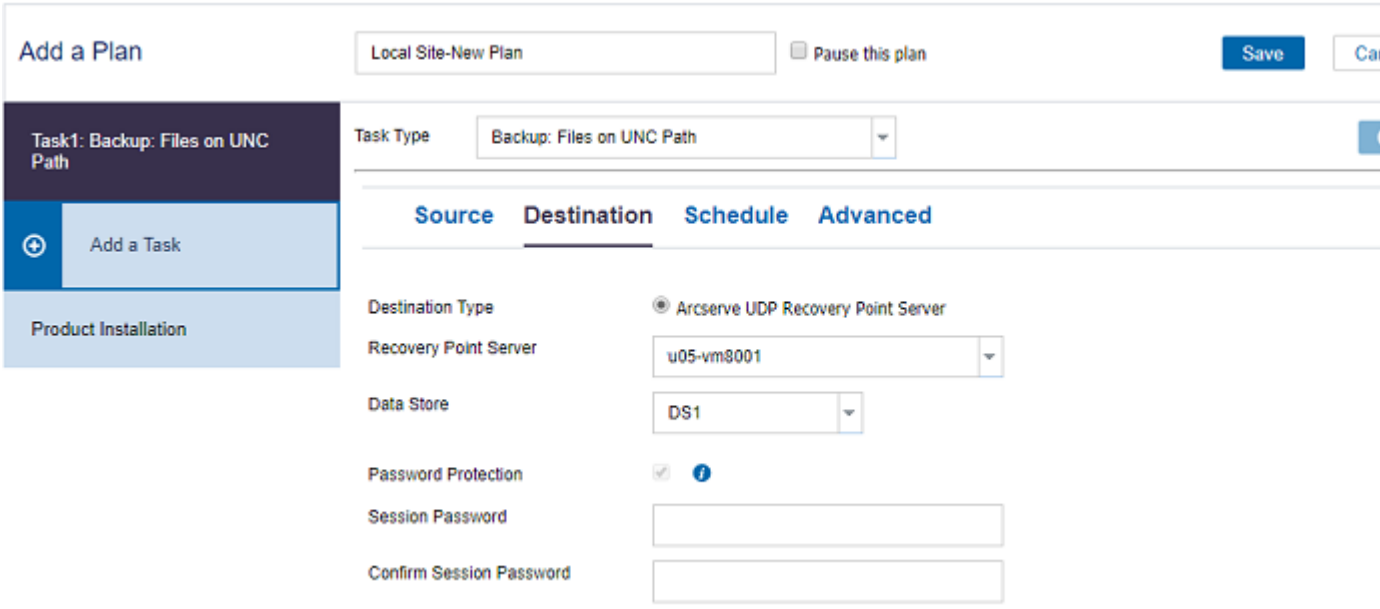
d. Click **Save**.

Note: If the connection to the Nutanix File Server fails, the Error dialog/a notification window displays. Click **OK** to close the Error dialog.

The source is specified.

Specify the Destination

The destination is a location where you store the backup data. You must at least specify the destination to save the plan.



Specify the following details:

- Destination Type: Verify if the **Arcserve UDP Recovery Point Server** option is selected by default.
- Recovery Point Server: Select a recovery point server from the drop-down list.
- Data Store: Select a data store from the drop-down list.

  The list displays all data stores that are created at the specified recovery point server.
- Password Protection: Provide a session password, and then retype the session password to confirm.

  Note: The session password is optional when the backup destination is an unencrypted RPS data store.

  The destination is specified.

Specify the Schedule

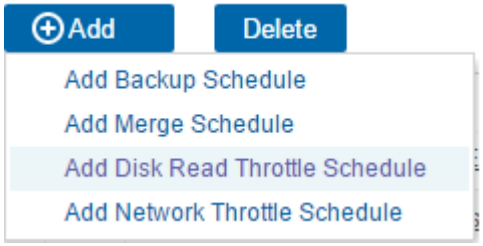
The Schedule page lets you define a schedule for Backup, Merge, and Throttle functions to repeat at specific intervals. After you define a schedule, the jobs run automatically as per the schedule. You can add multiple schedules and can provide retention settings.

A Backup Schedule refers to regular schedule that is repeated multiple times a day based on the number of hours or minutes you select. Besides the regular schedule, a backup schedule also provides options to add daily, weekly, and monthly schedules.

Note: For more information on scheduling and retention settings, see Understanding Advanced Scheduling and Retention.

Follow these steps:

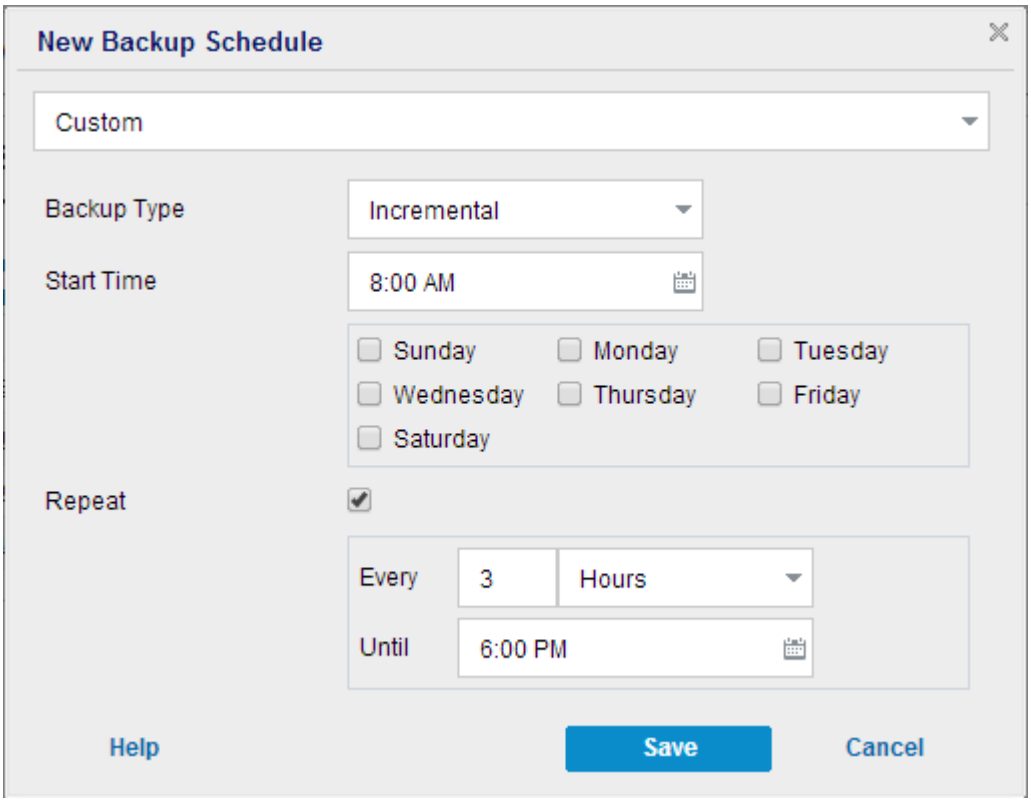
1. Add backup, merge, Disk Read throttle, and Network Throttle schedules.



Add Backup Schedule

a. Click **Add** and select **Add Backup Schedule**.

The **New Backup Schedule** dialog opens.



b. Select one of the following options:

Custom

Specifies the backup schedule that repeats multiple times a day.

Daily

Specifies the backup schedule that occurs once a day. By default, all the days of the week are selected for Daily backup. If you do not want to run the backup job on a specific day, clear the check box for that day of the week.

Weekly

Specifies the backup schedule that occurs once a week.

Monthly

Specifies the backup schedule that occurs once a month.

c. Select the backup type.

Full

Determines the backup schedule for Full Backups. As scheduled, Arcserve UDP performs a full backup of all used blocks from the source machine. A full backup typically consumes time depending on the backup size.

Verify

Determines the backup schedule for Verify Backups.

Arcserve UDP verifies that the protected data is valid and complete by performing a confidence check of the stored backup image to the backup source. If necessary, the image is resynchronized. A Verify Backup looks at the most recent backup of each individual block and compares the content and information to the source. This comparison verifies that the latest backed up blocks represent the corresponding information at the source. If the backup image for any block does not match the source (possibly because of changes in the system since the last backup), Arcserve UDP refreshes (resynchronizes) the backup of the block that does not match. You can also use a Verify Backup (infrequently) to get the guarantee of full backup without using the space required for a full backup.

Advantages: Produces a small backup image when compared to full backup because only the changed blocks (blocks that do not match the last backup) are backed up.

Disadvantages: Backup time is long because all source blocks are compared with the blocks of the last backup.

Incremental

Determines the backup schedule for Incremental Backups.

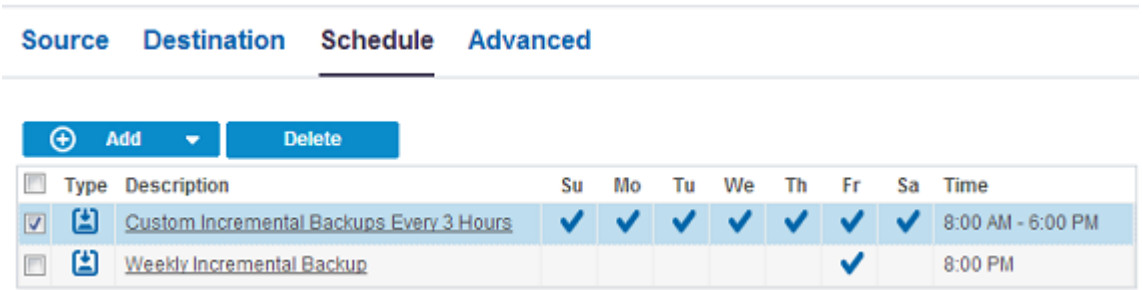
As scheduled, Arcserve UDP incrementally backs up only those blocks that have changed after the last successful backup. The advantages of Incremental Backups are that it is a fast backup and it produces a small backup image. This is the most optimal way to perform a backup.

d. Specify the backup start time.

e. (Optional) Select the **Repeat** check box and specify the repeat schedule.

f. Click **Save**.

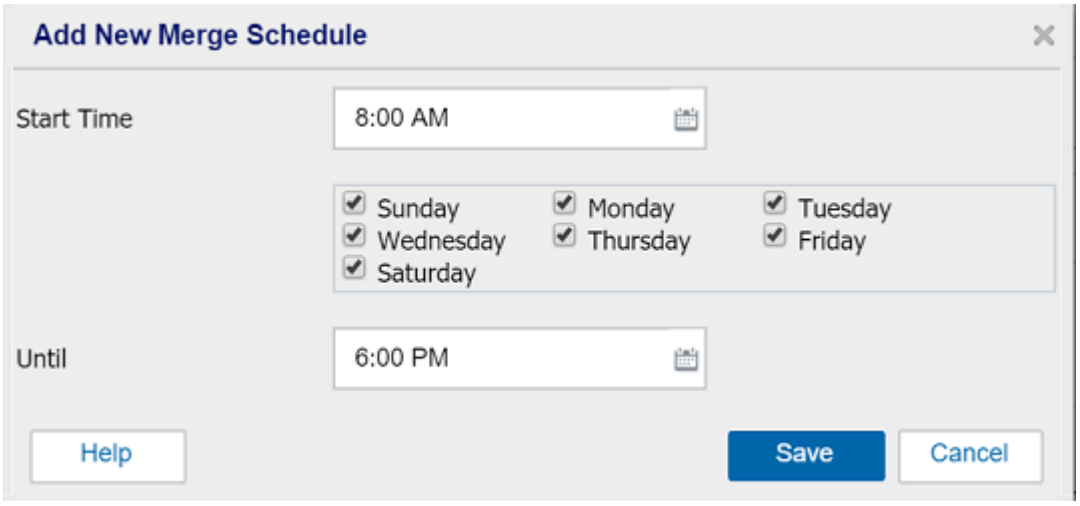
The Backup Schedule is specified and appears on the **Schedule** page.



Add Merge Schedule

a. Click **Add** and select **Add Merge Schedule**.

   The **Add New Merge Schedule** dialog opens.



b. Specify the start time to start the merge job.

c. Specify **Until** to specify an end time for the merge job.

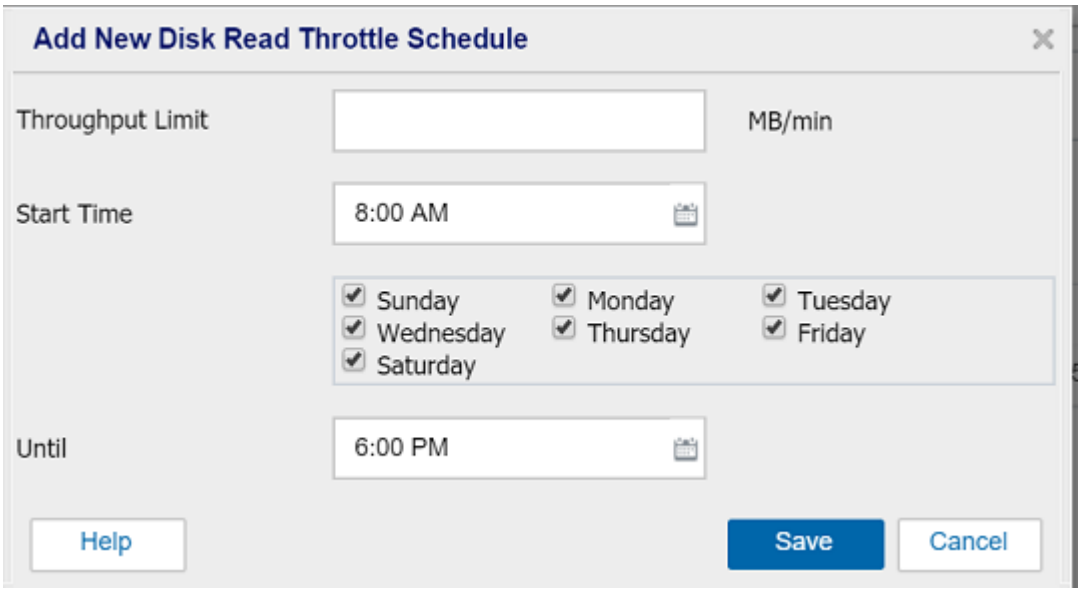
d. Click **Save**.

The Merge Schedule is specified and it is displayed on the **Schedule** page.

Add Disk Read Throttle Schedule

a. Click **Add** and select **Add Disk Read Throttle Schedule**.

   The **Add New Disk Read Throttle Schedule** dialog opens.



b. Specify the throughput limit in MB per minute unit.

c. Specify the start time to start the backup throughput job.

d. Specify **Until** to specify an end time for the throughput job.

e. Click **Save**.

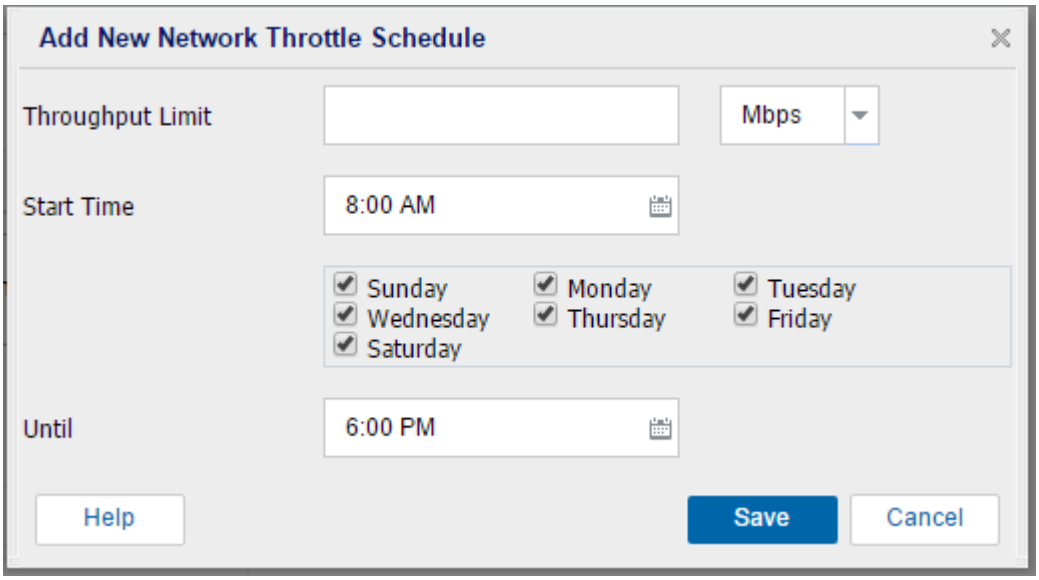
The Disk Read Throttle Schedule is specified and appears on the **Schedule** page.

Add Network Throttle Schedule

Note: Network Throttle Schedule appears only for Windows Agent based backup when you define a Deduplication enabled Data Store as the destination for the plan.

a. Click **Add** and select **Add Network Throttle Schedule**.

The **Add New Network Throttle Schedule** dialog opens.



b. Specify the throughput limit in Mbps or Kbps unit.

Note: Default minimum value: 500 kbps. To change the default value perform the following steps:

i. From the registry path SOFTWARE\Arcserve\Unified Data Protection\Management\Console, add a key MinNetworkThrottleValueInKpbs, type is REG_SZ, and set the value.

ii. Restart the Arcserve UDP Management service.

iii. Modify plan or create new plan.

The custom value takes effect.

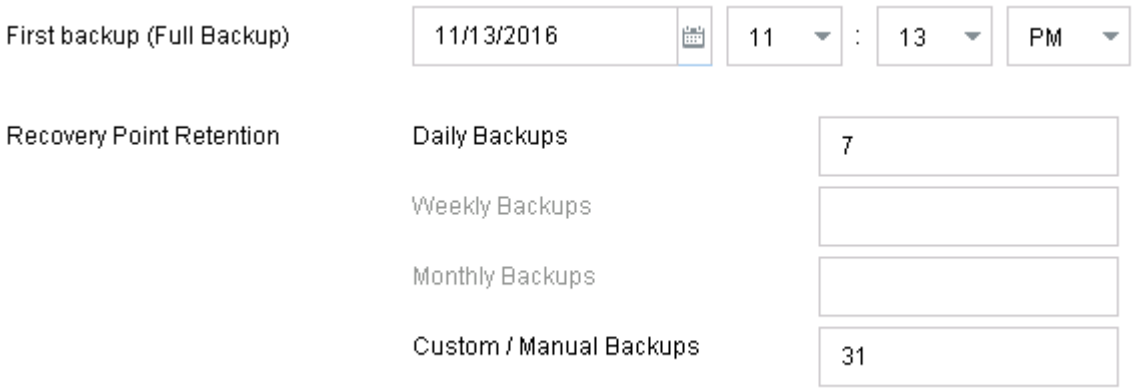
c. Specify the start time to start the backup throughput job.

d. Specify **Until** to specify an end time for the throughput job.

e. Click **Save**.

The Network Throttle schedule is specified and appears on the **Schedule** page.

2. Specify the start time for the scheduled backup.



3. Specify the recovery points retention settings for Custom, Daily, Weekly, and Monthly schedule.

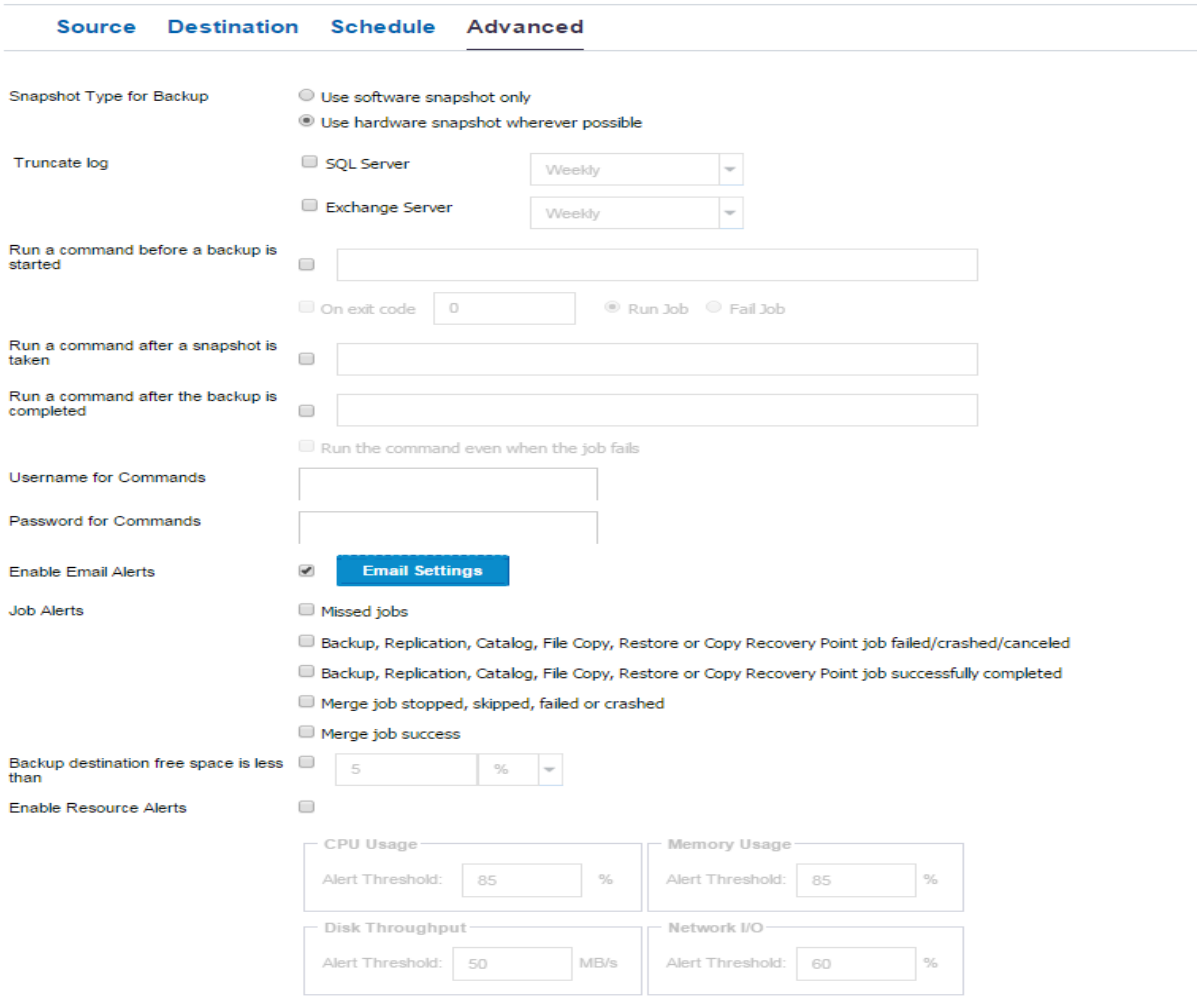
   These options are enabled if you have added the corresponding backup schedule. If you modify the retention settings on this page, the changes are reflected on the Backup Schedule dialog.

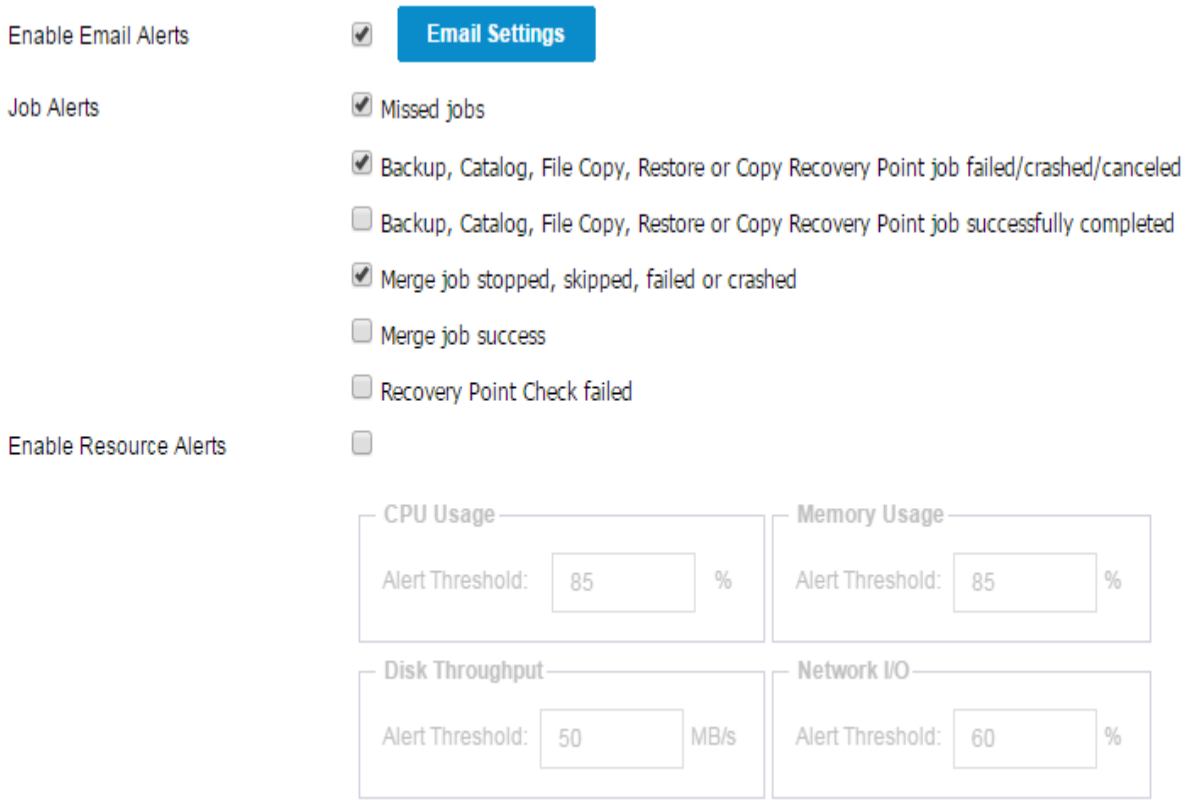
   The schedule is specified.

Specify the Advanced Settings

The **Advanced** tab lets you specify some advanced settings for the backup job. The advanced settings include providing truncate log settings, providing the location of any scripts, and email settings.

The following image displays the Advanced tab:


Source
Destination
Schedule
Advanced
Snapshot Type for Backup
Use software snapshot only
Use hardware snapshot wherever possible
Truncate log
SQL Server
Weekly
Exchange Server
Weekly
Run a command before a backup is started
On exit code
0
Run Job
Fail Job
Run a command after a snapshot is taken
Run a command after the backup is completed
Run the command even when the job fails
Username for Commands
Password for Commands
Enable Email Alerts
Email Settings
Job Alerts
Missed jobs
Backup, Replication, Catalog, File Copy, Restore or Copy Recovery Point job failed/crashed/canceled
Backup, Replication, Catalog, File Copy, Restore or Copy Recovery Point job successfully completed
Merge job stopped, skipped, failed or crashed
Merge job success
Backup destination free space is less than
5
%
Enable Resource Alerts
CPU Usage
Alert Threshold:
85
%
Memory Usage
Alert Threshold:
85
%
Disk Throughput
Alert Threshold:
50
MB/s
Network I/O
Alert Threshold:
60
%



Follow these steps:

1. Specify the following details.

Snapshot Type for Backup

Select one of the following options for the backup snapshot.

Use software snapshot only

Specifies that the backup type uses only the software snapshot. Arcserve UDP will not check for hardware snapshot. The software snapshot utilizes less resources on the virtual machines. You can use this option if the server has lower configurations and processing speed.

Use hardware snapshot wherever possible

Specifies that the backup type first checks for a hardware snapshot. If all the criteria are met, the backup type uses hardware snapshot.

Note: For more information on the hardware snapshot criteria, see the pre-requisite.

Truncate Log

Lets you specify the schedule to truncate logs for SQL Server and Exchange Server. You can specify the schedule as **Daily**, **Weekly**, or **Monthly**.

Enable purging logs at hourly interval for SQL via a registry switch

- Configure the Plan Settings. Check the "SQL Server" option in the "Truncate Log" Section under the "Advanced" Tab, and then select "Daily".
- Set the registry key on the SQL Server machine where the UDP Agent hosts. "PurgeSqlLogPerHour" is the interval in hours for purging SQL Log.

  Path: HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine

  Value Name: PurgeSqlLogPerHour (indicates the interval in hours for purging SQL Log.)

  Value type: REG_DWORD

User Name

Lets you specify the user who is authorized to run a script.

Password

Lets you specify the password of the user who is authorized to run the script.

Run a command before backup is started

Lets you run a script before the backup job starts. Specify the complete path where the script is stored. Click **On exit code** and specify the exit code for **Run Job** or **Fail Job**. **Run Job** indicates that the backup job will continue when the script returns the exit code. **Fail Job** indicates that the backup job will stop when the script returns the exit code.

Run a command after snapshot is taken

Lets you run a script after the backup snapshot is taken. Specify the complete path where the script is stored.

To run the post snapshot script even during a snapshot failure, select **Run the command even when the snapshot fails**.

Run a command after backup is over

Lets you run a script after the backup job is completed. Specify the complete path where the script is stored.

Enable Email Alerts

Lets you enable email alerts. You can configure email settings and can specify the types of alerts that you want to receive in an email. When you select this option, the following options are enabled for your selection.

Email Settings

Lets you configure the email settings. Click **Email Settings** and configure the email server and proxy server details. For more information about how to configure Email Settings, see Configure Email and Alert .

Job Alerts

Lets you select the types of job emails you want to receive.

Enable Resource Alerts

Lets you specify a threshold for CPU Usage, Memory Usage, Disk Throughput, Network I/O. You can provide the value in percentage. You will receive an email when the Alert Threshold value exceeds.

2. Click **Save**.

Note: When you select a node as a backup source or backup proxy, Arcserve UDP checks whether the agent is installed on the node and if it is the latest version. Arcserve UDP then displays a verification dialog that lists all the nodes that either have an outdated version of the agent or does not have the agent installed. To install/upgrade the agent on these nodes, select the installation method and click Save.

The changes are saved and a green check mark is displayed next to the task name. The plan page closes.

Note: If you have to add another task, you must select the plan from the **Resources** tab and modify the plan. To modify the plan, click the plan from the center pane. The plan opens and you can modify it.

The backup plan is created and automatically deployed to the source node. The backup runs per the schedule that you have configured in the **Schedule** tab. You can also perform a manual backup at any time.

(Optional) Perform a Manual Backup

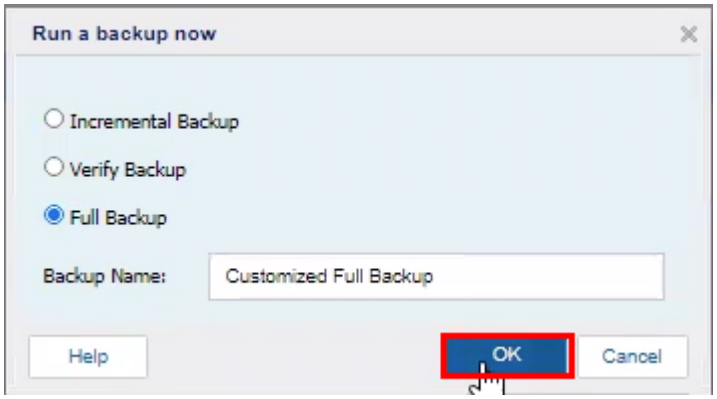
Typically, backups are performed automatically and are controlled by the schedule settings. In addition to the scheduled backup, a manual backup provides you the option to back up your nodes on a need basis. For example, if you have a repeat schedule for Full, Incremental, and Verify backups and you want to make major changes to your machine, you must perform an immediate manual backup without waiting for the next scheduled backup to occur.

Follow these steps:

1. Log into Arcserve UDP Console.
2. Navigate to **resources** > **Plan** > **All Plans**.

   The list of all plans is displayed on the center pane.
3. Right-click the plan, and then select **Backup Now**.

   The Run a backup now dialog opens.
4. On the Run a backup now dialog, select the type of backup you want to perform, and then specify the backup name if required.



The available options are:

Incremental Backup

Initiates an Incremental Backup of your machine. An Incremental Backup backs up only those blocks that have changed after the previous backup.

Note: If you are performing a backup for the first time, the full backup gets initiated automatically even though you selected the Incremental Backup option.

The advantages of Incremental Backups are that it is a fast backup and produces a small backup image. This is the most optimal way to perform backups.

Verify Backup

Initiates a Verify Backup of your machine by examining the most recent backup of each individual block and comparing the content and information to the original source. This comparison verifies that the latest backed up blocks represent the corresponding information at the source. If the backup image for any block does not match the source, Arcserve UDP Agent (Windows) refreshes (resynchronizes) the backup of the block that does not match.

Advantage: A very small backup image is produced when compared to a Full Backup because only the changed blocks (blocks that do not match the last backup) are backed up.

Disadvantage: The backup time is slow because all the source disk blocks are compared with the blocks of the last backup.

Full Backup

Initiates a Full Backup of your entire machine or the selected volumes.

Notes:

- If you add a new volume to the backup source, the newly added volume is fully backed up regardless of the overall backup method selected.
- If no backup name is specified, by default, it is automatically named as Customized Full/Incremental/Verify Backup.

5. Click **OK**.

   The Progress Information dialog appears.

6. Wait until the backup job completes, and then click **OK**.

7. To monitor the backup jobs, navigate to **jobs** > **Jobs in Progress**.

   Wait until the backup job is finished.

   Note: Only one job can be run at a time. If you manually attempt to launch a backup job when another job is currently running, an alert message informs you that another job is running and requests that you try again at a later time.

   The manual backup is successfully performed.

Other tasks on UNC/NFS Plan

You can follow the UNC/NFS Path plan task with other tasks. You can create follow-up tasks such as copy recovery point, Copy to Tape, Replicate to Arcserve Cloud, Replicate and replicate to remotely-managed RPS.

Copy Recovery Points Task

Lets you copy the recovery points to a local or shared folder or Cloud.

Copy to Tape

Lets you store the recovery point to tape by integrating with Arcserve Backup.

Replicate Task

Lets you create a task to replicate backup data from a recovery point server to another recovery point server.

Replicate to a remotely-managed RPS

Lets you create a task to replicate or send data to a remote recovery point server.

Replicate to Arcserve Cloud

Lets you create a task to replicate or send data to a Cloud recovery point server.

Assured Recovery Test

Lets you verify accessibility and assure recovery of the data.

Verify the Backup

To verify your backup job, confirm that you have successfully created the backup plan. After you verify that the plan is created successfully, confirm whether the backup job is running as scheduled. You can verify the status of backup jobs from the **Jobs** tab.

Follow these steps to verify plans:

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and then click **All Nodes**.

   A list of all nodes is displayed on the center pane.
3. Verify that plans are mapped with nodes.

Follow one of these methods to verify backup jobs:

Method 1

1. Click the **jobs** tab.
2. From the left pane, click **All Jobs Completed**.

   The status of each job is listed on the center pane.
3. Verify that the backup job is successful.

   The backup job is verified.

Method 2

1. Click the **jobs** tab.
2. From the left pane, click **Jobs Successfully Completed**.

   The list of backup jobs completed appear on the center pane.
3. Click the job you want to verify.

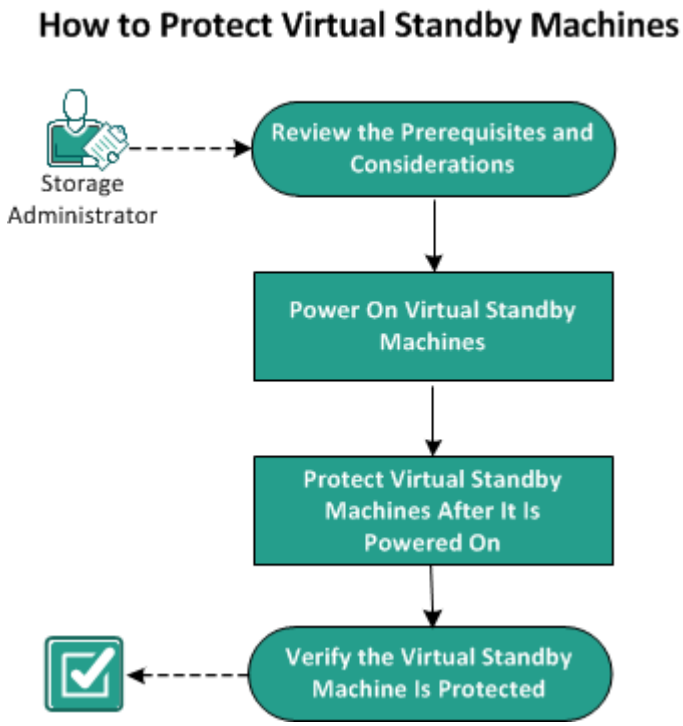
   The right pane refreshes.
4. Under Job Details, click the **View Logs** hyperlink.
5. Verify that the backup job is successful.

   The backup job is verified.

How to Protect Virtual Standby Machines

You can back up virtual standby machines and protect the data from getting corrupted. Before you protect the machine, you have to power on the machine.

The following diagram illustrates the process to protect virtual standby machines:



What To Do Next?

- Review the Prerequisites and Considerations
- Power On Virtual Standby Machines
- Protect Virtual Standby Machines After Power On
- Verify the Virtual Standby Machine Is Protected

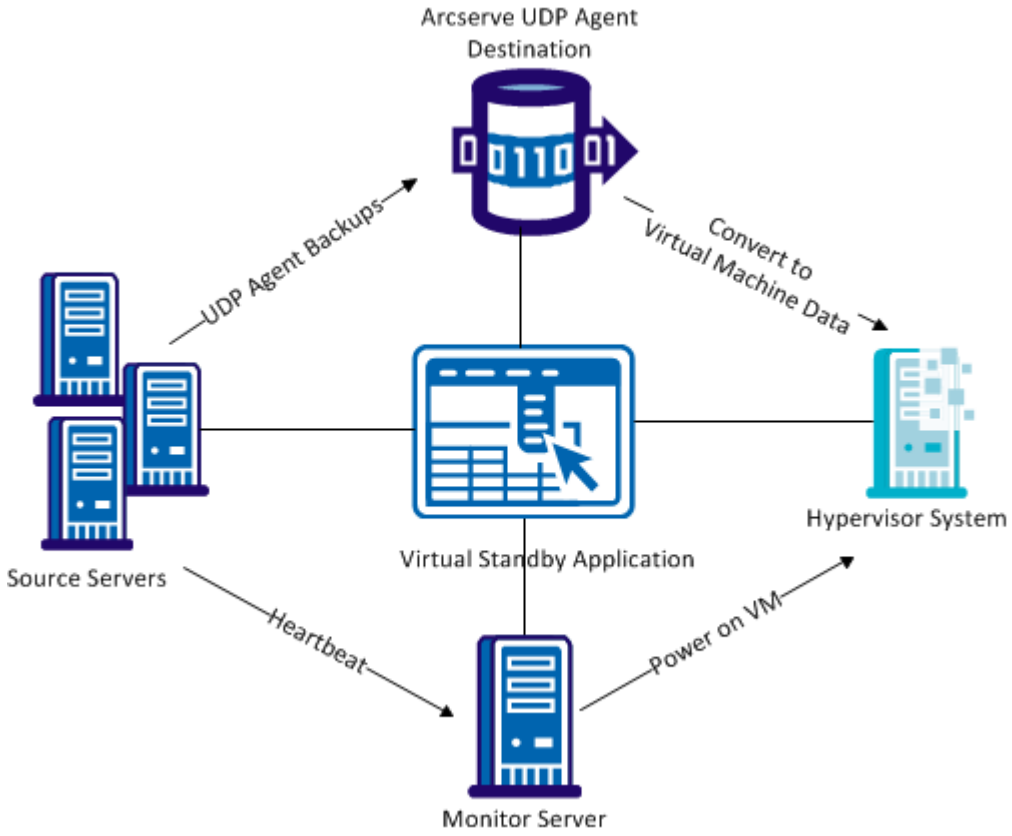
Review the Prerequisites and Considerations

Verify that you have completed the following prerequisite tasks:

- Logged into the Console
- Have a virtual standby machine ready.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Power On Virtual Standby Machines

You can power on virtual standby machines and protect the virtual machines after the machines are powered on. The following diagram describes the process flow to power on the virtual machines:



Power On Virtual Standby Machines from Recovery Point Snapshots

Virtual standby can be configured to power on virtual standby machines from recovery point snapshots automatically when the monitoring server does not detect a heartbeat from the source server. Optionally, you can power on virtual standby machines from recovery point snapshots manually in the event a source server fails, an emergency occurs, or you want to offline a source node for maintenance.

Note: The following steps describe how to power on virtual standby machines from recovery point snapshots manually. For information about how to allow Virtual Standby to power on Recovery Point Snapshots automatically, see Add a Virtual Standby to EC2 Task to the Plan.

Follow these steps:

1. From the **resources** tab, navigate to the **Virtual Standby** node group.

   The virtual standby nodes are displayed on the center pane.

2. On the center pane, select the node and click **Standby VM**.

   The **Standby VM** dialog opens.

3. On the **Standby VM** dialog, perform the following tasks:

   - Select a date and time snapshot of the recovery point snapshot to power on the virtual machine.

     Note: If the standby virtual machine was not configured yet, the link "The standby virtual machine network is not configured." is displayed.

     a. Click this link to configure the network.

     b. Click **Save**. The settings are saved for the virtual standby virtual machine.

     c. Click **Close** and the **Recovery Point Snapshot** dialog appears.

   - Click **Power On VM**.

     The virtual machine is powered on using the data contained in the recovery point snapshot.

     Note: After the virtual machine is powered on, you can be prompted to restart the computer one or more times. This behavior occurs because Amazon PV driver is installed on the virtual machine.

     After you power on virtual standby machines from recovery point snapshots, you may need to complete the following tasks:

- Activate the Windows operating system that is running on the virtual machine.
- Start Arcserve Unified Data Protection Agent for Windows on the virtual machine.
- Update Arcserve UDP with the host name, IP address, and the login credentials for the virtual machine.
- Assign the node to a plan.

Note: This task is required only when you want to create recovery point snapshots for the virtual machine that was powered on.

Power on Virtual Standby Machines from Hyper-V Manager

When you want to power-on Virtual Standby virtual machines manually, the best practice is to power on the virtual machines from the Standby VM dialog of Arcserve Unified Data Protection server. For more information, see Power on Virtual Standby Virtual Machines from Recovery Point Snapshots. However, if you want to start the Virtual Standby virtual machines from the Hyper-V server, you can do so using Hyper-V Manager.

Note: The Hyper-V Manager lets you access the recovery point snapshots that virtual standby created to protect the node. You should not delete the snapshots. When you delete the snapshots, the relationship between the data contained in the snapshots becomes inconsistent the next time a Virtual Standby job runs. With inconsistent data, you cannot power on Virtual Standby virtual machines properly.

Follow these steps:

1. Log into the Hyper-V server that is monitoring the nodes that you are protecting.
2. Start Hyper-V Manager performing the following steps:
   a. Click Start, click All Programs, click Administrative Tools, and then click Hyper-V Manager.

      Hyper-V Manager opens.

   b. From the Hyper-V Manager directory tree, expand Hyper-V Manager and click the Hyper-V server containing the virtual machine that you want to power on.

      The virtual machines associated with the specified Hyper-V server display in the Virtual Machines list in the center pane.
3. Perform one of the following tasks:
   - **To power on the virtual machine using the latest snapshot:** In the Virtual Machines list, right-click the virtual machine that you want to power on and click Start on the pop-up menu.
   - **To power on the virtual machine using an older snapshot:**
     a. In the Virtual Machines list, click the virtual machine that you want to power on.

        The snapshots associated with the virtual machine display in the Snapshots list.

b. Right-click the snapshot that you want to use to power on the virtual machine and click Apply on the pop-up menu.

   The Apply Snapshot dialog opens.

c. Click Apply.

d. In the Virtual Machines list, right-click the virtual machine that you want to power on and click Start on the pop-up menu.

   The virtual standby machine is powered on.

If necessary, you can back up the virtual machines and create recovery point snapshots after you power on the virtual machine.

Power on Virtual Standby Machines from VMware vSphere Client

When you want to power-on virtual standby machines manually, the best practice is to power on the virtual machines from the Standby VM dialog of Arcserve UDP. For more information, see Power on Virtual Standby Virtual Machines from Recovery Point Snapshots. However, if you want to start the virtual standby machines from the ESX Server or the vCenter Server system, you can do so using VMware vSphere Client.

Note: The VMware vSphere Client lets you access the recovery point snapshots that virtual standby created to protect the node. You should not delete the snapshots. When you delete the snapshots, the relationship between the data contained in the snapshots becomes inconsistent the next time a virtual standby runs. With inconsistent data, you cannot power on virtual standby machines properly.

Follow these steps:

1. Open VMware vSphere Client and log in to the ESX Server or vCenter Server system that is monitoring the nodes that you are protecting.
2. From the directory tree, expand the ESX Server system or the vCenter Server system, locate, and click the virtual machine that you want to power on.
3. Perform one of the following tasks:

   To power on the virtual machine using the latest snapshot: Click the Getting Started tab and then click Power on the virtual machine located on the bottom the screen.

   To power on the virtual machine using an older snapshot:

   a. Click the Snapshot Manager button on the toolbar.

      The Snapshots for (virtual machine name) dialog opens to display a list of snapshots that are available for the virtual machine.

   b. From the list of snapshots, click the snapshot that you want to use to power on the virtual machine and then click Go to.

      The virtual standby machine is powered on.

   If necessary, you can back up the virtual machines and create recovery point snapshots after you power on the virtual machine.

Protect Virtual Standby Machines After it is Powered On

After a virtual standby machine is powered on (either manually or automatically), the Arcserve Unified Data Protection Agent for Windows backup job and the virtual standby job will not run as they were scheduled. You have to manually configure the virtual standby machine to protect it.

Follow these steps:

1. Modify the **VM Name Prefix** in the Virtual Standby task.

   When you power on virtual standby machines, the application defines the virtual machine names of the powered on virtual machines as the concatenation of the **VM Name Prefix** option specified in the Virtual Standby task and the host name of the source node.

   Example:

   - VM Name Prefix: AA_
   - Host name of the source node: Server1
   - Virtual machine name of the virtual standby machine: AA_Server1

   After the virtual standby machines are powered on, virtual machine name conflicts can occur when you do not modify the **VM Name Prefix** in the Virtual Standby task. Problems of this type occur when the source nodes and the virtual standby machines reside on the same hypervisor.

   If necessary, you can update other Virtual Standby task settings. Optionally, you can create a new Virtual Standby task to protect the Virtual Standby virtual machine.

2. After you deploy the plan to the virtual standby machine, resume the Virtual Standby job.

   For more information, see Pause and Resume Virtual Standby Jobs.

3. After you deploy the plan, log in to Arcserve Unified Data Protection Agent for Windows on the virtual standby machine and schedule a repeat method for the Arcserve Unified Data Protection Agent for Windows backup job.

   For more information, see the *Arcserve Unified Data Protection Agent for Windows User Guide*.

   Note: Pause and Resume Virtual Standby Jobs.

Verify if the Virtual Standby Machine is Protected

Verify if the virtual standby machines are protected by confirming that the valid recovery points are available at the backup destination.

Follow these steps:

1. Log in to the backup destination and navigate to the backup destination folder.
2. Verify that the backup of the virtual standby machine was successful and recovery points are available.

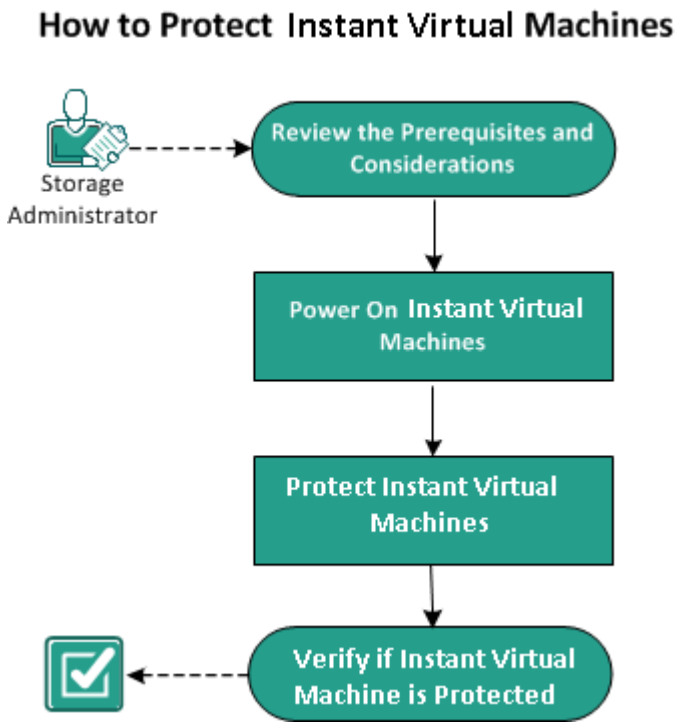
   The virtual standby machine is verified.

   The virtual standby machines are successfully protected.

How to Protect Instant Virtual Machines

You can back up instant virtual machines and protect the data from getting corrupted. Before you protect the machine, you may need to power on the machine.

The following diagram illustrates the process to protect instant virtual machines:



What To Do Next?

- Review the Prerequisites and Considerations
- Power On Instant Virtual Machines
- Protect Instant Virtual Machines
- Verify if the Instant Virtual Machine is Protected

Review the Prerequisites and Considerations

Verify that you have completed the following prerequisite tasks:

- Logged into the Console.
- Have an instant virtual machine ready.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Power On Instant Virtual Machines from Recovery Point

You can create instant virtual machines and protect the virtual machines manually from recovery point, after the machines are powered on.

You can select Boot Now or Boot Later when you create an Instant virtual machine. Select Boot Now option boots the Instant virtual machine immediately upon successful creation, otherwise you need to power on Instant VM manually.

You can power on Instant VMs from recovery point only when you want to create a recovery point for the virtual machine that was powered on.

Follow these steps:

1. From the **resources** tab, navigate to the **Infrastructure: Instant Virtual Machines** node group.

   The instant VM nodes appear on the center pane.

2. From the center pane, select the node, click **Actions**, and then click **Power on** from the drop-down list.

   The virtual machine is powered on using the data contained in the recovery point snapshot.

   Note: After the virtual machine is powered on, you may be asked to restart the computer one or more times. This behavior occurs because VMware installs VMware Tools on the virtual machine or Windows Hyper-V installs Integration Services on the virtual machine.

   After you power on instant virtual machines from recovery point snapshots, you may need to complete the following tasks:

   - Activate the Windows operating system that is running on the virtual machine.
   - Start Arcserve Unified Data Protection Agent for Windows on the virtual machine.
   - Update Arcserve UDP with the host name, IP address, and the login credentials for the virtual machine

Protect an Instant Virtual Machine when Powered On

After an Instant Virtual Machine is powered on (either manually or automatically), the Arcserve Unified Data Protection Agent for Windows backup job and the virtual standby job do not run as scheduled. You have to manually configure the instant virtual machine to protect.

Follow these steps:

1. Add an Instant Virtual Machine to the Console.

   Note: You can add a node by manually specifying the node details, or importing virtual machines from ESX/vCenter and Hyper-V servers.

   For more information, refer Add Nodes.

2. Add a destination.

   A destination could be a recovery point server, local folder, or remote shared folder.

3. Create a plan to protect the Instant Virtual Machine node.

   A plan is a group of tasks to manage backup, replication, and creation of virtual standby machines.

   Note: You can create a plan with Agent-Based Windows Backup task or Host-Based Agentless Backup task.

4. Perform jobs such as backup, create virtual standby, and replicate.

   For more information, see the *Arcserve UDP Agent (Windows) User Guide*.

Verify if the Instant Virtual Machine is Protected

Verify if the Instant Virtual machine is protected by confirming that the valid recovery points are available at the backup destination.

Follow these steps:

1. Log in to the backup destination and navigate to the backup destination folder.
2. Verify that the backup of the Instant Virtual machine is successful and recovery points are available.

   The Instant Virtual machine is verified.

   The Instant Virtual machine is successfully protected.

How to Replicate Data Between Data Stores Managed from a UDP Console

Using Arcserve UDP, you can replicate your backup data from one data store to another. These data stores are managed from the same UDP Console but are in different recovery point servers. You need to create a plan with two tasks--backup and replicate. The backup task will back up data based on the schedule and the replicate task will replicate the backed up data to the specified recovery point server. The replicate job runs per the schedule that you specify in the replicate task. You can create multiple replicate tasks in a plan.

If the replication job fails for some reasons (such as network problem), then the failed replication job resumes first before transferring any new session. The replication job resumes from the break point of the last failed replication job.

What To Do Next?

1. Review the Prerequisites and Considerations
2. Create a Plan with a Backup Task
3. Add a Replicate Task to the Plan
4. (Optional) Perform a Manual Replication
5. Verify the Plan

Review the Prerequisites and Considerations

Verify if you have completed the following prerequisites:

- Log into the Console.
- Install the server component and create Data Stores.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Create a Backup Task

A plan includes different types of tasks that you want to perform. Typically a plan includes a primary tasks followed by a secondary task. Typically, a primary task is a backup task or replicate from a remote Console task. The role of a backup task is to create a backup of the source nodes that you want to protect. You can back up data from Windows and Linux physical and virtual machines. You can then save the backup data to another location as an added precaution.

For more information on:

- How to back up a Windows node, see How to Create a Windows Backup Plan.
- How to back up virtual machines, see How to Create a Host-Based Virtual Machine Backup Plan.
- How to back up Linux nodes, see How to Create a Linux Backup Plan.

Add a Replicate Task to the Plan

Create a replicate task to further protect your data by copying your backup data from one recovery point server to another recovery point server. You may also copy your backup data from one data store to another data store in the same recovery point server. The replication destination must be a data store in the recovery point server. You can create multiple replicate task to perform multiple replications.

Notes:

- **For merge job:** When a replication task is configured and you run an on-demand merge job from the source data store, the job does not check whether the sessions are replicated or not. As a result, the merged sessions cannot be replicated to target data store and you end up replicating more data. For example, consider there are five sessions, s1, s2, s3, s4, and s5 respectively. s1 and s2 are replicated. Now, you run an on-demand merge job on the source side and retain two session. s4 and s5 are retained. s4 is a full session. So, when the next replication job starts, the job needs to replicate a full session.
- **For purge job:** When a replication task is configured and if you run purge job on the target data store, then the next replication job replicates all sessions to target data store.
- You cannot replicate from dedupe enabled datastore to non-dedupe enabled datastore.
- You cannot replicate from encryption enabled datastore to non-encryption enabled datastore.

Follow these steps:

1. Click **Add a Task** from the left pane.

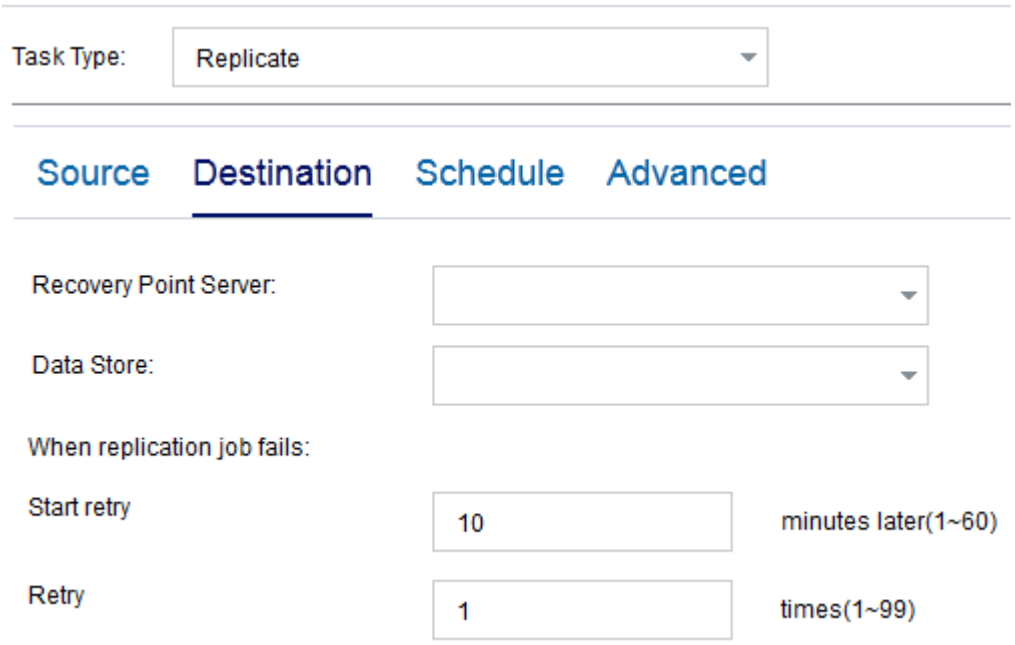
   A new task is added to the left pane.

2. From the **Task Type** drop-down menu, select **Replicate**.

   The Replicate task is added. You do not have to configure the **Source** tab in the Replicate task because it reflects the backup destination from the Backup task.

   Note: You cannot replicate a task without the primary task such as Backup to Recovery Point Server.

3. Click the **Destination** tab and enter the recovery point server details and retry schedule details.



Recovery Point Server

Select the recovery point server from the list.

Data Store

Select the data store from the list.

Start retry

Specify the time (in minutes) to restart the replicate job after the job fails. For example, if you specify 10 minutes, then the replicate job will restart after 10 minutes of its failure.

Limit: 1 to 60

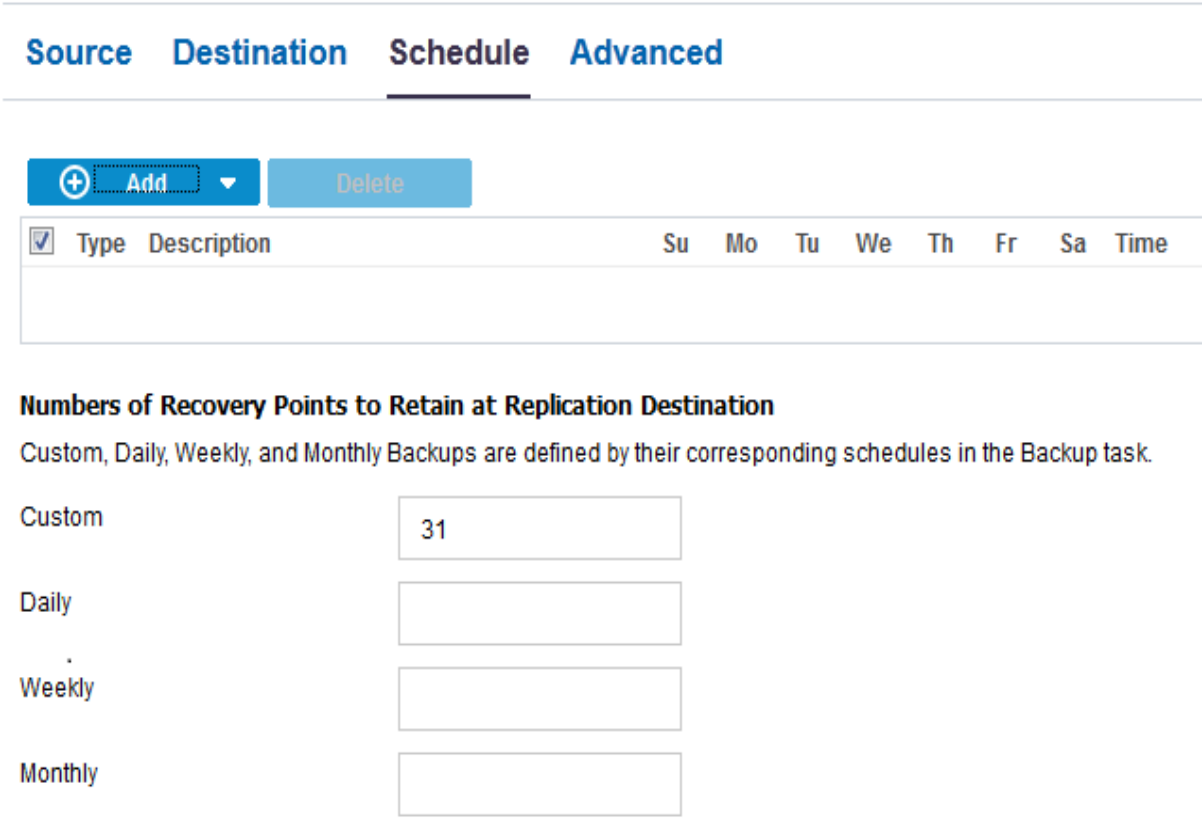
Retry

Specify the number of times you want to start the replicate job when the job fails. The replicate job runs until the job is successful, or until the limit is reached.

Limit: 1 to 99

4. Click the **Schedule** tab and add **Replication Job Schedule**, **Replication Throttle Schedule**, **Merge Schedule**, and **Retention Settings**.

   Note: The replication throttle quota is averagely shared by all the replication jobs started from all the nodes of a current plan.



5. Click the **Advanced** tab and enter the details.
6. Click **Save Changes** or **Add a Task**.

   If you have added a task, then you can create another replicate task to perform multiple levels of replication. You can add multiple replicate task in the plan.

   If you save the changes, then the plan is saved and the replication task is deployed to the replication destination.

   The replicate task is created.

   You have successfully created and automatically deployed a replication plan.

(Optional) Perform a Manual Replication

To manually run a replication job, you must have at least one successful backup data. If you have not set the replication schedule, the replication job will run immediately after the backup job, otherwise, it depends on your replication schedule setting.

Follow these steps:

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and click **All Nodes**.

   If you have added any plans, these plans will be displayed in the center pane.
3. Select the nodes that you want to backup and that has a plan assigned to it.
4. On the center pane, click **Actions**, **Replicate Now**.

   The **Replicate node** dialog opens.
5. Select the Source RPS and Target RPS for the job.
6. Click **OK**.

   The replication job runs.

   The manual replication is successfully performed.

Verify the Plan

To verify the replication feature, confirm that you have successfully created the replication plan. After you verify that the plan is created successfully, check whether the backup job is running as scheduled. After the backup job successfully completes, the replicate job runs. You can check the status of the backup job and the replicate job from the **jobs** tab.

Follow these steps: to verify plans

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and click **All Nodes**.

   A list of all nodes is displayed on the center pane.
3. Verify that plans are mapped with nodes.

Follow these steps: to verify replicate jobs

1. Click the **jobs** tab.
2. From the left pane, click **All Jobs**.

   The status of each job is listed on the center pane.
3. Verify that the backup job and replicate job is successful.

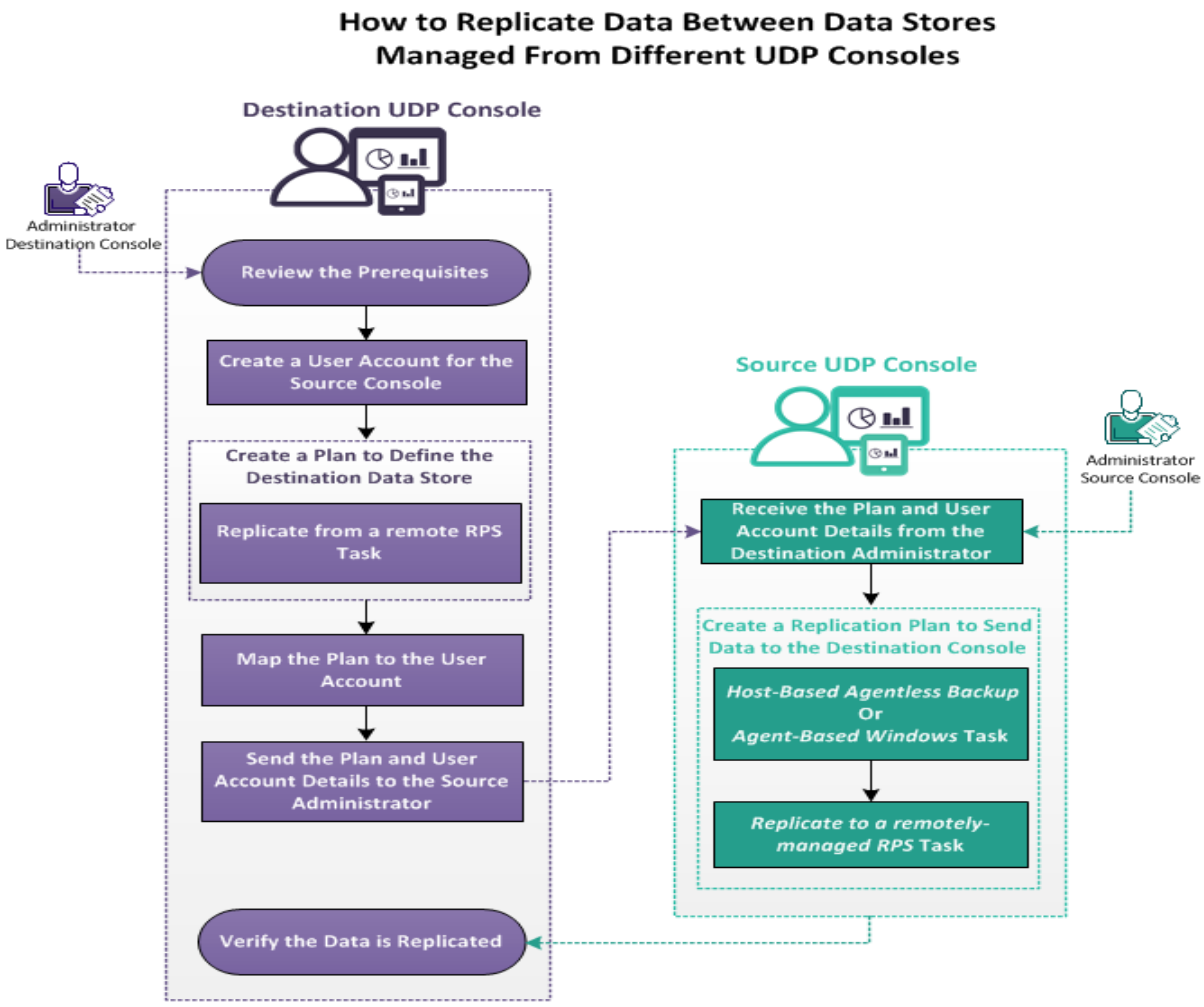
How to Replicate Data Between Data Stores Managed From Different UDP Consoles

To protect your data, you may have to replicate your backup data to another recovery point server that is managed from a different Arcserve UDP Console. For example, you can replicate your data to a service provider that offers its replication services to multiple customers. In this example, the data gets replicated from a source data store (on the source Console) to a destination data store (on the destination Console).

As the administrator of the destination Console, create a unique username, password, and a plan for the source Console. The plan defines the destination data store and the username and password helps the source administrator connect to your server and replicate data.

As the administrator of the source Console, create a plan to replicate data to the destination data store. While you create the plan, connect to the destination server and select the plan that is assigned to you by the destination administrator.

The following diagram illustrates how to replicate data to another data store that is managed from a different Console:



What To Do Next?

1. Review the Prerequisites
2. Create a User Account for the Source Console
3. Create a Plan to define the Destination Data Store
4. Map the Plan to the User Account
5. Send the Plan and User Account Details to the Source Administrator

6. Receive the Plan and User Account Details from the Destination Administrator
7. Create a Replication Plan to Send Data to the destination Console
8. Verify the Data is Replicated

Review the Prerequisites

Review the following prerequisites before replicating data:

- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Administrator—Destination Console

- Verify that you have installed Arcserve UDP on the destination server.
- Verify that you have full privileges to create Windows user accounts on the destination server.
- Verify that you have enabled the **Allow the retrieval of recovery points** check box in the Advanced settings during the Replicate from a remotely-managed RPS plan creation.

  Note: If you do not enable the **Allow the retrieval of recovery points** check box and initiate the Reverse Replication job for Remote Console, you are prompted to enable the Allow retrieval remote flag for Reverse Replication.

Administrator—Source UDP Console

- Verify that you have installed Arcserve UDP on the source server.
- Verify that you have at least completed one full backup on a data store.

Create a User Account for the Source Console

Destination Administrator

To identify and manage the replicated data on the destination server, create a Windows user account. If you are managing more than one source Console, then create a user account for each source Console.

The source Console administrator uses this account details to connect to the destination server.

To create a user account in a Windows operating system, use the User Accounts section in Windows Control Panel. For more information about creating user accounts in Microsoft Windows, see Microsoft documentation.

Create a Plan to Define the Destination Data Store

Destination Administrator

The source data is replicated to this destination data store. To define this destination data store, you create a plan. The plan lets you define the destination data store and the merge schedule.

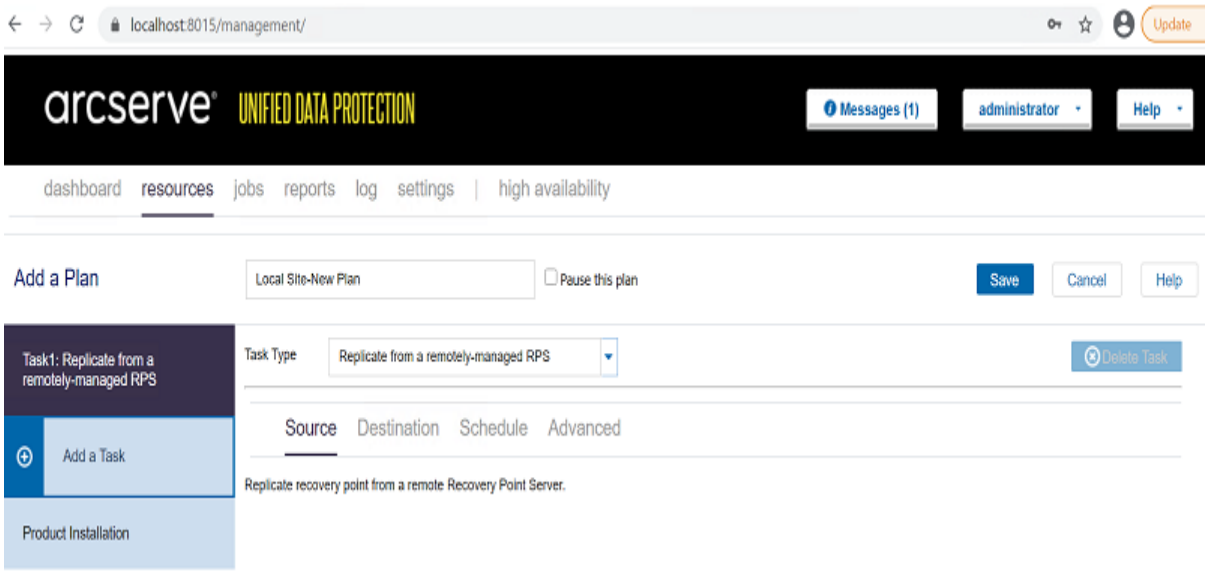
Follow these steps:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Plans**, and click **All Plans**.

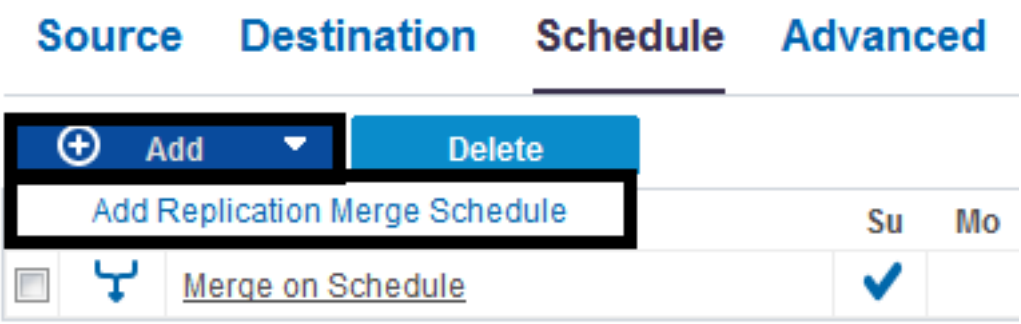
   If you have added any plans, these plans are on the center pane.
3. On the center pane, click **Add a Plan**.

   The **Add a Plan** page opens.
4. Enter a plan name in the **New Plan** field.
5. From the **Task Type** drop-down list, select **Replicate from a remotely-managed RPS**.

   The **Source** tab displays. You cannot provide any details on the Source tab. The source administrator at the source Console provides the source details.



6. Click the **Destination** tab and specify the recovery point server and the data store.
7. (Optional) Select the **Server is behind NAT router** check box and provide the server address and port number.
8. Click the **Schedule** tab.



9. Click **Add** and select **Add Replication Merge Schedule**.

   The **Add New Merge Schedule** dialog opens.

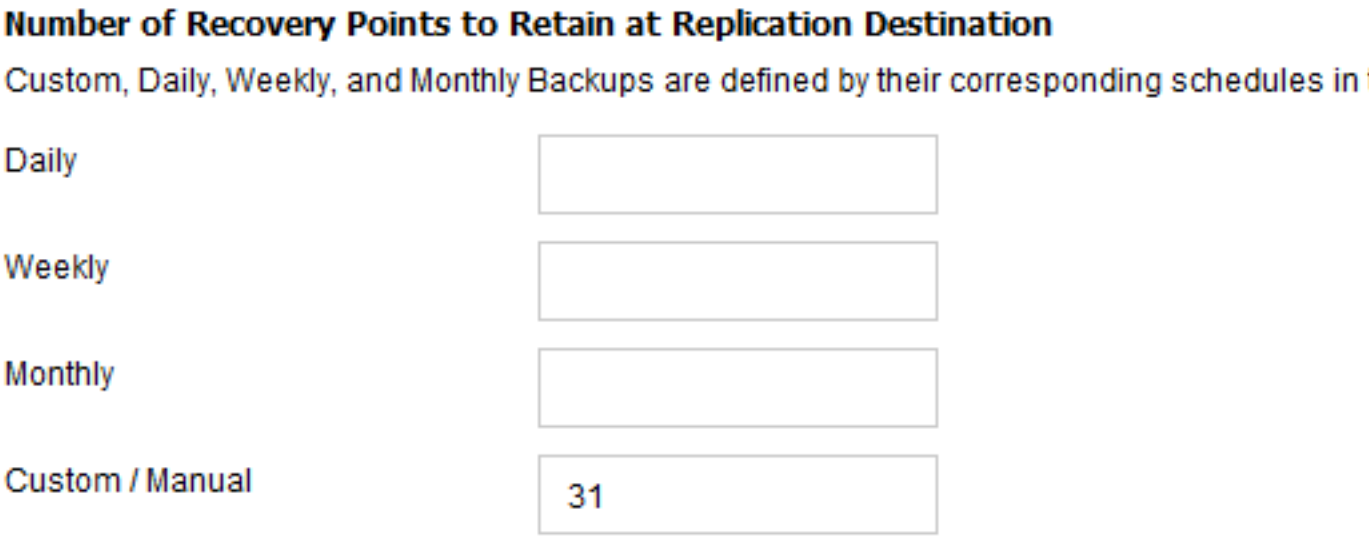
10. Enter the merge schedule.

    Note: To know more about the schedules, see Understanding Advanced Scheduling and Retention.

11. Click **Save**.

    The **Add New Merge Schedule** dialog closes.

12. Enter the recovery points retention details.



13. Click the **Advanced** tab and provide the following details.

    Retrieval Configuration

    Lets you perform the retrieval of recovery points. To enable, select the **Allow the retrieval of recovery points** checkbox.

    Enable Email Alerts

    Lets you enable email alerts. You can configure email settings and can specify the types of alerts that you want to receive in an email. When you select this option, the following options are enabled for your selection.

Email Settings

Lets you configure the email settings. Click **Email Settings** and configure the email server and proxy server details.

Job Alerts

Lets you select the types of job alerts that you want to receive.

14. Click **Save**.

The changes are saved and the plan is created.

The replication plan is successfully created. You can also add Replicate tasks, Replicate to a remotely-managed RPS task and Virtual Standby tasks to the plan.

Map the Plan to the User Account

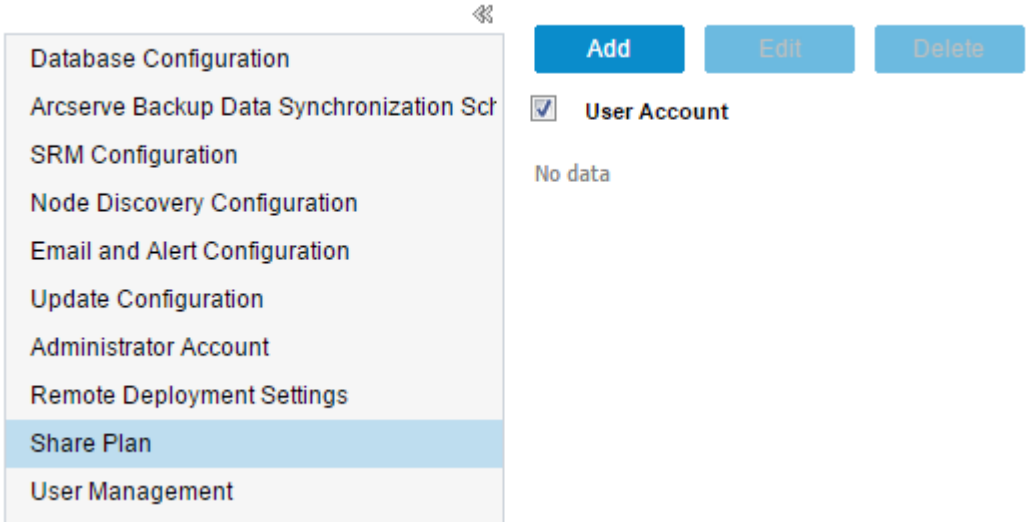
Destination Administrator

You have already created a user account and a plan for a source Console. To identify and manage replicated data, assign the plan to the user account.

Note: You can assign more than one plan to a user account but two different accounts cannot share a plan. However, we recommend assigning a single plan to a user account so that you can easily identify and manage the replicated data.

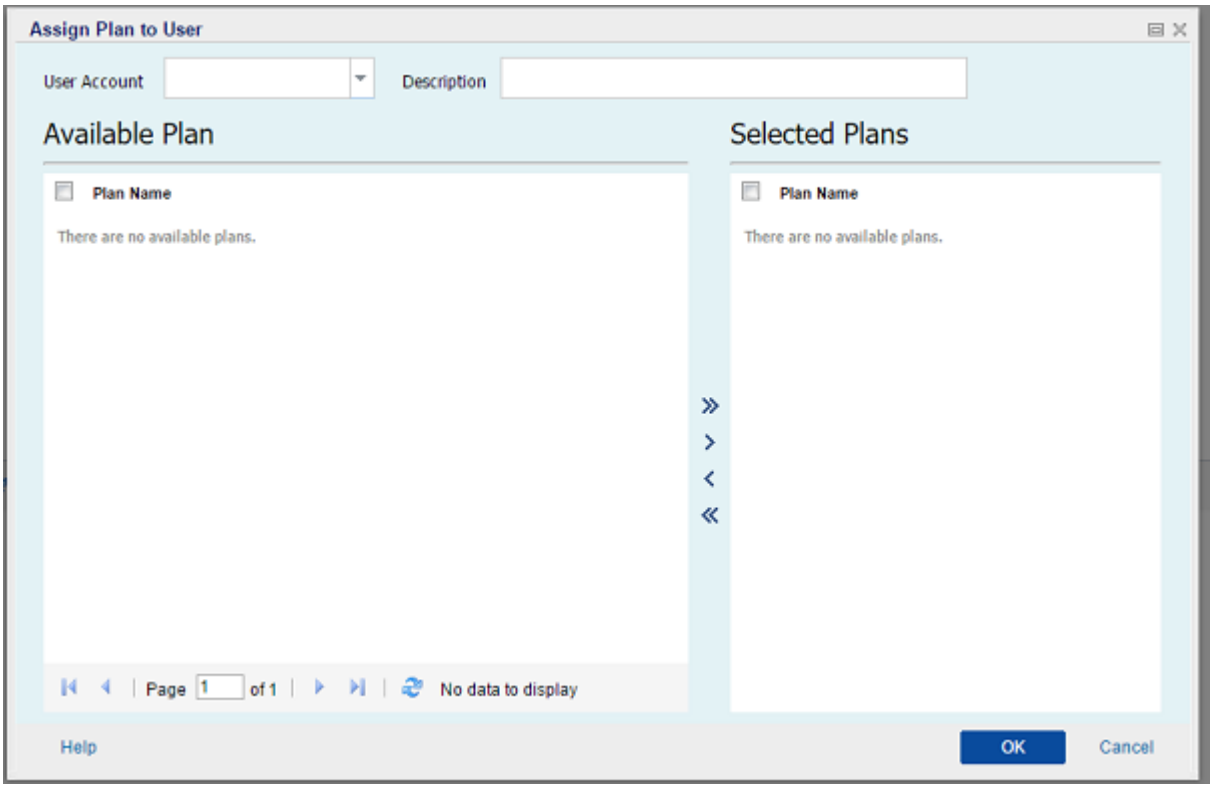
Follow these steps:

1. From the Console, click the **settings** tab.
2. From the left pane, click **Share Plan**.



3. From the center pane, click **Add**.

   The **Assign Plan to User** dialog opens.



4. Select the **User Account**.
5. Select a plan from the **Available Plan** column.

   Note: If a plan is already added to a user name, that plan is not displayed in the **Available Plan** column.
6. Click **Add all plans** or **Add selected plans** to add the plans in the **Selected Plans** column.
7. Click **OK**.

   The **Assign Plan to User** dialog closes. The user name and the associated plans are displayed on the **Share Plan** page.

   The user account is mapped to the plan created for the source Console.

   You can use **Edit** to modify the user configuration or **Delete** to remove the user account from the list.

Send the Plan and User Account Details to the Source Administrator

Destination Administrator

After the plan is associated with the user account, send the plan and user account details to the source administrator. The source administrator uses these details to connect to the destination Console.

As a destination administrator, you have completed all your tasks.

Receive the Plan and User Account Details from the Destination Administrator

Source Administrator

To replicate data to the source Console, you need the destination server, plan, and user account details from the destination administrator. You receive the details from the destination administrator. Understand the details and get your questions clarified from the destination administrator before you start creating replication plans.

Create a Replication Plan to Send Data to the Destination Console

Source Administrator

To replicate your backup data to the destination recovery point server that is managed from a different console, create a replication plan. This replication plan includes a backup task and a remotely managed replication task. In the replication task, specify the remote server and plan details and connect to the remote server. If the connection is successful, select the plan that the destination administrator created for you.

Follow these steps:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Plans** and click **All Plans**.
3. Click **Add a Plan**.

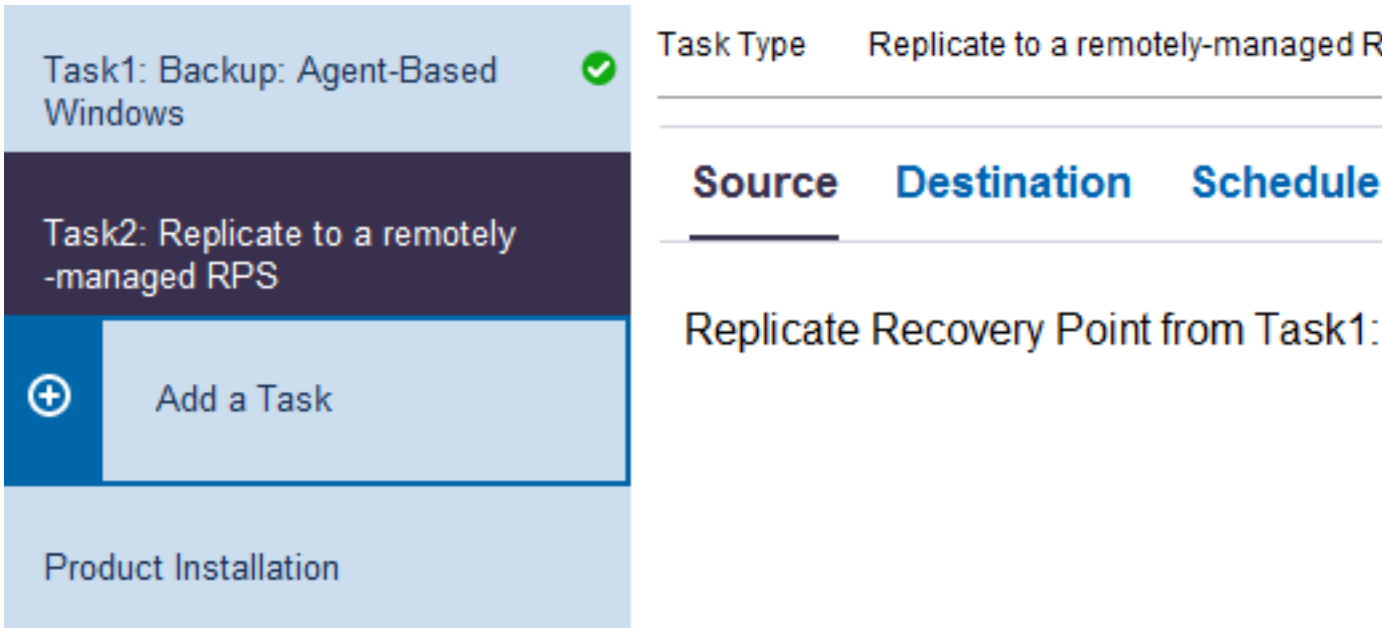
   The **Add a Plan** page opens.
4. Enter a plan name and select one of the following backup tasks and create the task:
   - **Backup: Agent-Based Windows**
   - **Backup: Host-Based Agentless**
   - **Backup: Agent-Based Linux**

     Note: For more information about creating a backup task, see the following topics:

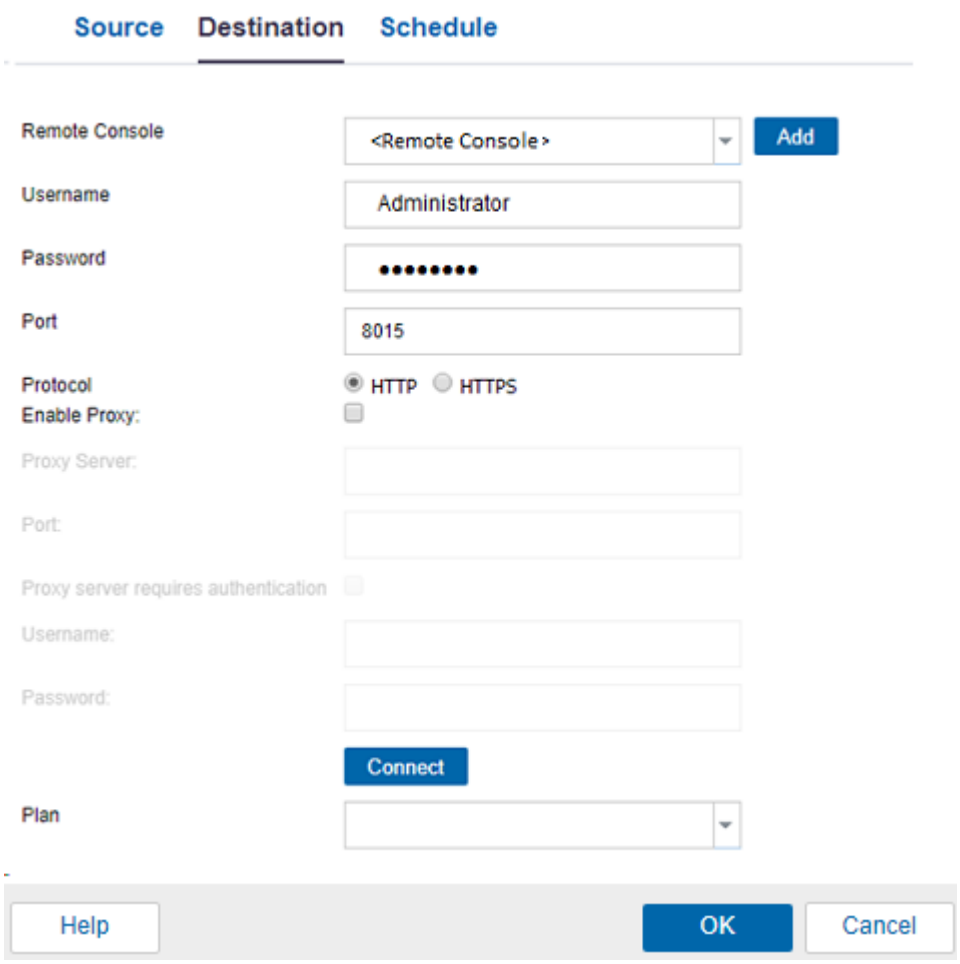
     - How to Create a Windows Backup Plan
     - How to Create a Host-Based Virtual Machine Backup Plan
     - How to Create Linux Backup Plan
5. On the left pane, click **Add a Task**.

   A new task is added to the left pane.
6. From the **Task Type** drop-down list, select **Replicate to a remotely-managed RPS**.

   The Replicate task is added and the **Source** page opens. For the **Source** tab, the destination of the backup task (for example, Backup: Agent-Based Windows) is the source for the **Replicate to a remotely-managed RPS** task.



7. Click the **Destination** tab and enter the following details.



Remote Console

Select a remote console account from the drop-down list or add a new remote console account by clicking the **Add** button.

For more information, click Add remote console.

Username

Specify the user name created by the destination administrator. The destination administrator provides you the username.

Password

Specify the password created by the destination administrator. The destination administrator provides you the password.

Port

Specify the port number of the destination Console. The destination administrator provides you the port number of the destination Console.

Protocol

Specify the protocol used by the destination administrator to connect to the destination Console.

Enable Proxy

Select the check box to enable the proxy server selection.

Proxy Server

Specify the address of the proxy server.

Port

Specify the port number of the proxy server.

Proxy server requires authentication

Select the check box to enable the authentication fields for the proxy server.

Username

Specify the username to connect to the proxy server.

Password

Specify the password to authenticate the proxy server connection.

Connect

Verifies the connection between the source Console and the destination Console. If the connection is successful, then you can see the plan name in the **Plan** field. This plan name is assigned to this Console by the destination administrator.

Plan

Specify the plan that the destination administrator has created for you. If there are multiple plans in the list, then contact the destination administrator to know the correct plan.

Start retry

Reruns the replication job after the specified time if there is a failure. Enter a value from 1 to 60 and the time is defined in minutes.

Retry

Specify the number of retries that you want to perform if there is a job failure. After the number of retries is over, the replication job will run only at the next scheduled time. Enter a value from 1 to 99.

8. Click the **Schedule** tab and provide the replication job schedule and replication throttle schedule.

   Replication Job Schedule

   Specify the date and time to start the replication jobs. You can edit or delete a replication job schedule.

   Replication Throttle Schedule

   Specify the maximum speed (Mbps) at which the replication is done. You can throttle the replication speed to reduce the CPU or network usage. For a replication job, the **jobs** tab displays the average Read and Write speed of the job in progress and the configured throttle speed limit.

   You can edit or delete a replication throttle schedule.

9. Click **Save**.

   The plan is saved and runs per the schedule.

   You have successfully created and automatically deployed a replication plan. When the plan runs, the data gets replicated from the source location to the destination data location over a network.

   Note: After the replication process is complete, the replicated node details are automatically added to the destination Console.

   You have successfully replicated data between two data stores managed from different UDP Consoles.

Verify the Data is Replicated

Destination Administrator

After data is replicated, you can verify whether the replication is successful.

Follow these steps:

1. On the destination Console, navigate to the destination data store on the recovery point server.
2. Verify that the replicated data size matches the source data.

   You have successfully replicated data between two data stores managed from different UDP Consoles.

Applying Best Practices

Configure Multi-Stream Parameters

Configure Multi-Stream Parameters

Replication over WAN related settings are saved at the following registry key:

[HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\Network]

The following list displays the Registry key and their default value:

- "WAN_EnableAutoTunning"=dword:00000001
- "WAN_ChunkSizeByte"=dword:00001000
- "WAN_NumberofStreams"=dword:00000005
- "WAN_MultiStreamsMaxCacheSize"=dword:01000000
- "WAN_SendCommandFragDataMerged"=dword:00000000
- "WAN_RTT_Threshold"=dword:00000032

The following description describes the registry key settings:

WAN_EnableAutoTunning

Specifies the switch to enable or disable multiple streaming. If the value is 0, multi-stream is disabled. For other values, multi-stream is enabled. The default value to enable multi-stream is 1.

WAN_ChunkSizeByte

Specifies the data chunk size for each packet. Packet size affects the throughput. If the WAN bandwidth is higher, the data chunk size can be increased higher.

The default value is 4k byte. The range is limited from 512 byte to 1M byte in code.

WAN_NumberofStreams

Specifies the number of streams that needs to be created on WAN, when the latency is more than the WAN_RTT_Threshhold number. The default stream number is 5. The stream range is from 1 to 10.

WAN_RTT_Threshold

When RTT is greater than WAN_RTT_Threshold, multiple sockets are created. The unit of WAN_RTT_Threshold is millisecond (ms). The default value is 50 millisecond. The range is limited from 20 ms to 600 ms.

WAN_MultiStreamsMaxCacheSize

Specifies that the memory size will be allocated when the multi-stream is enabled. This memory buffer will be used to cached received fragged memory.

The range is from 16MB to 64MB. Default value is 16MB. If the value is zero, the value will be set to 64MB. The unit of this value is BYTE.

WAN_SendCommandFragDataMerged

Specifies that if the value is not zero, the communication library groups small files and send them in one chunk. If the value is zero, then small files are sent individually. The default value is zero.

Notes:

- In a replication job, the socket connection number may not be consistent with the WAN_NumberofStreams registry.

  Replication job from non-GDD to non-GDD

  When RTT is more than WAN_RTT_Threshold, the socket connection number is equal to WAN_NumberofStreams.

  Replication job from non-GDD to GDD or GDD to GDD

  There are four types of connections. Only the data block connection works with the multi-stream feature. So, when RTT is more than WAN_RTT_Threshold, the total socket connection is 3+WAN_NumberofStreams.

- Replication job detects the network status to determine whether the communication is on WAN or not. If the network status is weak, LAN may be accepted as WAN.

How to Perform an Offline Data Replication Using RPS Jumpstart

Replicating a large data store to another recovery point server (managed from a different UDP Console) is time consuming over a network (LAN, WAN, Internet). To replicate a large data store quickly, Arcserve UDP provides an offline data replication method. This method is named RPS Jumpstart.

RPS Jumpstart is an offline replication method that uses an external storage device such as a USB flash drive to replicate a data store. This replication is between two data stores that are managed from different UDP Consoles. For example, consider a service provider that offers its replication services to multiple customers. The customer replicates the data to a storage device and sends the storage device to the service provider. The service provider replicates data from the storage device to the destination server. Both the service provider and the customer must have Arcserve UDP that is installed at their locations.

The offline replication requires both the administrators (the Source and Destination administrators) to complete the following steps at their respective location.

Important! If you are replicating from a shared folder to a data store selected on Recovery Point Server, see How to Migrate r16.5 Recovery Point to Arcserve UDP.

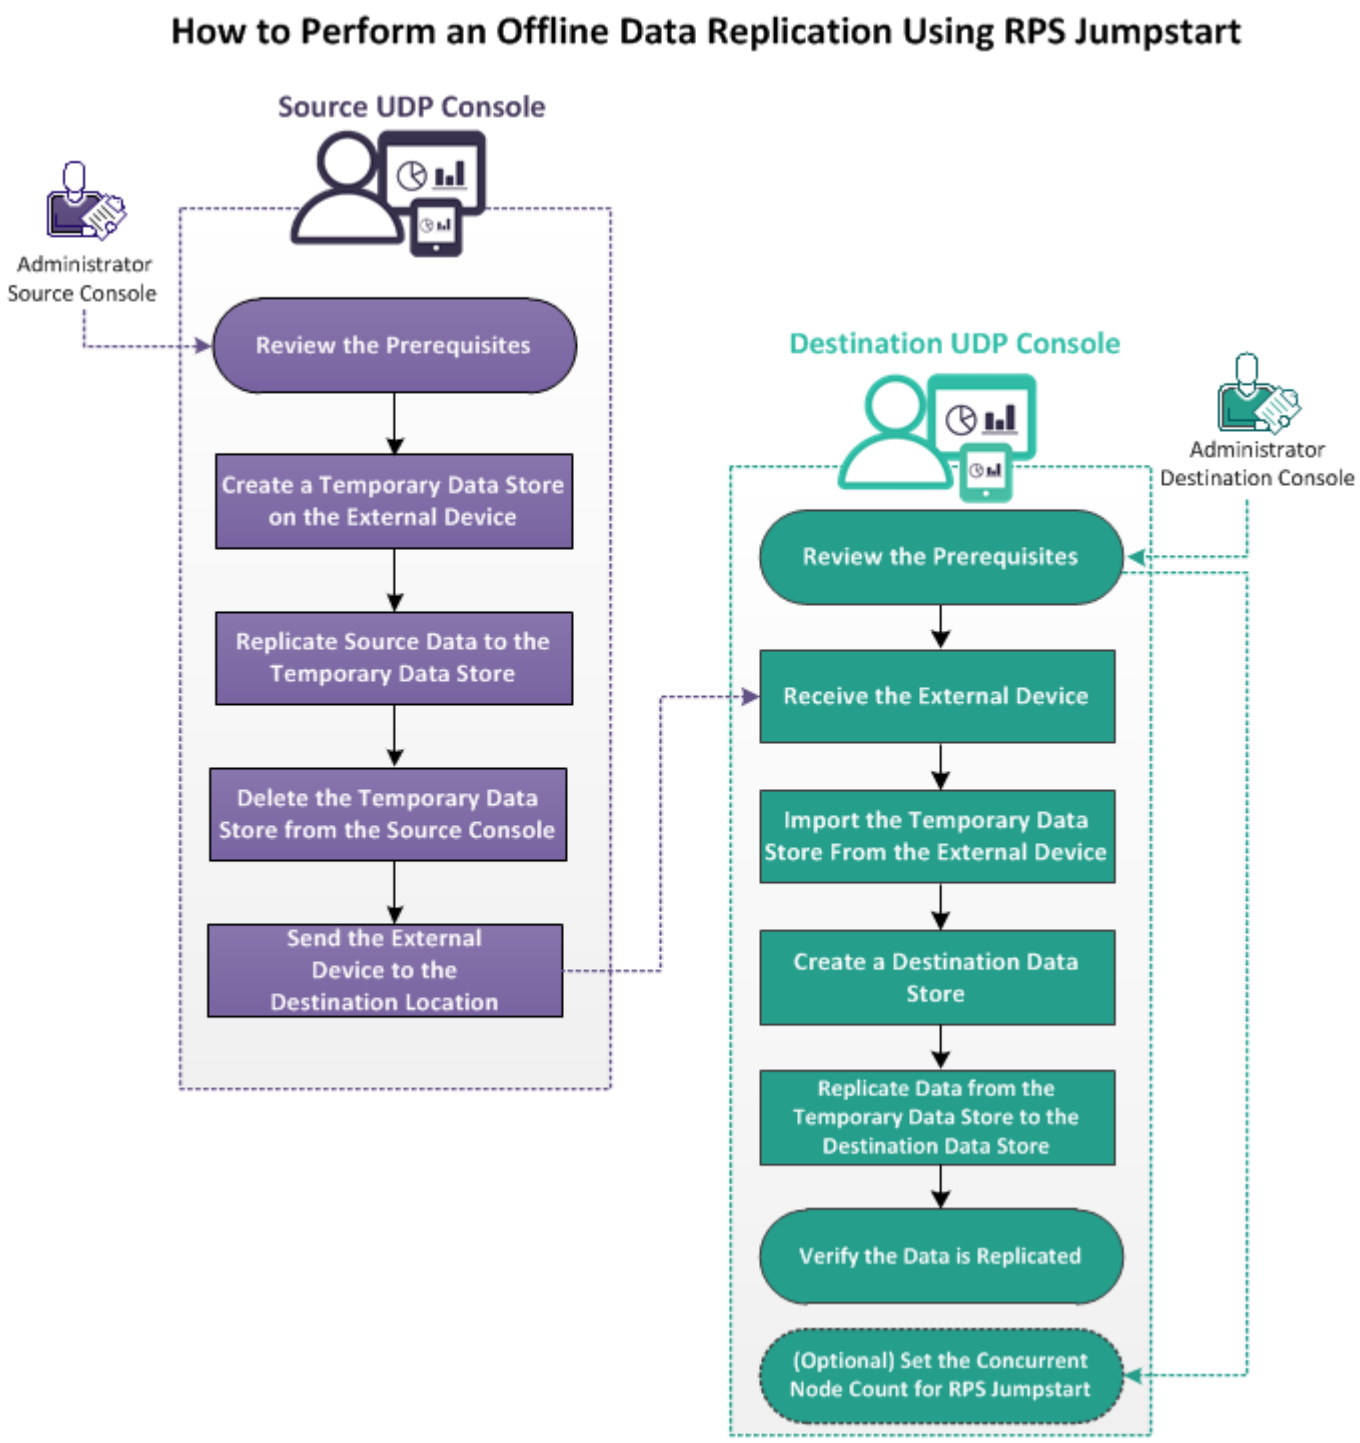
Source Administrator

1. Replicate the source data store to an external device.
2. Send the external device to the destination location.

Destination Administrator

1. Receive the external device.
2. Replicate the source data store from the external device to the destination recovery point server.

The following diagram illustrates how to perform an offline data replication using RPS Jumpstart.



What To Do Next?

- Review the Prerequisites
- Create a Temporary Data Store on an External Device
- Replicate Source Data to the Temporary Data Store
- Delete the Temporary Data Store from the Source Console

- Send the External Device to the Destination Location
- Receive the External Device
- Import the Temporary Data from the External Device
- Create a Destination Data Store
- Replicate Data from the Temporary Data Store to the Destination Data Store
- Verify that the Data is Replicated
- (Optional) Set the Concurrent Node Count for RPS Jumpstart

Review the Prerequisites

Review the following prerequisites before you perform an offline data replication:

- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.
- If the source is a deduplication-enabled data store, then the target must be also deduplication-enabled data store.
- If the source is encryption-enabled data store, then the target also should be encryption-enabled data store.

Administrator—Source Console

- Verify that you have created the source data store.
- Verify that you have at least completed one backup on a data store.
- (Optional) Verify that you have configured the concurrent node count for RPS Jumpstart. For more information about configuring the concurrent node count, see Set the Concurrent Node Count for RPS Jumpstart.

Administrator—Destination Console

- Verify that you have adequate available space for replication.
- Verify that you have the required privileges on the external device.

Create a Temporary Data Store on an External Device

Source Administrator

To import data from an existing data store to an external device, you first create a temporary data store on the external device. To create the temporary data store, connect the external device to the computer.

Follow these steps:

1. Log in to the UDP Console.
2. Navigate to **Destinations**, **Recovery Point Server**.
3. Select the Recovery Point Server.
4. Right-click and select **Add a Data Store**.
5. Enter the details on the **Add a Data Store** page.

   Note: Make sure that the backup destination folders are on the external device.
6. Save the data store.

   The temporary data store is created on the external device.

Replicate Source Data to the Temporary Data Store

Source Administrator

After creating the temporary data store on the external device, you can replicate the source data to the external device using RPS Jumpstart.

Note: Before you begin the RPS Jumpstart process, pause the related plan. Pausing the plan ensures that any scheduled replication job does not start when the Jumpstart process is in progress.

Follow these steps:

1. Click **Actions** and then click **RPS Jumpstart**.

   The **RPS Jumpstart Wizard** opens.

2. Select whether you want to migrate from the same data store or from a shared location.

3. Select the source recovery point server, source data store, and plan.

   The nodes that belong to the plan are displayed.

4. Select the nodes that you want to migrate.

5. Click **Next**.

   The **Select Target Data Store** page opens. If the source data store is encrypted, only the encrypted data stores are displayed in the drop-down list.

6. Select the target data store. The target data store should be on the external device.

7. Click **Finish**.

   The **Recent Events** section on the right pane displays the replication progress.

   After the replication process is complete, the data is replicated to the temporary data store. You can verify the size of both the data store from the **Destinations: Recovery Point Server** page.

Delete the Temporary Data Store from the Source Console

Source Administrator

To maintain data integrity on the external device, delete the temporary data store from the UDP Console before removing the external device.

Note: Deleting the temporary data store from the source UDP Console does not delete the data store files from the external device.

Follow these steps:

1. Right-click the temporary data store and click **Stop**.

   The data store stops.

2. Right-click the temporary data store and select **Delete**.

   A confirmation dialog opens.

3. Click **Yes**.

   The data store is deleted.

   Now you can remove the external device from the computer.

Send the External Device to the Destination Location

Source Administrator

After you remove the external device, send the device to the destination location.

Receive the External Device

Destination Administrator

Receive the external device that includes source data. Now, connect this external device to the destination server.

Import the Temporary Data Store from the External Device

Destination Administrator

Before you can replicate the source data to the destination data store, import the temporary data store to the destination recovery point server.

Follow these steps:

1. Navigate to the **resources** tab and select the recovery point server where you want to import the data store.
2. Right-click the recovery point server and select **Import Data Store**.

   The **Import a Data Store** dialog opens.
3. Select the backup destination folder from the external device.
4. Click **Next**.

   The temporary data store details are displayed. If required, then change the Data, Index, and Hash path.
5. Click **Save**.

   The data store is imported and you can see the data store on the destination Console.

Create a Destination Data Store

Destination Administrator

To replicate data from the temporary data store, first create a destination data store. For more information, see How to Add a Data Store.

Note: You can also use an existing data store as a destination data store.

Replicate Data from the Temporary Data Store to the Destination Data Store

After you create the destination data store, replicate data from the temporary data store to the destination data store. After data is replicated to the destination data store, you can delete the temporary data store.

Follow these steps:

1. Click **Actions** and then click **RPS Jumpstart**.

   The **RPS Jumpstart Wizard** opens.

2. Select the source recovery point server and source data store. Which plan needs to be selected here.

   The nodes are displayed.

3. Select the nodes that you want to migrate.

4. Click **Next**.

   The **Select Target Data Store** page opens. If the source data store is encrypted, only the encrypted data stores are displayed in the drop-down list.

5. Select the target data store. The target data store should be on the external device.

6. Click **Finish**.

   The **Recent Events** section on the right pane displays the replication progress.

   After the replication process is complete, the data is replicated to the temporary data store. You can verify the size of both the data store from the **Destinations: Recovery Point Server** page.

   The data is replicated to the destination data store.

Verify that the Data is Replicated

Destination Administrator

After data is replicated, you can verify whether the replication is successful.

Follow these steps:

1. On the destination Console, navigate to the destination data store on the recovery point server.
2. Verify that the replicated data size matches the source data.

   You have successfully replicated data between two data stores managed from different UDP Consoles.

(Optional) Set the Concurrent Node Count for RPS Jumpstart

Source Administrator

When you start an RPS jumpstart job, the concurrent node value for data store is 4, by default. To specify the concurrent node count, create a key and manually add a DWORD to set the count.

Follow these steps:

1. Log in to the recovery point server.
2. Navigate to the following location:

   HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine
3. Create a key in the Engine directory and name the key as *RPS Jumpstart*.
4. Add the following DWORD in the RPS Jumpstart key:

   JumpStartConCurrencyCount
5. Provide a value for the DWORD.

   Example: If you want to limit to ten nodes per RPS Jumpstart job, then add the following value as DWORD:

   JumpStartConCurrencyCount=10

   The concurrent node count is set for RPS Jumpstart.

How to Perform Reverse Replication

Using this feature, you can replicate data (recovery point) from Cloud, MSP, or Remotely Managed RPS to local UDP or on-premises RPS server. To retrieve your replicated data back to the local UDP or the on-premises server from a recovery point server that is managed from a different Arcserve UDP Console such as Cloud, MSP, or remotely managed RPS, you need to perform reverse replication. In reverse replication, the data gets replicated from a destination data store (on the destination Console) to a source data store (on the source Console).

This section contains the following topics:

Review the Prerequisites

Review the following prerequisites before replicating data:

- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Administrator—Destination Console

- Verify that you have installed Arcserve UDP on the destination server.
- Verify that you have full privileges to create Windows user accounts on the destination server.
- Verify that you have enabled the **Allow the retrieval of recovery points** check box in the Advanced settings during the Replicate from a remotely-managed RPS plan creation.

  Note: If you do not enable the **Allow the retrieval of recovery points** check box and initiate the Reverse Replication job for Remote Console, you are prompted to enable the Allow retrieval remote flag for Reverse Replication.

Administrator—Source UDP Console

- Verify that you have installed Arcserve UDP on the source server.
- Verify that you have at least completed one full backup on a data store.

Perform Ad-hoc Replication

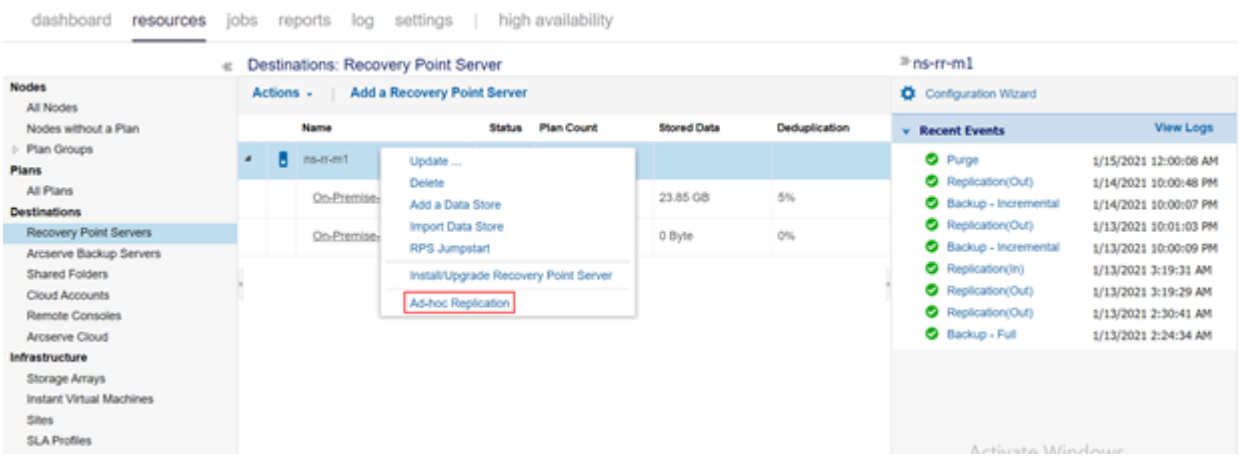
This section explains about how to initiate Ad-hoc replication.

Note: Before performing the Ad-hoc replication, make sure the forward replication is executed successfully. To verify whether the forward replication is successful, see Verify the Data is Replicated.

Follow these steps:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Destinations** > **Recovery Point Servers**.

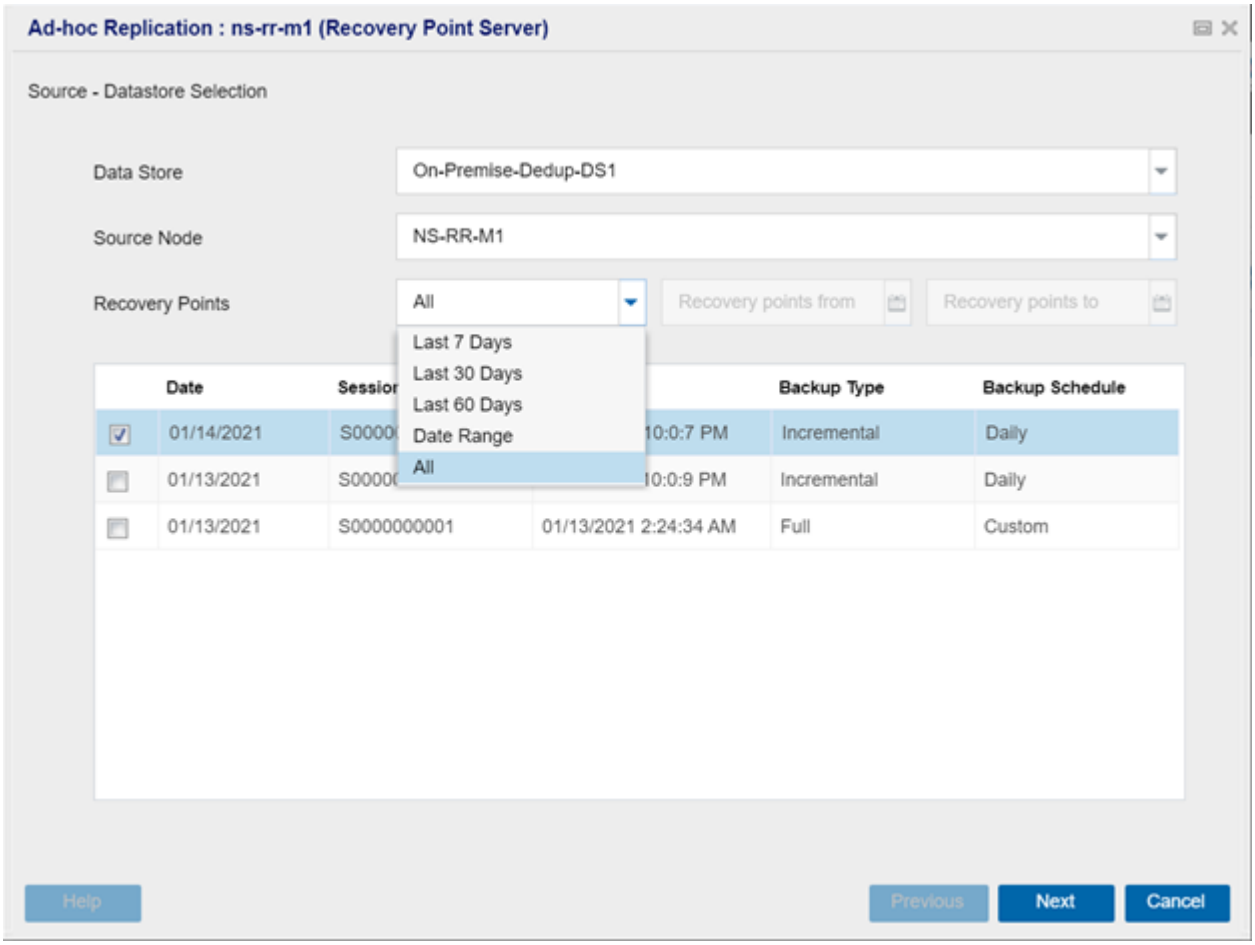
   The list of recovery point servers appears in the center pane.
3. Right-click the recovery point server, and then select **Ad-hoc Replication**.



Wait until you get connected to the recovery point server.

The Ad-hoc Replication : xx-xx-xx (Recovery Point Server) dialog opens.

4. On the Source - Datastore Selection page, do the following:



- Data Store – Select the source data store from the drop-down list.
- Source Node – Select the source node from the drop-down list.
- Recovery Points – Do one of the following:

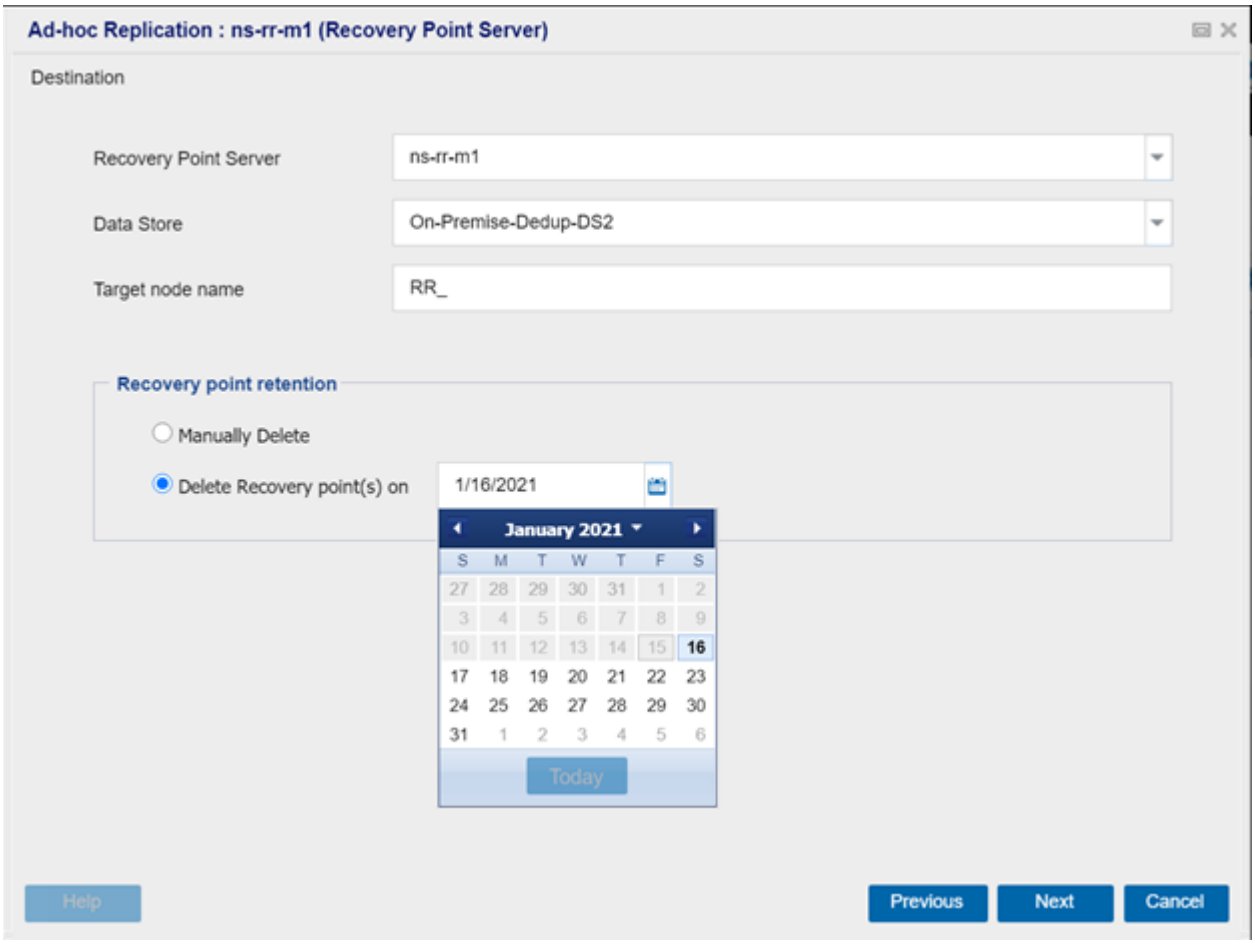
- If you want all the recovery points you backed up to display, select **All** from the drop-down list.
- If you want the recovery points to display for a specific duration, select one of the following from the drop-down list:
  - Last 7 days
  - Last 30 days
  - Last 60 days
  - Date Range

  Note: If you selected the **Date Range**, set the date on **Recovery points from** and **Recovery point to** fields.

All the associated recovery points between the specified duration are displayed.

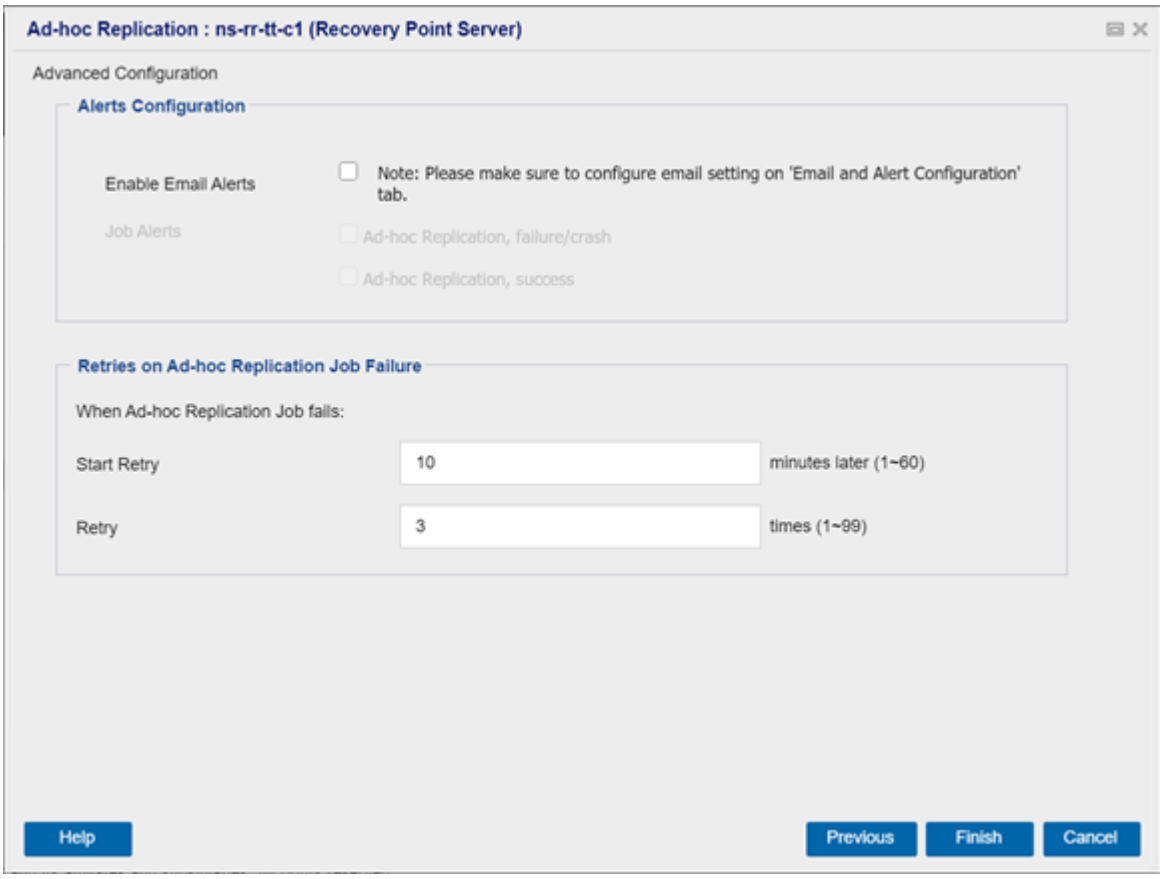
- From the table, click the checkboxes to select one or more recovery points.

5. Click **Next**.
6. On the Destination page, do the following:



- Recovery Point Server – Select the destination recovery point server from the drop-down list.
- Data Store – Select the destination data store from the drop-down list.
- Target node name – Type a target node name of your choice.
- Recovery point retention – Delete the recovery points in one of the following ways:
  - To delete recovery points manually, select **Manually delete** option.
  - If you want the recovery points to get deleted automatically, select the **Delete Recovery point(s) on** option, and then specify the date. All the recovery points corresponding to that date get deleted automatically.

7. Click **Next**.
8. On the Advanced Configuration page, do the following:



- For Alerts Configuration

  Enable Email Alerts

  Lets you enable email alerts. You can configure email settings and specify the types of alerts that you want to receive in an email. When you select the **Note: Please make sure to configure email setting on 'Email and Alert Configuration' tab.** check box, the following option is enabled for your selection:

  Job Alerts

  Lets you select the type of job alert emails that you want to receive.

  - If you want to receive email alerts for Ad-hoc replication failure or crash, select the **Ad-hoc Replication, failure/crash** check box.

    Note: This email alert is sent with a high importance. The email alerts that have a high importance level setting display a visual indicator of an exclamation point in their Inbox.

  - If you want to receive email alerts for Ad-hoc replication success, select the **Ad-hoc Replication, Success** check box.

- For Retries on Ad-hoc Replication Job Failure

  When the Ad-hoc Replication job fails, do the following:

  Start Retry

  Lets you specify the time (in minutes) to restart the replication job after the job fails. For example, if you specify 10 minutes, then the replication job restarts after 10 minutes of its failure.

  Limit: 1 to 60

  Retry

  Lets you specify the number of retries that you want to perform if the replication job fails. The replication job runs until the job is successful, or until the limit is reached. After the number of retries is over, the replication job runs only at the next scheduled time.

  Limit: 1 to 99

9. Click **Finish**.

   Note: If you want to change the settings of the previous screens, click **Previous**, and then change the setting as needed.

   The Information dialog box appears and displays the *Ad-hoc Replication Job Submitted Successfully!* message.

10. Click **OK** to close the Information dialog box.

    The Ad-hoc replication is performed successfully.

Monitor the Status of Ad-hoc Replication

Source Administrator

After Ad-hoc replication is successful, you can monitor the status of the Ad-hoc replication job.

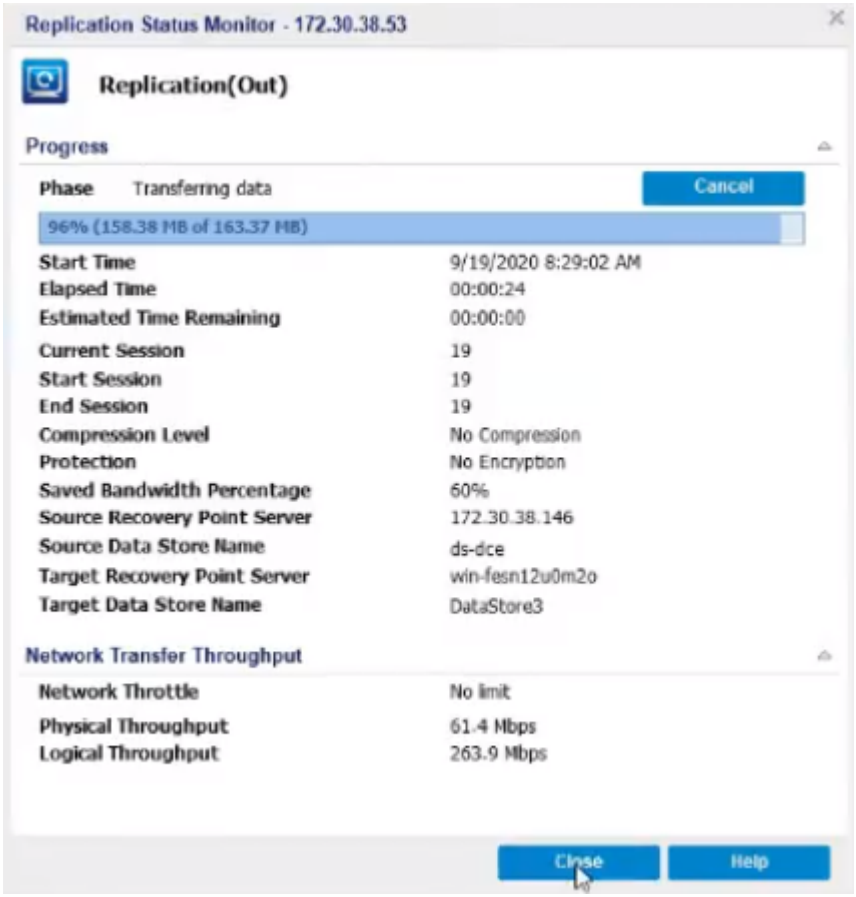
Follow these steps:

1. Navigate to **jobs** > **Jobs in Progress**.

   The Replication(Out) and Replication(In) jobs that are in progress appear in the center pane.

2. To monitor the status of a specific job, click **Details** next to the Progress bar corresponding to that job.

   The Replication Status Monitor – xx.xx.xx.xx dialog appears.



3. Wait until the data transfer is 100%, and then click **Close**.

You have successfully monitored the status of Ad-hoc replication.

Verify whether the Ad-hoc Replication is Successful

Source Administrator

After data is replicated, you can verify whether Ad-hoc replication is successful.

Follow these steps:

1. On the Source Console, navigate to **jobs** > **Jobs Successfully Completed**.

   A list of jobs that are completed appear in the center pane.
2. Click the finished reverse replication job. The right pane refreshes.
3. Under Job Details, click the **View Logs** hyperlink.
4. Verify that the replication of data is successful.

You have successfully verified the Ad-hoc replication.

Perform Reverse Replication from Remote Console to On-Premises RPS Server

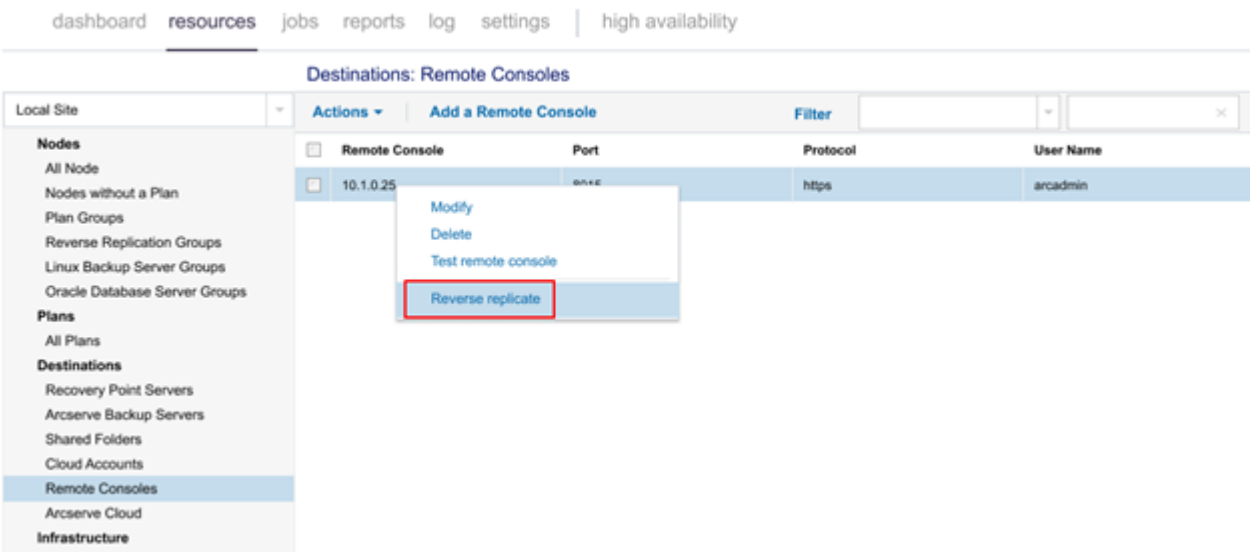
This section explains about how to perform reverse replication from Remote Console to local UDP or on-premises RPS server.

Note: Before performing the reverse replication, make sure the forward replication is executed successfully. To verify whether the forward replication is successful, see Verify the Data is Replicated.

Follow these steps:

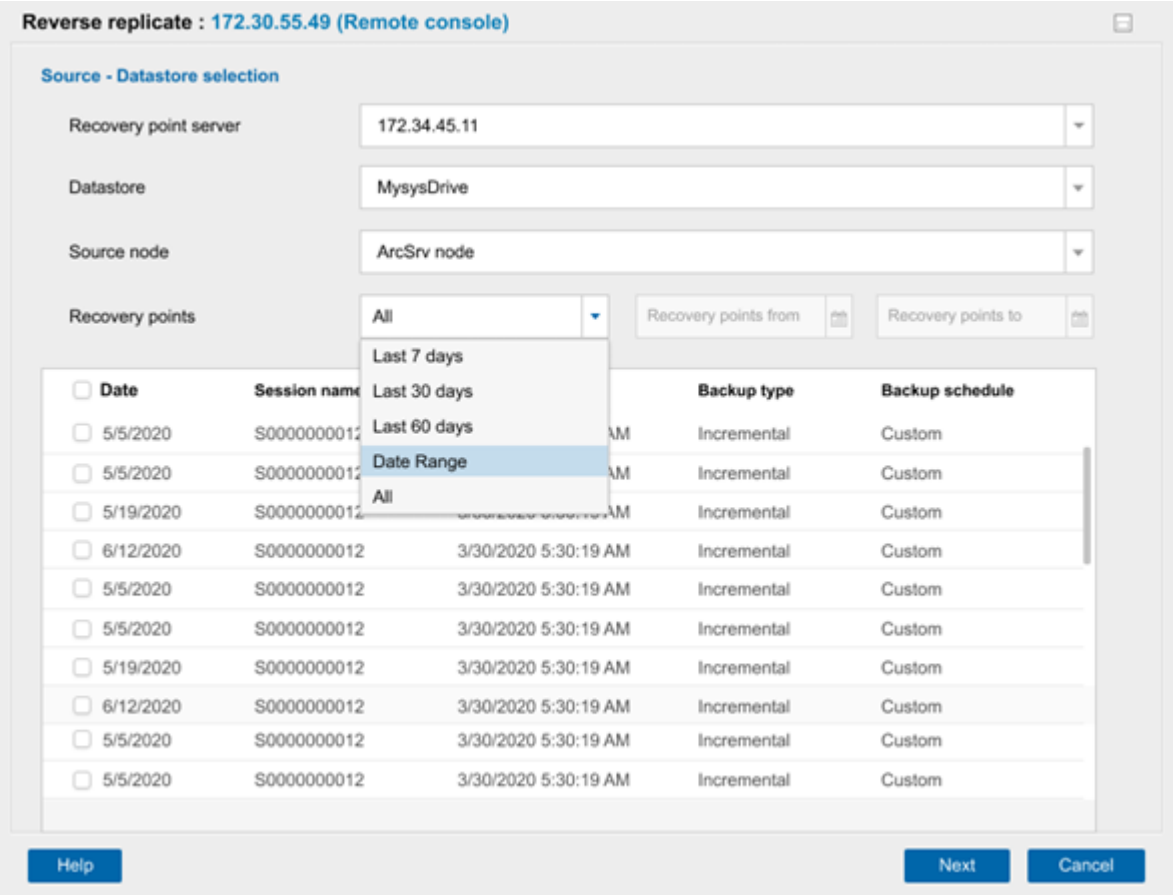
1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Destinations** > **Remote Consoles**.

   The list of remote consoles appears in the center pane.
3. Right-click the remote console, and then select **Reverse replicate**.



   Wait until you get connected to the remote console.

   The Reverse Replicate: xxx.xx.xx.xxx (Remote Console) dialog opens.
4. On the Source - Datastore Selection page, do the following:



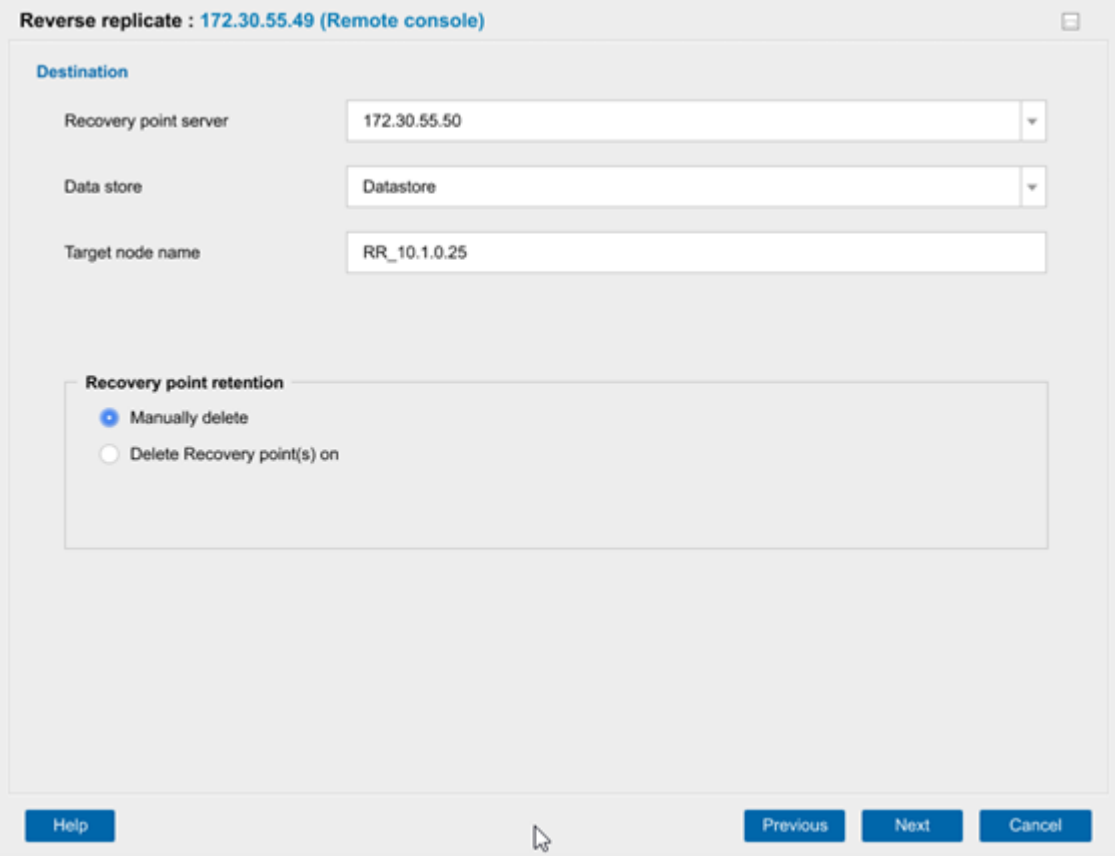
- Recovery point server – Select the source recovery point server from the drop-down list.
- Datastore – Select the source datastore from the drop-down list.
- Source node – Select the source node from the drop-down list.
- Recovery points – Do one of the following:
  - If you want all the recovery points you backed up to display, select **All** from the drop-down list.
  - If you want the recovery points to display for a specific duration, select one of the following from the drop-down list:
    - Last 7 days
    - Last 30 days
    - Last 60 days
    - Date Range

    Note: If you selected the **Date Range**, set the date on **Recovery points from** and **Recovery point to** fields.

All the associated recovery points between the specified duration are displayed.

- From the table, click the checkboxes to select one or more recovery points.

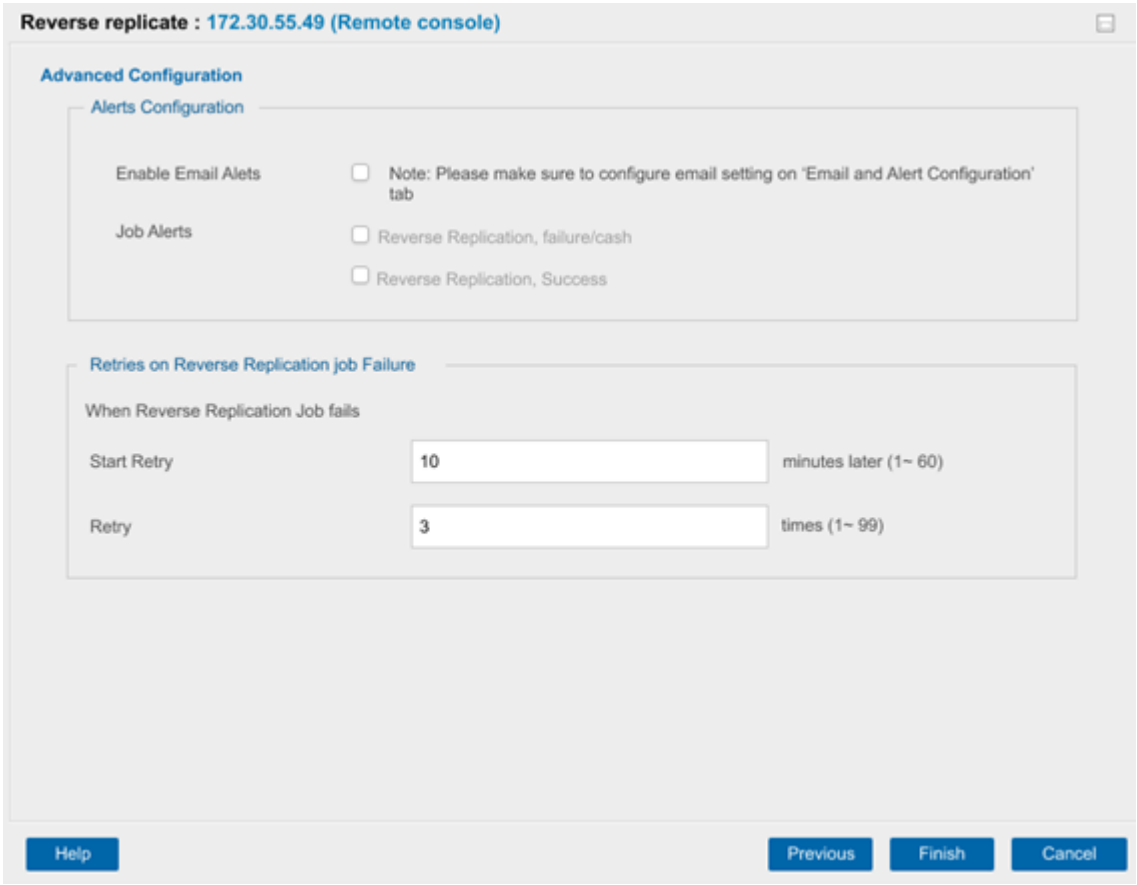
5. Click **Next**.
6. On the Destination page, do the following:



- Recovery point server – Select the destination recovery point server from the drop-down list.
- Data store – Select the destination data store from the drop-down list.
- Target node name – Type a target node name of your choice.
- Recovery Point Retention – Delete the recovery points in one of the following ways:
  - To delete recovery points manually, select **Manually delete** option.
  - If you want the recovery points to get deleted automatically, select the **Delete Recovery point(s) on** option, and then specify

the date. All the recovery points corresponding to that date get deleted automatically.

7. Click **Next**.
8. On the Advanced Configuration page, do the following:



- For Alerts Configuration

  Enable Email Alerts

  Lets you enable email alerts. You can configure email settings and specify the types of alerts that you want to receive in an email. When you select the **Note: Please make sure to configure email setting on 'Email and Alert Configuration' tab.** check box, the following option is enabled for your selection:

  Job Alerts

  Lets you select the type of job alert emails that you want to receive.

  - If you want to receive email alerts for reverse replication failure or crash, select the **Reverse Replication, failure/crash** check box.

Note: This email alert is sent with a high importance. The email alerts that have a high importance level setting display a visual indicator of an exclamation point in their Inbox.

- If you want to receive email alerts for reverse replication success, select the **Reverse Replication, Success** check box.

- For Retries on Reverse Replication Job Failure

  When the reverse replication job fails, do the following:

  Start Retry

  Lets you specify the time (in minutes) to restart the replication job after the job fails. For example, if you specify 10 minutes, then the replication job restarts after 10 minutes of its failure.

  Limit: 1 to 60

  Retry

  Lets you specify the number of retries that you want to perform if the replication job fails. The replication job runs until the job is successful, or until the limit is reached. After the number of retries is over, the replication job runs only at the next scheduled time.

  Limit: 1 to 99

9. Click **Finish**.

   Note: If you want to change the settings of the previous screens, click **Previous**, and then change the setting as needed.

   The Information dialog box appears, which displays the *Reverse Replication Job Submitted Successfully!* message.

10. Click **OK** to close the Information dialog box.

    Reverse replication is completed from Remote Console to local UDP or on-premises server successfully.

Monitor the Status of Reverse Replication

Source Administrator

After reverse replication is successful, you can monitor the status of the reverse replication job.

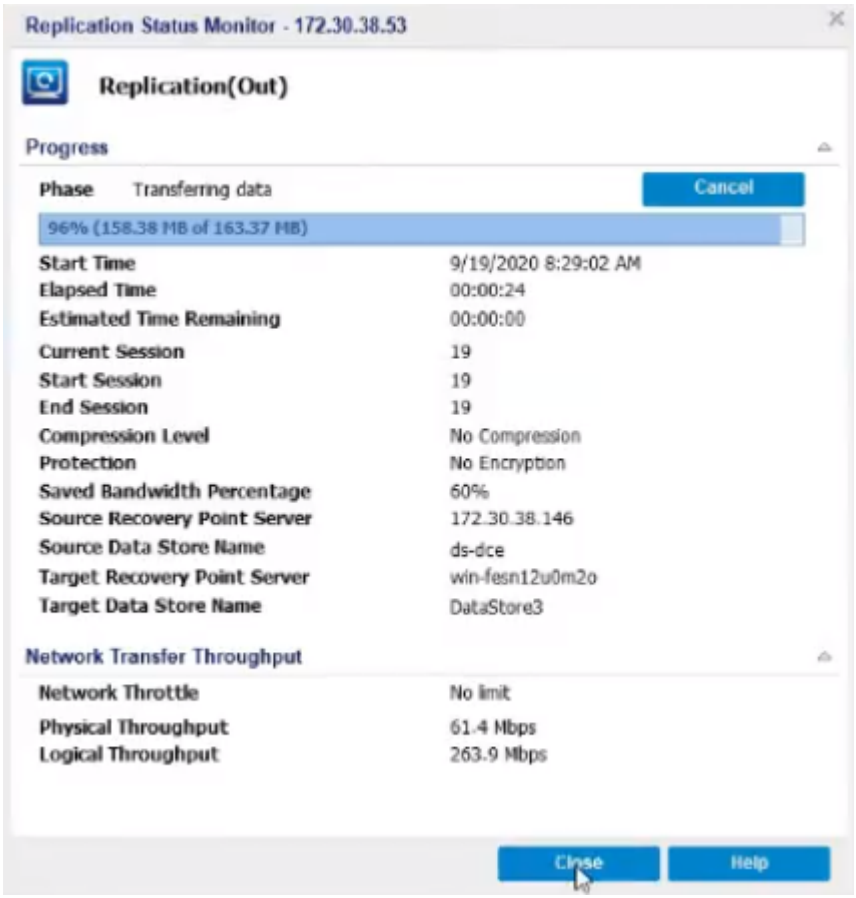
Follow these steps:

1. Navigate to **jobs** > **Jobs in Progress**.

   The Replication(Out) and Replication(In) jobs that are in progress appear in the center pane.

2. To monitor the status of a specific job, click **Details** next to the Progress bar corresponding to that job.

   The Replication Status Monitor – xx.xx.xx.xx dialog appears.



3. Wait until the data transfer is 100%, and then click **Close**.

You have successfully monitored the status of reverse replication.

Verify whether the Reverse Replication is Successful

Source Administrator

After data is replicated, you can verify whether reverse replication is successful.

Follow these steps:

1. On the Source Console, navigate to **jobs** > **Jobs Successfully Completed**.

   A list of jobs that are completed appear in the center pane.

2. Click the finished reverse replication job. The right pane refreshes.

3. Under Job Details, click the **View Logs** hyperlink.
4. Verify that the replication of data is successful.

You have successfully verified the reverse replication of data from Remote Console to local UDP or on-premises RPS server.

How to Create a Copy Recovery Points Plan

Using Arcserve Unified Data Protection, you can copy the recovery points to Cloud or a shared folder or local volume to protect the recovery points. This process helps ensure that you have an additional copy of the recovery points if your original recovery points are accidentally deleted. The copy recovery points task copies the recovery points from the backup destination to Cloud or a shared folder or a local volume only. You cannot copy the recovery point to a recovery point server.

You can add only one Copy Recovery Points task in a plan.

Notes:

- In the current version, copy recovery Point jobs are not supported if **Backup: Agent-Based Linux** is created as Task1.
- Copy recovery Point jobs always runs on Agent, even if the backup is configured on RPS.

  For agentless VM backup, the UDP agent proxy used in Task1 processes the Copy to Recovery Point job.

Troubleshooting: Bandwidth Congestion with Copy Recovery Point to Cloud Jobs

What To Do Next?

- Review the Prerequisites and Considerations
- Create a Plan with a Backup Task
- Add a Copy Recovery Points Task to the Plan
- Verify the Plan

Review the Prerequisites and Considerations

Verify if you have completed the following prerequisites:

- Log into the Console.
- Install the server component and create Data Stores if you want to store the backup data to recovery point servers.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Create a Plan with a Backup Task

A plan includes different types of tasks that you want to perform. To create a copy recovery points task, you must first have a valid recovery point. To get a valid recovery point, you have to create a backup task.

The backup task performs a backup of the source nodes and stores the data to the specified destination. Copy Recovery Points is supported for both Agent-based Windows and Host-based agentless backup. The following procedure explains the steps to create the agent-based Windows backup task. You cannot perform copy recovery point for Linux Agent-based backup plans.

Note: For more information about host-based agentless backup, see How to Create a Host-Based Virtual Machine Backup Plan.

For more information about UNC path backup, see How to Create a UNC Path Backup Plan.

For more information about Exchange Online backup, see How to Create an Exchange Online Backup Plan.

Follow these steps:

1. Click the **resources** tab on the Console.
2. From the left pane, navigate to **Plans**, and click **All Plans**.

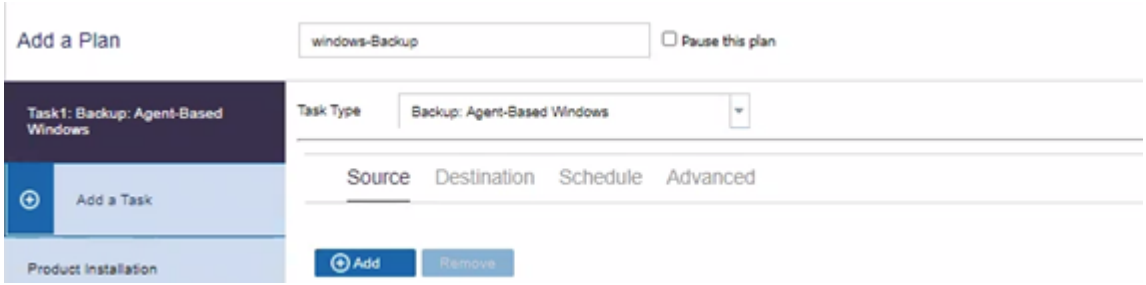
   If you have created plans earlier, those plans are displayed on the center pane
3. On the center pane, click **Add a Plan**.

   Add a Plan opens.
4. Enter a plan name.
5. (Optional) Select **Pause this plan** check box to pause the plan.

   The plan will not run until you clear the check box to resume the plan.

   Note: If a plan is paused, then any in-progress job is not paused but all corresponding scheduled jobs associated with that plan are paused. However, you can manually run a job. For example, you can manually run backup job and replication job for a node even if the respective plan is paused. In such case, the following task to the on-demand (manual) job does not run. For example, there is a replication task after an on-demand backup job, the replication job does not run for the on-demand backup job. You need to manually run the replication job. When you resume the plan, the pending jobs do not resume immediately. After you resume the plan, the pending jobs run from the next scheduled time.

6. From the **Task Type** drop-down list, select **Backup, Agent-Based Windows**.



Now, specify the Source, Destination, Schedule, and Advanced details.

Specify the Source

The Source page lets you specify the source nodes that you want to protect. You can select more than one node in a plan. If you have not added any nodes to the Console, you can add nodes when you create or modify a plan from the Source page. You can also save a plan without adding any source nodes. The plan gets deployed only after you add source nodes.

Follow these steps:

1. On the Source tab, click **Add**.
2. Select one of the following options:

   Select Nodes to Protect in Arcserve UDP

   Opens the Select Nodes to Protect in Arcserve UDP dialog. Select the nodes from the displayed list. Select this option if you have already added the nodes to the Console.

   Add Windows Node

   Opens the Add nodes to a plan dialog. Select this option if you have not added the nodes and you want to manually add the nodes to protect.

   Discover Nodes from Active Directory

   Opens the Add Nodes to Arcserve UDP Console dialog. Select this option if you want to discover and add nodes from the Active Directory.

3. If you selected the **Select Nodes to Protect in Arcserve UDP** option, on the Select Nodes to Protect dialog, do the following:

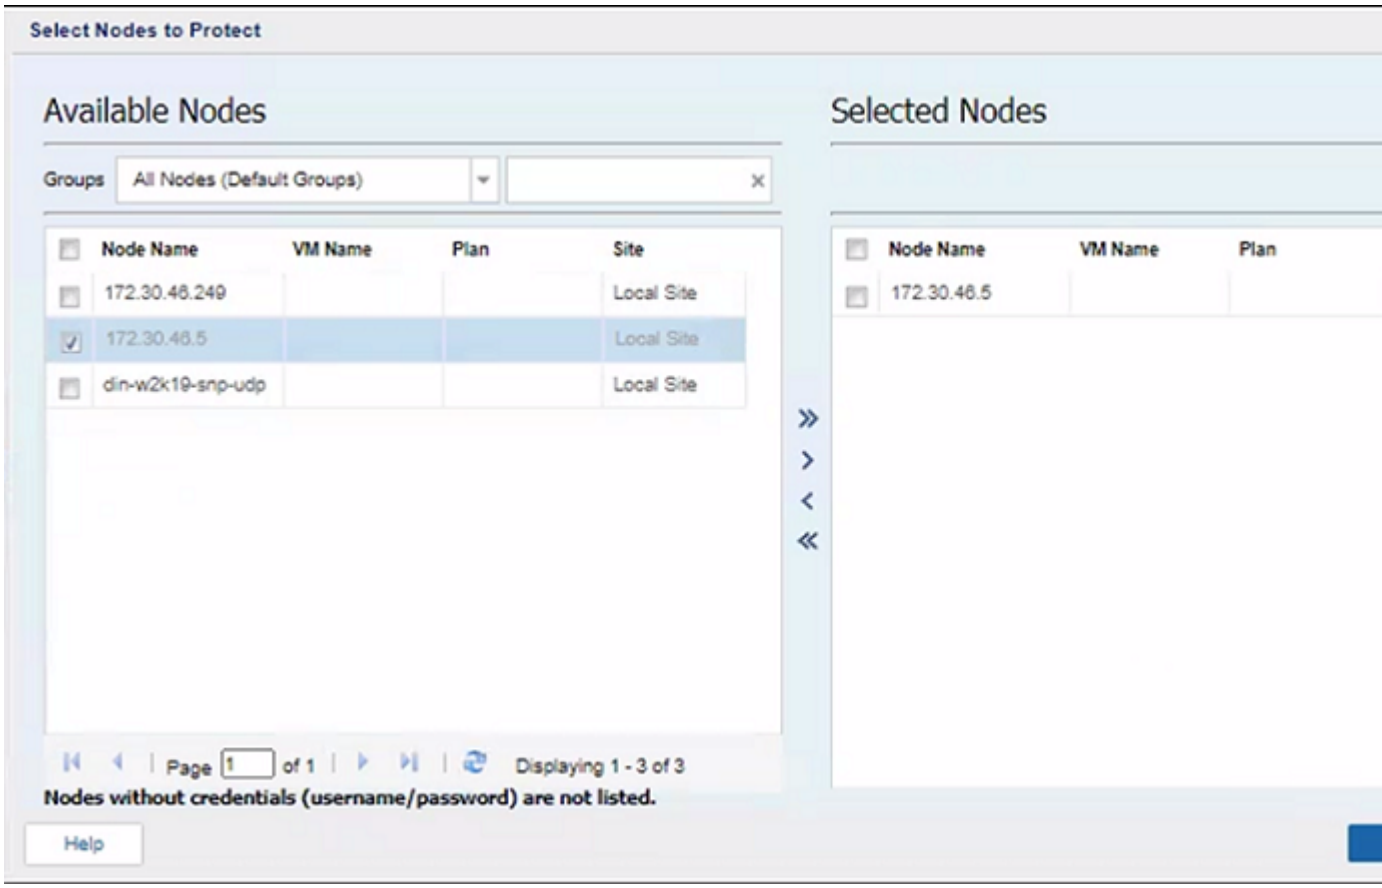
   a. (Optional) Select a filter from the Groups drop-down list to filter nodes. You can enter keywords to further filter your nodes.

      The nodes are displayed on the Available Nodes column.

   b. Select the nodes from the Available Nodes column, and then click the single right arrow (>) or double right arrow (>>) to move the selected nodes or all the nodes respectively to the Selected Nodes column.

      The nodes are displayed in the Selected Nodes column.

   c. Click **OK** to close the dialog.



4. To choose the Protection Type, select one of the following options:

   Back up all volumes

   Prepares a backup snapshot of all the volumes.

   Back up selected volumes

   Prepares a backup snapshot of the selected volume.

   The source is specified.

Specify the Destination

The destination is a location where you store the backup data. You must at least specify the destination to save the plan.

Follow these steps:

1. Select one of the following **Destination Type**:

   Local disk or shared folder

   Specifies that the backup destination is either a local destination or a shared folder. If you select this option, you can save data as either recovery points or recovery sets. The recovery points and recovery sets options are available on the **Schedule** tab.

   Arcserve UDP Recovery Point Server

   Specifies that the backup destination is a recovery point server. If you select this option, then data is stored as recovery points. You cannot store data as recovery sets.

2. If you have selected **Arcserve UDP Recovery Point Server**, then provide the following details:
   a. Select a recovery point server
   b. Select a data store. The list displays all data stores that are created at the specified recovery point server.
   c. Provide a session password.
   d. Confirm the session password.
3. If you have selected **Local disk or shared folder**, then provide the following details:
   a. Provide the full path of the local or network destination. For the network destination, specify the credentials with the write access.
   b. Select the encryption algorithm. For more information, see Encryption Settings.
   c. Optionally, provide an encryption password.
   d. Confirm the encryption password.
   e. Select a type of compression. For more information, see Compression Type.

   Notes:

   - Copy Recovery Point (CRP) doesn't support custom backups even when the CRP schedule is added. Additionally, the CRP task does not support the Custom backup schedule.

- The CRP to cloud job supports only from a console. It doesn't support a standalone agent.
- Both *Local disk/share folders* and *backups stored on RPS* support the CRP tasks.

The destination is specified.

Specify the Schedule

The Schedule page lets you define a schedule for Backup. Copy Recovery Points supports Daily, Weekly and Monthly backups from console. After you define a schedule, the jobs run automatically per the schedule. You can add Daily, Weekly and Monthly schedules as well as can provide retention settings.

Note: For more information on scheduling and retention settings, see Understanding Advanced Scheduling and Retention.

Follow these steps:

1. (Optional) Select the option to manage recovery points. This option is visible only if you have selected Local or shared folder as your backup destination.

   Retain by Recovery Points

   The backup data is stored as recovery points.

   Retain by Recovery Sets

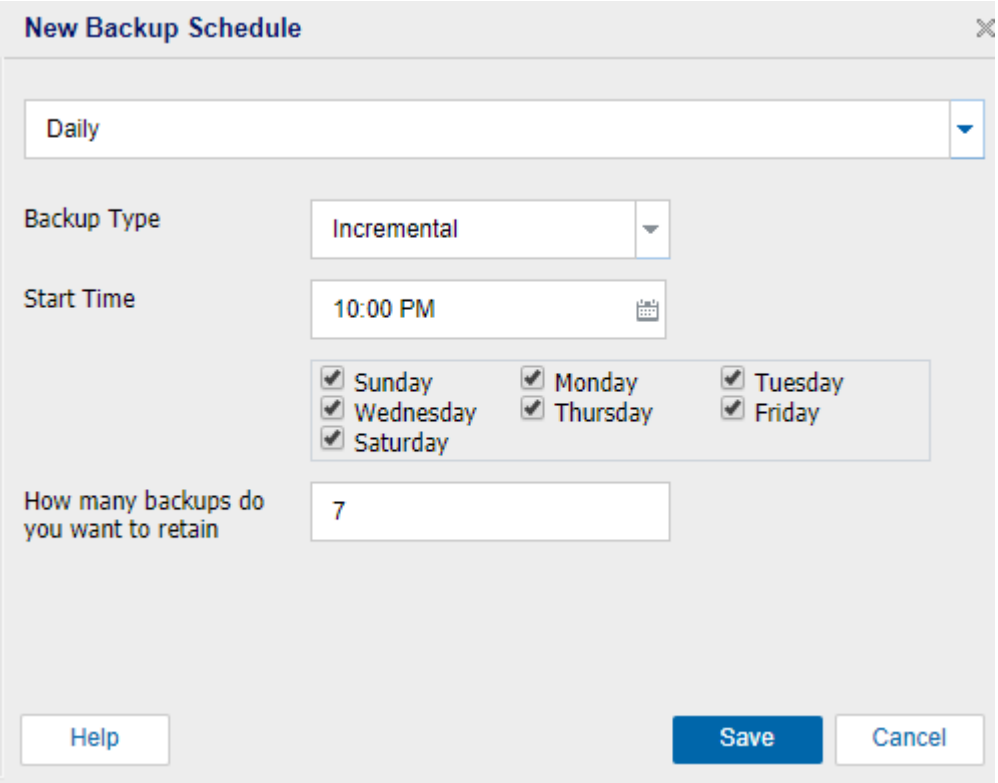
   The backup data is stored as recovery sets.

   Note: Retain by Recovery Sets is not supported for Copy Recovery Points.

2. Add backup schedule.

   a. Click **Add** and select **Add Backup Schedule**.

      The **New Backup Schedule** dialog opens.



b. Select one of the following options:

Daily

Specifies the backup schedule that occurs once a day. By default, all the days of the week are selected for Daily backup. If you do not want to run the backup job on a specific day, clear the check box for that day of the week.

Weekly

Specifies the backup schedule that occurs once a week.

Monthly

Specifies the backup schedule that occurs once a month.

c. Select the backup type.

Full

Determines the backup schedule for Full Backups. As scheduled, Arcserve UDP performs a full backup of all used blocks from the source machine. A full backup typically consumes time depending on the backup size.

Incremental

Determines the backup schedule for Incremental Backups.

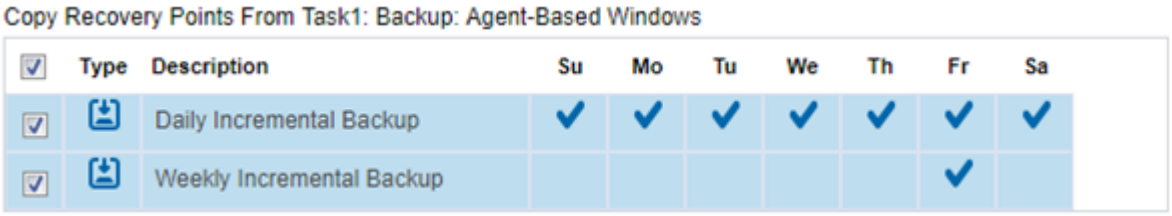
As scheduled, Arcserve UDP incrementally backs up only those blocks that have changed since the last successful backup. The advantages of Incremental Backups are that it is a fast backup and it produces a small backup image. This is the most optimal way to perform a backup.

d. Specify the backup start time.

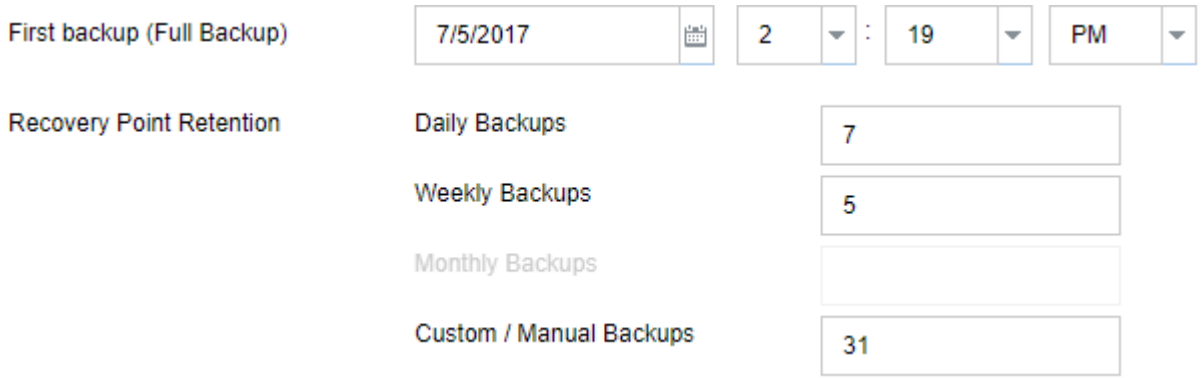
e. (Optional) Select the **Repeat** check box and specify the repeat schedule.

f. Click **Save**.

The Backup Schedule is specified and displayed on the **Schedule** page.



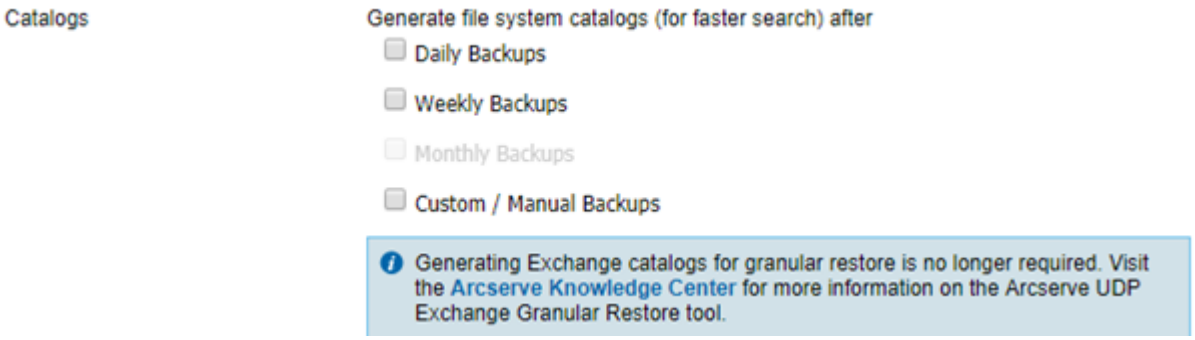
3. Specify the start time for the scheduled backup.



4. Specify the recovery points retention settings for Daily, Weekly, and Monthly schedule.

   These options are enabled if you have added the corresponding backup schedule. If you modify the retention settings on this page, the changes are reflected on the Backup Schedule dialog.

5. Specify the catalog details.



Catalogs let you generate the File System catalog. The File System catalog is required to perform faster and better search. If you select the catalog check boxes, the catalogs are enabled depending on the type of backup that you have specified. Clear the check box to disable generating the catalog.

The schedule is specified.

Understanding Advanced Scheduling and Retention

The scheduling option lets you specify a Custom schedule or a Daily / Weekly / Monthly schedule, or both the schedules. In the Custom schedule, you can configure the backup schedule for each day of the week and you can add up to four backup schedules each day. You can select a specific day of a week and create a time window to define when to run backup and at what frequency.

Schedule	Supported Job	Comments
Backup	Backup job	Define time windows to run backup jobs.
Backup throttling	Backup job	Define time windows to control the backup speed.
Merge	Merge job	Define when to run merge jobs.
Daily schedule	Backup job	Define when to run daily backup jobs.
Weekly schedule	Backup job	Define when to run weekly backup jobs.
Monthly schedule	Backup job	Define when to run monthly backup jobs.

You can also specify the retention settings for the recovery points.

Note: Set the retention settings within each plan to control how data for the nodes assigned to that plan are retained at the target data store.

Schedules for Daily / Weekly / Monthly backups are independent to the Custom schedule, and each other. You can configure only to run Daily backup, Weekly backup or Monthly backup, without configuring the Custom schedule.

Backup Job Schedule

You can add four time windows per day in your backup schedule. A valid time window is from 12:00 AM until 11:59 PM. You cannot specify a time window such as 6:00 PM to 6:00 AM. In such cases, you must manually specify two different time windows.

For each time window, the start time is inclusive, and the end time is exclusive. For example, you have configured to run Incremental Backup every one hour between 6:00 AM and 9:00 AM and the backup will start at 6:00 AM. This means the backup runs at 6:00 AM, 7:00 AM, 8:00 AM, but NOT at 9:00 AM.

Note: If you want to run the backup job repeatedly until the end of day, set the schedule until 12:00 AM. For example, to run backup job every 15 minutes for the entire day, set the schedule from 12:00 AM to 12:00 AM, every 15 minutes.

Backup Throttle Schedule

Backup throttle schedule lets you control the backup throughput speed that in turn controls the resource usage (disk I/O, CPU, network bandwidth) of the server being backed up. This is useful if you do not want to affect the server

performance during business hours. You can add four time windows per day in your backup throttle schedule. For each time window, you can specify a value, in MB per minute. This value is used to control the backup throughput. Valid values are from 1 MB/minutes to 99999 MB/minutes.

If a backup job extends its specified time, then the throttle limit adjusts according to the specified time window. For example, you have defined the backup throttle limit as 500 MB/minute from 8:00 AM to 8:00 PM, and 2500 MB/minute from 8:00 PM to 10:00 PM. If a backup job starts at 7:00 PM and it runs for three hours, then, from 7:00 PM to 8:00 PM the throttle limit is 500 MB/minute and from 8:00 PM to 10:00 PM the throttle limit is 2500 MB/minute.

If you do not define any backup schedule and backup throughput schedule, the backup runs as fast as it can.

Merge Schedule

Lets you merge recovery points based on the provided schedule.

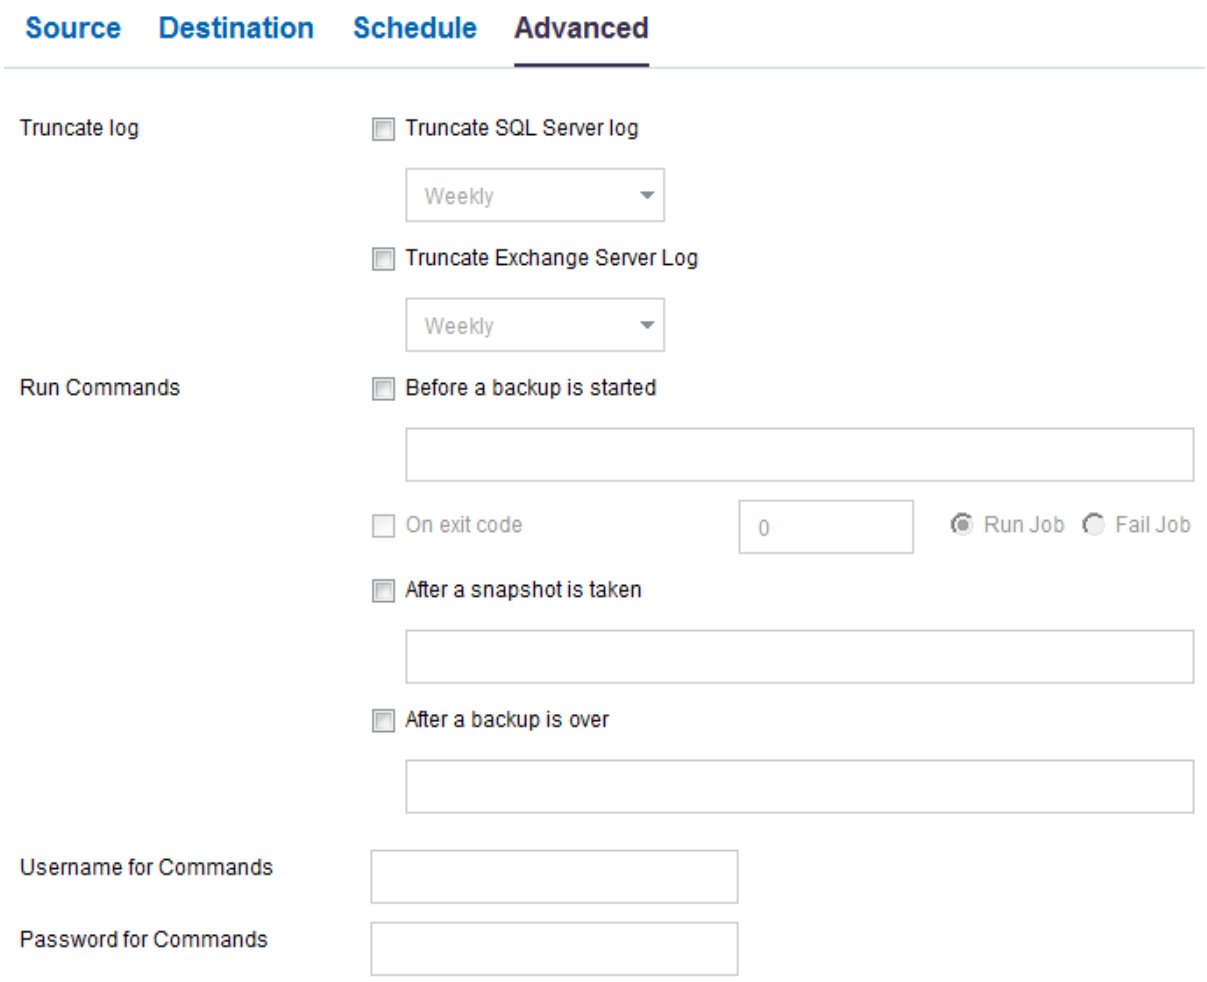
Consider the following points for the merge job:

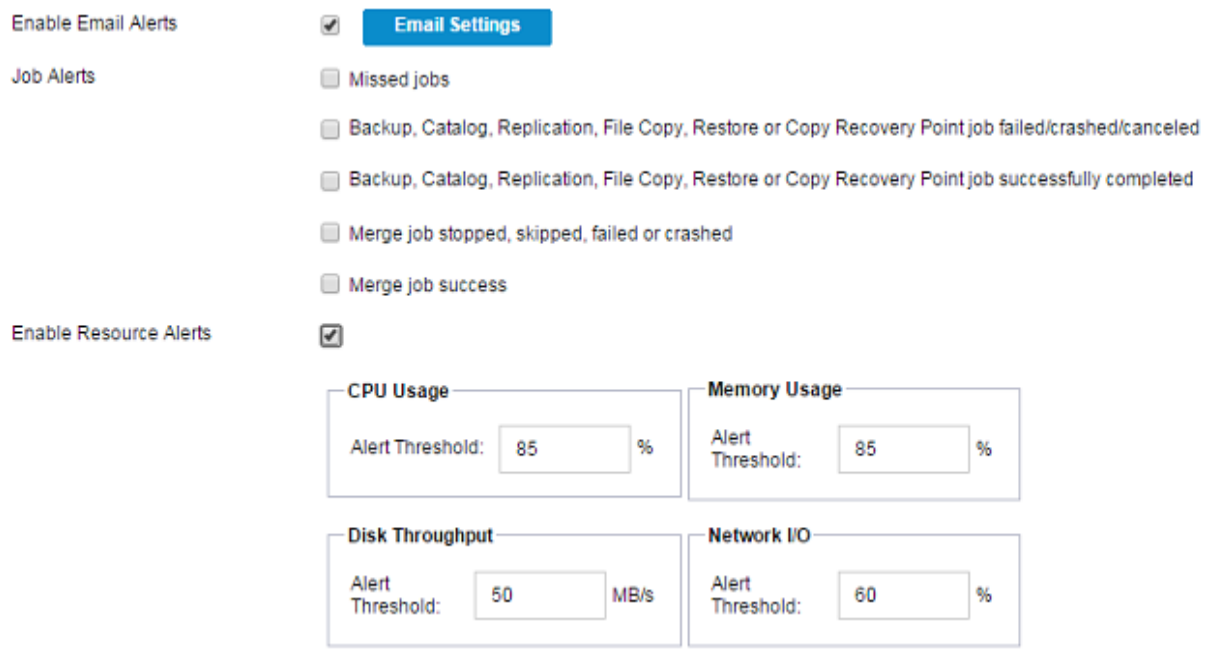
- At any given time only one merge job can run for a node.
- If a merge job starts, it has to complete before the next merge job can start. This means, if one or more sets of recovery points are merged, new recovery point cannot be added to this merge process, until the merge process of the current set of recovery point completes.
- If a merge job is processing more than one set of recovery points (for example set [1~4], set [5~11], and set [12~14]; they are three sets), recovery point server processes these sets one by one.
- If a merge job is resumed after a pause, the job detects at which point it is paused and resumes the merge from the break-point.

Specify the Advanced Settings

The **Advanced** tab lets you specify some advanced settings for the backup job. The advanced settings include providing truncate log settings, providing the location of any scripts, and email settings.

The following image displays the Advanced tab:





Follow these steps:

1. Specify the following details.

Truncate Log

Lets you specify the schedule to truncate logs for SQL Server and Exchange Server. You can specify the schedule as **Daily**, **Weekly**, or **Monthly**.

User Name

Lets you specify the user who is authorized to run a script.

Password

Lets you specify the password of the user who is authorized to run the script.

Run a command before backup is started

Lets you run a script before the backup job starts. Specify the path where the script is stored. Click **On exit code** and specify the exit code for **Run Job** or **Fail Job**. **Run Job** indicates that the backup job will continue when the script returns the exit code. **Fail Job** indicates that the backup job will stop when the script returns the exit code.

Run a command after snapshot is taken

Lets you run a script after the backup snapshot is taken. Specify the path where the script is stored.

Run a command after backup is over

Lets you run a script after the backup job is completed. Specify the path where the script is stored.

Enable Email Alerts

Lets you enable email alerts. You can configure email settings and can specify the types of alerts that you want to receive in an email. When you select this option, the following options are enabled for your selection.

Email Settings

Lets you configure the email settings. Click **Email Settings** and configure the email server and proxy server details.

Job Alerts

Lets you select the types of job emails you want to receive.

Enable Resource Alerts

Lets you specify a threshold for CPU Usage, Memory Usage, Disk Throughput, Network I/O. You can provide the value in percentage. You will receive an email when the Alert Threshold value exceeds.

2. Click **Save**.

Note: When you select a node as a backup source or backup proxy, Arcserve UDP checks whether the agent is installed on the node and if it is the latest version. Arcserve UDP then displays a verification dialog that lists all the nodes that either have an outdated version of the agent or does not have the agent installed. To install/upgrade the agent on these nodes, select the installation method and click Save.

The changes are saved and a green checkmark is displayed next to the task name. The plan page closes.

Note: If you have to add another task, you must select the plan from the **Resources** tab and modify the plan. To modify the plan, click the plan from the center pane. The plan opens and you can modify it.

The backup plan is created and automatically deployed to the source node. The backup runs per the schedule that you have configured in the **Schedule** tab. You can also perform a manual backup at any time.

Add a Copy Recovery Points Task to the Plan

The copy recovery task copies the recovery points from the backup destination to Cloud or a shared folder or local volume.

Note: If a backup job is in progress and you pause the plan, the backup job will get over and the copy recovery points job will not start. When you resume the plan again, the copy recovery points job is not resumed automatically. You need to manually run another backup job to start the copy recovery points job.

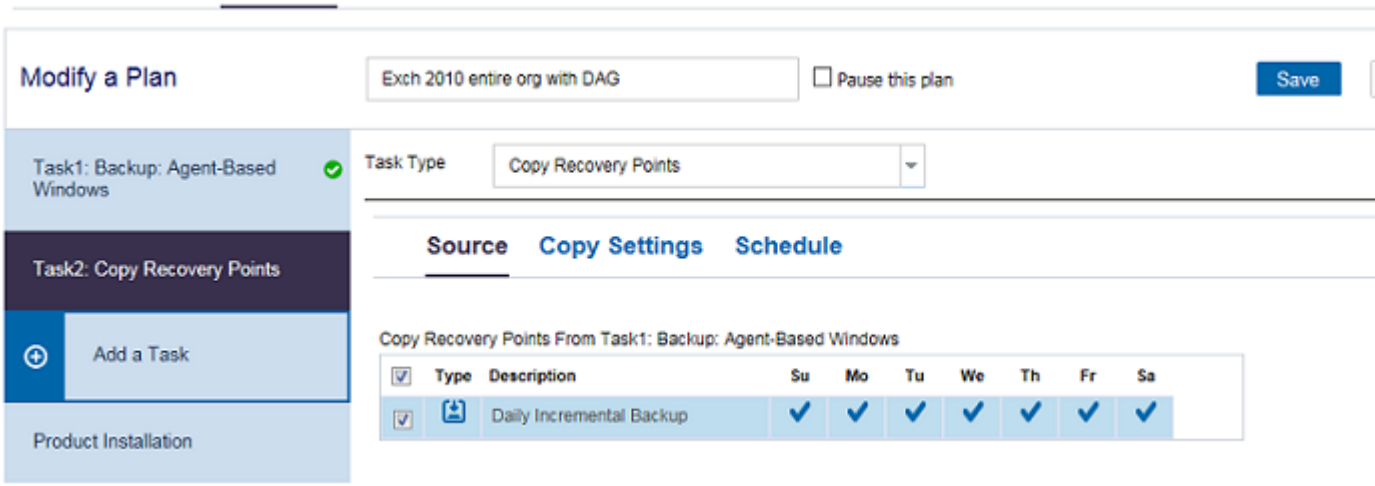
Follow these steps:

1. From the left pane, click **Add a Task**.

   A new task is added to the left pane.

2. From the **Task Type** drop-down list, select **Copy Recovery Points**.

   The Copy Recovery Points task is added. You do not have to configure the **Source** tab in the Copy Recovery Points (CRP) task. Additionally, you can view the backup destination from the Backup task.



   Note: You can configure Copy Recovery Point path to a customized location by using the below registry key.

   Path: *HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFCopySession*

   Key Name: *LocalTempPath*

   Key Type: *REG_SZ (String)*

   Key Value: *"Local_disk_path"*

3. On the Source tab, select a schedule.
4. On the Copy Settings tab, provide the following details:

Destination Type

Specifies the type of destination. You can select one of the options such as Cloud storage or a local disk or shared folder. You cannot opt for Recovery point server as destination type.

Local disk or shared folder

You need to enter the details for destination.

Destination

Specifies the destination where you want to keep the copy recovery points.

Note: Click arrow to validate the provided destination. The arrow is visible when you enter destination.

Cloud Storage

You need to select the cloud service and add a cloud storage.

For detailed information about the directory structure, view Recovery Points Directory Structure in S3 cloud bucket.

Cloud Service

Lets you select one of the options from the multiple available storage service.

Cloud Storage

Lets you select the cloud account of the selected storage service. If the drop-down list does not display any account then click **Add** to add an account.

If you select **Amazon S3** as Cloud Service, the **Enable Object Lock** check-box appears on the screen.

To configure Object Lock Settings, follow these steps:

a. Select the Enable Object Lock checkbox.

   The Object Lock Settings button is enabled.

b. Click the Object Lock Settings button.

   The Object Lock Settings dialog appears.

c. On the Object Lock Settings dialog, perform the following:

Default Retention Mode

Lets you enable different levels of protection on object versions. To enable, select one of the following:

- **Governance:** If you have special IAM permissions during the retention period, you can overwrite or delete the protected object versions, or change the lock settings.
- **Compliance:** You cannot overwrite or delete protected object versions even if you are a root user in the AWS account.

Default Retention Period

Lets you specify fixed amount of time to protect object versions. You can provide retention period in days or years.

d. Click **OK**.

Notes:

- Object Lock gets enabled only for new bucket.
- Object Lock does not support other cloud accounts.

Compression

Specifies to select a compression level for the recovery point copies. Compression is typically performed to decrease your disk space usage, but also has an inverse impact on your backup speed due to the increased CPU usage. The available options are:

No Compression - Compression is not performed. Files are pure VHD. This option has the lowest CPU usage (fastest speed), but also has the highest disk space usage for your backup image.

No Compression - VHD - Compression is not performed. Files are converted to .vhd format directly, without the need for manual operations. This option has the lowest CPU usage (fastest speed), but also has the highest disk space usage for your backup image.

Standard Compression - Some compression is performed. This option provides a good balance between CPU usage and disk space usage. This setting is the default setting.

Maximum Compression - Maximum compression is performed. This option provides the highest CPU usage (lowest speed), but also has the lowest disk space usage for your backup image.

Note: If your backup image contains incompressible data (such as JPG images or ZIP files), additional storage space can be allocated to handle such data. As a result, if you select any compression option and you have incompressible data in your backup, it can actually result in an increase in your disk space usage.

Encryption Algorithm

Specifies the type of encryption algorithm that is used for the recovery point copies. The available format options are No Encryption, AES-128, AES-192, and AES-256.

Encryption Password

Specifies to provide an encryption password that will be used to encrypt the destination session. When you restore from the copy recovery points, you must provide this password to confirm your authentication.

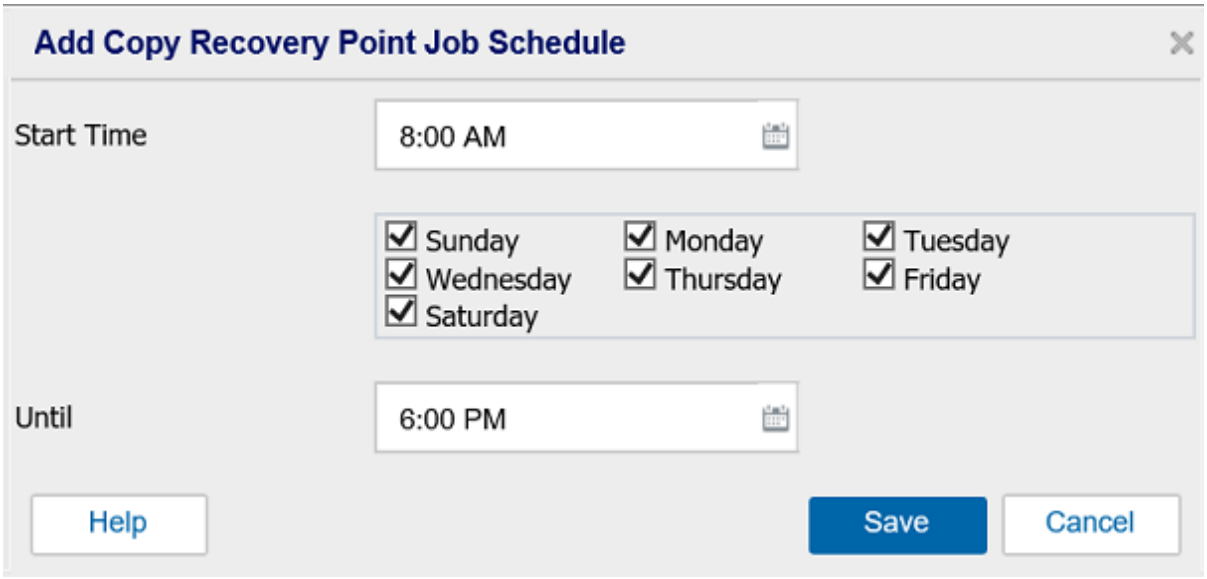
Confirm Encryption Password

Specifies to reenter the password.

5. Click the **Schedule** tab and specify the job schedule.

   Using the schedule option, you can add multiple schedules for a recovery point. Also consider the following details:

   - If no schedule is configured, then the Copy Recovery Points job runs immediately after completion of backup job.
   - Now, you can configure the Copy Recovery Points for start time and end time.



   - CRP always tries to check the session available in given time.

- During the specified time if any backup session configured in the Source tab is available, then CRP copies the session to respective destination.
- You can also configure retention points for Daily, Weekly, and Monthly.

Note: Number of successful backups is counted for any daily, weekly, and monthly backups that are configured.

6. Click **Save Changes**.

   The changes are saved and the copy recovery points task is automatically deployed to the node.

   You have successfully created and deployed the copy recovery points plan.

Recovery Points Directory Structure in Cloud bucket/Container

All recovery points of nodes with a single CRP plan and one cloud account (created in Arcserve Management console) as a destination are stored inside one cloud bucket.

Following is the recovery points directory structure in Cloud bucket/Container:

arcserve-{random-generated-id}-<BucketName>

ca_root_arcserve-recovery-points_<NodeName1>

Set0

<YYYY-MM-DD_HH-MM-SS_<ScheduleType>>

<NodeName1>

<YYYY-MM-DD_HH-MM-SS_<ScheduleType>>

<NodeName1>

.....

<YYYY-MM-DD_HH-MM-SS_<ScheduleType>>

<NodeName1>

ca_root_arcserve-recovery-points_<NodeName2>

.....

ca_root_arcserve-recovery-points_<NodeNameN>

Notes:

- The new arcserve-{random-generated-id}- prefix changes have no effect on existing plans.
- arcserve-{random-generated-id}- prefix is added to bucket name configured in cloud account.
- Node name is similar to what appears in the backup destination.
- YYYY-MM-DD: Date format (Y – Year, M – Month, D – Day)
- HH-MM-SS: Time format (H – Hour, M – Minute, S – Second)
- <ScheduleType> is the backup schedule type and refers to one of the following options:

- Daily: Daily backup recovery point
- Weekly: Weekly backup recovery point
- Monthly: Monthly backup recovery point
- Custom: Recovery point uploaded using the Upload Recovery Point to Cloud option (Ad hoc copy recovery point)

- Recovery points of nodes that share same name are stored inside the same directory if they are destined to the same Cloud account. As a result, we recommend to use different cloud account for nodes sharing the same name to help them store in different buckets.

Verify the Plan

To verify the copy recovery points feature, confirm that you have successfully created the plan. After you verify that the plan is created successfully, check whether the backup job is running as scheduled. After the backup job successfully completes, the copy recovery points job runs. You can check the status of the backup job and the copy recovery points job from the **jobs** tab.

Follow these steps to verify plans:

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and click **All Nodes**.

   A list of all nodes is displayed on the center pane.
3. Verify that plans are mapped with nodes.

Follow these steps to verify copy recovery points jobs:

1. Click the **jobs** tab.
2. From the left pane, click **All Jobs**.

   The status of each job is listed on the center pane.
3. Verify that the backup job and copy recovery points job is successful.

How to Create a File Copy Plan

Using Arcserve Unified Data Protection, you can copy or move selected source files to a destination. The destination can be a cloud storage, or shared network. The source file must be from the same volume that you have already backed up. For example, you have backed up the entire D:\ volume of your source node. Now you want to copy a specific file from the D:\ volume of the source node. You can create a file copy task to perform this operation.

File Copy can be used for copying critical data to secondary locations and can also be used as an archiving solution.

The advantages of copying files are:

- Improve Efficiency - Helps you to speed backup and recovery processes by copying and moving unchanged data and reduce the amount of real data being backed up and stored to tape or disk.
- Meet Regulatory Compliance - Helps you to preserve important documents, emails, and other critical data, as necessary to comply with internal rules and external regulations.
- Reduce Storage Cost - Helps you to reclaim storage capacity by migrating older or infrequently accessed data from your primary systems to more cost-effective archival storage locations.
- Maintain Multiple File Versions - Helps you to roll back to previous versions of backed-up files (if necessary) or maintain multiple versions of the same files at different destinations.

What To Do Next?

- Review the Prerequisites
- Create a Plan with a Backup Task
- Add a File Copy Task to the Plan
- (Optional) Perform a Manual File Copy
- Verify the Plan

Review the Prerequisites and Considerations

Verify if you have completed the following prerequisites:

- Log into the Console.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Considerations:

- The following table provides the limit on the file name length for a file copy job. The files exceeding the limits are skipped.

Destination	Limit
Network Share	240
Cloud	245

- If the File Copy is running from a replicated session and the Replicate *source* has multiple backup sessions, then the File Copy job runs for each session separately. For example, if there are **five** *backup sessions* in the backup destination and you add a Replicate task, then the Replicate task replicates all the sessions in *a single job*. Now, if you add a File Copy task and the File Copy source is the Replicate destination, then **five** *File Copy job* runs to replicate each session.

Create a Plan with a Backup Task

A plan includes different types of tasks that you want to perform. To create a file copy task, you must first have a valid recovery point. To get a valid recovery point, you have to create a backup task.

The backup task performs a backup of the source nodes and stores the data to the specified destination. File Copy is supported only for the Agent-based Windows backup. The following procedure explains the steps to create the agent-based Windows backup task.

Follow these steps:

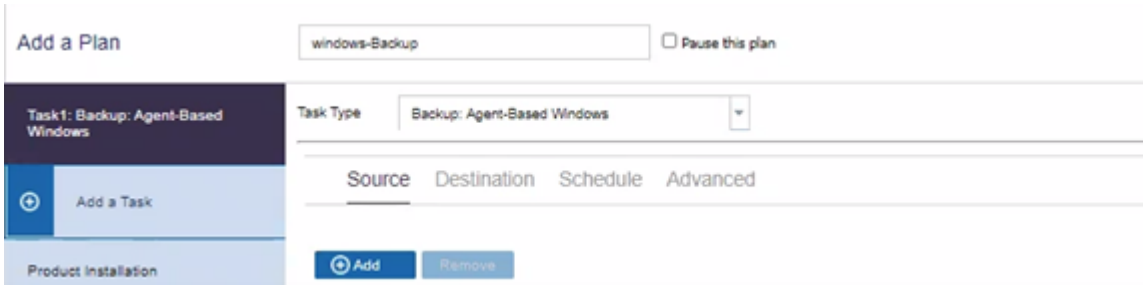
1. Click the **resources** tab on the Console.
2. From the left pane, navigate to **Plans**, and click **All Plans**.

   If you have created plans earlier, those plans are displayed on the center pane.
3. On the center pane, click **Add a Plan**.

   Add a Plan opens.
4. Enter a plan name.
5. (Optional) Select **Pause this plan** check box.

   The plan will not run until you clear the check box to resume the plan.

   Note: If a plan is paused, then any in-progress job is not paused but all corresponding scheduled jobs associated with that plan are paused. However, you can manually run a job. For example, you can manually run backup job and replication job for a node even if the respective plan is paused. In such case, the following task to the on-demand (manual) job does not run. For example, there is a replication task after an on-demand backup job, the replication job does not run for the on-demand backup job. You need to manually run the replication job. When you resume the plan, the pending jobs do not resume immediately. After you resume the plan, the pending jobs run from the next scheduled time.
6. From the **Task Type** drop-down list, select **Backup, Agent-Based Windows**.



Now, specify the Source, Destination, Schedule, and Advanced settings.

Specify the Source

The Source page lets you specify the source nodes that you want to protect. You can select more than one node in a plan. If you have not added any nodes to the Console, you can add nodes when you create or modify a plan from the Source page. You can also save a plan without adding any source nodes. The plan gets deployed only after you add source nodes.

Follow these steps:

1. On the Source tab, click **Add**.
2. Select one of the following options:

   Select Nodes to Protect in Arcserve UDP

   Opens the Select Nodes to Protect in Arcserve UDP dialog. Select the nodes from the displayed list. Select this option if you have already added the nodes to the Console.

   Add Windows Node

   Opens the Add nodes to a plan dialog. Select this option if you have not added the nodes and you want to manually add the nodes to protect.

   Discover Nodes from Active Directory

   Opens the Add Nodes to Arcserve UDP Console dialog. Select this option if you want to discover and add nodes from the Active Directory.

3. If you selected the **Select Nodes to Protect in Arcserve UDP** option, on the Select Nodes to Protect dialog, do the following:

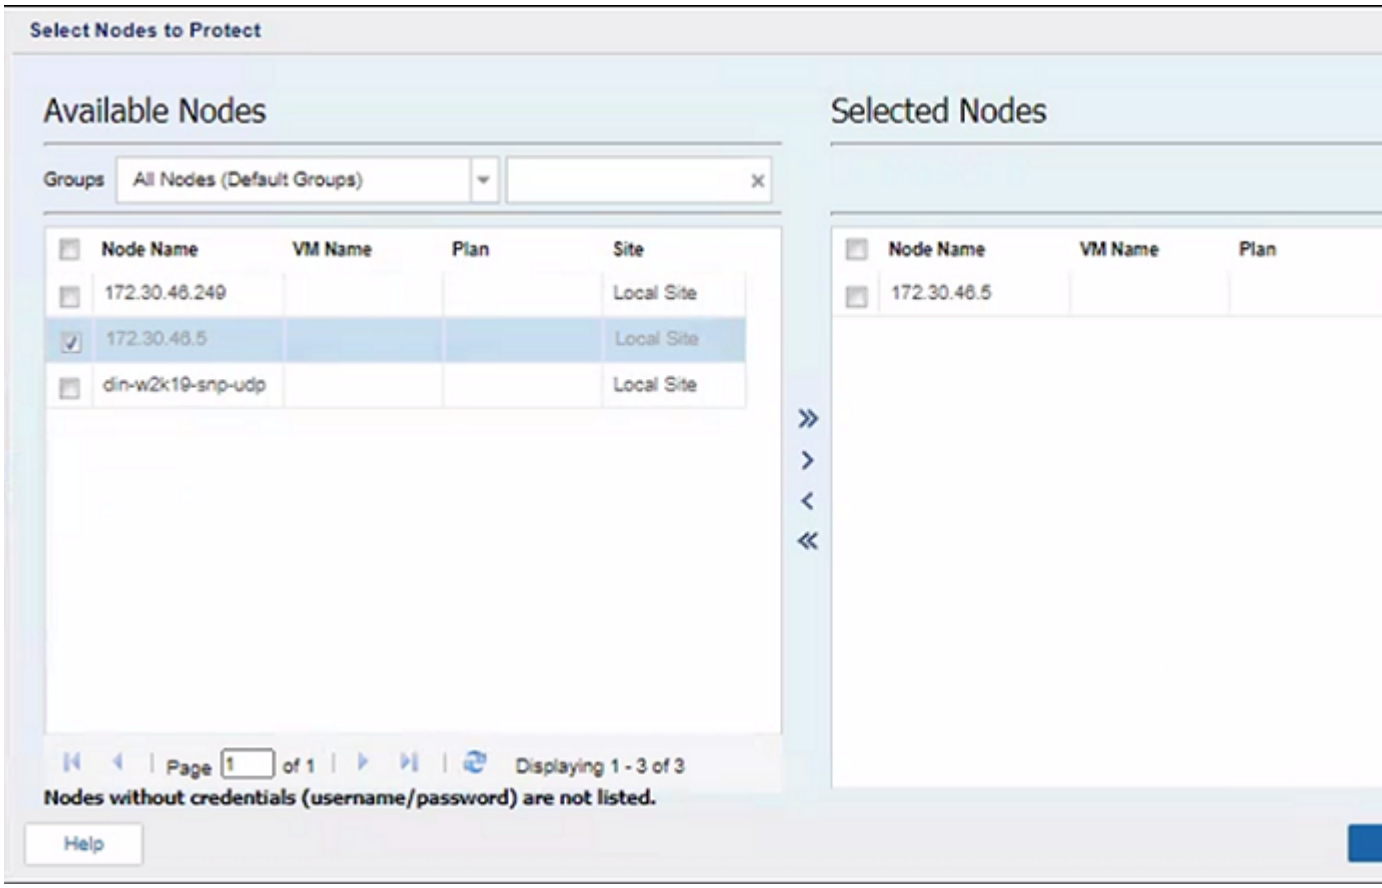
   a. (Optional) Select a filter from the Groups drop-down list to filter nodes. You can enter keywords to further filter your nodes.

      The nodes are displayed on the Available Nodes column.

   b. Select the nodes from the Available Nodes column, and then click the single right arrow (>) or double right arrow (>>) to move the selected nodes or all the nodes respectively to the Selected Nodes column.

      The nodes are displayed in the Selected Nodes column.

   c. Click **OK** to close the dialog.



4. To choose the Protection Type, select one of the following options:

   Back up all volumes

   Prepares a backup snapshot of all the volumes.

   Back up selected volumes

   Prepares a backup snapshot of the selected volume.

   The source is specified.

Specify the Destination

The destination is a location where you store the backup data. You must at least specify the destination to save the plan.

Follow these steps:

1. Select one of the following **Destination Type**:

   Local disk or shared folder

   Specifies that the backup destination is either a local destination or a shared folder. If you select this option, you can save data as either recovery points or recovery sets. The recovery points and recovery sets options are available on the **Schedule** tab.

   Arcserve UDP Recovery Point Server

   Specifies that the backup destination is a recovery point server. If you select this option, then data is stored as recovery points. You cannot store data as recovery sets.

2. If you have selected **Arcserve UDP Recovery Point Server**, then provide the following details:
   a. Select a recovery point server
   b. Select a data store. The list displays all data stores that are created at the specified recovery point server.
   c. Provide a session password.
   d. Confirm the session password.
3. If you have selected **Local disk or shared folder**, then provide the following details:
   a. Provide the full path of the local or network destination. For the network destination, specify the credentials with the write access.
   b. Select the encryption algorithm. For more information, see Encryption Settings.
   c. Optionally, provide an encryption password.
   d. Confirm the encryption password.
   e. Select a type of compression. For more information, see Compression Type.

Note: If you store the data to a local disk or shared folder, you cannot replicate the data to another recovery point server. Replication is supported only if you store the data to a recovery point server.

The destination is specified.

Specify the Schedule

The Schedule page lets you define a schedule for Backup, Merge, and Throttle functions to repeat at specific intervals. After you define a schedule, the jobs run automatically per the schedule. You can add multiple schedules and can provide retention settings.

A Backup Schedule refers to regular schedule that is repeated multiple times a day based on the number of hours or minutes you select. Besides the regular schedule, a backup schedule also provides options to add daily, weekly, and monthly schedules.

Note: For more information on scheduling and retention settings, see Understanding Advanced Scheduling and Retention.

Follow these steps:

1. (Optional) Select the option to manage recovery points. This option is visible only if you have selected Local or shared folder as your backup destination.

   Retain by Recovery Points

   The backup data is stored as recovery points.

   Retain by Recovery Sets

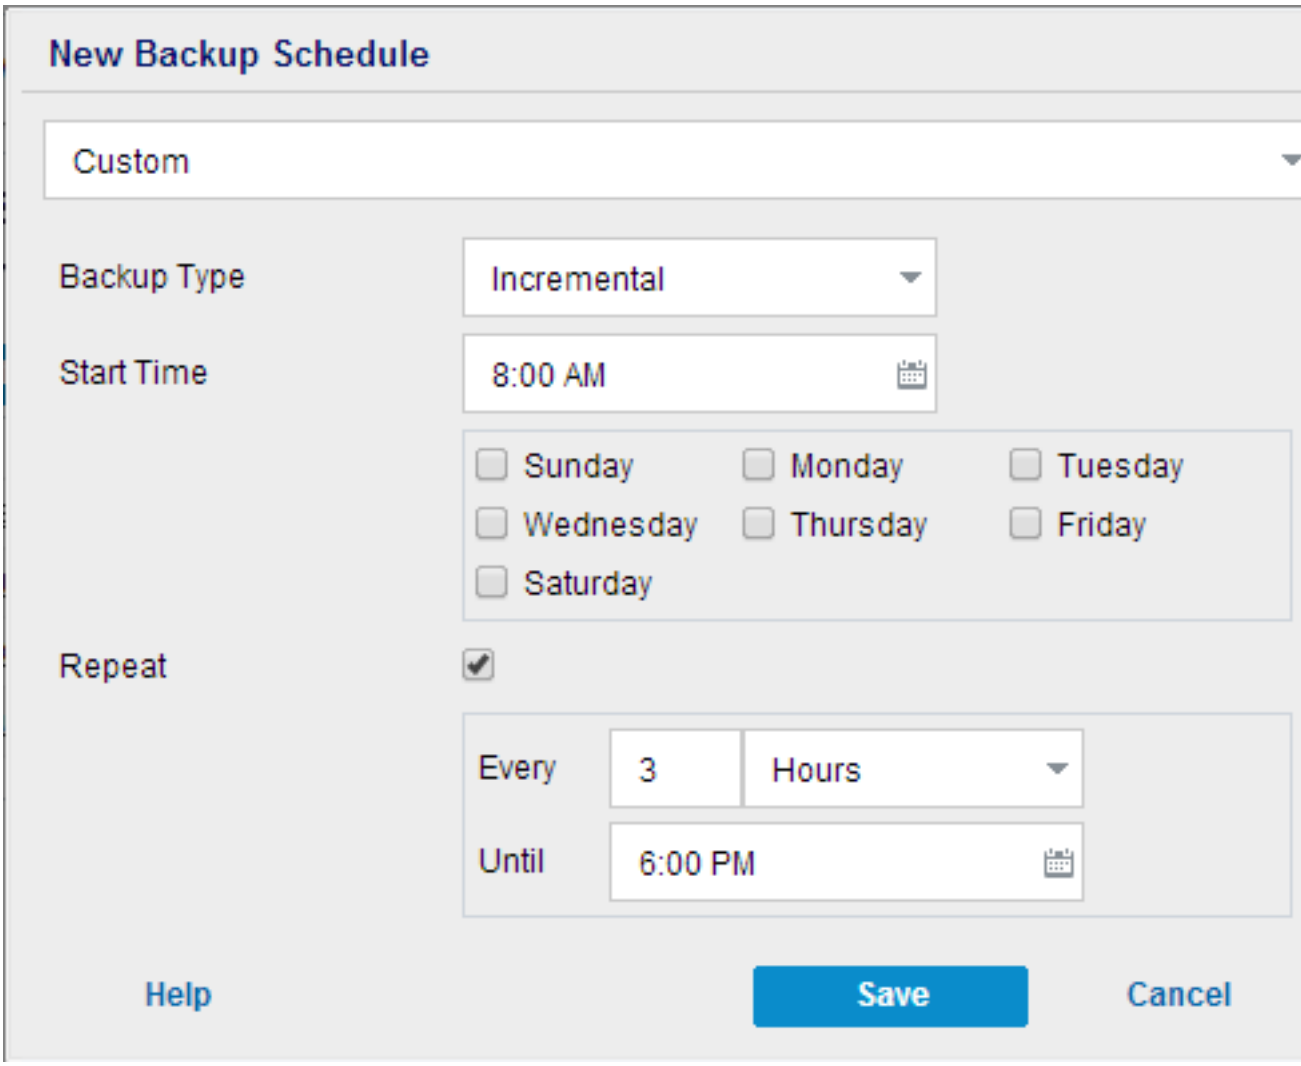
   The backup data is stored as recovery sets.

2. Add backup, merge, and throttle schedules.

   Add Backup Schedule

   a. Click **Add** and select **Add Backup Schedule**.

      The **New Backup Schedule** dialog opens.



b. Select one of the following options:

Custom

Specifies the backup schedule that repeats multiple times a day.

Daily

Specifies the backup schedule that occurs once a day. By default, all the days of the week are selected for Daily backup. If you do not want to run the backup job on a specific day, clear the check box for that day of the week.

Weekly

Specifies the backup schedule that occurs once a week.

Monthly

Specifies the backup schedule that occurs once a month.

c. Select the backup type.

Full

Determines the backup schedule for Full Backups. As scheduled, Arcserve UDP performs a full backup of all used blocks from the source machine. A full backup typically consumes time depending on the backup size.

Verify

Determines the backup schedule for Verify Backups.

Arcserve UDP verifies that the protected data is valid and complete by performing a confidence check of the stored backup image to the backup source. If necessary, the image is resynchronized. A Verify Backup looks at the most recent backup of each individual block and compares the content and information to the source. This comparison verifies that the latest backed up blocks represent the corresponding information at the source. If the backup image for any block does not match the source (possibly because of changes in the system since the last backup), Arcserve UDP refreshes (resynchronizes) the backup of the block that does not match. You can also use a Verify Backup (infrequently) to get the guarantee of full backup without using the space required for a full backup.

Advantages: Produces a small backup image when compared to full backup because only the changed blocks (blocks that do not match the last backup) are backed up.

Disadvantages: Backup time is long because all source blocks are compared with the blocks of the last backup.

Incremental

Determines the backup schedule for Incremental Backups.

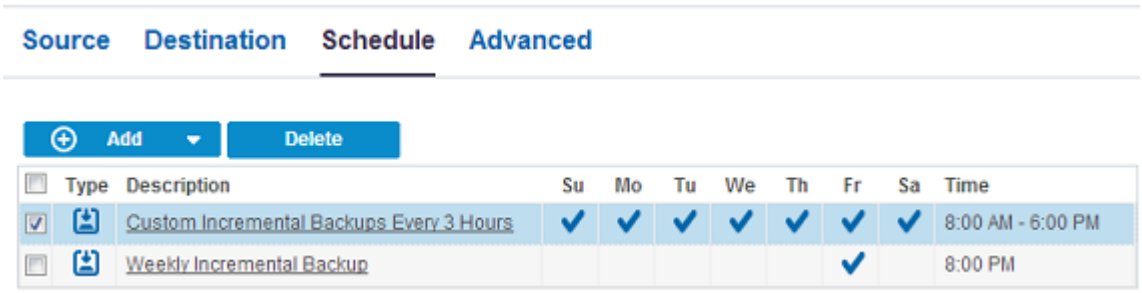
As scheduled, Arcserve UDP incrementally backs up only those blocks that have changed since the last successful backup. The advantages of Incremental Backups are that it is a fast backup and it produces a small backup image. This is the most optimal way to perform a backup.

d. Specify the backup start time.

e. (Optional) Select the **Repeat** check box and specify the repeat schedule.

f. Click **Save**.

The Backup Schedule is specified and displayed on the **Schedule** page.



Add Merge Schedule

a. Click Add and select Add Merge Schedule.

   The **Add New Merge Schedule** dialog opens.

b. Specify the start time to start the merge job.

c. Specify **Until** to specify an end time for the merge job.

d. Click **Save**.

   The Merge Schedule is specified and displayed on the **Schedule** page.

Add Throttle Schedule

a. Click **Add** and select **Add Throttle Schedule**.

   The Add New Throttle Schedule dialog opens.

b. Specify the throughput limit in MB per minute unit.

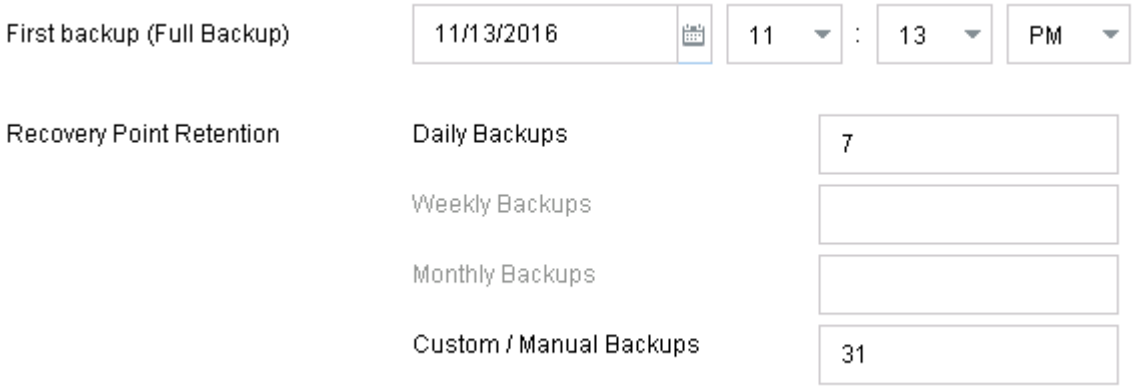
c. Specify the start time to start the backup throughput job.

d. Specify **Until** to specify an end time for the throughput job.

e. Click **Save**.

   The Throttle Schedule is specified and displayed on the **Schedule** page.

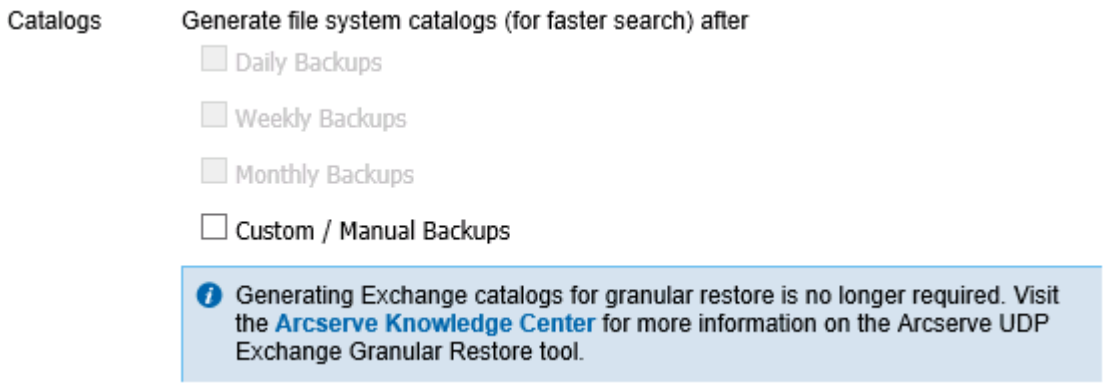
3. Specify the start time for the scheduled backup.



4. Specify the recovery points retention settings for Custom, Daily, Weekly, and Monthly schedule.

   These options are enabled if you have added the corresponding backup schedule. If you modify the retention settings on this page, the changes are reflected on the Backup Schedule dialog.

5. Specify the catalog details.



   Catalogs let you generate the File System catalog. The File System catalog is required to perform faster and better search. If you select the catalog check boxes, the catalogs are enabled depending on the type of backup that you have specified. Clear the check box to disable generating the catalog.

   The schedule is specified.

Understanding Advanced Scheduling and Retention

The scheduling option lets you specify a Custom schedule or a Daily / Weekly / Monthly schedule, or both the schedules. In the Custom schedule, you can configure the backup schedule for each day of the week and you can add up to four backup schedules each day. You can select a specific day of a week and create a time window to define when to run backup and at what frequency.

Schedule	Supported Job	Comments
Backup	Backup job	Define time windows to run backup jobs.
Backup throttling	Backup job	Define time windows to control the backup speed.
Merge	Merge job	Define when to run merge jobs.
Daily schedule	Backup job	Define when to run daily backup jobs.
Weekly schedule	Backup job	Define when to run weekly backup jobs.
Monthly schedule	Backup job	Define when to run monthly backup jobs.

You can also specify the retention settings for the recovery points.

Note: Set the retention settings within each plan to control how data for the nodes assigned to that plan are retained at the target data store.

Schedules for Daily / Weekly / Monthly backups are independent to the Custom schedule, and each other. You can configure only to run Daily backup, Weekly backup or Monthly backup, without configuring the Custom schedule.

Backup Job Schedule

You can add four time windows per day in your backup schedule. A valid time window is from 12:00 AM until 11:59 PM. You cannot specify a time window such as 6:00 PM to 6:00 AM. In such cases, you must manually specify two different time windows.

For each time window, the start time is inclusive, and the end time is exclusive. For example, you have configured to run Incremental Backup every one hour between 6:00 AM and 9:00 AM and the backup will start at 6:00 AM. This means the backup runs at 6:00 AM, 7:00 AM, 8:00 AM, but NOT at 9:00 AM.

Note: If you want to run the backup job repeatedly until the end of day, set the schedule until 12:00 AM. For example, to run backup job every 15 minutes for the entire day, set the schedule from 12:00 AM to 12:00 AM, every 15 minutes.

Backup Throttle Schedule

Backup throttle schedule lets you control the backup throughput speed that in turn controls the resource usage (disk I/O, CPU, network bandwidth) of the server being backed up. This is useful if you do not want to affect the server

performance during business hours. You can add four time windows per day in your backup throttle schedule. For each time window, you can specify a value, in MB per minute. This value is used to control the backup throughput. Valid values are from 1 MB/minutes to 99999 MB/minutes.

If a backup job extends its specified time, then the throttle limit adjusts according to the specified time window. For example, you have defined the backup throttle limit as 500 MB/minute from 8:00 AM to 8:00 PM, and 2500 MB/minute from 8:00 PM to 10:00 PM. If a backup job starts at 7:00 PM and it runs for three hours, then, from 7:00 PM to 8:00 PM the throttle limit is 500 MB/minute and from 8:00 PM to 10:00 PM the throttle limit is 2500 MB/minute.

If you do not define any backup schedule and backup throughput schedule, the backup runs as fast as it can.

Merge Schedule

Lets you merge recovery points based on the provided schedule.

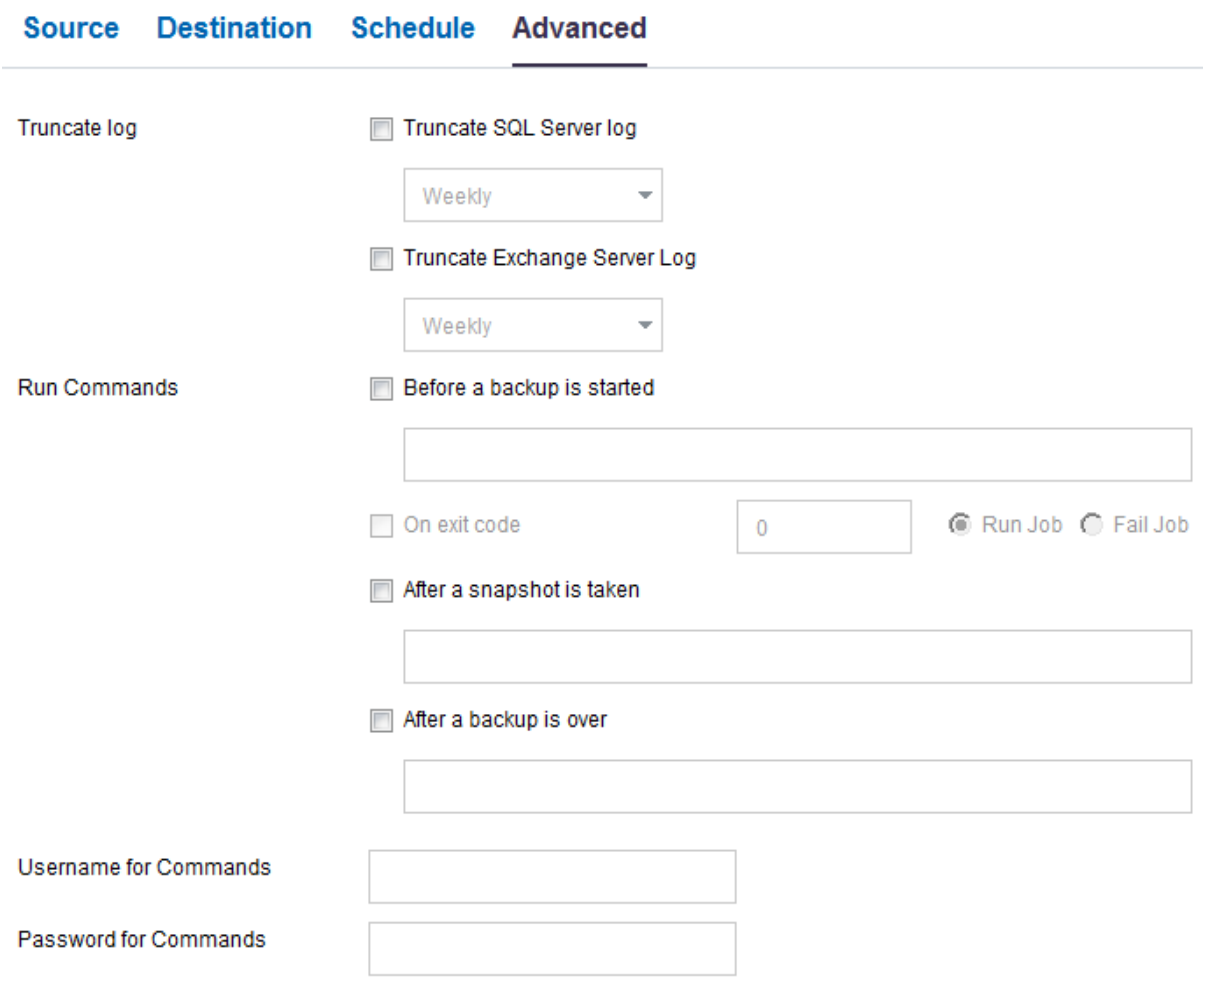
Consider the following points for the merge job:

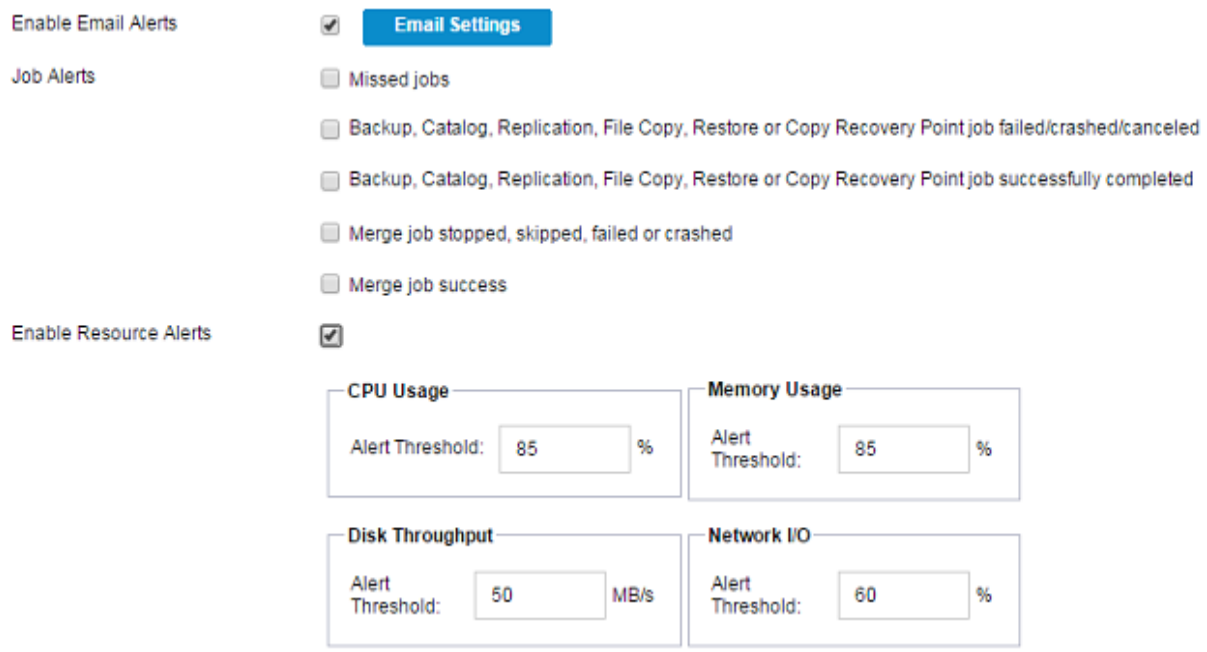
- At any given time only one merge job can run for a node.
- If a merge job starts, it has to complete before the next merge job can start. This means, if one or more sets of recovery points are merged, new recovery point cannot be added to this merge process, until the merge process of the current set of recovery point completes.
- If a merge job is processing more than one set of recovery points (for example set [1~4], set [5~11], and set [12~14]; they are three sets), recovery point server processes these sets one by one.
- If a merge job is resumed after a pause, the job detects at which point it is paused and resumes the merge from the break-point.

Specify the Advanced Settings

The **Advanced** tab lets you specify some advanced settings for the backup job. The advanced settings include providing truncate log settings, providing the location of any scripts, and email settings.

The following image displays the Advanced tab:





Follow these steps:

1. Specify the following details.

Truncate Log

Lets you specify the schedule to truncate logs for SQL Server and Exchange Server. You can specify the schedule as **Daily**, **Weekly**, or **Monthly**.

User Name

Lets you specify the user who is authorized to run a script.

Password

Lets you specify the password of the user who is authorized to run the script.

Run a command before backup is started

Lets you run a script before the backup job starts. Specify the path where the script is stored. Click **On exit code** and specify the exit code for **Run Job** or **Fail Job**. **Run Job** indicates that the backup job will continue when the script returns the exit code. **Fail Job** indicates that the backup job will stop when the script returns the exit code.

Run a command after snapshot is taken

Lets you run a script after the backup snapshot is taken. Specify the path where the script is stored.

Run a command after backup is over

Lets you run a script after the backup job is completed. Specify the path where the script is stored.

Enable Email Alerts

Lets you enable email alerts. You can configure email settings and can specify the types of alerts that you want to receive in an email. When you select this option, the following options are enabled for your selection.

Email Settings

Lets you configure the email settings. Click **Email Settings** and configure the email server and proxy server details.

Job Alerts

Lets you select the types of job emails you want to receive.

Enable Resource Alerts

Lets you specify a threshold for CPU Usage, Memory Usage, Disk Throughput, Network I/O. You can provide the value in percentage. You will receive an email when the Alert Threshold value exceeds.

2. Click **Save**.

Note: When you select a node as a backup source or backup proxy, Arcserve UDP checks whether the agent is installed on the node and if it is the latest version. Arcserve UDP then displays a verification dialog that lists all the nodes that either have an outdated version of the agent or does not have the agent installed. To install/upgrade the agent on these nodes, select the installation method and click Save.

The changes are saved and a green checkmark is displayed next to the task name. The plan page closes.

Note: If you have to add another task, you must select the plan from the **Resources** tab and modify the plan. To modify the plan, click the plan from the center pane. The plan opens and you can modify it.

The backup plan is created and automatically deployed to the source node. The backup runs per the schedule that you have configured in the **Schedule** tab. You can also perform a manual backup at any time.

Add a File Copy Task to the Plan

The file copy task lets you copy individual files to the specified destination. The original copy is retained after you have copied the files to the specified destination. File copy is not dependent on catalog jobs. If the backup destination is a deduplication or non-deduplication data store, the Catalog folder is moved to the backup destination. The catalog job is a part of the file copy job. The file copy job runs on recovery point server and by doing so Arcserve UDP offloads the task from agents.

Arcserve UDP supports file copy from a replication data store. ;

Pre-flight Check for File Copy Nodes: You can also perform Pre-flight Check (PFC) for file copy. Only those node that are added for file copy are eligible for PFC. To perform PFC for file copy, right-click the node from All Nodes and select Filecopy Preflight check.

Note: If a backup job is in progress and you pause the plan, the backup job will get over and the file copy job will not start. When you resume the plan again, the file copy job is not resumed automatically. You have to manually run another backup job to start the file copy job.

Follow these steps:

1. Click **Add a Task** from the left pane.

   A new task is added to the left pane.

2. From the **Task Type** drop-down menu, select **File Copy**.

   The File Copy task is added.

3. Click the **Source** tab and specify the details.

   Recovery Points Source for File Copy

   Let you select the source of recovery points. If there is only one source, the source is displayed automatically. If there are more than one source, you have to select the correct source from the drop-down list.

   Types of Recovery Points

   Lets you select the recovery points for file copy. You can select either of the following two options:

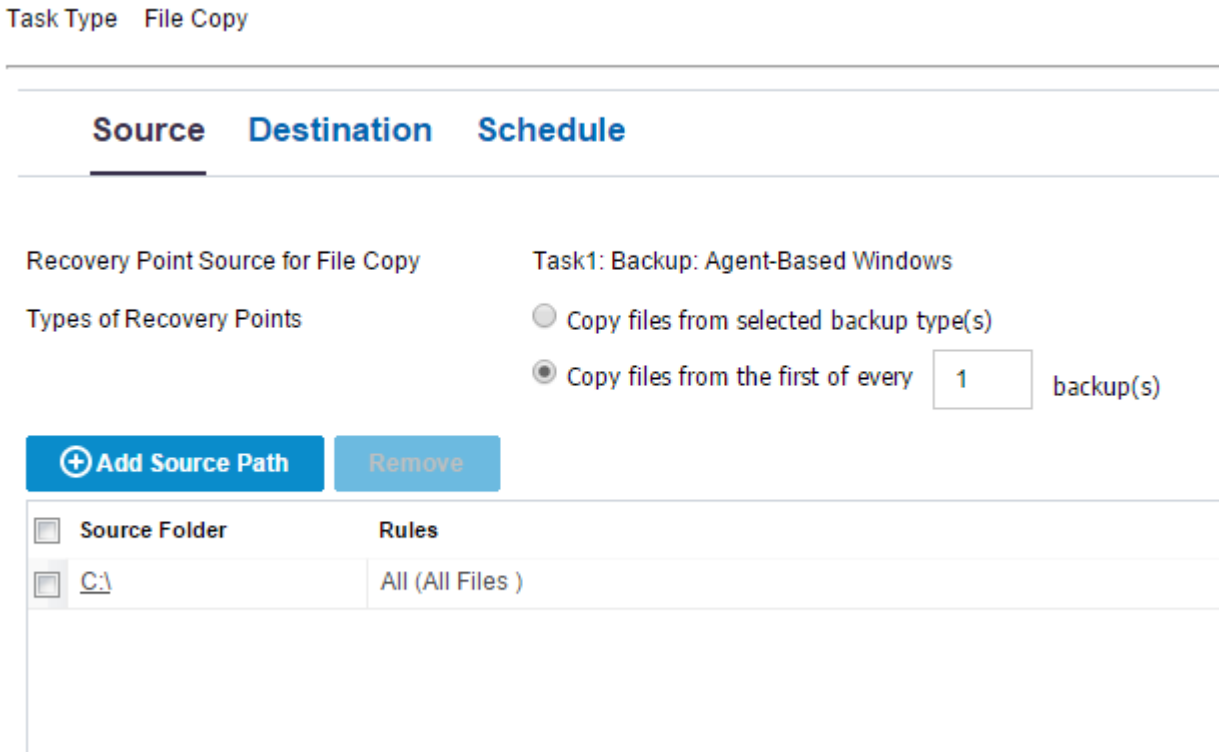
   Copy files from selected backup types(s)

   You can select either **Daily Backups**, **Weekly Backups**, **Monthly Backups**. The daily, weekly, or monthly options are enabled depending on the schedule of the source backup. For example, if the source backup has only monthly

backup scheduled, then the **Copy files from selected backup types(s)** option has only **Monthly Backups** enabled.

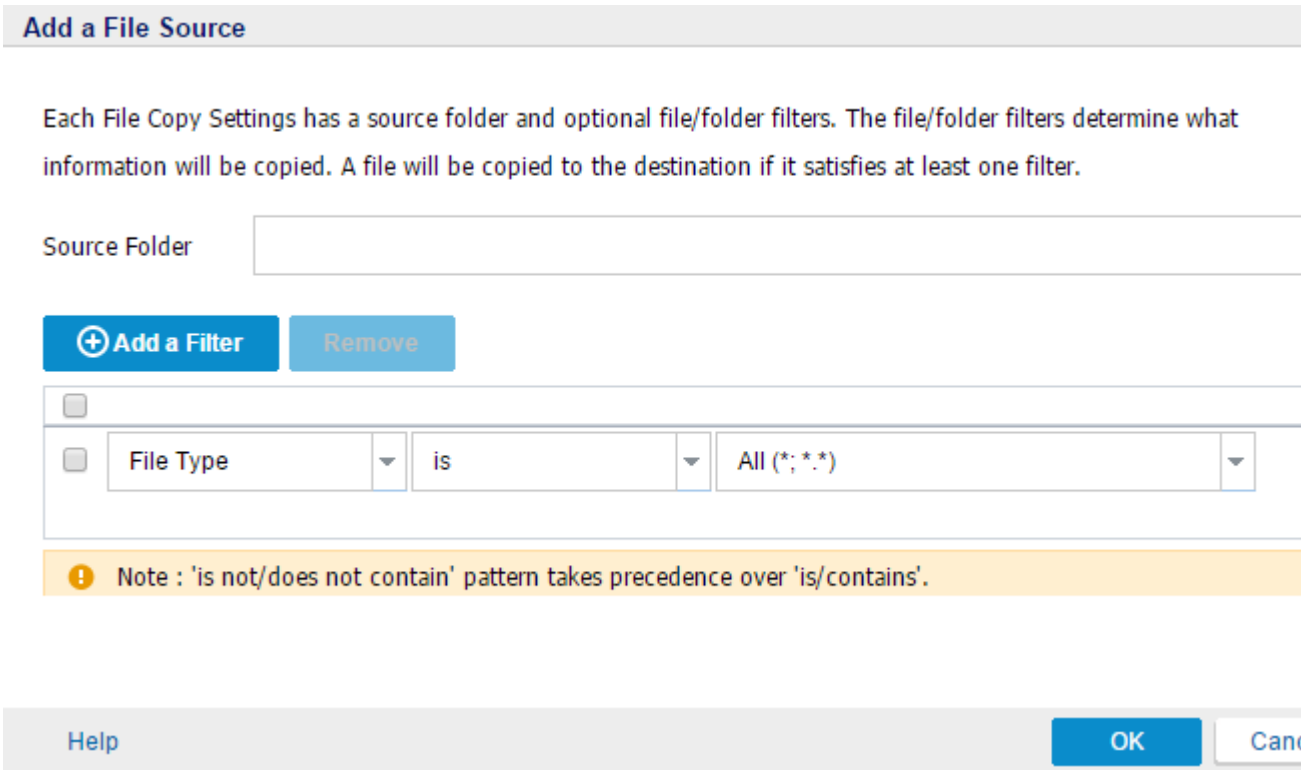
Copy file from the first of every backup(s)

You can specify the backup number that you want to copy. For example, if you specify 3, then the third backup will be copied. File copy supports up to 700 recovery points to copy from.



4. Click **Add Source Path**.

   The **Add a File Source** dialog opens.



You can avoid skipping Windows System (C:\Windows) and program files (C:\Program Files, C:\Program Files (x86)) directories using the following configuration:

Add the following XML tag, if not present already or update in the FileCopyDebugSetting.xml file present under $UDPHome\Engine\Configuration folder:

<SkipWindowsFolders>0</SkipWindowsFolders>

The XML file appears as below:

<?xml version="1.0" encoding="UTF-8"?>

<HKLM>

<AFArchiveDLL>

.....

<SkipWindowsFolders>0</SkipWindowsFolders>

</AFArchiveDLL>

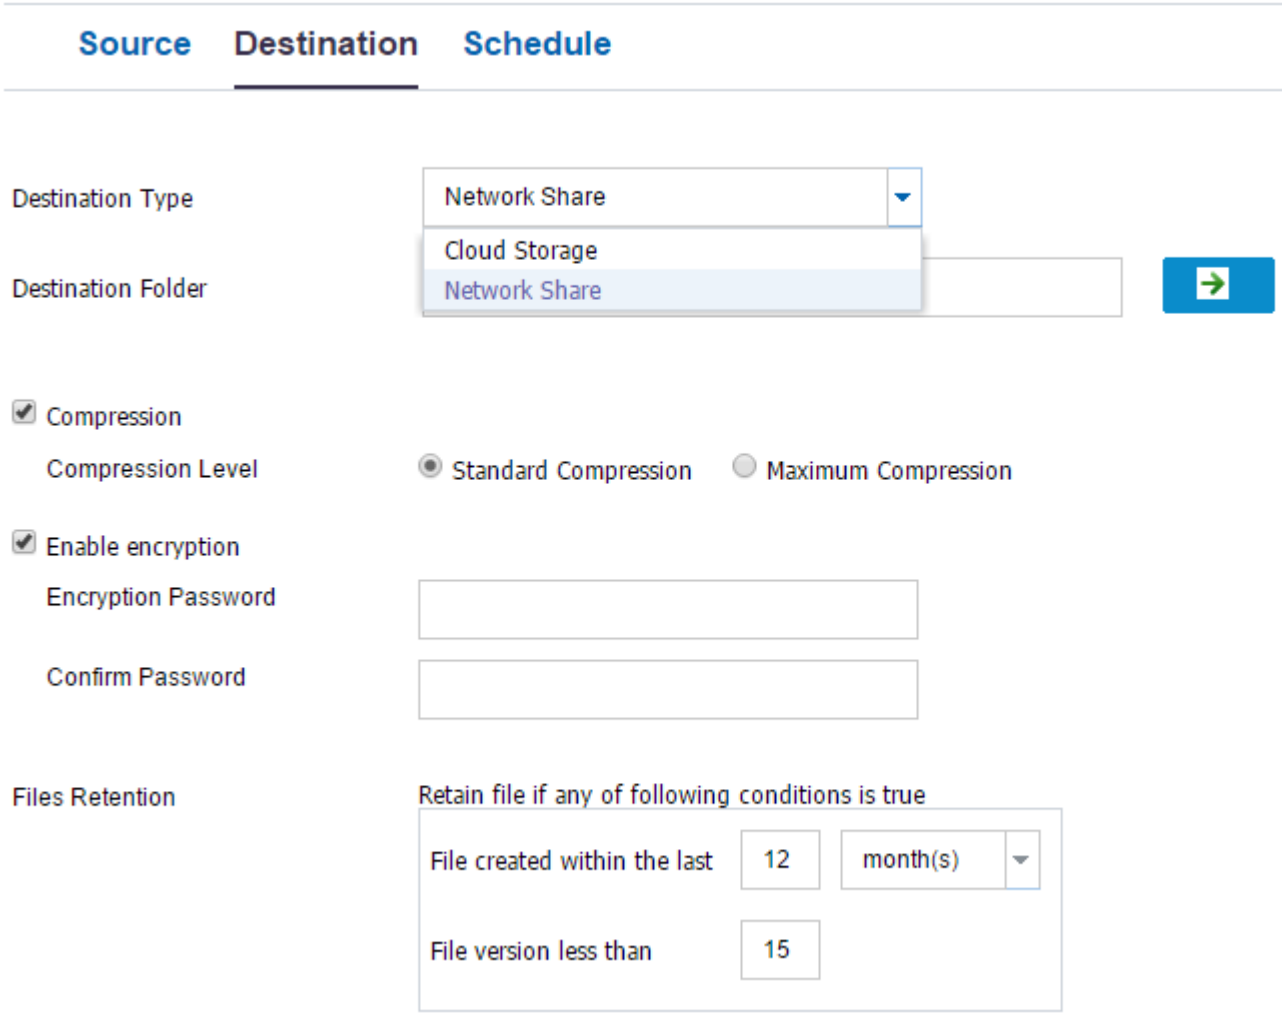
</HKLM>

5. Specify the path of the source folder that you want to copy.
6. Click **Add a Filter**.

   The filter is added below the **Add a Filter** button. You can add multiple filters and can also remove the filters. For more information, see Add File Copy Filters.
7. Select the filter from the list and click **OK**.

The **Add a File Source** dialog closes.

8. Click the **Destination** tab and specify the destination details.



Destination Type

Specifies that the destination types is a network share or a cloud storage. For either destination option, if the connection to the specified destination is lost or broken, Arcserve UDP makes several attempts to continue the file copy job. If these reattempts are not successful, a makeup job is then performed from the point where the failure occurred. In addition, the activity log is updated with a corresponding error message and an email notification is sent (if configured).

Network Share

Specifies that the destination is a shared folder. When selected, lets you specify the full path of the location where you want to move or copy the source files/-folders.

Destination Folder

Specifies the destination where the copied files are stored. The destination can be any local volume or folder or a file share accessible by any uniform naming convention (UNC) path. This field is available when you select Network Share as the destination type. You can also browse the destination folder.

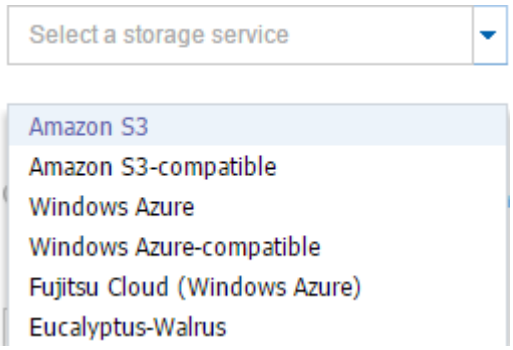
Cloud Storage

Specifies that the copied files are stored in a cloud environment. Arcserve UDP currently supports file copying to multiple cloud vendors, such as Amazon S3 (Simple Storage Service), Amazon S3-compatible, Windows Azure, Windows Azure-compatible, Eucalyptus-Walrus, and FUJITSU Hybrid IT Service FJcloud-O. These cloud vendors are publicly available web services that let you store and retrieve any amount of data, at any time from anywhere on the web in a safe and secure environment.

Note: To eliminate any potential clock skew error when attempting to connect to the cloud, verify that your machine has the correct time zone set and the clock is in sync with the global time. Always check the time of your machine against the GMT time. If the time of your machine is not synchronized with the correct global clock time (within 5 to 10 minutes), your cloud connection may not work. If necessary, reset the correct time for your machine and rerun your file copy job.

Storage Device

Select the device type from the drop down list.



Cloud Storage

Select the cloud storage path from the drop-down list. The drop-down list is available if you have specified your cloud storage details. If you are specifying the cloud storage account for the first time, click Add to add your cloud account. When you select Cloud Storage from the next time, the account will be displayed in the Cloud Storage drop-down list.

Note: For more information on adding a cloud account, see Add a Cloud Account.

Compression

Specifies the type of compression that is used for the File Copy jobs.

Compression is performed to decrease your storage space at the File Copy destination, but also has an inverse impact on your file copy speed due to the increased CPU usage.

Note: For a compressed File Copy job, the Activity log displays only the uncompressed size.

The available options are:

Standard Compression

Some compression is performed. This option provides a good balance between CPU usage and storage space requirement. This is the default setting.

Maximum Compression

Maximum compression is performed. This option provides the highest CPU usage (lowest speed), but also has the lowest storage space requirement for your file copy.

Enable Encryption

Specifies to use encryption for file copying.

Data encryption is the translation of data into a form that is unintelligible without a deciphering mechanism. Arcserve UDP data protection uses secure, AES-256 (Advanced Encryption Standard) encryption algorithms to achieve maximum security and privacy of your specified data. When an encryption is selected, you must provide (and confirm) an encryption password.

Note: When you modify a file copy task, the Encryption or Compression cannot be modified.

File Retention

Retains the files in the file copy destination if the specified criteria is met.

File created within the last

Specifies the amount of time (years, months, days) that the stored data is retained at the destination location. At the end of the specified retention time period, the stored data is purged from the destination.

Important! At the end of the specified retention time when the data is purged from the destination, all of this purged data is no longer stored or saved.

Note: The Retention Time purge process is only triggered if the File Copy Schedule option is enabled.

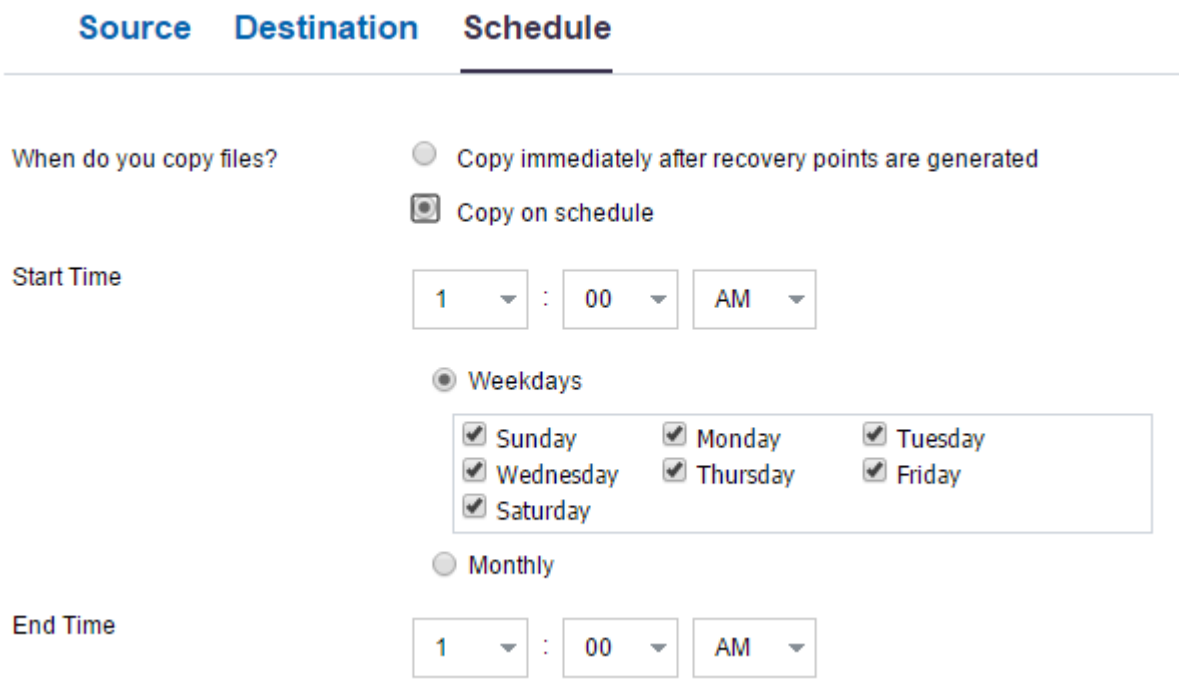
File version less than

Specifies the number of copies retained and stored at the destination location. After this number is exceeded, the earliest (oldest) version will be discarded. This cycle of discarding the oldest stored version repeats as newer versions are added to the destination, allowing you to always maintain the specified number of stored versions.

For example, if your specified File Versions retention count is set to 5 and you perform five file copies at times t1, t2, t3, t4, and t5, these file copies become the five file copy versions retained and available to recover. After the sixth file copy is performed (new version is saved), Arcserve UDP will remove the t1 copy and the five available versions to recover are now t2, t3, t4, t5, and t6.

By default, the number of copies retained at the destination location before discarding is 15.

9. Click the **Schedule** tab and specify the file copy schedule.



If a file copy job runs beyond the end time, the job will continue to run until it is complete. The next file copy job does not run until the previous file copy job is complete even if a scheduled job overlaps the running job.

The File Copy job runs as specified in the schedule.

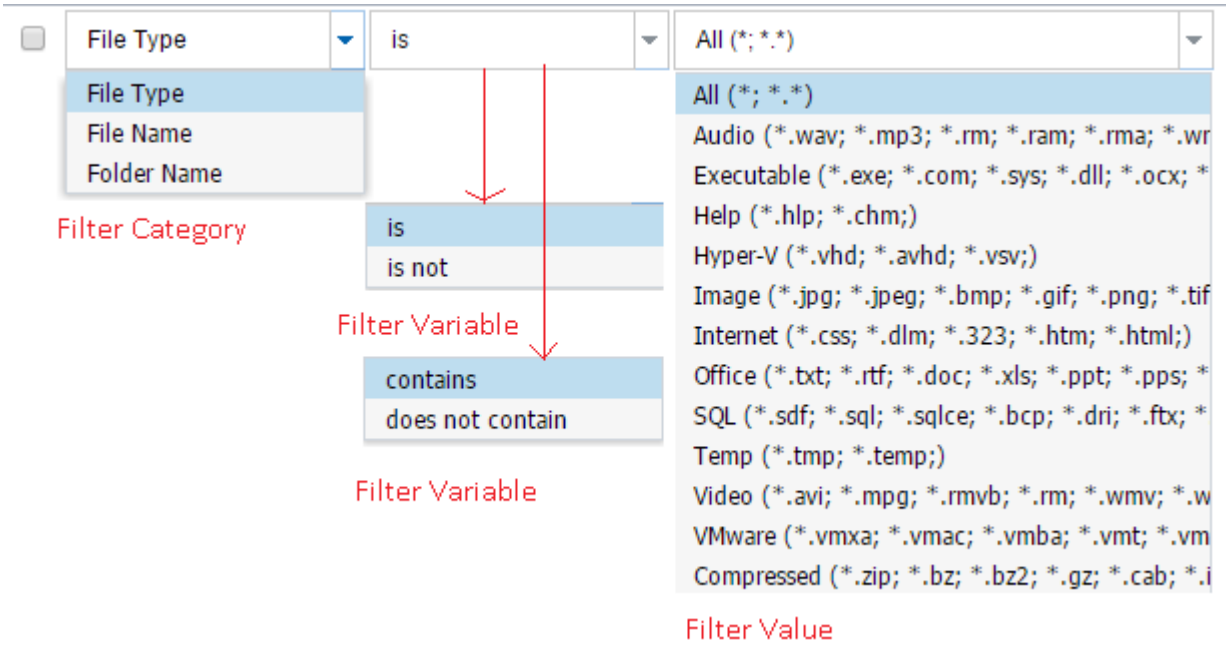
10. Click **Save**.

The changes are saved and the file copy task is automatically deployed to the node.

Add File Copy Filters

Add a Filter

Lets you add a filter. Filters let you limit the objects to be file copied by certain specified types and values.



Filter Category

There are three categories of filters available: **File Type**, **File Name**, and **Folder Name**. The Filter Variable and Filter Value field changes depending on Filter category.

Filter Variable

If the Filter Category is **File Type**, the Filter Variable options are **is** and **is not**. If the Filter Category is **File Name or Folder Name**, the Filter Variable options are **contains** or **does not contain**.

You can specify multiple filters within the same file copy request. You can specify the same Filter Category but different Filter Variable.

Note: When the Filter Variable conflict for the same Filter Category, the **is not** or **does not contain** variable is always a higher priority and is enforced.

Filter Value

The filter value lets you limit the information that is file copied by selecting only the parameter information that you specify, such as .txt files.

Arcserve UDP supports the use of wildcard characters to help select multiple objects to file copy with a single request. A wildcard character is a special

character that can be used as a substitute to represent either a single character or a string of text.

The wildcard characters asterisk and question mark are supported in the Value field. If you do not know the complete file/folder pattern value, you can simplify the results of the filter by specifying a wildcard character.

"*" - Use the asterisk to substitute zero or more characters in the value.

"?" - Use the question mark to substitute a single character in the value.

For example, you can enter *.txt to exclude all files with a .txt extension if you do not know the specific file name. You can provide as much of the file name as you know, then use wildcards to fill in the blanks.

Note: When you select **File Type** as the filter type, a drop-down list of predefined filters for many commonly used files is available (MS-Office files, Image files, Executable files, Temp files, etc.). After choosing any of the predefined filters, you can still append or modify the corresponding values.

Protect System folders from Skipping

File copy job skips the system folders by default when the job is performed on the volumes containing system folders. In order to protect those folders and copied in the destination, you can change the default text by adding a configuration entry.

Add the following XML tag entry in the FileCopyDebugSetting.xml file from the $UDPHome\Engine\Configuration directory:

FileName : FileCopyDebugSetting.xml

TagName: SkipWindowsFolders

DefaultValue: 1

To Protect change the value to: 0

Note: This Option is limited to File Copy Job only. File Archive Skips all the system folders irrespective of the TagValue configured.

(Optional) Perform a Manual File Copy

Typically, a File Copy is performed automatically and it is controlled by the File Copy schedule settings. In addition to the scheduled File Copy, a manual File Copy provides you the option to copy your important files on a need basis.

When you run a manual File Copy, the File Copy job runs only for the first backup session, which is qualified for a File Copy. (Backup sessions qualify for File Copy per the backup schedule and if the sessions are in a queue. For example, if you specify to run File Copy for every second backup, then every second backup only is qualified for File Copy, not all the backups are qualified for File Copy.) After the File Copy is complete, the first session is removed and the second session in line becomes the first session. For example, if there are three backup sessions (S1, S2, S3 respectively) and you run a manual File Copy, then the File Copy job runs only for S1. The File Copy job does not run for S2 and S3. When you run the manual File Copy job again, S2 is copied.

File Copy can be run manually from the Nodes view and Plans view by clicking **Actions** menu or context menu.

Follow these steps to run the File Copy manually from the Nodes view:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and click **All Nodes**.

   Nodes are displayed in the center pane.
3. Select the nodes for a File Copy job. The nodes must have a File Copy plan assigned to it.
4. On the center pane, click **Actions**, **File copy Now**.

   The **Run File copy now** dialog opens.
5. Click **OK**.

   The File Copy job runs.

   The manual File Copy is successfully performed.

Verify the Plan

To verify the file copy plan, confirm that you have successfully created the plan. After you verify that the plan is created successfully, check whether the backup job is running as scheduled. After the backup job successfully completes, the file copy job runs as scheduled. You can check the status of the backup job and the file copy job from the **jobs** tab.

Follow these steps to verify plans:

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and click **All Nodes**.

   A list of all nodes is displayed on the center pane.
3. Verify that plans are mapped with nodes.

Follow these steps to verify file copy jobs:

1. Click the **jobs** tab.
2. From the left pane, click **All Jobs Completed**.

   The status of each job is listed on the center pane.
3. Verify that the backup job and file copy job is successful.

How to Create a File Archive Plan

Using Arcserve Unified Data Protection, you can archive selected source files to a destination. The destination can be a cloud account or a shared network. The source file must be from the same volume that you have already backed up. For example, you have backed up the entire D:\ volume of your source node. Now you want to copy a specific type of file (for example, .htm) from the D:\ volume of the source node. After you copy the file, you want to delete that file from the source node. You can create a file archive plan to perform this operation.

File Archive allows you to safely and securely delete the source data after it has been copied to an off-site or secondary storage repository.

The advantages of archiving files are:

- Improve Efficiency - Helps you to speed backup and recovery processes by archiving unchanged data and reduce the amount of real data being backed up and stored to tape or disk.
- Meet Regulatory Compliance - Helps you to preserve important documents, emails, and other critical data, as necessary to comply with internal rules and external regulations.
- Reduce Storage Cost - Helps you to reclaim storage capacity by migrating older or infrequently accessed data from your primary systems to more cost-effective archival storage locations.
- Maintain Multiple File Versions - Helps you to roll back to previous versions of backed-up files (if necessary) or maintain multiple versions of the same files at different destinations.

What To Do Next?

Review the Prerequisites and Considerations

Verify if you have completed the following prerequisites:

- Log into the Console.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Be aware of the following considerations:

- File archive job runs from the latest available backup session.
- File archive job runs only once a day per the schedule.
- File archive delete job runs as a new job to delete the source files.
- The following table provides the limit on the file name length for a file archive job. The files exceeding the limits are skipped.

Destination	Limit
Network Share	240
Cloud	245

Create a Plan with a Backup Task

A plan includes different types of tasks that you want to perform. To create a file archive task, you must first have a valid recovery point. To get a valid recovery point, you have to create a backup task.

The backup task performs a backup of the source nodes and stores the data to the specified destination. File Archive is supported only for the Agent-based Windows backup. The following procedure explains the steps to create the agent-based Windows backup task.

Follow these steps:

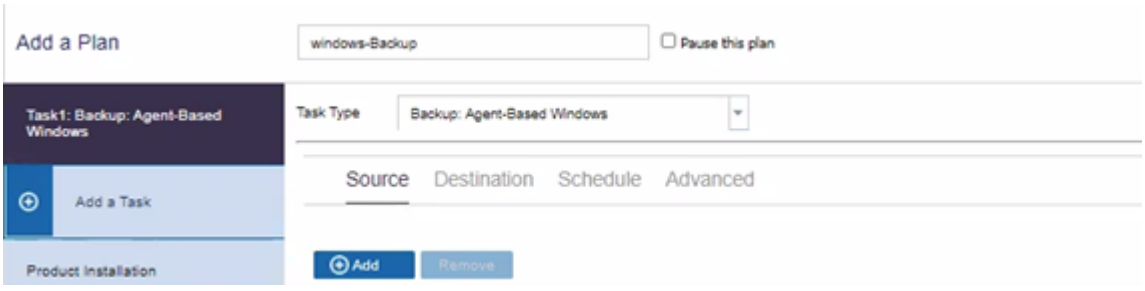
1. Click the **resources** tab on the Console.
2. From the left pane, navigate to **Plans**, and click **All Plans**.

   If you have created plans earlier, those plans are displayed on the center pane.
3. On the center pane, click **Add a Plan**.

   Add a Plan opens.
4. Enter a plan name.
5. (Optional) Select **Pause this plan** check box.

   The plan will not run until you clear the check box to resume the plan.

   Note: If a plan is paused, then any in-progress job is not paused but all corresponding scheduled jobs associated with that plan are paused. However, you can manually run a job. For example, you can manually run backup job and replication job for a node even if the respective plan is paused. In such case, the following task to the on-demand (manual) job does not run. For example, there is a replication task after an on-demand backup job, the replication job does not run for the on-demand backup job. You need to manually run the replication job. When you resume the plan, the pending jobs do not resume immediately. After you resume the plan, the pending jobs run from the next scheduled time.
6. From the **Task Type** drop-down list, select **Backup, Agent-Based Windows**.



Now, specify the Source, Destination, Schedule, and Advanced settings.

Specify the Source

The Source page lets you specify the source nodes that you want to protect. You can select more than one node in a plan. If you have not added any nodes to the Console, you can add nodes when you create or modify a plan from the Source page. You can also save a plan without adding any source nodes. The plan gets deployed only after you add source nodes.

Follow these steps:

1. On the Source tab, click **Add**.
2. Select one of the following options:

   Select Nodes to Protect in Arcserve UDP

   Opens the Select Nodes to Protect in Arcserve UDP dialog. Select the nodes from the displayed list. Select this option if you have already added the nodes to the Console.

   Add Windows Node

   Opens the Add nodes to a plan dialog. Select this option if you have not added the nodes and you want to manually add the nodes to protect.

   Discover Nodes from Active Directory

   Opens the Add Nodes to Arcserve UDP Console dialog. Select this option if you want to discover and add nodes from the Active Directory.

3. If you selected the **Select Nodes to Protect in Arcserve UDP** option, on the Select Nodes to Protect dialog, do the following:

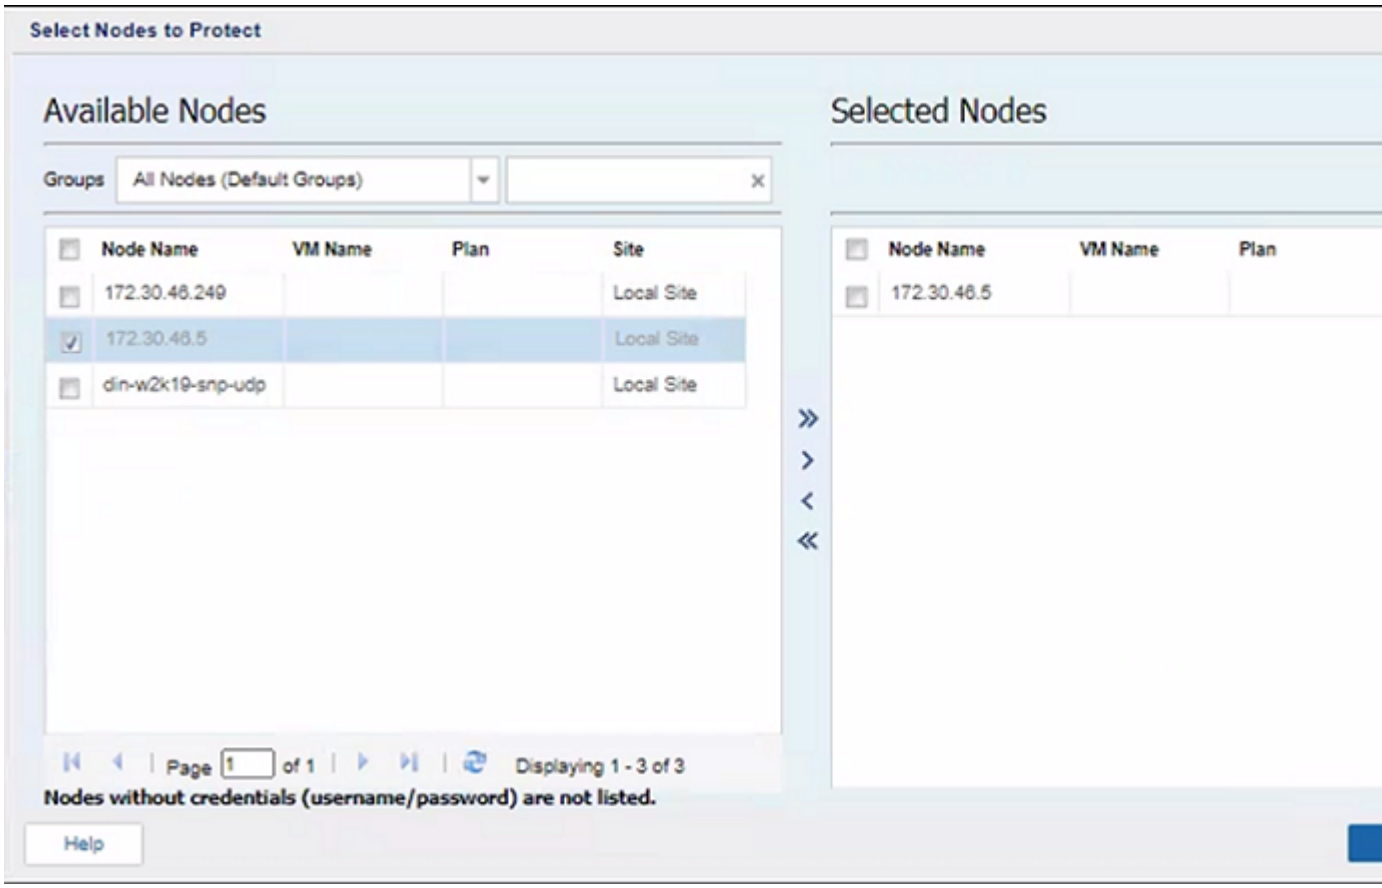
   a. (Optional) Select a filter from the Groups drop-down list to filter nodes. You can enter keywords to further filter your nodes.

      The nodes are displayed on the Available Nodes column.

   b. Select the nodes from the Available Nodes column, and then click the single right arrow (>) or double right arrow (>>) to move the selected nodes or all the nodes respectively to the Selected Nodes column.

      The nodes are displayed in the Selected Nodes column.

   c. Click **OK** to close the dialog.



4. To choose the Protection Type, select one of the following options:

   Back up all volumes

   Prepares a backup snapshot of all the volumes.

   Back up selected volumes

   Prepares a backup snapshot of the selected volume.

   The source is specified.

Specify the Destination

The destination is a location where you store the backup data. You must at least specify the destination to save the plan.

Follow these steps:

1. Select one of the following **Destination Type**:

   Local disk or shared folder

   Specifies that the backup destination is either a local destination or a shared folder. If you select this option, you can save data as either recovery points or recovery sets. The recovery points and recovery sets options are available on the **Schedule** tab.

   Arcserve UDP Recovery Point Server

   Specifies that the backup destination is a recovery point server. If you select this option, then data is stored as recovery points. You cannot store data as recovery sets.

2. If you have selected **Arcserve UDP Recovery Point Server**, then provide the following details:
   a. Select a recovery point server
   b. Select a data store. The list displays all data stores that are created at the specified recovery point server.
   c. Provide a session password.
   d. Confirm the session password.
3. If you have selected **Local disk or shared folder**, then provide the following details:
   a. Provide the full path of the local or network destination. For the network destination, specify the credentials with the write access.
   b. Select the encryption algorithm. For more information, see Encryption Settings.
   c. Optionally, provide an encryption password.
   d. Confirm the encryption password.
   e. Select a type of compression. For more information, see Compression Type.

Note: If you store the data to a local disk or shared folder, you cannot replicate the data to another recovery point server. Replication is supported only if you store the data to a recovery point server.

The destination is specified.

Specify the Schedule

The Schedule page lets you define a schedule for Backup, Merge, and Throttle functions to repeat at specific intervals. After you define a schedule, the jobs run automatically per the schedule. You can add multiple schedules and can provide retention settings.

A Backup Schedule refers to regular schedule that is repeated multiple times a day based on the number of hours or minutes you select. Besides the regular schedule, a backup schedule also provides options to add daily, weekly, and monthly schedules.

Note: For more information on scheduling and retention settings, see Understanding Advanced Scheduling and Retention.

Follow these steps:

1. (Optional) Select the option to manage recovery points. This option is visible only if you have selected Local or shared folder as your backup destination.

   Retain by Recovery Points

   The backup data is stored as recovery points.

   Retain by Recovery Sets

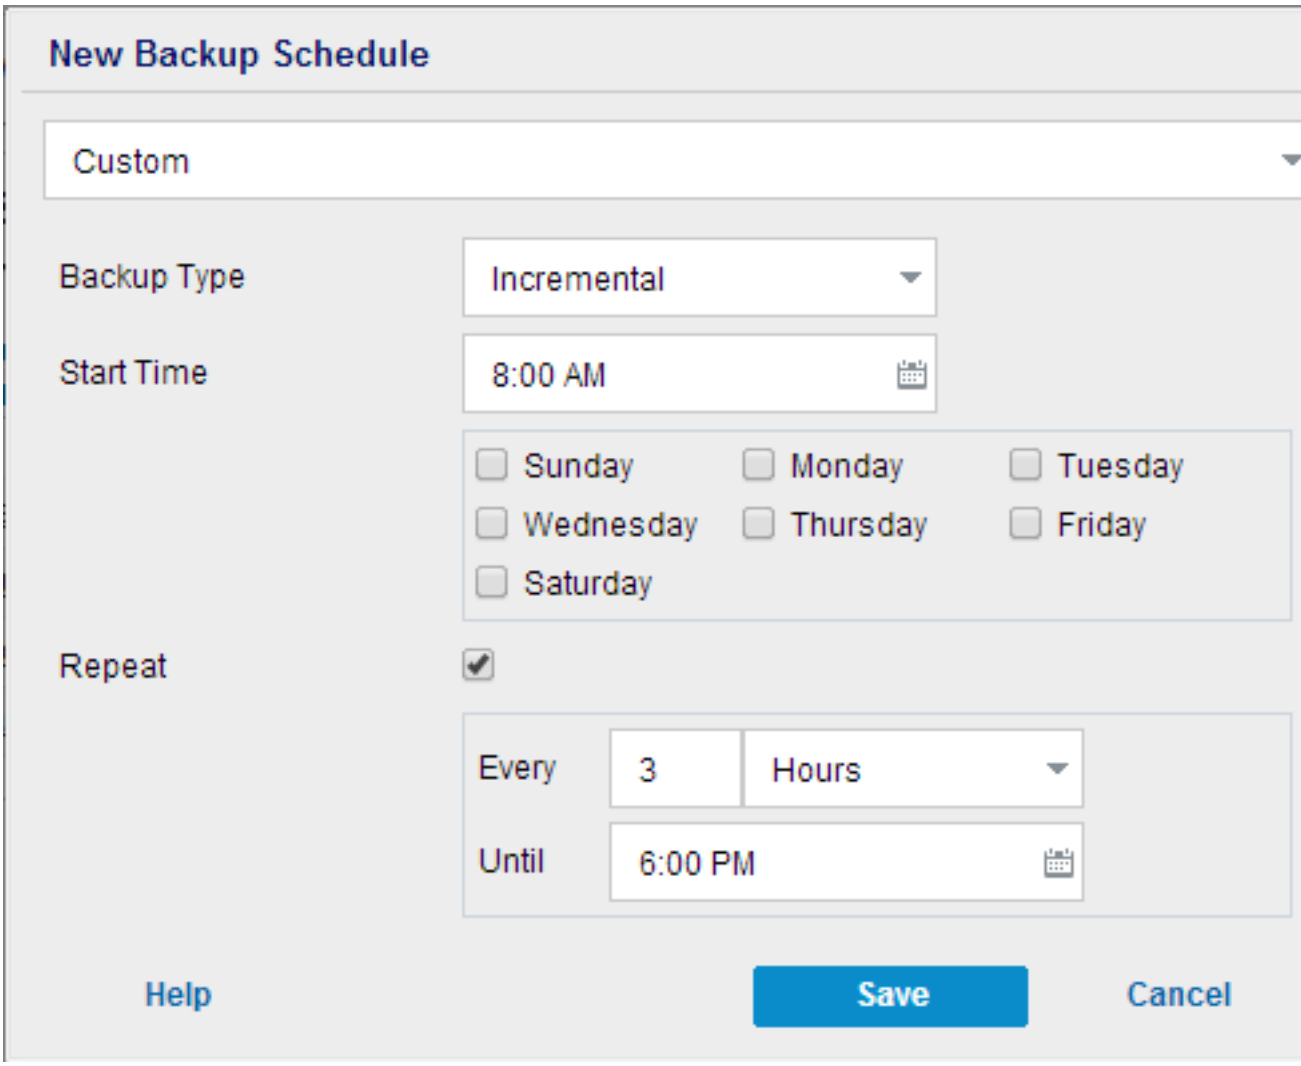
   The backup data is stored as recovery sets.

2. Add backup, merge, and throttle schedules.

   Add Backup Schedule

   a. Click **Add** and select **Add Backup Schedule**.

      The **New Backup Schedule** dialog opens.



b. Select one of the following options:

Custom

Specifies the backup schedule that repeats multiple times a day.

Daily

Specifies the backup schedule that occurs once a day. By default, all the days of the week are selected for Daily backup. If you do not want to run the backup job on a specific day, clear the check box for that day of the week.

Weekly

Specifies the backup schedule that occurs once a week.

Monthly

Specifies the backup schedule that occurs once a month.

c. Select the backup type.

Full

Determines the backup schedule for Full Backups. As scheduled, Arcserve UDP performs a full backup of all used blocks from the source machine. A full backup typically consumes time depending on the backup size.

Verify

Determines the backup schedule for Verify Backups.

Arcserve UDP verifies that the protected data is valid and complete by performing a confidence check of the stored backup image to the backup source. If necessary, the image is resynchronized. A Verify Backup looks at the most recent backup of each individual block and compares the content and information to the source. This comparison verifies that the latest backed up blocks represent the corresponding information at the source. If the backup image for any block does not match the source (possibly because of changes in the system since the last backup), Arcserve UDP refreshes (resynchronizes) the backup of the block that does not match. You can also use a Verify Backup (infrequently) to get the guarantee of full backup without using the space required for a full backup.

Advantages: Produces a small backup image when compared to full backup because only the changed blocks (blocks that do not match the last backup) are backed up.

Disadvantages: Backup time is long because all source blocks are compared with the blocks of the last backup.

Incremental

Determines the backup schedule for Incremental Backups.

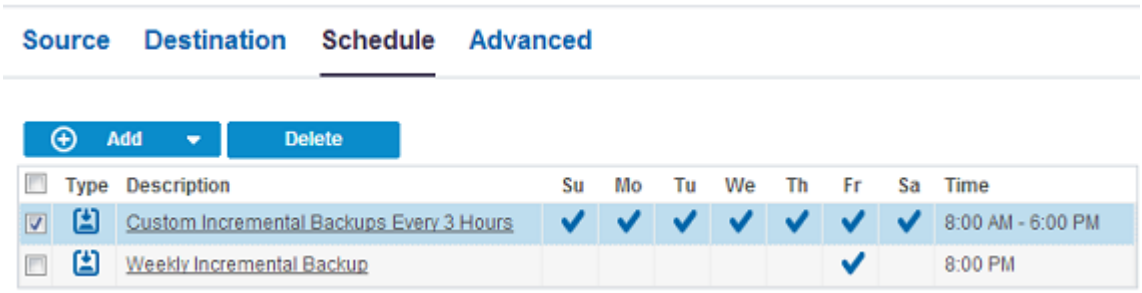
As scheduled, Arcserve UDP incrementally backs up only those blocks that have changed since the last successful backup. The advantages of Incremental Backups are that it is a fast backup and it produces a small backup image. This is the most optimal way to perform a backup.

d. Specify the backup start time.

e. (Optional) Select the **Repeat** check box and specify the repeat schedule.

f. Click **Save**.

The Backup Schedule is specified and displayed on the **Schedule** page.



Add Merge Schedule

a. Click Add and select Add Merge Schedule.

   The **Add New Merge Schedule** dialog opens.

b. Specify the start time to start the merge job.

c. Specify **Until** to specify an end time for the merge job.

d. Click **Save**.

   The Merge Schedule is specified and displayed on the **Schedule** page.

Add Throttle Schedule

a. Click **Add** and select **Add Throttle Schedule**.

   The Add New Throttle Schedule dialog opens.

b. Specify the throughput limit in MB per minute unit.

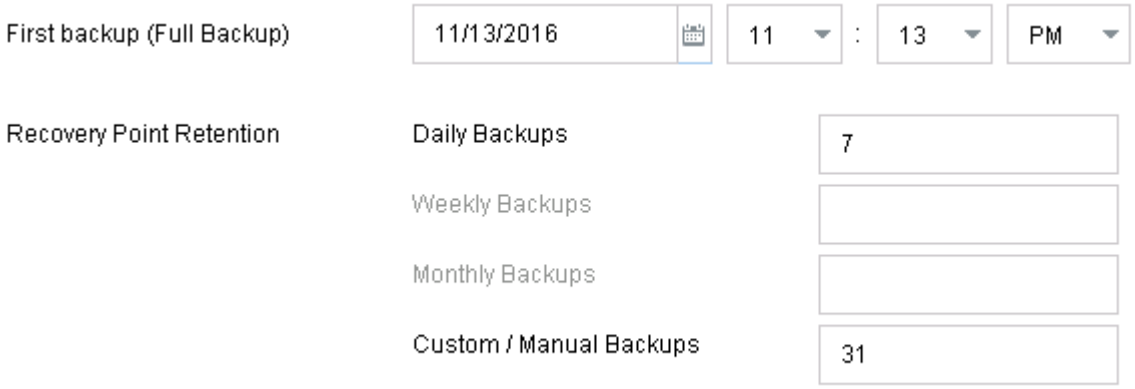
c. Specify the start time to start the backup throughput job.

d. Specify **Until** to specify an end time for the throughput job.

e. Click **Save**.

   The Throttle Schedule is specified and displayed on the **Schedule** page.

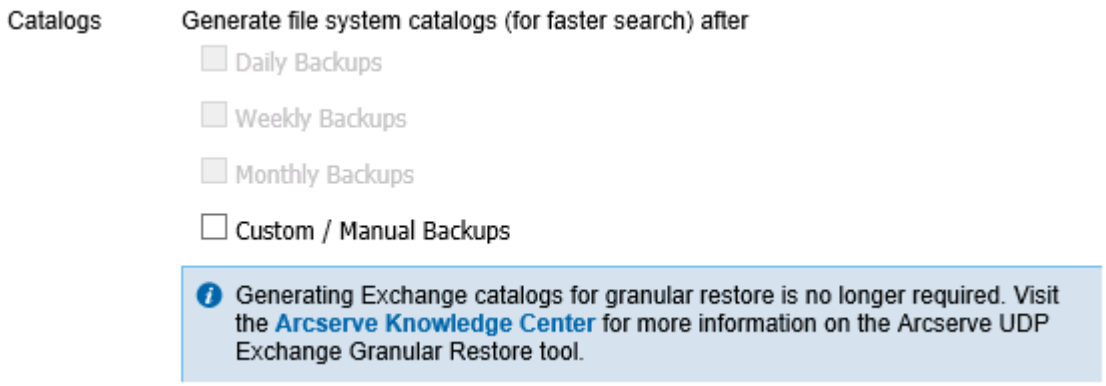
3. Specify the start time for the scheduled backup.



4. Specify the recovery points retention settings for Custom, Daily, Weekly, and Monthly schedule.

   These options are enabled if you have added the corresponding backup schedule. If you modify the retention settings on this page, the changes are reflected on the Backup Schedule dialog.

5. Specify the catalog details.



   Catalogs let you generate the File System catalog. The File System catalog is required to perform faster and better search. If you select the catalog check boxes, the catalogs are enabled depending on the type of backup that you have specified. Clear the check box to disable generating the catalog.

   The schedule is specified.

Understanding Advanced Scheduling and Retention

The scheduling option lets you specify a Custom schedule or a Daily / Weekly / Monthly schedule, or both the schedules. In the Custom schedule, you can configure the backup schedule for each day of the week and you can add up to four backup schedules each day. You can select a specific day of a week and create a time window to define when to run backup and at what frequency.

Schedule	Supported Job	Comments
Backup	Backup job	Define time windows to run backup jobs.
Backup throttling	Backup job	Define time windows to control the backup speed.
Merge	Merge job	Define when to run merge jobs.
Daily schedule	Backup job	Define when to run daily backup jobs.
Weekly schedule	Backup job	Define when to run weekly backup jobs.
Monthly schedule	Backup job	Define when to run monthly backup jobs.

You can also specify the retention settings for the recovery points.

Note: Set the retention settings within each plan to control how data for the nodes assigned to that plan are retained at the target data store.

Schedules for Daily / Weekly / Monthly backups are independent to the Custom schedule, and each other. You can configure only to run Daily backup, Weekly backup or Monthly backup, without configuring the Custom schedule.

Backup Job Schedule

You can add four time windows per day in your backup schedule. A valid time window is from 12:00 AM until 11:59 PM. You cannot specify a time window such as 6:00 PM to 6:00 AM. In such cases, you must manually specify two different time windows.

For each time window, the start time is inclusive, and the end time is exclusive. For example, you have configured to run Incremental Backup every one hour between 6:00 AM and 9:00 AM and the backup will start at 6:00 AM. This means the backup runs at 6:00 AM, 7:00 AM, 8:00 AM, but NOT at 9:00 AM.

Note: If you want to run the backup job repeatedly until the end of day, set the schedule until 12:00 AM. For example, to run backup job every 15 minutes for the entire day, set the schedule from 12:00 AM to 12:00 AM, every 15 minutes.

Backup Throttle Schedule

Backup throttle schedule lets you control the backup throughput speed that in turn controls the resource usage (disk I/O, CPU, network bandwidth) of the server being backed up. This is useful if you do not want to affect the server

performance during business hours. You can add four time windows per day in your backup throttle schedule. For each time window, you can specify a value, in MB per minute. This value is used to control the backup throughput. Valid values are from 1 MB/minutes to 99999 MB/minutes.

If a backup job extends its specified time, then the throttle limit adjusts according to the specified time window. For example, you have defined the backup throttle limit as 500 MB/minute from 8:00 AM to 8:00 PM, and 2500 MB/minute from 8:00 PM to 10:00 PM. If a backup job starts at 7:00 PM and it runs for three hours, then, from 7:00 PM to 8:00 PM the throttle limit is 500 MB/minute and from 8:00 PM to 10:00 PM the throttle limit is 2500 MB/minute.

If you do not define any backup schedule and backup throughput schedule, the backup runs as fast as it can.

Merge Schedule

Lets you merge recovery points based on the provided schedule.

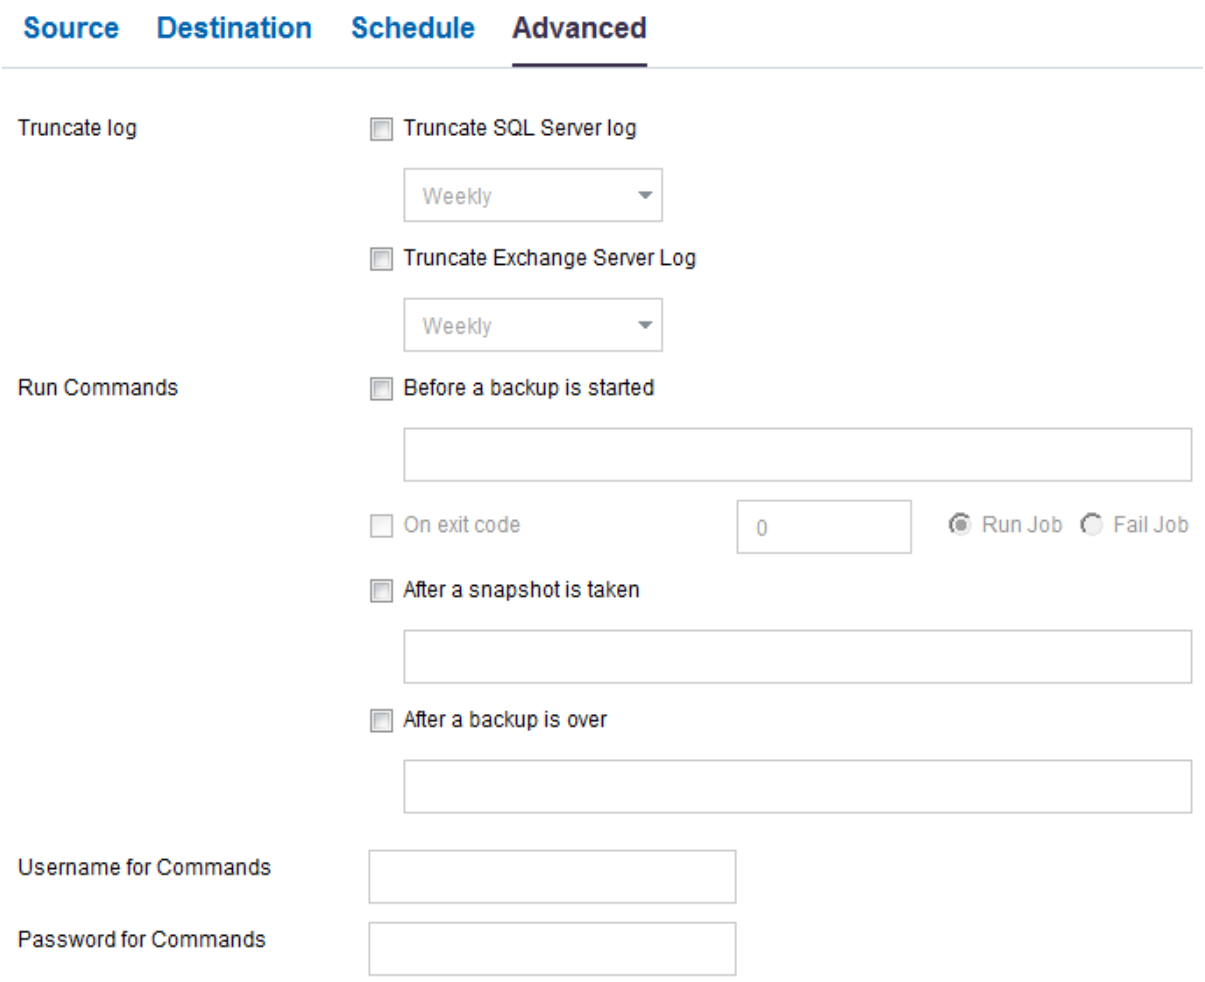
Consider the following points for the merge job:

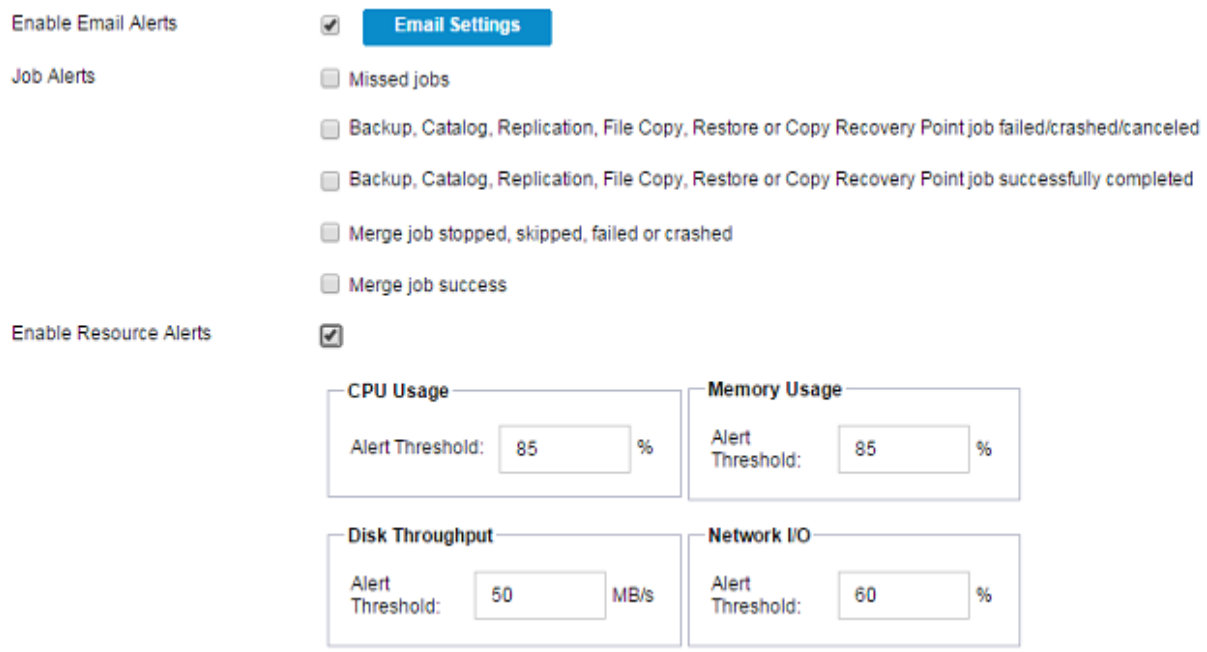
- At any given time only one merge job can run for a node.
- If a merge job starts, it has to complete before the next merge job can start. This means, if one or more sets of recovery points are merged, new recovery point cannot be added to this merge process, until the merge process of the current set of recovery point completes.
- If a merge job is processing more than one set of recovery points (for example set [1~4], set [5~11], and set [12~14]; they are three sets), recovery point server processes these sets one by one.
- If a merge job is resumed after a pause, the job detects at which point it is paused and resumes the merge from the break-point.

Specify the Advanced Settings

The **Advanced** tab lets you specify some advanced settings for the backup job. The advanced settings include providing truncate log settings, providing the location of any scripts, and email settings.

The following image displays the Advanced tab:





Follow these steps:

1. Specify the following details.

Truncate Log

Lets you specify the schedule to truncate logs for SQL Server and Exchange Server. You can specify the schedule as **Daily**, **Weekly**, or **Monthly**.

User Name

Lets you specify the user who is authorized to run a script.

Password

Lets you specify the password of the user who is authorized to run the script.

Run a command before backup is started

Lets you run a script before the backup job starts. Specify the path where the script is stored. Click **On exit code** and specify the exit code for **Run Job** or **Fail Job**. **Run Job** indicates that the backup job will continue when the script returns the exit code. **Fail Job** indicates that the backup job will stop when the script returns the exit code.

Run a command after snapshot is taken

Lets you run a script after the backup snapshot is taken. Specify the path where the script is stored.

Run a command after backup is over

Lets you run a script after the backup job is completed. Specify the path where the script is stored.

Enable Email Alerts

Lets you enable email alerts. You can configure email settings and can specify the types of alerts that you want to receive in an email. When you select this option, the following options are enabled for your selection.

Email Settings

Lets you configure the email settings. Click **Email Settings** and configure the email server and proxy server details.

Job Alerts

Lets you select the types of job emails you want to receive.

Enable Resource Alerts

Lets you specify a threshold for CPU Usage, Memory Usage, Disk Throughput, Network I/O. You can provide the value in percentage. You will receive an email when the Alert Threshold value exceeds.

2. Click **Save**.

Note: When you select a node as a backup source or backup proxy, Arcserve UDP checks whether the agent is installed on the node and if it is the latest version. Arcserve UDP then displays a verification dialog that lists all the nodes that either have an outdated version of the agent or does not have the agent installed. To install/upgrade the agent on these nodes, select the installation method and click Save.

The changes are saved and a green checkmark is displayed next to the task name. The plan page closes.

Note: If you have to add another task, you must select the plan from the **Resources** tab and modify the plan. To modify the plan, click the plan from the center pane. The plan opens and you can modify it.

The backup plan is created and automatically deployed to the source node. The backup runs per the schedule that you have configured in the **Schedule** tab. You can also perform a manual backup at any time.

Add a File Archive Task to the Plan

The file archive task lets you archive individual files to the specified destination. The original files are deleted from the source after you have copied the files to the specified destination and provides more available free space at your source. The file archive job is not dependent on the catalog job.

Note: If a backup job is in progress and you pause the plan, the backup job will get over and the file archive job will not start. When you resume the plan again, the file archive job is not resumed automatically. You have to manually run another backup job to start the file archive job.

For files copied using File Archive, Arcserve UDP leaves a stub file with the "UDP.txt" extension. For more information on updating the stub file, see Update Stub Files.

Follow these steps:

1. Click **Add a Task** from the left pane.

   A new task is added to the left pane.

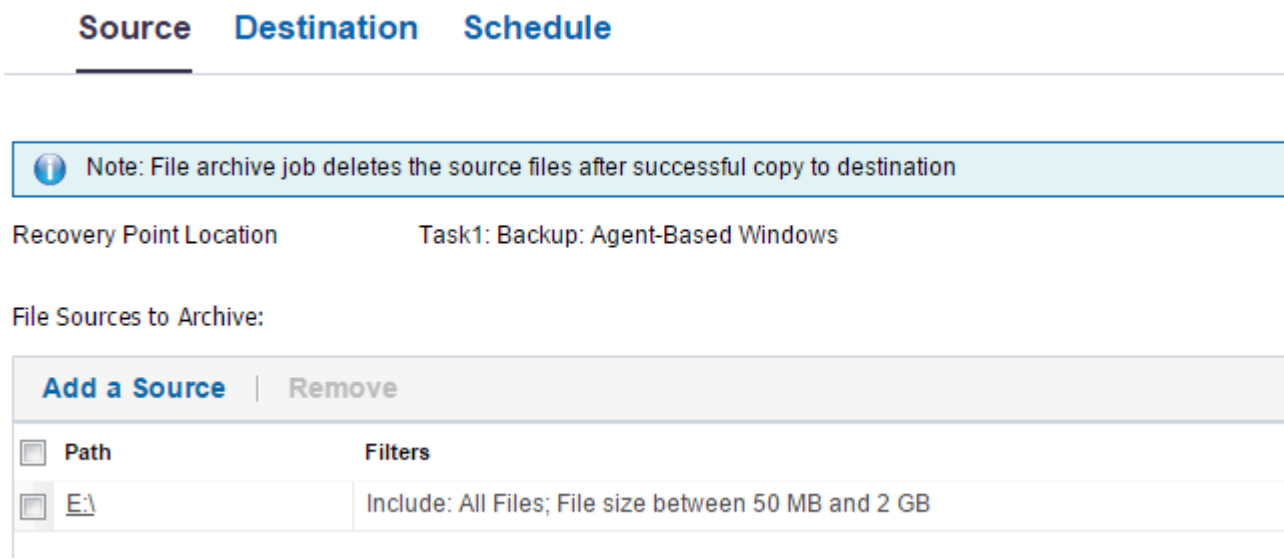
2. From the **Task Type** drop-down menu, select **File Archive**.

   The File Archive task is added.

3. Click the **Source** tab and specify the details.

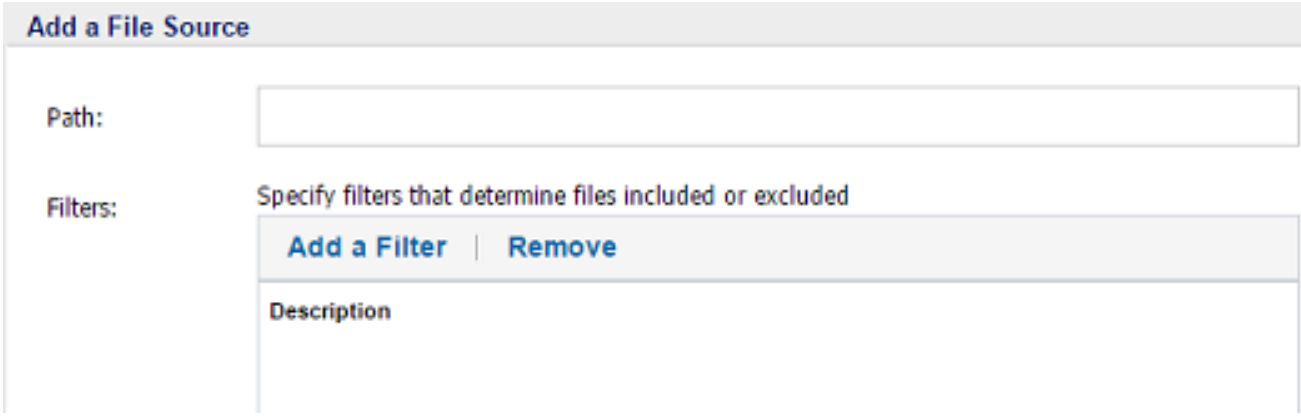
   Recovery Point Location

   Specifies the location of the recovery points that will be archived. This field is pre-selected.

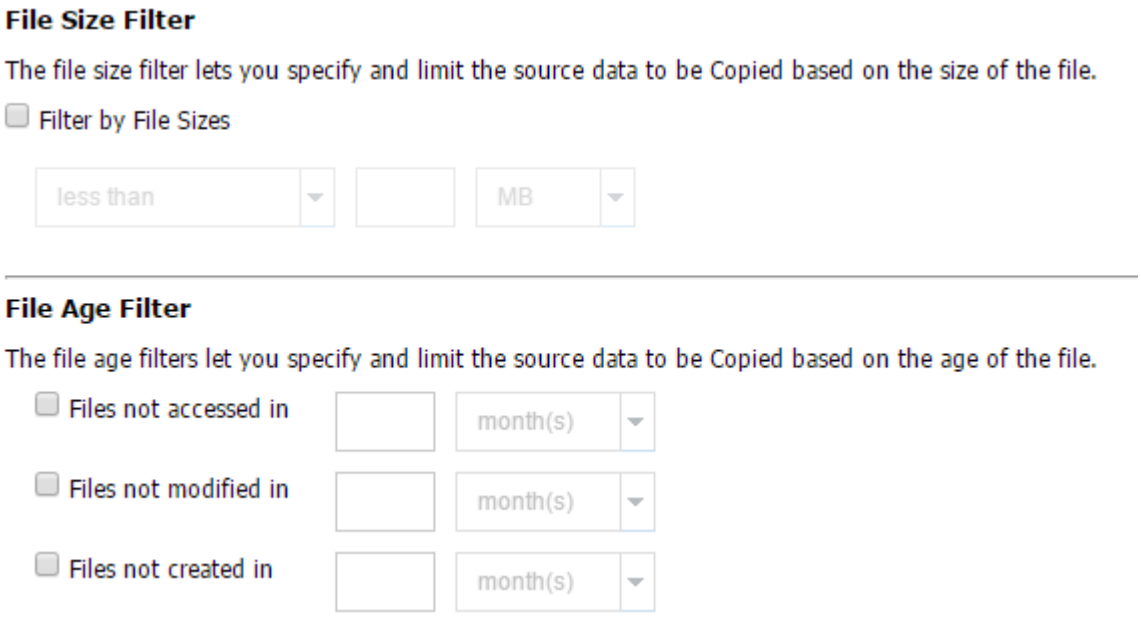


4. Click **Add a Source**.

The **Add a File Source** dialog opens.

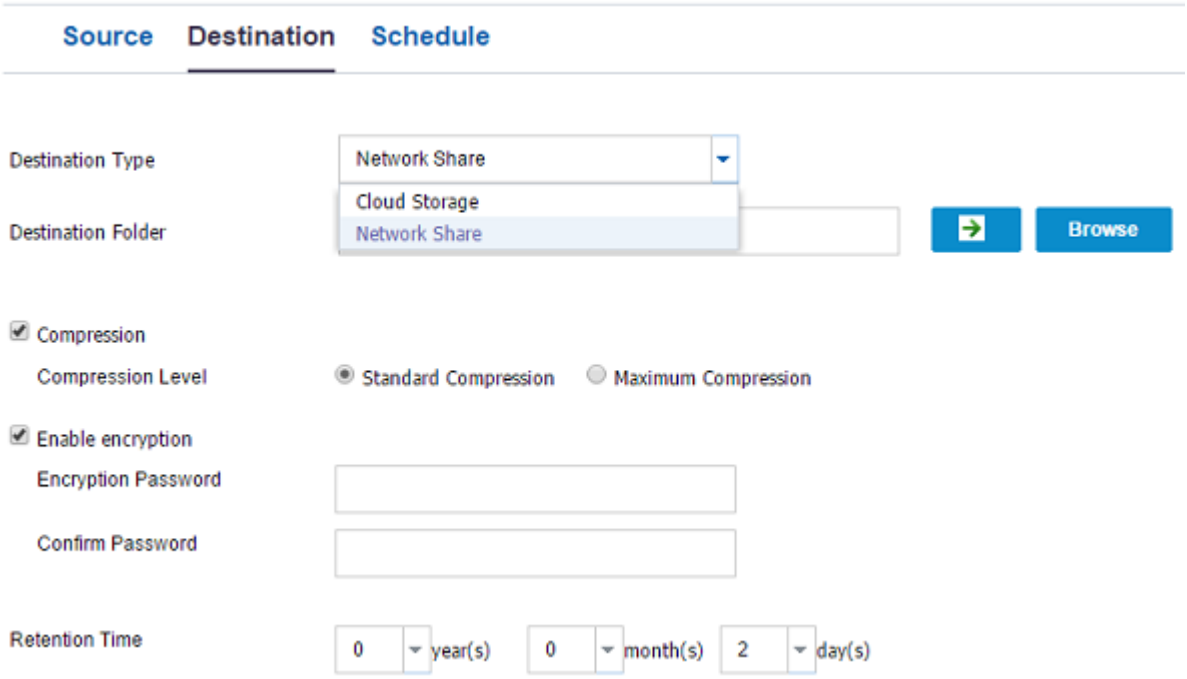


5. Specify the file path of the source path that you want to copy.
6. Specify the **File Size Filter** and **File Age Filter**.



7. Click **Add a Filter**.
8. Select the filter from the list and click **Apply**.
9. Click **OK**.

   The **Add a File Source** dialog closes.
10. Click the **Destinations** tab and specify the destination details.



Destination Type

Specifies that the destination types is a network share or cloud storage. For either destination option, if the connection to the specified destination is lost or broken, Arcserve UDP makes several attempts to continue the file archive job. If these reattempts are not successful, a makeup job is then performed from the point where the failure occurred. In addition, the activity log is updated with a corresponding error message and an email notification is sent (if configured).

Network Share

Specifies that the destination is a shared folder. When selected, lets you specify the full path of the location where you want to move the source files/folders.

Destination Folder

Specifies the destination where the archived files are stored. The destination can be any local volume or folder or a file share accessible by any uniform naming convention (UNC) path. This field is available when you select Network Share or Volume on a Protected Node as the destination type. You can also browse the destination folder.

Cloud Storage

Specifies that the copied files are stored in a cloud environment. Arcserve UDP currently supports file copying to multiple cloud vendors, such as Amazon S3 (Simple Storage Service), Amazon S3-compatible, Windows Azure, Windows

Azure-compatible, Eucalyptus-Walrus, and FUJITSU Hybrid IT Service FJcloud-O. These cloud vendors are publicly available web services that let you store and retrieve any amount of data at any time from anywhere on the web in a safe and secure environment.

Note: To eliminate any potential clock skew error when attempting to connect to the cloud, verify that your machine has the correct time zone set and the clock is in sync with the global time. Always check the time of your machine against the GMT time. If the time of your machine is not synchronized with the correct global clock time (within 5 to 10 minutes), your cloud connection may not work. If necessary, reset the correct time for your machine and rerun your file copy job.

Storage Device

Select the device type from the drop down list.

Cloud Storage

Select the cloud storage path from the drop-down list. The drop-down list is available if you have specified your cloud storage details. If you are specifying the cloud storage account for the first time, click Add to add your cloud account. When you select Cloud Storage from the next time, the account will be displayed in the Cloud Storage drop-down list.

Note: For more information on adding a cloud account, see Add a Cloud Account.

Compression

Specifies the type of compression that is used for the File Archive jobs.

Compression is performed to decrease your storage space at the File Archive destination, but also has an inverse impact on your file archive speed due to the increased CPU usage.

Note: For a compressed File Archive job, the Activity log displays only the uncompressed size.

The available options are:

Standard Compression

Some compression is performed. This option provides a good balance between CPU usage and storage space requirement. This is the default setting.

Maximum Compression

Maximum compression is performed. This option provides the highest CPU usage (lowest speed), but also has the lowest storage space requirement for your file copy.

Enable Encryption

Specifies to use encryption for file archiving.

Data encryption is the translation of data into a form that is unintelligible without a deciphering mechanism. Arcserve UDP data protection uses secure, AES-256 (Advanced Encryption Standard) encryption algorithms to achieve maximum security and privacy of your specified data. When an encryption is selected, you must provide (and confirm) an encryption password.

Retention Time

Specifies the amount of time (years, months, days) that the stored data is retained at the destination location. At the end of the specified retention time period, the stored data is purged from the destination.

The retention time calculations are based upon a month being 30 days and a year being 365 days. For example: If you specify a retention time of 2 years, 2 months, and 5 days, then the total retention time for your file copied data is 795 days (365 + 365 + 30 + 30 + 5).

Important! At the end of the specified retention time when the data is purged from the destination, all of this moved data is no longer stored or saved.

Note: The Retention Time purge process is only triggered if the File Copy Schedule option is enabled.

File version less than

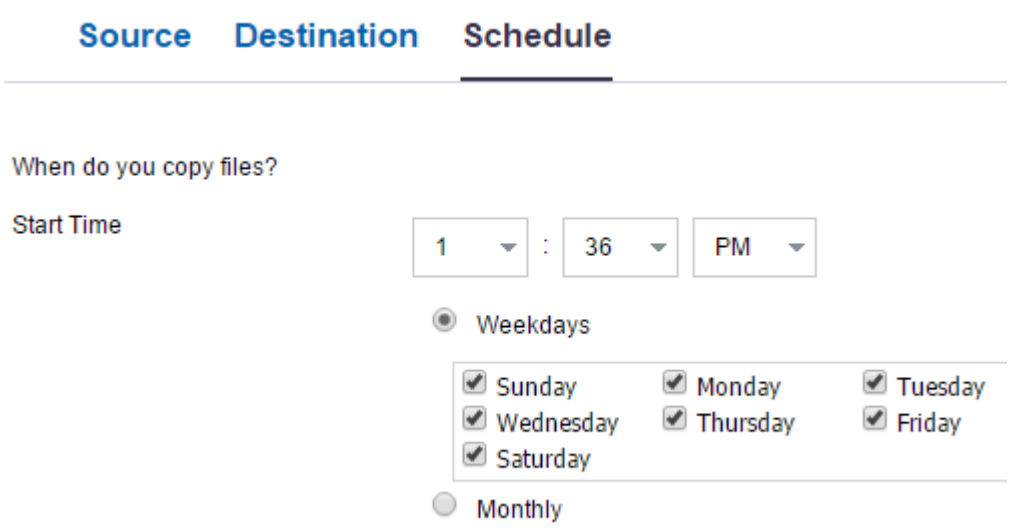
This setting only applies to copied data that is retained (not copied data that is moved).

Specifies the number of copies retained and stored at the destination location (cloud or disk). After this number is exceeded, the earliest (oldest) version will be discarded. This cycle of discarding the oldest stored version repeats as newer versions are added to the destination, allowing you to always maintain the specified number of stored versions.

For example, if your specified File Versions retention count is set to 5 and you perform five file copies at times t1, t2, t3, t4, and t5, these file copies become the five file copy versions retained and available to recover. After the sixth file copy is performed (new version is saved), Arcserve UDP will remove the t1 copy and the five available versions to recover are now t2, t3, t4, t5, and t6.

By default, the number of copies retained at the destination location before discarding is 15.

11. Click the **Schedule** tab and specify the time to run the file archive job.



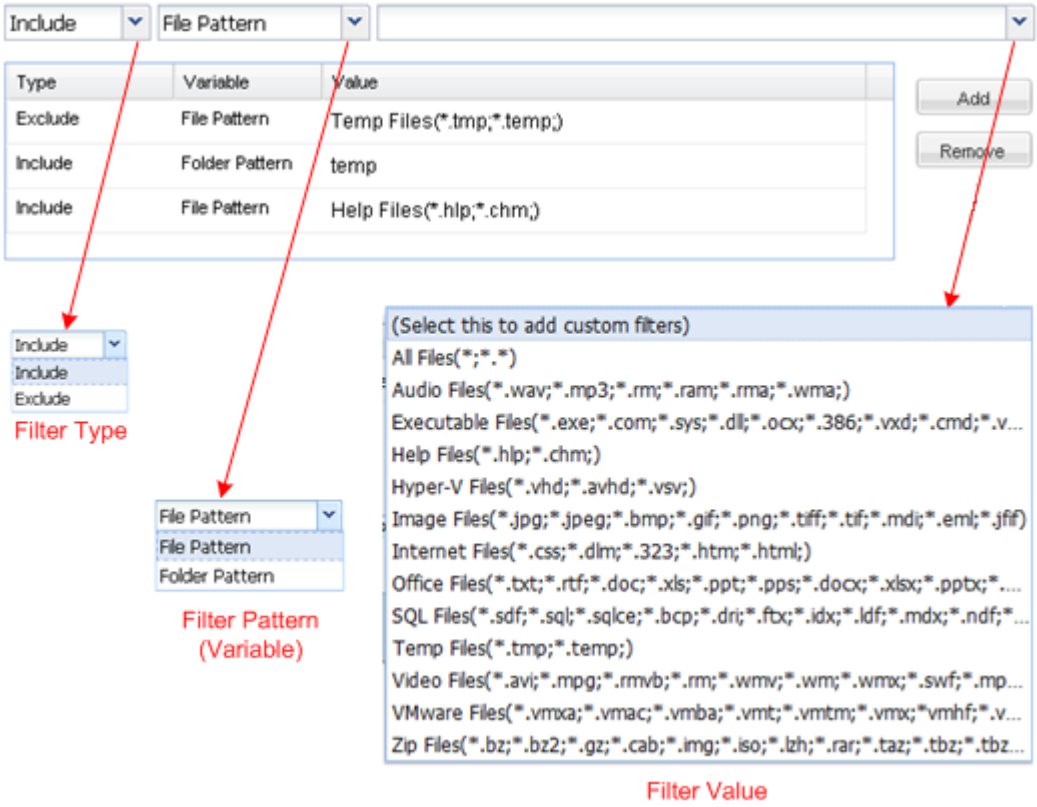
12. Click **Save**.

    The changes are saved and the file archive task is automatically deployed to the node.

Add File Archive Filters

Add a Filter

Lets you add a filter. Filters let you limit the objects to be file copied by certain specified types and values.



Filter Type

There are two types of filters: Include and Exclude.

An Include filter copies only those objects from the file copy source that match the specified value.

An Exclude filter copies all objects from the file copy source except those that match the specified value.

You can specify multiple filters within the same file copy request by separating each filter value with a comma.

- If you specify multiple Include filters, the data is included in the file copy if any one of those Include filters matches.
- If you specify multiple Exclude filters, the data is excluded from the file copy if any one of those Exclude filters matches.
- You can mix both Include and Exclude filters in the same file copy request.

Note: When the specified parameters of Exclude and Include filters conflict, the Exclude filter is always a higher priority and is enforced. An Include filter can never file copy an object that was also Excluded.

Filter Variable (Pattern)

There are two types of variable pattern filters: File Pattern and Folder Pattern.

You can use a File Pattern filter or Folder Pattern filter to include or exclude certain objects from the file copy.

Filter Value

The filter value lets you limit the information that is file copied by selecting only the parameter information that you specify, such as .txt files.

Arcserve UDP supports the use of wildcard characters to help select multiple objects to file copy with a single request. A wildcard character is a special character that can be used as a substitute to represent either a single character or a string of text.

The wildcard characters asterisk and question mark are supported in the Value field. If you do not know the complete file/folder pattern value, you can simplify the results of the filter by specifying a wildcard character.

"*" - Use the asterisk to substitute zero or more characters in the value.

"?" - Use the question mark to substitute a single character in the value.

For example, you can enter *.txt to exclude all files with a .txt extension if you do not know the specific file name. You can provide as much of the file name as you know, then use wildcards to fill in the blanks.

Note: When you select File Pattern as the filter type, a drop-down list of predefined filters for many commonly used files is available (MS-Office files, Image files, Executable files, Temp files, etc.). After choosing any of the predefined filters, you can still append or modify the corresponding values.

File Size Filter

File size filters let you limit the source objects to be file copied based upon the size of the file. When you enable the file size filter, the parameters that you specify become the filter for which objects will and will not be included in the file copy. You can select the range (Equal to or Greater Than, Equal to or Less Than, or Between) and then enter a value for the size.

For example, if you specify Equal to or Greater Than 10MB, then Arcserve UDP only file copies objects that meet this criteria. All other objects that do not meet this file size criteria are not file copied.

File Age Filter

File age filters let you automatically include source objects to be file copied based upon certain dates for the file. You can select a parameter (Files not accessed in, Files not modified in, and/or Files not created in) and then enter a value for the number of days, months, or years for the file age filter. You can select multiple file age filters for automatic file copying.

For example, if you specify Files not modified in 180 days, then Arcserve UDP automatically copies all files that meet this criteria (have not been modified during the last 180 days).

Important! If you specify both File Size and File Age filters (or multiple File Age filters), then only the files which meet all of the specified filter parameters are copied. Files which do not meet any one of these specified parameters are not copied.

Update Stub Files

For files copied using the File Archive, Arcserve UDP leaves a stub file with the "UDP.txt" extension. The stub file contains information about the destination where the files were moved and some additional information. If a file is restored to the original location and then gets moved again to the specified destination, then the stub file is updated with this move information. If necessary, these file copy stub files can be safely disabled or deleted without any negative impact. (Existing stub files are not deleted when the registry key is changed to no longer create stub files).

The following information is present in the stub file by default:

Please contact your IT department to restore this file.

You can change the default text by adding a configuration entry. Add the following XML tag entry in the FileCopyDebugSetting.xml file from the $UDPHome\Engine\Configuration directory:

<ArchiveStubFileText> New Text can be added here to display in stub file

</ArchiveStubFileText>

Example: The FileCopyDebugSetting.xml file would look like

```
<?xml version="1.0" encoding="UTF-8"?>
<HKLM>
<AFArchiveDLL>
<ArchiveStubFileText> New Text can be added here to
display in stub file </ArchiveStubFileText>.
</AFArchiveDLL>
</HKLM>
```

If the FileCopyDebugSetting.xml file is not present under the $UDPHome\Engine\Configuration directory, create the XML file.

If you want to disable the stub file creation, add the following XML tag entry in the FileCopyDebugSetting.xml file from the $UDPHome\Engine\Configuration directory:

<CreateStubFile>0</CreateStubFile>

Example: The FileCopyDebugSetting.xml file would look like

```
<?xml version="1.0" encoding="UTF-8"?>
<HKLM>

<AFArchiveDLL>
```

```
<CreateStubFile>0</CreateStubFile>
</AFArchiveDLL>
</HKLM>
```

Note: If you disable or delete the file copy stub files, you can no longer track the status and location of moved files.

(Optional) Perform a Manual File Archive

Typically, a File Archive is performed automatically and it is controlled by the File Archive schedule settings. In addition to the scheduled File Archive, a manual File Archive provides you the option to copy your important files on a need basis. When you run a manual File Archive, the job archives all the sessions in the File Archive source.

File Archive can be run manually from the Nodes view and Plans view by clicking **Actions** menu or context menu.

Follow these steps to run the File Archive manually from the Nodes view:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and click **All Nodes**.

   Nodes are displayed in the center pane.
3. Select the nodes for a File Archive job. The nodes must have a File Archive plan assigned to it.
4. On the center pane, click **Actions**, **File archive Now**.

   The Run File archive now dialog opens.
5. Click **OK**.

   The File Archive job runs.

   The manual File Archive is successfully performed.

Verify the Plan

To verify the file copy plan, confirm that you have successfully created the plan. After you verify that the plan is created successfully, check whether the backup job is running as scheduled. After the backup job successfully completes, the file copy job runs as scheduled. You can check the status of the backup job and the file copy job from the **jobs** tab.

Follow these steps to verify plans:

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and click **All Nodes**.

   A list of all nodes is displayed on the center pane.
3. Verify that plans are mapped with nodes.

Follow these steps to verify file copy jobs:

1. Click the **jobs** tab.
2. From the left pane, click **All Jobs Completed**.

   The status of each job is listed on the center pane.
3. Verify that the backup job and file copy job is successful.

How to Create a Copy to Tape Plan

Arcserve UDP integrates with Arcserve Backup to copy backup data to a tape media destination. You can create a plan to backup source data and copy the backup data to the tape. You can create and manage the plan from the Console.

The advantages of archiving recovery points to a tape media are:

- Meet Regulatory Compliance - Helps you to preserve important documents, emails, and other critical data, as necessary to comply with internal rules and external regulations.
- Reduce Storage Cost - Helps you to reclaim storage capacity by migrating older or infrequently accessed data from your primary systems to more cost-effective archival storage locations.
- Maintain Multiple File Versions - Helps you to roll back to previous versions of backed-up files (if necessary) or maintain multiple versions of the same files at different destinations.

Supported Scenarios

- If the Destination of Task 1 is **Arcserve UDP Recovery Point Server**, install the client agent on the RPS node.
- For Agent-based plan with Destination as **Local/Remote share**, install the client agent on all the Arcserve UDP agent nodes.
- For Host-Based Agentless plan, install the client agent on the Arcserve UDP Proxy Node.

What To Do Next?

- Review the Prerequisites
- Create a Plan with a Backup Task
- Add a Copy to Tape Plan
- Verify the Plan

Review the Prerequisites and Considerations

Verify if you have completed the following prerequisites:

- Log into the Console.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.
- Added the Arcserve Backup Server to the Console. For more information on adding Arcserve Backup Server to the Console, see Add Arcserve Backup Servers.
- The Recovery point server that you specify in the primary task of a plan must have the Arcserve Backup Server agent installed.

Create a Plan with a Backup Task

A plan includes different types of tasks that you want to perform. Typically a plan includes a primary tasks followed by a secondary task. Typically, a primary task is a backup task or replicate from a remote Console task. The role of a backup task is to create a backup of the source nodes that you want to protect. You can back up data from Windows and Linux physical and virtual machines. You can then save the backup data to another location as an added precaution.

For more information on:

- How to back up a Windows node, see How to Create a Windows Backup Plan.
- How to back up virtual machines, see How to Create a Host-Based Virtual Machine Backup Plan.
- How to back up Linux nodes, see How to Create a Linux Backup Plan.
- Replicating a backup data to a remote destination, see How to Replicate Data Between Data Stores Managed from Different UDP Consoles.

Add a Copy to Tape Task to the Plan

The Copy to Tape task lets you copy your data to a tape. The tape media is identified from the Arcserve Backup Server that you have added.

Follow these steps:

1. Click **Add a Task** from the left pane.

   A new task is added to the left pane.

2. From the **Task Type** drop-down menu, select **Copy to Tape**.

   The Copy to Tape task is added.

   Note: The custom backups does not trigger the Copy to tape jobs.

   Now specify the Source, Destination, Schedule, and Advanced details.

Specify the Source

The source file is typically a backup destination or a replication destination.

Follow these steps:

1. Specify the following details in the **Source** tab.

   Source of Copy to Tape

   Specifies the location of the recovery points that will be copied. If you have only one source, then this field is pre-selected. For example, if your plan has only one backup task and then you add Copy to Tape task, then the destination of the backup task becomes the source of copy to tape. If you have multiple source for copy to tape, you can add a Copy to Tape for each of those source. For example, if the first task is a backup task, the second and third tasks are replicate tasks, then you can add three Copy to Tape task to the plan.

   Types of Recovery Points

   Specifies the types of recovery points. The available options are Daily Backups, Weekly Backups, and Monthly Backups.

2. Click the **Destination** tab.

   The **Destination** page opens.

   You have successfully provided the source details.

Specify the Destination

The destination is a tape media group in your Arcserve Backup Server. You must add the Arcserve Backup Server to Console before you create this task.

Follow these steps:

1. Specify the following details for the backup destination.

   Arcserve Backup Server

   Select the Arcserve Backup server from the drop-down list.

   Media Group

   Select the media group from the drop-down list. The media group depends on the Arcserve Backup server. By default, the <ASBU> disk-based device are not listed in the media group. Also, the tape raid group is not listed in the media group.

   To migrate Arcserve UDP recovery points to an Arcserve Backup disk-based device, the incremental recovery points must be converted to full recovery points. The recommended way to migrate Arcserve UDP recovery point to a disk-based device is to perform a replication from Arcserve UDP data store to another data store.

   Only tape groups can be configured as the destination for a Copy to Tape job. You can modify a setting key in ConsoleConfiguration.xml file available at the UDP Management Configuration path to display the ASBU FSD group as the Copy to Tape destination:

   C:\Program Files\Arcserve\Unified Data Protection\Management\Configuration\ConsoleConfiguration.xml

   <displayASBUFSDGroup>0</displayASBUFSDGroup>

   When the value is 0, the ASBU FSD group is not displayed.

   <displayASBUFSDGroup>1</displayASBUFSDGroup>

   When the value is 1, the ASBU FSD group is displayed.

   After changing the settings, restart Arcserve UDP Management Service to make effective.

   The list of the media group is queried from Arcserve Backup Server. It is corresponding to the device group list in the Arcserve Backup Server.

   Multiplexing

Select the check box to enable multiplexing. Specify the maximum number of streams that can write to a tape at the same time. The default number of streams is 4 and the supported range is from 2 through 32.

Encryption

Specifies to use encryption for copy to tape.

Enabling encryption ensures that data is encrypted on tape. When an encryption is selected, you must provide (and confirm) an encryption password.

Compression

Enabling compression ensures that the data is compressed on tape.

2. Verify the Media Group Details for the selected Arcserve Backup server and media group.
3. Click the **Schedule** tab.

   The Schedule page opens.

   You have specified the destination.

Specify the Schedule

You can specify the schedule to start your copy to tape job. You can also decide the media retention policy and tape usage mode.

Follow these steps:

1. Click **Add**, **Add Copy to Tape Schedule**.

   The **Add New Copy to Tape Schedule** dialog opens.

2. Specify a tape schedule.

   A schedule defines the time range to start a copy to tape job. If you add a schedule, then the copy to tape job runs only during the defined time schedule. If you do not specify a schedule, then the copy to tape job runs within 30 minutes after a qualified recovery point is ready on a data store.

3. Click **Save**.

   The **Add New Copy to Tape Schedule** dialog closes.

4. Select the **Medial Pool Name** from the drop-down list.

   A default media pool name is selected based on the plan name.

   You can also select an existing media pool name from the drop-down list. In that case, the media retention policy and tape usage mode associated with that media pool would be copied into this task. You can share the tapes across multiple Arcserve UDP plans by specifying the same media pool in all the Arcserve UDP plans.

   You can also specify a different pool name. A maximum of 13 characters is accepted for a media pool name.

5. Specify a **Recovery Point Retention** policy.

   A recovery point retention policy allows you to retain recovery points on a daily, weekly, or monthly basis. You can specify a different retention time for different types of recovery points. For example, if you select **Daily Backups** and **Weekly Backups** from the **Source** tab, then you can specify different recovery point retention period for both these types of backup.

6. Select one of the **Tape Usage** options.

   Append to Existing Tapes

   Indicates that all recovery points generated within the specified retention period are copied to the same tape. For example, if you have specified the retention time for daily backup as 7, then all the recovery points from Day 1 to Day 7 are copied to the same tape. Then, all the recovery points from the next 7 days (Day 8 to Day 14) are copied to a different tape, and so on.

The recovery points from the first week (Day 1 to Day 7)) are retained for the next 7 days (Day 8 to Day 14). From Day 15, the recovery points are again copied to Tape 1 because the retention policy for first week recovery points is expired.

The following list shows the default retention time for Append to Existing Tape:

- Daily - 7 days
- Weekly - 5 weeks
- Monthly - 12 months

Seven daily recovery points are copied to the same tape, five weekly recovery points are copied to the same tape, and 12 monthly recovery points are copied to the same tape.

Copy to Separate Tapes

Indicates that recovery points of each day are copied to a separate tape. For example, if you have specified the retention time for daily backup as 7, then recovery points from Day 1 are copied to Tape 1, recovery point from Day 2 are copied to Tape 2, recovery points from Day 3 are copied to Tape 3, and so on.

The recovery points from Day 1 is retained for 7 days. On Day 8, the recovery points are copied to Tape 1 because the retention policy is expired for Day 1 recovery points.

The following list shows the default retention time for Copy to Separate Tapes:

- Daily - 7 days
- Weekly - 5 weeks
- Monthly - 12 months

Each of the seven daily recovery points are copied to separate tapes, each of the five weekly recovery points are copied to separate tapes, and each of the 12 monthly recovery points are copied to separate tapes.

7. Click the **Advanced** tab.

   The **Advanced** page opens.

   You have specified the schedule.

Specify the Advanced Settings

The advanced settings lets you configure some additional settings for the copy to tape task.

1. Specify the following details.

Media Eject

Specifies that the media is ejected from the drive after the job finishes. This helps prevent any other job from overwriting information on this media.

Backup Verification

Specifies that Arcserve Backup verifies the reliability of the backup by checking the header of each file for readability. This option does not apply to multiplexed backups.

Run a command before a copy to tape job is started

Lets you run a script before the backup job starts. Specify the path where the script is stored. Click On exit code and specify the exit code for Run Job or Fail Job. Run Job indicates that the backup job will continue when the script returns the exit code. Fail Job indicates that the backup job will stop when the script returns the exit code.

Run a command after a copy to tape job is over

Lets you run a script after the backup job is completed. Specify the path where the script is stored.

Username for commands

Lets you specify username to run the script.

Password for commands

Lets you specify the password to run the script.

Enable Email Alerts

Lets you enable email alerts. You can configure email settings and can specify the types of alerts that you want to receive in an email. When you select this option, the following options are enabled for your selection.

Email Settings

Lets you configure the email settings. Click **Email Settings** and configure the email server and proxy server details.

Job Alerts

Lets you select the types of job emails you want to receive.

2. Click **Save**.

   The changes are saved and a green checkmark is displayed next to the task name. The plan page closes.

   The copy to tape task is created and it runs per the schedule.

   Note: When you run the **Copy to Tape** job, the job monitor values in Arcserve UDP console are different from those in the Arcserve Backup.

   Important! After you copy the recovery points to a tape, you cannot restore the data in tape from Arcserve UDP Console. You have to restore the data from Arcserve Backup Manager. For more information on restoring the tape data, see Backing Up and Recovering D2D/UDP Data in Arcserve Backup Administrator Guide.

How to Create and Manage an Instant Virtual Machine on Microsoft Azure

Instant virtual machine (Instant VM) supports creating a virtual machine on Microsoft Azure instantly The IVM on Microsoft Azure has following advantages:

- Provides immediate access to data and applications present in the Arcserve UDP backup sessions.
- Eliminates the downtime associated with a traditional restore or conversion of the backup session to a virtual machine.
- Provides an alternative to create virtual machine on cloud, rather than local.

You can create an Instant VM from the following backup sessions:

- Agent-based Linux backup
- Host-based Agentless backup for Linux VM

What To Do Next?

- Best Practices
- Review the Prerequisites
- Create an Instant Virtual Machine Plan on Microsoft Azure
- Manage an Instant Virtual Machine Plan on Microsoft Azure

Best Practices for an Instant Virtual Machine on Microsoft Azure

To protect the nodes in On-premise network, you must install UDP Console on the host.

- Select HTTPS as protocol when installing the UDP components.
- The easiest way to create a resource group is to create at least one test virtual machine. Azure will run you through the steps to create all of the resources for the test Virtual Machine that you can use for standby Virtual Machine.
- (Optional) Create an RPS in Azure.

  Follow these steps:

1. Open TCP ports 8014 and 8015 inbound.
2. If accessing the RPS from a remote web browser, resolve the name of the RPS server to the public IP.
3. Use shared plan task “Replicate to remotely managed RPS” to replicate.

Review the Prerequisites for an Instant Virtual Machine on Microsoft Azure

Complete the following prerequisites before creating an Instant VM:

- From the Compatibility Matrix, verify if the VM is supported by Microsoft Azure and UDP.
- Add a Microsoft Azure Cloud Account
- Verify that you have at least one Recovery Point Server on local machine, as the backup destination.
- Verify that you have at least one Recovery Point Server on Microsoft Azure, as replication destination.
- Verify that you have at least one Linux backup server locally, for the backup job.
- Verify that you have at least one Linux backup server on Microsoft Azure, for the instant VM job.
- Verify that a Microsoft Azure account has been added.

Limitation

- Linux machines booted with UEFI are NOT supported.
- Linux machines having Btrfs file system across multiple disks are NOT supported.

Create an Instant Virtual Machine on Microsoft Azure

To create an IVM on Microsoft Azure plan, we recommend to perform one of the following options:

- Backup local protected node to local Recovery Point Server data store, then replicate to Recovery Point Server on Microsoft Azure. For information about adding node, see How to Add Nodes to the Console.
- Backup local protected node to local CIFS (NFS) share location, then copy to CIFS (NFS) share location on Microsoft Azure. For information about backing up protected node, see How to Create a Linux Backup Plan.

Creating an Instant VM involves the following five broad steps:

1. Open the Instant VM wizard
2. Select the recovery point
3. Select the VM location
4. Select the recovery server
5. Specify the Instant VM details
6. Submit the Instant VM job

Open the Instant Virtual Machine Wizard

You can configure and create an Instant VM from the Instant VM wizard. There are three ways to open the Instant VM wizard:

- From the Node Management
- From the Destination Management: Recovery Point Server
- From the Destination Management: Shared Folder

Open the Wizard from the Node management view

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes** and click **All Nodes**.
3. All the nodes are displayed on the center pane.
4. Right-click a node and select **Create an Instant VM**.

   The Instant VM wizard opens.

   Note: If a node is not associated with any plan, such node does not have the **Create an Instant VM** option.

Open the Wizard from the Destination management view

From Destinations: Recovery Point Server

1. Click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Recovery Point Servers**.

   If you have added any data stores, these data stores are displayed in the center pane.
3. Click the data store.

   If you have already backed up data to the RPS, all the source node are listed in the pane.
4. Right-click a node and select **Create an Instant VM**.

   The Instant VM wizard opens.

From Destinations: Shared Folder

1. Click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Shared Folders**.
3. If you have added any shared folders, these shared folders are displayed on the center pane.

4. Click a shared folder and select **Recovery Point Browser**.

   If you have already backed up data to the shared folder, all the source nodes are listed in the pane.

5. Right-click a node, and select **Create an Instant VM**.

   The Instant VM wizard opens.

Select a Recovery Point

The Select a Recovery Point page displays the location of the recovery point and lets you select a recovery point. The recovery point can be in a shared location or in a data store in RPS.

The Console automatically recognizes the location of the recovery point and pre-selects the **Location Type**, **Recovery Point Server**, and **Data Store** or **Network Share Folder** fields.

Note: Select Recovery Point Server (Shared folder) session as replication destination to save network bandwidth and time.

Expand the Date list, select the required recovery point from the list, and click **Next**.

The **VM Location** page opens.

Note: If you create an Instant Virtual Machine from the Node management view, then only those Recovery Point Servers are listed that are in the same Site as the source node. If you want to create an Instant Virtual Machine from the Recovery Point Server managed from a different Site, then navigate to that Site and open the Wizard from the Destination management view.

Select a VM Location

Specify the location of the virtual machine where you want to create the Instant VM.

Follow these steps:

1. Select Microsoft Azure.
2. Select an Account name.

   Note: Select corresponding account, if not added, refer to Add a Microsoft Azure Cloud Account.
3. Click **Next**.

   The **Recovery Server** page opens.

Select a Recovery Server

The recovery server hosts the core module of the Instant VM. The default recovery server is the RPS. You can also assign a node as a recovery server.

For Linux backup sessions, the recovery server is a Linux Backup Server.

Follow these steps:

1. Select a Linux Backup Server from the node list.

   Notes: Select Recovery Point Server (Shared folder) session as replication destination to save network bandwidth and time.

2. Click **Next**.

   The Instant VM Details page opens.

   Note: Ensure your select the Recovery Point Server that is located at Microsoft Azure.

Specify the Instant Virtual Machine Details

Specify the details for the Instant VM.

Follow these steps:

1. Specify the name and description of the Instant VM.

VM Name

Specifies the name of the Instant VM. The name of the source node prefix with "UDPIVM_" becomes the default name of the Instant VM. You can modify the name. Some special characters, such as '@', \ and so on, are not allowed in the name.

Description

(Optional) Specifies the description for the Instant VM.

Location

Specifies the location of your Microsoft Azure account.

Disk Type

Specifies the type of disk.

Azure VM Size

Specifies the supported VM size.

Network

Specifies the network existing on the Microsoft Azure account.

Subnet

Specifies the subnet existing on the Microsoft Azure account.

Auto-assign Public IP

Specifies the public IP assigned to the VM. .

Primary IP

Specifies the primary IP of VM. Specifies automatically, if you do not specify.

Select a Security Group

Specifies the security group. You can select multiple groups. if you do not select, automatically creates a new security group.

Advance: Change Host Name

Specifies the host name of new VM.

Advance: Recover data automatically after Instant VM is started

Specifies if data recovery happens automatically, after creating the Instant VM.

2. Click **Finish**.

   Now you can submit the job.

Submit the Instant Virtual Machine Job

To create the Instant VM, submit the Instant VM job. After the job is complete, you can see the Instant VM in **resources**, **Infrastructure**, **Instant Virtual Machine**.

Follow these steps:

1. Click **Create VM**.

   The **Create VM** dialog opens.

2. Select one of the following options:

   Boot Now

   Submits a job to create the Instant VM. After the VM is created, it automatically starts the VM.

   Boot Later

   Creates an Instant VM. You have to manually start the VM. You can start the VM after the Instant VM job is complete.

   Cancel

   Closes the Create VM dialog without creating any VM. You are returned to the Create VM page.

   The Instant VM job is successfully created.

Manage an Instant Virtual Machine

You can manage the Instant VM from Console. You can power-on or power-off an Instant VM from the Console. Also, you can delete any Instant VM.

Note: The Console displays only those Instant VM that are created from the recovery points managed from the selected Site.

- Start or Stop an Instant Virtual Machine
- Delete an Instant Virtual Machine

Start or Stop an Instant Virtual Machine

You can start or stop an Instant VM after it is created. The start or stop button displays depending on the status of the VM.

Follow these steps:

1. From the Console, click **resources**.
2. Navigate to **Infrastructures** and click **Instant Virtual Machines**.
3. Select the virtual machine from the center pane and click **Actions**.
4. Select **Power On** or **Power Off** depending on the status of the virtual machine.

   The virtual machine successfully starts or stops.

Delete an Instant Virtual Machine

You can delete any Instant VM that you do not require anymore.

Follow these steps:

1. From the Console, click **resources**.
2. Navigate to **Infrastructures** and click **Instant Virtual Machines**.
3. Select the virtual machine from the center pane and click **Actions**.
4. Click **Delete**.

   A confirmation dialog opens.
5. Click **OK**.

   The virtual machine is successfully deleted.

How to Create and Manage an Instant Virtual Machine on Hyper-V and VMware ESX Servers

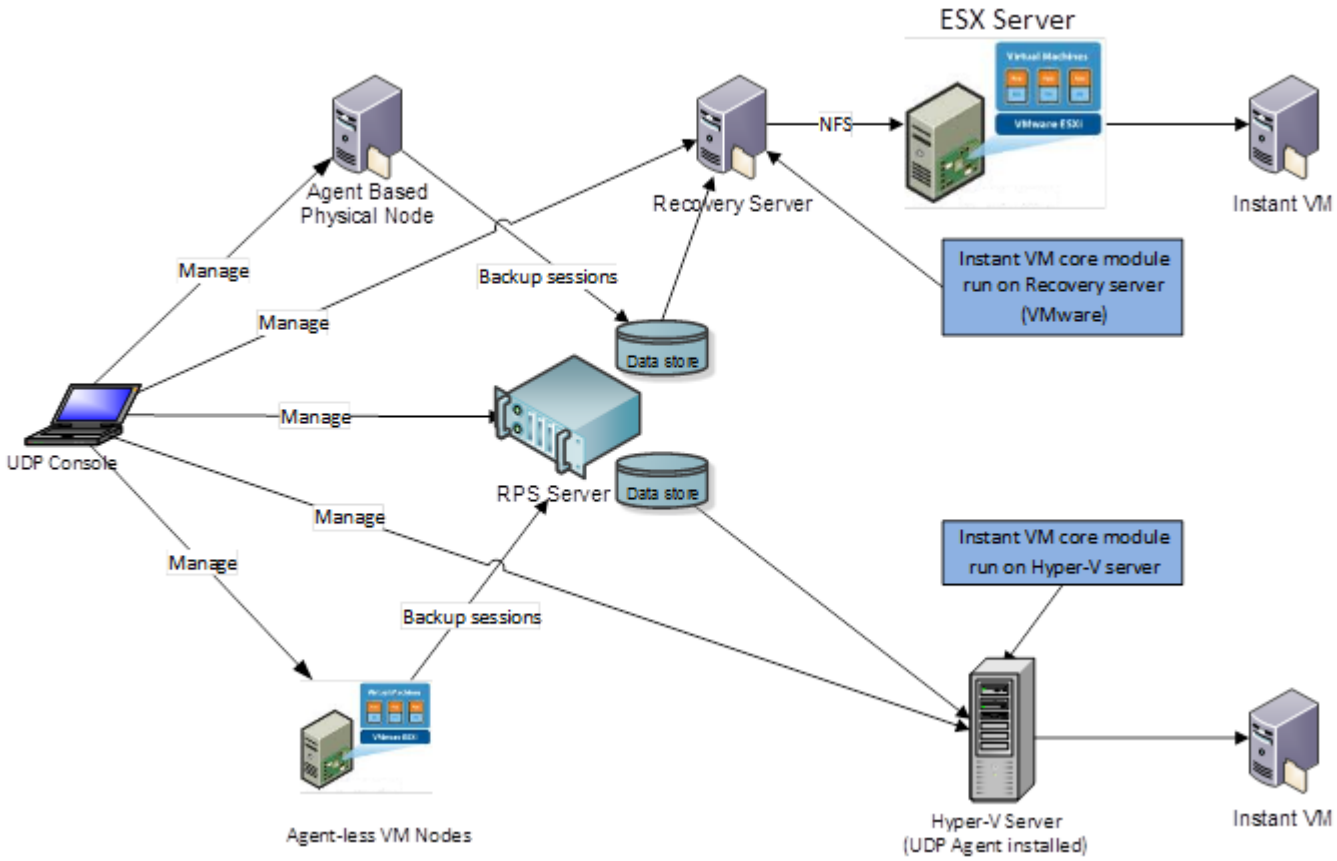
Instant virtual machine (Instant VM) creates a virtual machine in the hypervisor and runs the backup session inside the virtual machine without any prior conversion. The advantage of an Instant virtual machine is that it provides an immediate access to data and applications present in the Arcserve UDP backup sessions. Instant VM eliminates the downtime associated with a traditional restore or conversion of the backup session to a physical or virtual machine.

You can create an Instant VM from the following backup sessions:

- Agent-based Windows backup
- Agent-based Linux backup
- Host-based agentless backup

You can choose VMware vCenter/ESX(i) server or Windows Hyper-V server as the hypervisor.

The following diagram explains the architecture of an Instant VM:



What To Do Next?

- Review the Prerequisites for an Instant Virtual Machine
- Create an Instant Virtual Machine
- Manage an Instant Virtual Machine

Review the Prerequisites for an Instant Virtual Machine

Complete the following prerequisites before creating an Instant VM:

- Verify that you have at least one Arcserve UDP backup.
- Verify that the Arcserve UDP Agents are already installed on the Recovery Server.
- Verify that the NFS feature is installed on the Recovery Server, if the destination hypervisor is a VMware vCenter/ESX(i) server.
- Verify that the operating system of the Recovery Server is 64-bit Windows Server 2008 R2 or later.
- Verify that the Recovery Server has enough space for the Instant VM.
- Verify if you have the minimum permission to perform required Instant VM tasks. For more information, see Minimum Permission of VMware Required for IVM Tasks.

Note: Machine can boot up. NIC(s) are configured according to the user input on UI.

Consideration

- The Migrate InstantVM between nodes feature is not supported when an Instant VM is started in a Hyper-V cluster.
- When you create an Instant VM from a Linux agent backup, the settings of a virtual machine created by Instant VM cannot be modified using a vSphere client. You must use the vSphere web client to modify the virtual machine settings.
- If the number of NFS datastore reaches the maximum number of NFS mounts on an ESXi/ESX host and you create an Instant VM, then Arcserve UDP fails to create the NFS datastore. To increase the maximum number of NFS mounts on the ESXi/ESX host, see the VMware KB article.
- If recovery point is from agentless backup, and destination hypervisor is vsphere ESX/VC, the NFS server is required on the Linux Backup Server machine.
- If target hypervisor is MS Hyper-V, command net must be available on the Linux Backup Server. This command may be installed by samba client packages.
- The virtual machine does not work in the following situations:

- Recovery server reboots.
- Recovery server crashes.
- The network connection between recovery server and backup destination (Data store or share folder) disconnects.

Minimum Permission of VMware Required for IVM Tasks

The table displays list of minimum permission of VMware required to perform all the instant VM tasks.

Note: Global permissions are set at vCenter level.

Tasks	**Permission**
Datastore	Allocate space
Global	Disable methods
	Enable methods
	Licenses
Host>Configuration	Storage partition configuration
Network	Assign network
Resource	Assign virtual machine to resource pool
Virtual machine > Configuration	Add existing disk
	Advanced
Virtual machine > Interaction	Power off
	Power on
	Reset
	Console Interaction
Virtual machine > Inventory	Create new
	Remove
Virtual machine > Provisioning	Allow disk access
	Allow read-only disk access
	Allow virtual machine download
Virtual machine > Snapshot management	Create snapshot
	Remove snapshot
	Revert to snapshot
Virtual machine > Guest Operations	Guest Operation Queries

Create an Instant Virtual Machine

Creating an Instant VM involves the following five broad steps:

1. Open the Instant VM wizard
2. Select a Recovery Point
3. Specify the VM location
4. Specify the recovery server
5. Specify the Instant VM details
6. Submit the Instant VM job

On the successful completion of the job, an Instant VM is created.

Open the Instant Virtual Machine Wizard

You can configure and create an Instant VM from the Instant VM wizard. There are three ways to open the Instant VM wizard:

- From the Node Management
- From the Destination Management: Recovery Point Server
- From the Destination Management: Shared Folder

Open the Wizard from the Node management view

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes** and click **All Nodes**.
3. All the nodes are displayed on the center pane.
4. Right-click a node and select **Create an Instant VM**.

   The Instant VM wizard opens.

   Note: If a node is not associated with any plan, such node does not have the **Create an Instant VM** option.

Open the Wizard from the Destination management view

From Destinations: Recovery Point Server

1. Click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Recovery Point Servers**.

   If you have added any data stores, these data stores are displayed in the center pane.
3. Click the data store.

   If you have already backed up data to the RPS, all the source node are listed in the pane.
4. Right-click a node and select **Create an Instant VM**.

   The Instant VM wizard opens.

From Destinations: Shared Folder

1. Click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Shared Folders**.
3. If you have added any shared folders, these shared folders are displayed on the center pane.

4. Click a shared folder and select **Recovery Point Browser**.

   If you have already backed up data to the shared folder, all the source nodes are listed in the pane.

5. Right-click a node, and select **Create an Instant VM**.

   The Instant VM wizard opens.

Select a Recovery Point

The Select a Recovery Point page displays the location of the recovery point and lets you select a recovery point. The recovery point can be in a shared location or in a data store in RPS.

The Console automatically recognizes the location of the recovery point and pre-selects the **Location Type**, **Recovery Point Server**, and **Data Store** or **Network Share Folder** fields.

Expand the Date list, select the required recovery point from the list, and click **Next**.

The **VM Location** page opens.

Note: If you create an Instant Virtual Machine from the Node management view, then only those Recovery Point Servers are listed that are in the same Site as the source node. If you want to create an Instant Virtual Machine from the Recovery Point Server managed from a different Site, then navigate to that Site and open the Wizard from the Destination management view.

Select a VM Location

Specify the location of the virtual machine where you want to create the Instant VM. You can specify either VMware or a Microsoft Hyper-V virtual machine.

Follow these steps:

1. Select a Hypervisor Type.

   VMware vSphere

   a. Select **VMware vSphere**.

   b. If you have already added a VMware node to the selected Site in the Console, select the node from the **Host** drop-down list.

   c. If you have not added any VMware nodes, from the Host drop-down list, click **+Add New Hypervisor**.

      The Add Hypervisor dialog opens.

   d. Specify the VMware vCenter or ESX(i) server details and click **OK**.

      The Add Hypervisor dialog closes and you see the **VM Location** page again. All the ESX(i) or resource pools are displayed on the central pane.

   e. Select either the ESX(i), cluster, resource pool, virtual app as the location.

   The VMware vSphere machine is specified.

   Microsoft Hyper-V

   a. Select **Microsoft Hyper-V**.

   b. If you have already added a Hyper-V node to the selected Site in the Console, select the node from the **Hyper-V Server/Cluster** drop-down list.

   c. If you have not added any Hyper-V nodes, then click **Add**.

      The **Specify the VM Destination** dialog opens.

      Note: When you connect to the Hyper-V Instant VM using a local non-built-in administrator account, the remote UAC needs to be disabled. For more information on how to disable the remote UAC for non-built-in administrator, see How to disable a remote UAC for a non-built-in administrator.

   d. Specify the Hyper-V server details and click **OK**.

      The Hyper-V virtual machine is specified.

      Note: The Instant VM helper cannot install the Integration Service in Instant Virtual machines with Microsoft Hyper-V 2016, if the source node is of Windows 2008 or lower versions.

2. Click **Next**.

   The **Recovery Server** page opens.

How to Disable a Remote UAC for a non-built-in Administrator

Additional administrative account refers to those accounts that are not default administrators. Such accounts are also referred as non-built-in administrative accounts. To import virtual machine from a Hyper-V host, you can either use the built-in administrator account of the Hyper-V host, or a domain account which is in the local administrators group of the Hyper-V host, or a non-built-in administrative user.

The user with additional administrative account can use the procedures to disable UAC remote access.

Notes:

- This procedure is not similar to disabling UAC. Using this procedure, you can disable some of the functionalities of UAC.
- Considering that remote Windows Management Instrumentation (WMI) technology is used for import, ensure that WMI is not blocked by the firewall.

Follow these steps:

1. Click Start, type *regedit* in the Search programs and files field, and then press Enter.

   The Windows Registry Editor opens.

   Note: You may need to provide administrative credentials to open Windows Registry Editor.

2. Locate and click the following registry key:

   HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System

3. From the Edit menu, click New, and then click DWORD (32-bit) Value.
4. Specify *LocalAccountTokenFilterPolicy* as the name for the new entry and then press Enter.
5. Right-click *LocalAccountTokenFilterPolicy* and then click Modify.
6. Specify 1 in the Value data field and then click OK.
7. Exit the Registry Editor.

   For more information about the Windows behavior, see the Microsoft documentation.

Select a Recovery Server

The recovery server hosts the core module of the Instant VM. The default recovery server is the RPS. You can also assign a node as a recovery server.

For Linux backup sessions, the recovery server is a Linux Backup Server.

Follow these steps:

1. Select one of the following as recovery server:

   Use Recovery Point Server

   Specifies that the RPS is used as a recovery server.

   Select a Windows node from the node list

   Specifies that the recovery server is a Windows node. You can select the node from the list. Only those nodes are displayed that are managed by the selected Site.

   Notes:

   - You have to select a Recovery Server only when the hypervisor is VMware vSphere.
   - When the hypervisor is VMware vSphere, the Windows Network File System (NFS) role must be installed on the Recovery Point Server. It could be automatically installed by Instant VM process. To manually install the Network File System, see How to manually install Network File System on a Windows Server.
   - If Arcserve Backup is installed on the Recovery Server, the Windows Network File System (NFS) service may fail to start. It is because the default port number of the Windows NFS service is 111 and it is used by Arcserve Backup service **Remote Procedure Call Server**. To change the default port number of Arcserve Backup service **Remote Procedure Call Server**, see Modify the Ports Configuration File and Primary Server and Member Server Communication Ports in Arcserve Backup documentation

2. Click **Next**.

   The Instant VM Details page opens.

How to Manually Install Network File System on a Windows Server

You can manually install Network File System (NFS) on a Windows Server using the Server Manager.

Follow these steps:

1. Open the Server Manager and click Manage, Add Roles and Features Wizard.

   The Add Roles and Features Wizard opens.
2. Click Server Roles and select *File and Storage Services*.
3. Expand *File and iSCSI Services*.
4. Select *File Server* and *Server for NFS*.
5. Click Add Features to include the selected NFS features.
6. Click Install to install the NFS components on the server.

   You have manually installed Network File System on the Windows Server.

Specify the Instant Virtual Machine Details

Specify the details for the Instant VM. The options may vary depending on the hypervisor.

Follow these steps:

1. Specify the name and description of the Instant VM.

   VM Name

   Specifies the name of the Instant VM. The name of the source node with a prefix is the default name of the Instant VM. Some special characters are not allowed in the name, such as '@', '\' and so on.

   Description

   (Optional) Specifies the description for the Instant VM.

2. Specify the folder location of the Instant VM on the recovery server.

   You can browse the volume information of the recovery server.

   Note: If you select **VMware vSphere** as the VM location, the selected folder mounts as NFS Datastore to VMware. A shared icon appears on this folder in the local machine. By default, the NFS Datastore is mounted on the same folder location and you can only manually unmount it from VMware ESX(i) server.

3. Specify the Instant VM settings.

   CPU Count

   Specifies the number of CPU that you would require in the Instant VM.

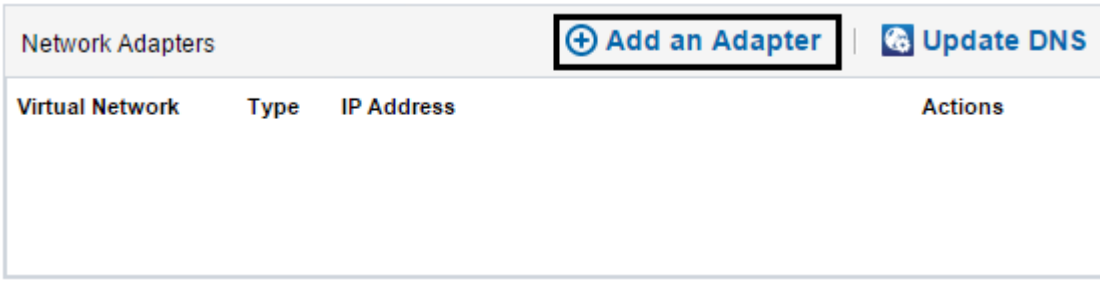
   Memory Size

   Specifies the size of memory that you would require in the Instant VM.

4. Add network adapters.

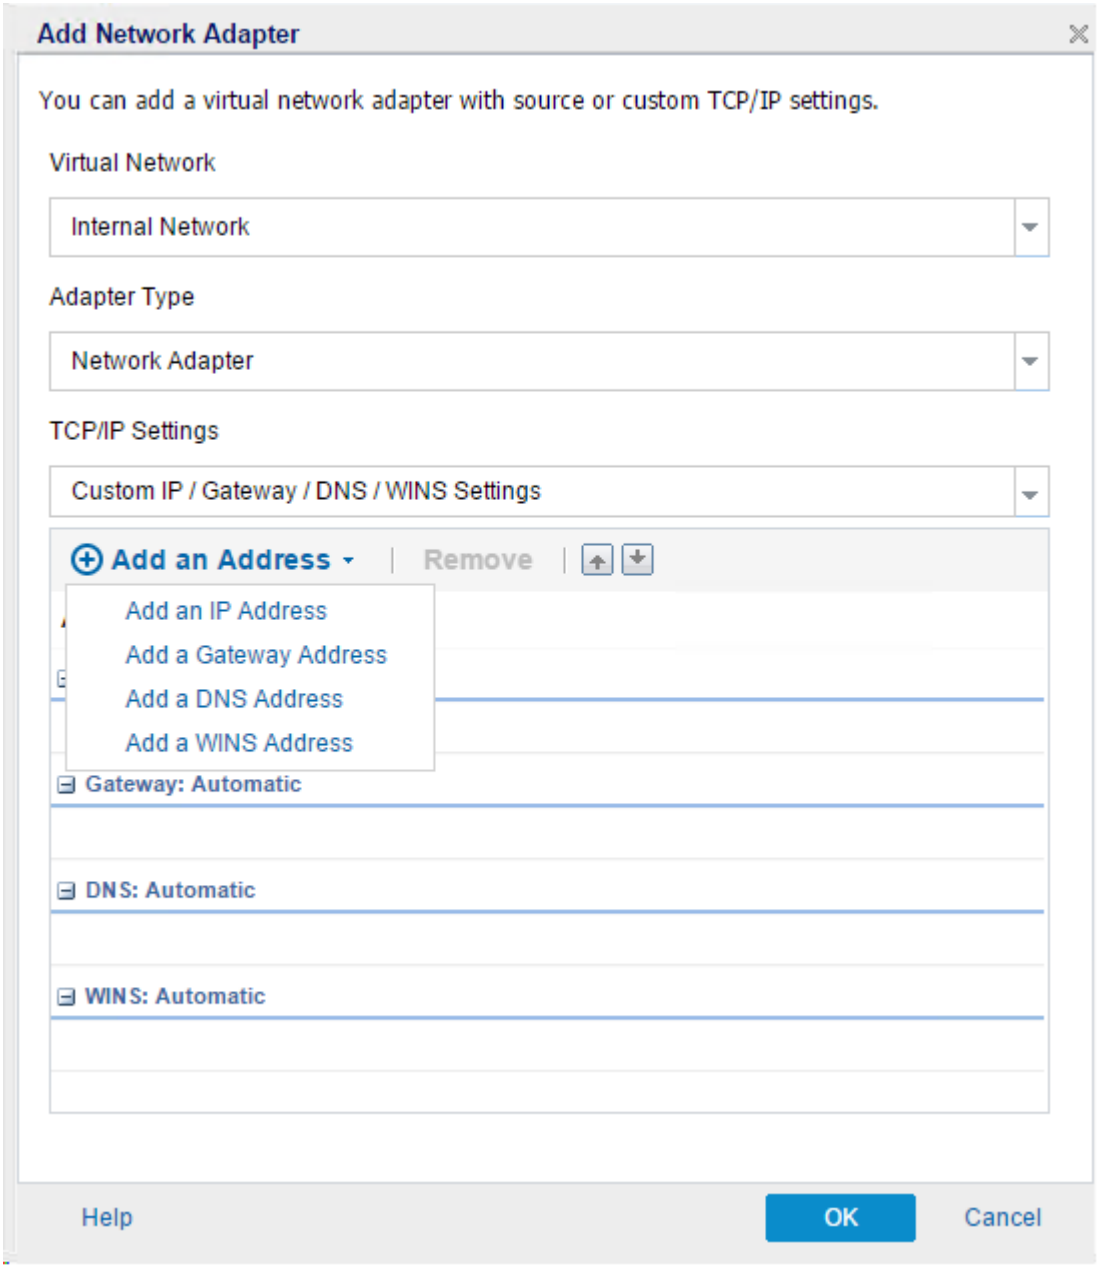
   Note: For a Linux Instant VM, at least one virtual NIC can be used to connect to Linux Backup Server.

   a. Click the **Add an Adapter** button to add an adapter and specify the details of the network.



You can add multiple network adapters. After adding a network adapter, you can edit and delete the network adapter from the **Actions** column.

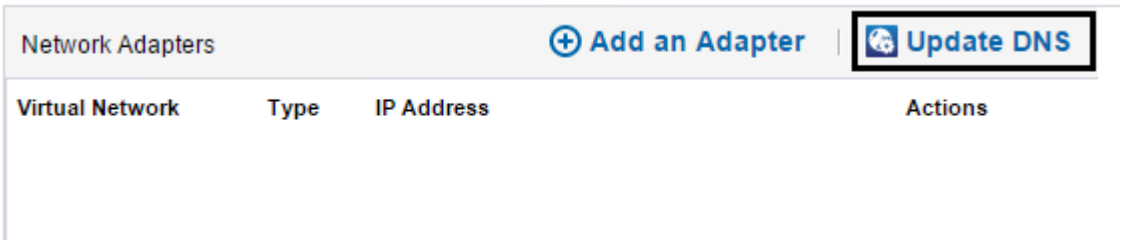
b. Specify the **Virtual Network**, **Adapter Type**, and **TCP/IP Settings**. If you want to specify the IP address of the Instant VM, you can click **Add an Address** and select the address you want to configure.



5. Update DNS.

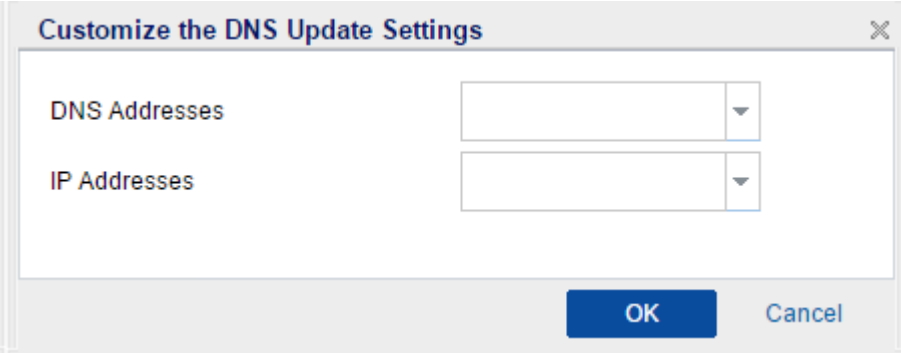
   Note: This feature is only available for Windows Instant VM.

   a. Click the **Update DNS** button to specify DNS details.



Note: You can specify the Update DNS detail information if you have specified an IP address and DNS address in the Network Adapter and the source machine is in a domain.

b. Click **Add a DNS Address** to add a DNS update record. Click **Remove** to remove DNS update record. Click Up and Down button to adjust the sequence of the records.



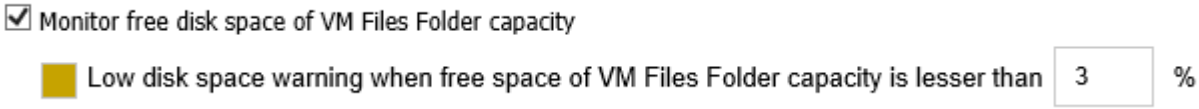
c. Select a DNS Address and an IP Address from the drop-down list and then click **OK**.

d. Specify the **Time to Live** (TTL).

e. Specify the DNS Authentication.

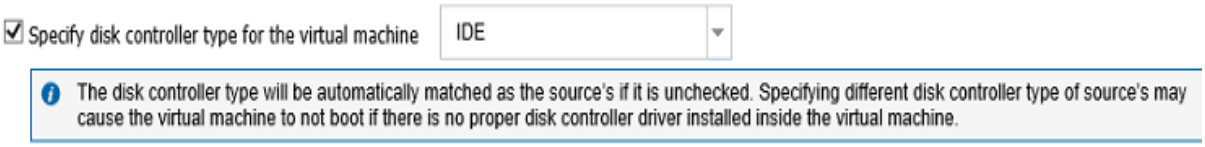
For Microsoft DNS Server, enter the username and password. For Bind Server, you need to specify the full path, including the file name of the key file in the Recovery Server.

6. Verify the free disk space of VM Files Folder capacity.



Note: The **Monitor free disk space of VM Files Folder capacity** check box is selected by default. The capacity bar appears in yellow in Instant Virtual Machines page, if the free space of VM Files Folder capacity is lesser than the threshold value. The default threshold value is 3%. You may change the value, if required.

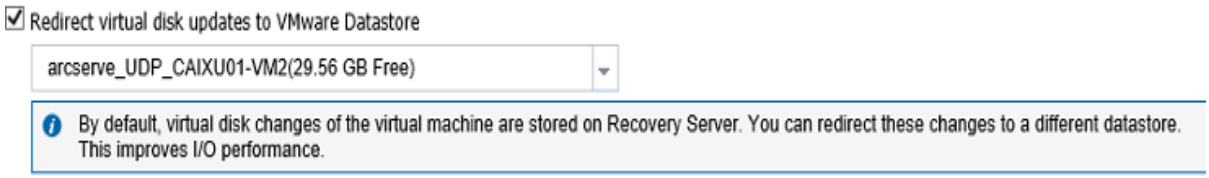
7. Select the **Specify disk controller type for the virtual machine** check box and select the type of the disk controller for the virtual machine from the drop-down.



Note: The **Specify disk controller type for the virtual machine** check box is available if you have selected **VMware vSphere** as the VM location.

The instant virtual machine is created applying the specified disk controller in the VMware.

8. To redirect the virtual disk updates to the VMware Datastore, follow these steps:
   a. Select the **Redirect virtual disk updates to VMware Datastore** check box.

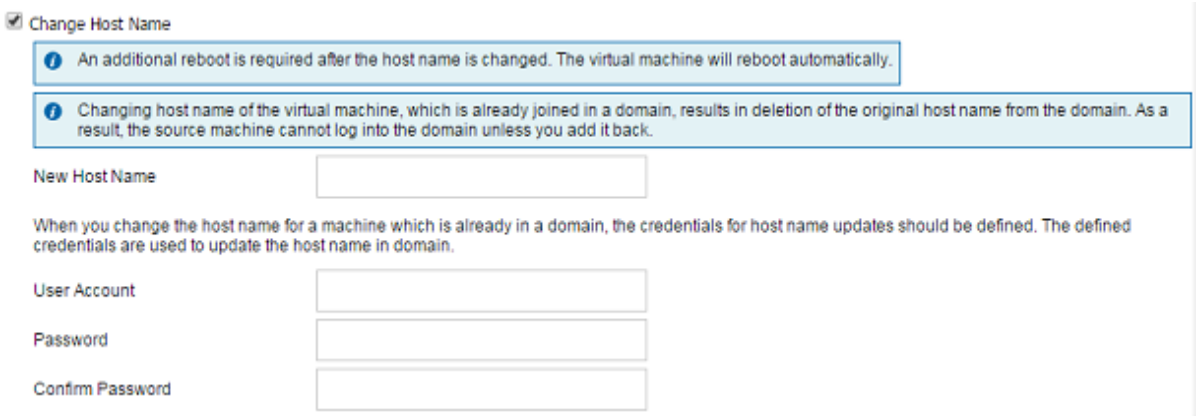


Note: The **Redirect virtual disk updates to VMware Datastore** check box is available if you have selected **VMware vSphere** as VM location.

   b. Select the required VMware Datastore from the drop-down.

The updates of the virtual disk are redirected to the selected VMware datastore.

9. To change the Instant Virtual Machine hostname, follow these steps:
   a. Click the **Change Host Name** check box to update the hostname of the Instant Virtual Machine.



   b. Specify a New Host Name for the Instant Virtual Machine. If the source machine is in a domain, provide the User Account and password.

Note: If the source machine is in a domain, the account should have the permission to change the hostname in domain.

10. (Optional) For **Linux Instant VM**, select the **Recover data automatically after Instant VM is started** option to enable the automatic recovery of data when the instant VM target is started.

    The default behavior of **Linux Instant VM** is to recover the necessary data first and start the VM. If the option is not selected, then when the VM starts the remaining data is not recovered even if used a normal VM. If the option is enabled, the remaining data is recovered at the backend when you are using the VM. You can also preserve the Instant VM target permanently when the data recovery is finished.

    Note: When status of the Linux Instant VM target is Power Off, Instant VM job fails. If that recovery point is merged, the Linux Instant VM fails to Power ON.

    Now, you can submit the job.

Submit the Instant Virtual Machine Job

To create the Instant VM, submit the Instant VM job. After the job is complete, you can see the Instant VM in **resources**, **Infrastructure**, **Instant Virtual Machine**.

Follow these steps:

1. Click **Finish**.

   The **Boot VM** dialog opens.

2. Select one of the following options:

   Boot Now

   Submits a job to create the Instant VM. After the VM is created, it automatically starts the VM.

   Boot Later

   Creates an Instant VM. You have to manually start the VM. You can start the VM after the Instant VM job is complete.

   Cancel

   Closes the Create VM dialog without creating any VM. You are returned to the Create VM page.

   The Instant VM job is successfully created.

Manage an Instant Virtual Machine

You can manage the Instant VM from Console. You can power-on or power-off an Instant VM from the Console. Also, you can delete any Instant VM.

Note: The Console displays only those Instant VM that are created from the recovery points managed from the selected Site.

- Start or Stop an Instant Virtual Machine
- Restart an Instant Virtual Machine
- Delete an Instant Virtual Machine
- Convert the Linux Instant Virtual Machine to an Independent Virtual Machine
- Migrate the Linux Instant Virtual Machine to a Physical Machine

Start or Stop an Instant Virtual Machine

You can start or stop an Instant VM after it is created. The start or stop button displays depending on the status of the VM.

Follow these steps:

1. From the Console, click **resources**.
2. Navigate to **Infrastructures** and click **Instant Virtual Machines**.
3. Select the virtual machine from the center pane and click **Actions**.
4. Select **Power On** or **Power Off** depending on the status of the virtual machine.

   The virtual machine successfully starts or stops.

Restart an Instant Virtual Machine

You can restart an Instant VM after you create.

Note: You can restart an Instant VM only if it is in a **Failed/Job Crash** status.

Follow these steps:

1. From the Console, click **resources**.
2. Navigate to **Infrastructure** and click **Instant Virtual Machines**.
3. Select the virtual machine in **Failed/Job Crash** status from the center pane, and click **Actions**.
4. Click **Restart**.

   The virtual machine restarts.

Delete an Instant Virtual Machine

You can delete any Instant VM that you do not require anymore.

Follow these steps:

1. From the Console, click **resources**.
2. Navigate to **Infrastructures** and click **Instant Virtual Machines**.
3. Select the virtual machine from the center pane and click **Actions**.
4. Click **Delete**.

   A confirmation dialog opens.
5. Click **OK**.

   The virtual machine is successfully deleted.

Convert the Linux Instant Virtual Machine to an Independent Virtual Machine

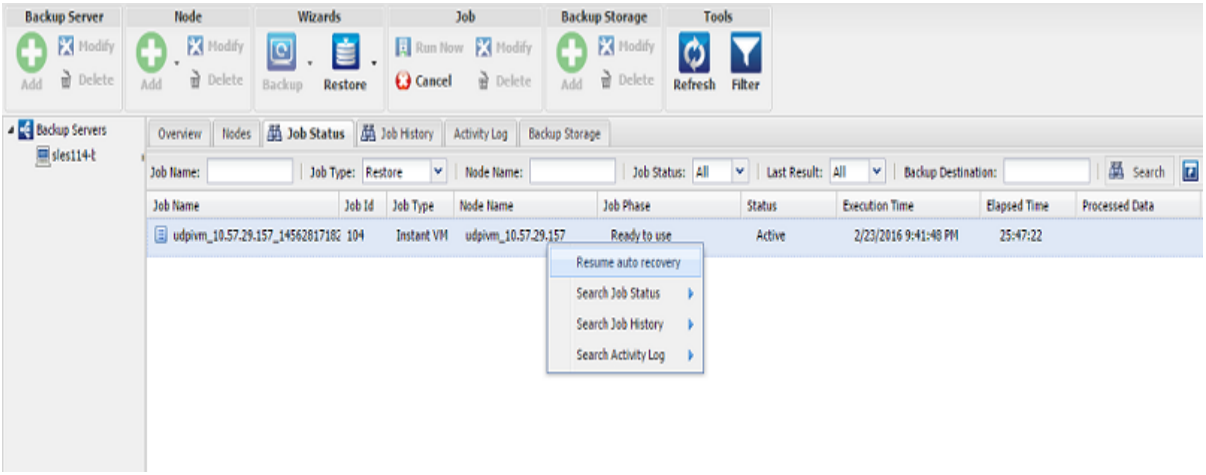
You can change a Linux Instant Virtual machine (IVM) to an Independent virtual machine (VM).

Notes:

- To continue you need a menu item that is available only after the IVM has run into “Ready to use” job phase.
- The menu item is unavailable if the recovery point is agentless backup, and the target hypervisor is vSphere ESX/VC. You need to use VMware Storage vMotion to convert the IVM into an independent VM.

Follow these steps:

1. Open Linux Backup Server UI.



2. From the Job Status tab, select the IVM job, right click for context menu.
3. Select Resume auto recovery.

   The IVM job is moved to Job History after the process finishes successfully.

Migrate the Linux Instant Virtual Machine to a Physical Machine

To migrate the Linux Instant Virtual Machine to a physical machine, refer to the **How to Perform a Migration BMR for Linux Machines** section in the *Agent for Linux User Guide*.

How to Create and Manage an Instant Virtual Machine on Amazon EC2

Instant virtual machine (Instant VM) supports creating a virtual machine on Amazon EC2 instantly The IVM on Amazon EC2 has following advantages:

- Provides immediate access to data and applications present in the Arcserve UDP backup sessions.
- Eliminates the downtime associated with a traditional restore or conversion of the backup session to a virtual machine.
- Provides an alternative to create virtual machine on cloud, rather than local.

You can create an Instant VM from the following backup sessions:

- Agent-based Linux backup
- Host-based Agentless backup

What To Do Next?

- Review the Prerequisites
- Create an Instant Virtual Machine Plan on Amazon EC2
- Manage an Instant Virtual Machine Plan on Amazon EC2

You can create an Instant VM from the following backup sessions: On the successful completion of the job, an Instant VM is created on Amazon EC2.

Review the Prerequisites for an Instant Virtual Machine on Amazon EC2

Complete the following prerequisites before creating an Instant VM:

- From the Compatibility Matrix, verify if the VM is supported by Amazon EC2 and UDP.
- Verify that you have at least one Recovery Point Server on local machine, as backup destination.
- Verify that you have at least one Recovery Point Server on Amazon EC2, as replication destination.
- Verify that you have at least one Linux backup server locally, for backup job.
- Verify that you have at least one Linux backup server on Amazon EC2, for the instant VM job.
- Verify that the Amazon EC2 account has enough running instance quota for the Instant VM.

Limitation

The Agent-based Windows and Host-based Agentless Windows Virtual Machine backup are NOT supported.

Create an Instant Virtual Machine on Amazon EC2

Creating an Instant VM involves the following five broad steps:

1. Open the Instant VM wizard
2. Select the recovery point
3. Select the VM location
4. Select the recovery server
5. Specify the Instant VM details
6. Submit the Instant VM job

Open the Instant Virtual Machine Wizard

You can configure and create an Instant VM from the Instant VM wizard. There are three ways to open the Instant VM wizard:

- From the Node Management
- From the Destination Management: Recovery Point Server
- From the Destination Management: Shared Folder

Open the Wizard from the Node management view

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes** and click **All Nodes**.
3. All the nodes are displayed on the center pane.
4. Right-click a node and select **Create an Instant VM**.

   The Instant VM wizard opens.

   Note: If a node is not associated with any plan, such node does not have the **Create an Instant VM** option.

Open the Wizard from the Destination management view

From Destinations: Recovery Point Server

1. Click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Recovery Point Servers**.

   If you have added any data stores, these data stores are displayed in the center pane.
3. Click the data store.

   If you have already backed up data to the RPS, all the source node are listed in the pane.
4. Right-click a node and select **Create an Instant VM**.

   The Instant VM wizard opens.

From Destinations: Shared Folder

1. Click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Shared Folders**.
3. If you have added any shared folders, these shared folders are displayed on the center pane.

4. Click a shared folder and select **Recovery Point Browser**.

   If you have already backed up data to the shared folder, all the source nodes are listed in the pane.

5. Right-click a node, and select **Create an Instant VM**.

   The Instant VM wizard opens.

Select a Recovery Point

The Select a Recovery Point page displays the location of the recovery point and lets you select a recovery point. The recovery point can be in a shared location or in a data store in RPS.

The Console automatically recognizes the location of the recovery point and pre-selects the **Location Type**, **Recovery Point Server**, and **Data Store** or **Network Share Folder** fields.

Note: Select Recovery Point Server (Shared folder) session as replication destination to save network bandwidth and time.

Expand the Date list, select the required recovery point from the list, and click **Next**.

The **VM Location** page opens.

Note: If you create an Instant Virtual Machine from the Node management view, then only those Recovery Point Servers are listed that are in the same Site as the source node. If you want to create an Instant Virtual Machine from the Recovery Point Server managed from a different Site, then navigate to that Site and open the Wizard from the Destination management view.

Select a VM Location

Specify the location of the virtual machine where you want to create the Instant VM. You can specify either Amazon EC2 or Amazon EC2 (China) according to your account type.

Follow these steps:

1. Select Amazon EC2 or Amazon EC2 (China).
2. Select an Account name.

   Note: Select corresponding account, if not added, refer to Add a Cloud Account.
3. Click **Next**.

   The **Recovery Server** page opens.

Select a Recovery Server

The recovery server hosts the core module of the Instant VM. The default recovery server is the RPS. You can also assign a node as a recovery server.

For Linux backup sessions, the recovery server is a Linux Backup Server.

Follow these steps:

1. Select a Linux Backup Server from the node list.

   Notes: Select Recovery Point Server (Shared folder) session as replication destination to save network bandwidth and time.

2. Click **Next**.

   The Instant VM Details page opens.

Specify the Instant Virtual Machine Details

Specify the details for the Instant VM.

Follow these steps:

1. Specify the name and description of the Instant VM.

VM Name

Specifies the name of the Instant VM. The name of the source node prefix with “UDPIVM_” becomes the default name of the Instant VM. You can modify the name. Some special characters, such as ‘@’, \ and so on, are not allowed in the name.

Description

(Optional) Specifies the description for the Instant VM.

Region

Specifies the region of your Amazon EC2 account..

Instance Type

Specifies the supported instance type.

Volume Type

Specifies the type of volume.

Network

Specifies the network existing on the Amazon EC2 account.

Subnet

Specifies the subnet existing on the Amazon EC2 account.

Auto-assign Public IP

Specifies the public IP assigned to the VM. .

Primary IP

Specifies the primary IP of VM. Specifies automatically, if you do not specify.

Select a Security Group

Specifies the security group. You can select multiple groups. if you do not select, automatically creates a new security group.

Advance: Change Host Name

Specifies the host name of new VM.

Advance: Recover data automatically after Instant VM is started

Specifies if data recovery happens automatically, after creating the Instant VM.

2. Click **Finish**.

   Now you can submit the job.

Submit the Instant Virtual Machine Job

To create the Instant VM, submit the Instant VM job. After the job is complete, you can see the Instant VM in **resources**, **Infrastructure**, **Instant Virtual Machine**.

Follow these steps:

1. Click **Create VM**.

   The **Create VM** dialog opens.

2. Select one of the following options:

   Boot Now

   Submits a job to create the Instant VM. After the VM is created, it automatically starts the VM.

   Boot Later

   Creates an Instant VM. You have to manually start the VM. You can start the VM after the Instant VM job is complete.

   Cancel

   Closes the Create VM dialog without creating any VM. You are returned to the Create VM page.

   The Instant VM job is successfully created.

Manage an Instant Virtual Machine

You can manage the Instant VM from Console. You can power-on or power-off an Instant VM from the Console. Also, you can delete any Instant VM.

Note: The Console displays only those Instant VM that are created from the recovery points managed from the selected Site.

- Start or Stop an Instant Virtual Machine
- Restart an Instant Virtual Machine
- Delete an Instant Virtual Machine
- Convert the Linux Instant Virtual Machine to an Independent Virtual Machine
- Migrate the Linux Instant Virtual Machine to a Physical Machine

Start or Stop an Instant Virtual Machine

You can start or stop an Instant VM after it is created. The start or stop button displays depending on the status of the VM.

Follow these steps:

1. From the Console, click **resources**.
2. Navigate to **Infrastructures** and click **Instant Virtual Machines**.
3. Select the virtual machine from the center pane and click **Actions**.
4. Select **Power On** or **Power Off** depending on the status of the virtual machine.

   The virtual machine successfully starts or stops.

Restart an Instant Virtual Machine

You can restart an Instant VM after you create.

Note: You can restart an Instant VM only if it is in a **Failed/Job Crash** status.

Follow these steps:

1. From the Console, click **resources**.
2. Navigate to **Infrastructure** and click **Instant Virtual Machines**.
3. Select the virtual machine in **Failed/Job Crash** status from the center pane, and click **Actions**.
4. Click **Restart**.

   The virtual machine restarts.

Delete an Instant Virtual Machine

You can delete any Instant VM that you do not require anymore.

Follow these steps:

1. From the Console, click **resources**.
2. Navigate to **Infrastructures** and click **Instant Virtual Machines**.
3. Select the virtual machine from the center pane and click **Actions**.
4. Click **Delete**.

   A confirmation dialog opens.
5. Click **OK**.

   The virtual machine is successfully deleted.

Convert the Linux Instant Virtual Machine to an Independent Virtual Machine

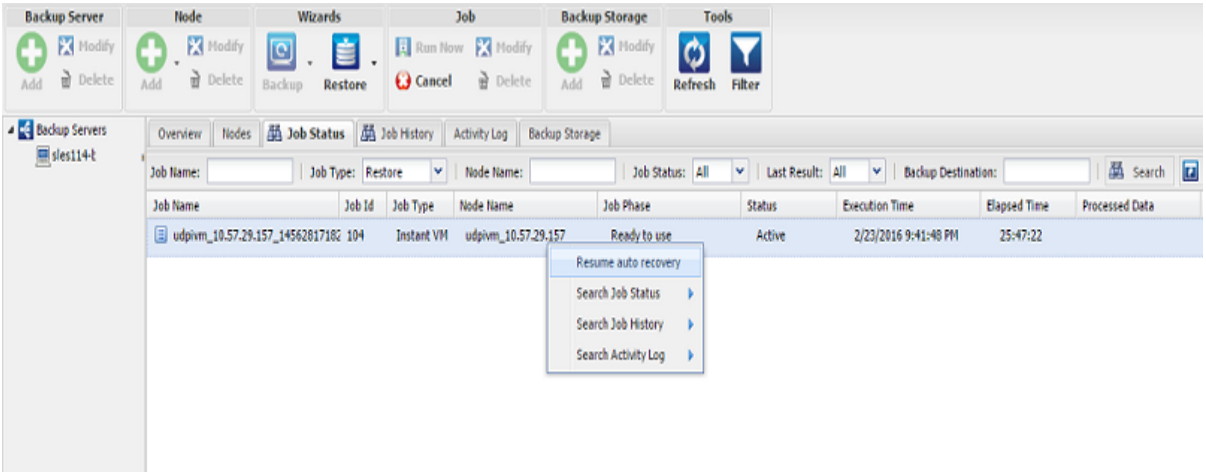
You can change a Linux Instant Virtual machine (IVM) to an Independent virtual machine (VM).

Notes:

- To continue you need a menu item that is available only after the IVM has run into “Ready to use” job phase.
- The menu item is unavailable if the recovery point is agentless backup, and the target hypervisor is vSphere ESX/VC. You need to use VMware Storage vMotion to convert the IVM into an independent VM.

Follow these steps:

1. Open Linux Backup Server UI.



2. From the Job Status tab, select the IVM job, right click for context menu.
3. Select Resume auto recovery.

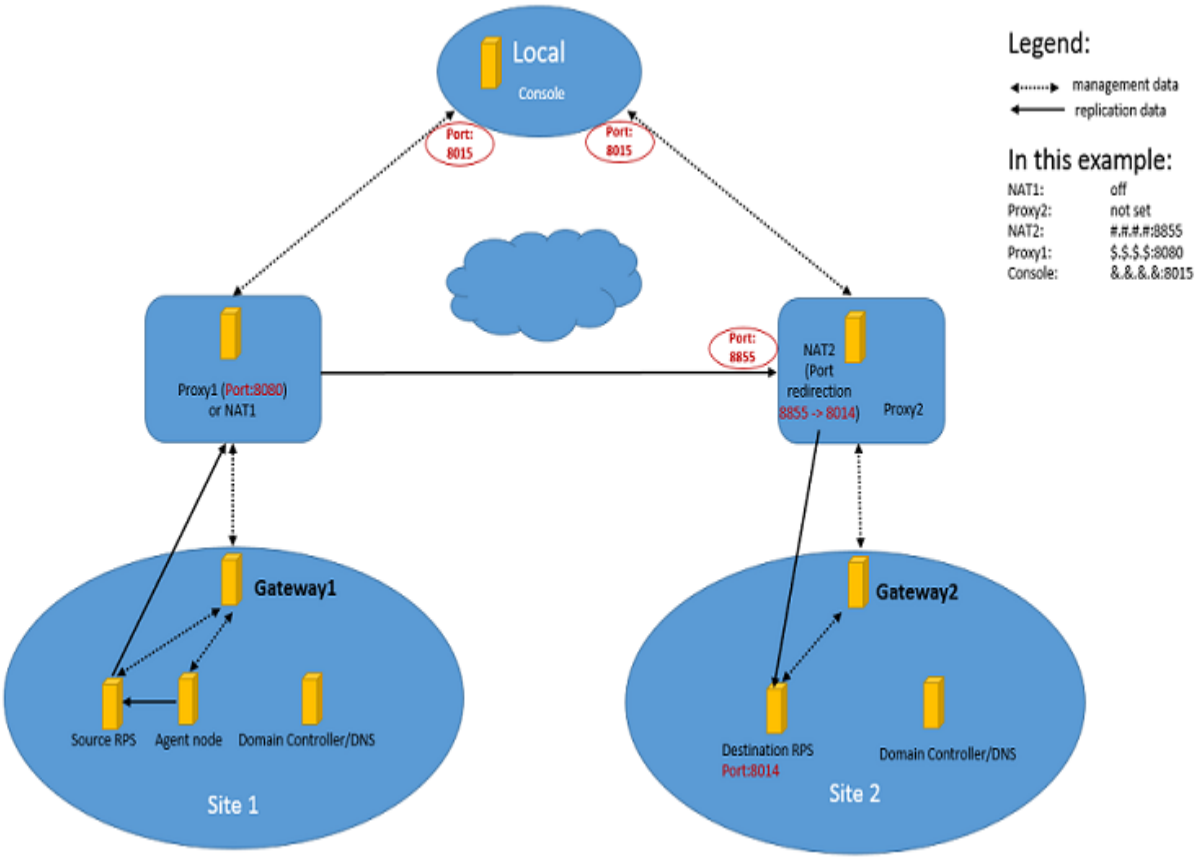
   The IVM job is moved to Job History after the process finishes successfully.

Migrate the Linux Instant Virtual Machine from Amazon EC2 to a Physical Machine

To migrate the Linux Instant Virtual Machine to a physical machine, refer to the **How to Perform a Migration BMR for Linux Machines from Amazon EC2 to local** section in the *Agent for Linux User Guide*.

How to Create a Plan for Replication Across Sites

The replication across site feature allows you to use the same Console to replicate data between different sites. The following diagram illustrates the connection among Console, Site 1 and Site 2:



Console (For example, in public network), Site 1 (For example, in a private network) and Site 2 (For example, in another private network) can be in different network segments. Site 1 and Site 2 are managed by console through the gateway.

The diagram gives an example to illustrate this feature of UDP, which is just a reference because the real environment that you are using may not be similar to the one described here. The example aims to replicate data from Site 1 source RPS to Site 2 destination RPS. Gateway in Site 1 and Site 2 can connect to the Console through the proxy or NAT. The proxy and NAT should be properly configured.

Description of diagram:

- Proxy1 or NAT1 has both public interface and private interface. For example: IP-1 as public IP is in the same segment with console and IP-2 as private IP is in the same segment with Site1. In the network setting of Gateway1 machine, default

gateway is set to private IP of Proxy1 or NAT1.

- Proxy2 and NAT2 also have the same setting with Proxy1/NAT1, besides that NAT2 needs extra configuration of a port redirection rule so that from the public network one can access the private service by this mapping. In this example, that is the source RPS through which NAT2 port redirection can connect to destination RPS.
- With windows server 2012R2 OS, both windows GUI and command could set port direction for NAT. Here is an example command, which means: input this address from source RPS web explorer https or http://<NAT2 IP-1>:<port number=8855>, and it will redirect to destination RPS https or http://<destination RPS IP=192.168.30.102>:8014, then you can execute the replication cross site.

netsh interface portproxy add v4tov4 listenport=8855 connectaddress="192.168.30.102" connectport=8014 protocol=tcp

Create a Plan for Replication Across Sites

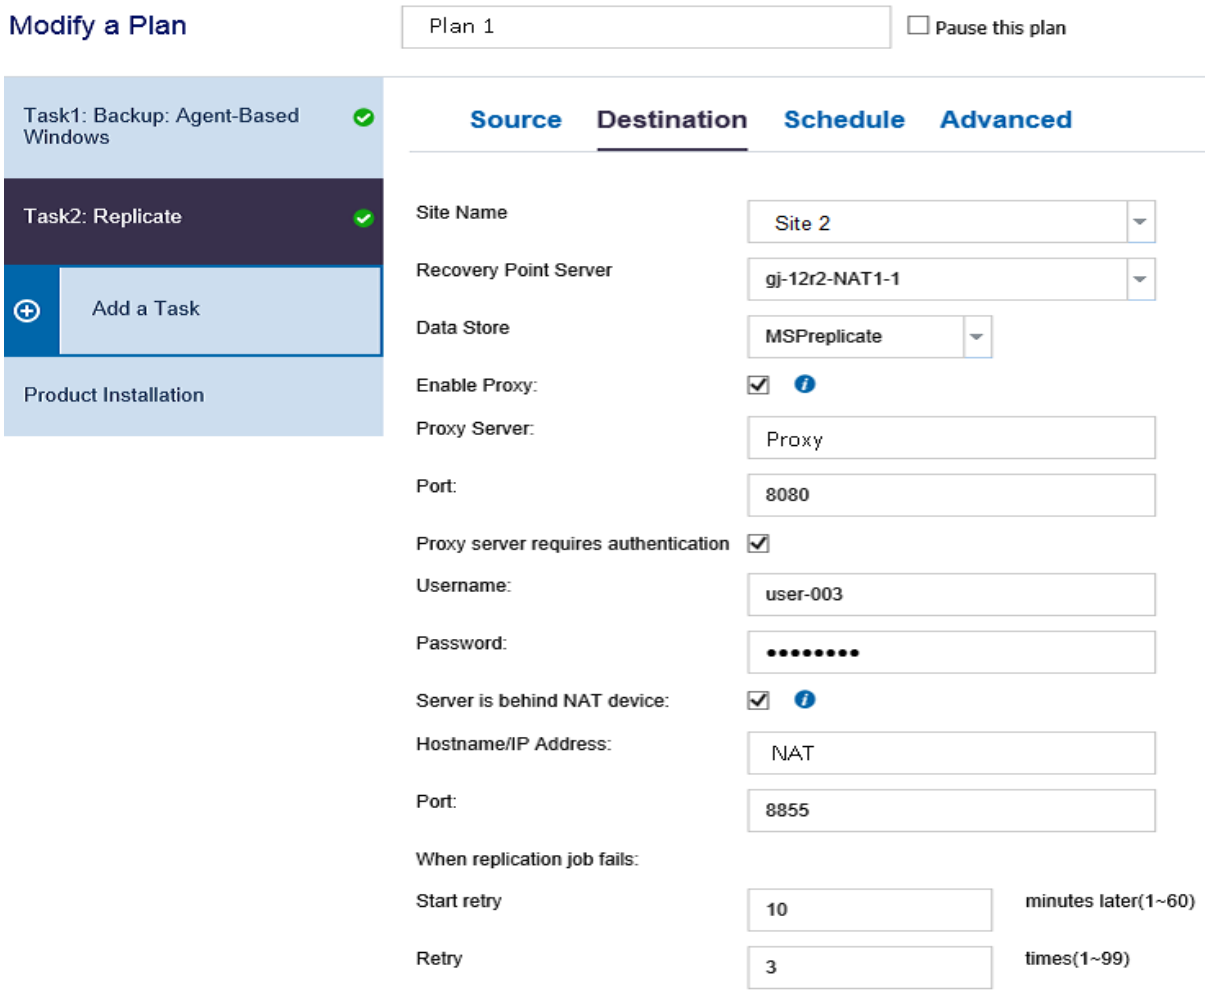
You can create a plan to replicate across sites.

Follow these steps:

1. In Site1, which is the source RPS, create one backup task as Task1.
2. Add a **Replicate** task to the same plan.
3. In the **Destination** tab, select the other site (Site 2 in the example) where the destination RPS or data store is located.
4. (Optional) Enable the Proxy details (server, port and authentication).

   Note: Before you enable the proxy details, you must configure the proxy server between Site 1 and the Console.
5. (Optional) Enable the NAT details.

   Note: Before you enable the NAT details, you must configure the NAT server and port redirection between the Console and Site 2.



6. Configure the other tabs such as **Schedule** and **Advanced** and save the plan.
7. When the job runs, verify the monitor and log status for backup, replication (out) and replication (in).

How to Protect Microsoft Office 365 Data

This section contains following topics:

Modern Authentication

This section provides the following information about Modern Authentication for Office 365 backups.

Review the Prerequisites and Considerations

Review the following prerequisites:

- Install Microsoft .NET Framework (version 4.7.1 or higher) and PowerShell (version 5.1 or higher) in the proxy server (64-bit computer).
- Before adding a node, you need to install Azure cmdlet in the proxy machine using the following relevant PowerShell CMD line: *Install-Module -Name AzureAD -Force*
- Before adding a node, you need to install the Exchange Online module in the proxy machine using the following relevant PowerShell CMD line: *Install-Module -Name ExchangeOnlineManagement -Force*

Note: To save the installation time, use the Powershell script, which installs the following modules automatically:

- Azure AD
- ExchnageOnlineManagement
- PackageManagement
- PowerShellGet

To download the Powershell script, click Script2InstallPrerequisites.ps1_. After downloading the Powershell script, do the following:

- Rename “Script2InstallPrerequisites.ps1_” to "Script2InstallPrerequisites.ps1".
- Run the Script2InstallPrerequisites.ps1 from the PowerShell window using the Run as Administrator option.

How to Use Self-Signed Certificate

You can use the self-signed certificate in one of the following ways:

Method 1: Generate and Download a new self signed certificate

To authenticate using a new self-signed certificate, follow these steps:

1. Under Step 1, to allow UDP to generate a new self signed certificate, select **Generate and download a new self signed certificate**.
2. Type and retype the certificate password as needed.
3. To download the certificate, click one of the following options:
   - **Public certificate** - Download the certificate to upload it on the Azure portal.
   - **Private certificate** - Download the certificate to save it on your computer for future use.

     Note: After you download the public certificate, the download option for the private certificate gets enabled.
4. To upload the certificate in Azure portal, see How to Upload Certificates in Azure.

Method 2: Use existing certificate

To authenticate using an existing certificate, follow these steps:

Notes:

- When this option is selected, the user is expected to have their own Self signed Certificate (.cer file) and Private key certificate files (.PFX file with password).
- If you are using Internet Explorer (IE 11), it blocks to proceed further when you select the Use existing certificate option. To resolve this, see Configure Security Settings of Internet Explorer (IE 11).

1. To upload and set up the application in the Azure portal, see How to Upload Certificates in Azure.

   Note: If the same certificate is used for multiple plan creations, it is enough to upload the certificate only once in the Azure portal.
2. Follow these steps from the Console window:
   a. Browse for the Private Key certificate [.PFX file] and enter the password.
   b. Click **Next** to proceed for Node registration.

How to Upload Certificates in Azure

This section provides information about how to upload certificates in Azure portal.

Follow these steps:

1. Log into the Azure portal.
2. To navigate to the App registrations page, do one of the following:
   - From the left pane, select **Azure Active Directory**, and then click **App registrations**.
   - Search for App registrations in the search bar, and then select **App registrations**.
3. On the Arcserve- App registrations page, click **New registration**.
4. On the Register an application page, do the following, and then click **Register**:
   - Name - Enter a name for the application.
   - Supported account types - Select the **Accounts in this organizational directory only (Arcserve only – Single tenant)** option.
5. Copy and save the **Application ID** for future use.
6. Navigate to **Certificates & Secrets**, click **Upload certificate**, browse your existing certificate [.cer file], and then click **Add**.
7. On the left pane, select **API permissions**, and then click **Grant Admin Consent for <tenant name>**.

   A confirmation message appears.
8. Click **Yes**.

   The admin consent is added successfully for the required permissions.
9. Assign the Global/Company Administrator role to the AzureAD App. For more information, see How to Assign Global/Company Administrator role to the AzureAD App.
10. Log out of the Azure portal.

How to Assign Global/Company Administrator Role to the AzureAD App

This section provides information about how to assign Global/Company Administrator role to the AzureAD app using the PowerShell cmdlets.

Follow these steps:

1. To install the MSOnline module, run the following command from the PowerShell Window using the Run as administrator option:

   Install-Module MSOnline

2. To initiate a connection with Azure Active Directory (AzureAD), run the following command:

   Connect-MSolService

3. Get the App ID using the following command:

   $sp = Get-MsolServicePrincipal -AppPrincipalID <APP ID GUID>

   Note: APP ID GUID is the application id of the UDP application.

4. Get the Global/Company Administrator role using the following command:

   $role = Get-MsolRole -RoleName "Company Administrator"

5. To add the Global/Company Administrator role to the AzureAD app, use the following command:

   Add-MsolRoleMember -RoleObjectId $role.ObjectId -RoleMemberType ServicePrincipal -RoleMemberObjectId $Sp.ObjectId

The Global/Company Administrator role is assigned to the AzureAD app.

How to Create an Exchange Online Backup Plan

Exchange Online is an email application hosted on Microsoft cloud. To protect your Exchange Online mail items (Mails, Calendar items, Contacts, and so on) from Microsoft cloud, you need to create a plan. The plan for Exchange Online consists of a backup task. This backup task lets you specify the Exchange Online nodes you want to protect, the backup destination, and the backup schedule.

What To Do Next?

1. Review the Prerequisites and Considerations
2. Create an Exchange Online Backup Plan
3. (Optional) Perform a Manual Backup
4. Configuration for Multi-Factor Authentication

Review the Prerequisites and Considerations

Prerequisites:

For Backup account:

- Use a backup service account with Global admin permissions.
- Add impersonation permission for backup user to the Exchange Online backup account to connect the Exchange Online organization, perform backups and restore.

  Note: If you do not add the backup account to the Discovery Management role group and do not assign the Application Impersonation permission, backup fails.

- Associate the Backup user account with one exchange online mailbox.

For Backup proxy:

- Log in to the Console.
- Install Arcserve UDP Agent on the proxy machine where you want to run backup/restore.
- Review the Prerequisites before you create a plan using Modern Authentication.

Considerations:

- While backing up a large set of users, you can use exchange groups for better backup and optimisation.
- Default setting for exchange online backup uses 4 backup threads (one thread per user) at a time. You can also modify the threads in configuration file of *Engine\BIN\Office365\Arcserve.Office365.Exchange.config* as follows:

  <!--#region for multi thread-->

  <!--MultiThreadEnable default value:0. if enable, set 1.-->

  <add key="MultiThreadEnable" value ="1"/>

  <!--set how many thread will be used to backup mailbox. default value is 4-->

  <add key="MaxDegreeOfParallelismForMailbox" value="4"/>

  <!--#endregion-->

  We recommend to set the value from 1-5. The maximum possible value is 10. But, we do not recommend setting the value from 6-10.

Note: UDP 8.0 supports the following:

- User mailbox
- Shared mailbox
- Mail-enabled Public folders
- Room mailboxes
- Equipment mailboxes
- Group mailboxes

Add the Required Role and Group to the Exchange Online Backup Account to Perform Backup and Restore

Add the backup account to the Discovery Management role group and assign Application Impersonation permission to it.

Follow these steps:

1. Add the required role and group using any one of the following ways:

   Using Office 365 portal

a. Log into Office 365 portal as an Administrator or with an account that has Global Admin permissions.

   The Exchange admin center page appears.

b. Go to **permissions** and from the **Add** drop-down, double-click **Discovery Management**.

   The Discovery Management dialog box appears.

   Note: Member of the Discovery Management role group can search mailboxes in the Exchange organization for data that meets specific criteria.

c. Under **Roles**, click **+** to add the ApplicationImpersonation role.

   The Select a Role dialog box appears.

d. From the **Display Name** drop-down, select **ApplicationImpersonation**.

   Note: The ApplicationImpersonation role enables applications to impersonate users in an organization to perform tasks on behalf of the user.

e. Click **Add** to add the ApplicationImpersonation role, and then click **OK**.

f. On the Discovery Management dialog box, under **Members**, click **+** to add the backup account as a member.

   The Select Members dialog box appears.

g. From the Name drop-down, select the backup account, click **Add** to add the member, and then click **OK**.

   The selected backup account is displayed under **Members** on the **Discovery Management** dialog.

h. Click **Save**.

Using Remote Powershell

a. To connect to the Exchange Online tenant using remote PowerShell, refer to the link.

b. Once connected, to add the backup account as a member of Discovery Management role group, use the:

`"Add-RoleGroupMember" cmdlet`

For example: Add-RoleGroupMember "discovery management" -member userName@domain.onmicrosoft.com.

c. To assign application impersonation role to the backup account, use the:

`"New-ManagementRoleAssignment" cmdlet`

For example:

New-ManagementRoleAssignment -Name: impersonationAssignmentName -Role:ApplicationImpersonation - User: "username@domain.onmicrosoft.com"

The ApplicationImpersonation role and Members group are added to the Exchange Online backup account.

Create an Exchange Online Backup Plan

A backup plan includes a backup task that performs a backup of Exchange Online mail data items (Mails, Calendar items, Contacts, so on) and stores data either at a non-deduplication data store or deduplication data store. Additionally, this backup plan allows UDP to discover and specify Exchange Online groups for which a backup task is performed. Each task consists of parameters that define the source, destination, schedule, and other backup details.

Follow these steps:

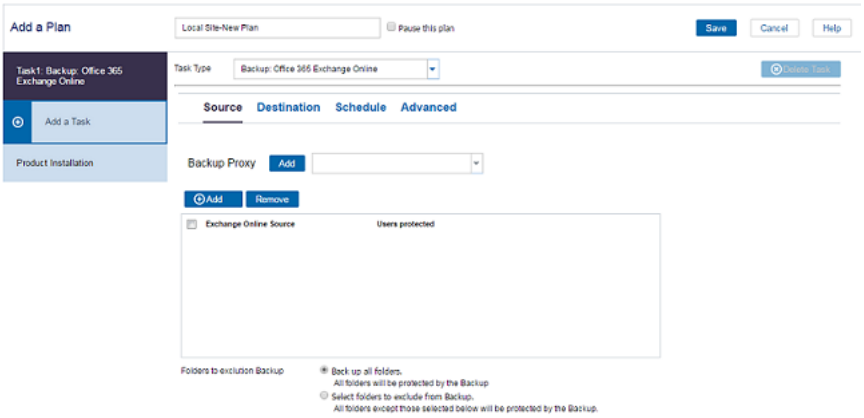
1. Click the **resources** tab on the Console.
2. From the left pane, navigate to **Plans**, and click **All Plans**.

   If you have created plans earlier, those plans are displayed on the center pane.
3. On the center pane, click **Add a Plan**.

   The **Add a Plan** dialog opens.
4. Enter a plan name.
5. (Optional) Select the **Pause this plan** check box.

   The plan does not run until you clear the check box to resume the plan.

   Note: If a plan is paused, then any in-progress job is not paused but all corresponding scheduled jobs associated with that plan are paused. However, you can manually run a job. For example, you can manually run backup job and replication job for a node even if the respective plan is paused. In such case, the following task to the on-demand (manual) job does not run. For example, there is a replication task after an on-demand backup job, the replication job does not run for the on-demand backup job. You need to manually run the replication job. When you resume the plan, the pending jobs do not resume immediately. After you resume the plan, the pending jobs run from the next scheduled time.
6. From the Task Type drop-down list, select **Backup: Office 365 Exchange Online**.



Now specify the Source, Destination, Schedule, and Advanced details.

Specify the Source

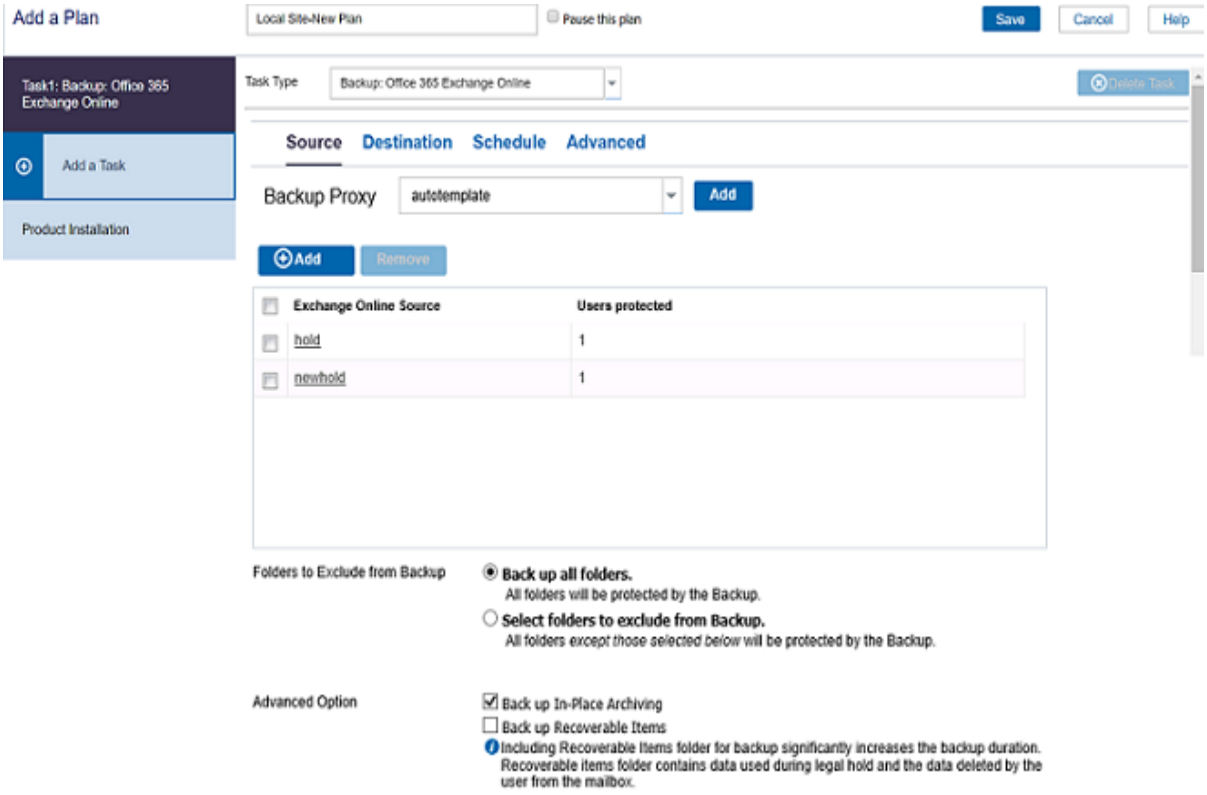
The Source page lets you specify the Exchange Online source nodes that you want to protect. You can select more than one Exchange Online source nodes in a plan. If you have not added any nodes to the Console, you can add Exchange Online source nodes from the Source page.

Note: You can save a plan without adding any source nodes but the plan is not deployed unless you add any nodes.

You can also manage Exchange Online nodes using Public folder Mailbox Support for Exchange Online Protection.

Follow these steps:

1. Click the **Source** tab.



2. For Backup Proxy, do one of the following:
   - From the drop-down list, select an existing backup proxy.
   - To add a new backup proxy, next to the drop-down list, click **Add**.
3. Add Exchange Online node using one of the following options:

   Method 1

a. Click **Add** and then click **Select Source to Protect in Arcserve UDP**.

   The Add Nodes to Plan dialog is displayed.

b. Select a node, and then click **Connect**.

   Note: You can also search for the Exchange Online nodes that you want to protect in **Search**.

c. Select the **Protect all of the Office 365 Exchange Sources** check box to protect all the Exchange Online accounts across all the pages.

   Note: To add all the Exchange Online accounts to the protected list, you may click the right (>) arrow.

   The Exchange Online accounts that you selected are added.

Method 2

a. Click **Add** and then click **Add Exchange online Source**.

   The Add nodes to a plan dialog appears.

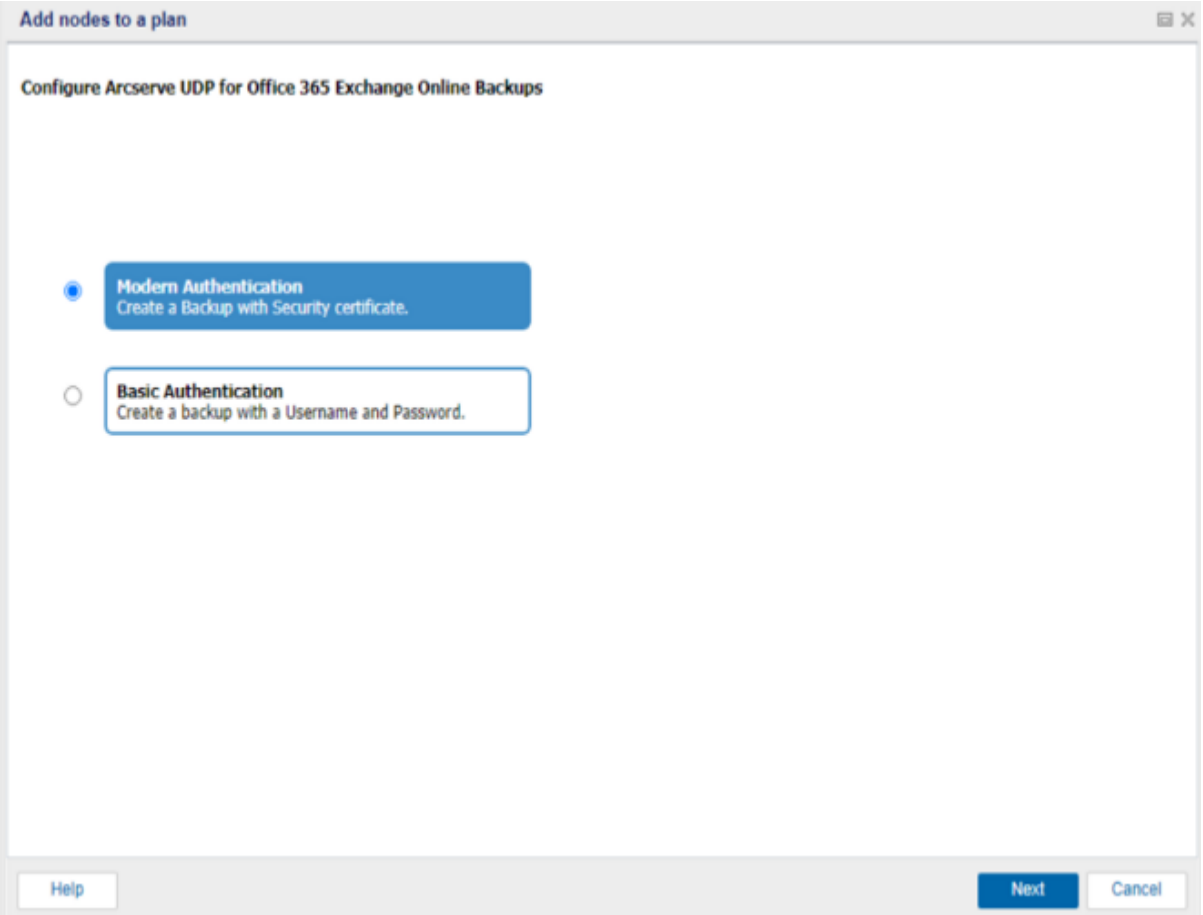
   Note: Unlike other nodes, you cannot add the Exchange Online node from **All Node**s page. You can add an Exchange Online node only in a plan or when you modify a plan.

   Multiple Exchange online nodes can use the same user account (service account) of Exchange online.

4. To apply authentication, do one of the following:

   Modern Authentication

   For modern authentication, you can run a backup using security certificate.



Use security certificate

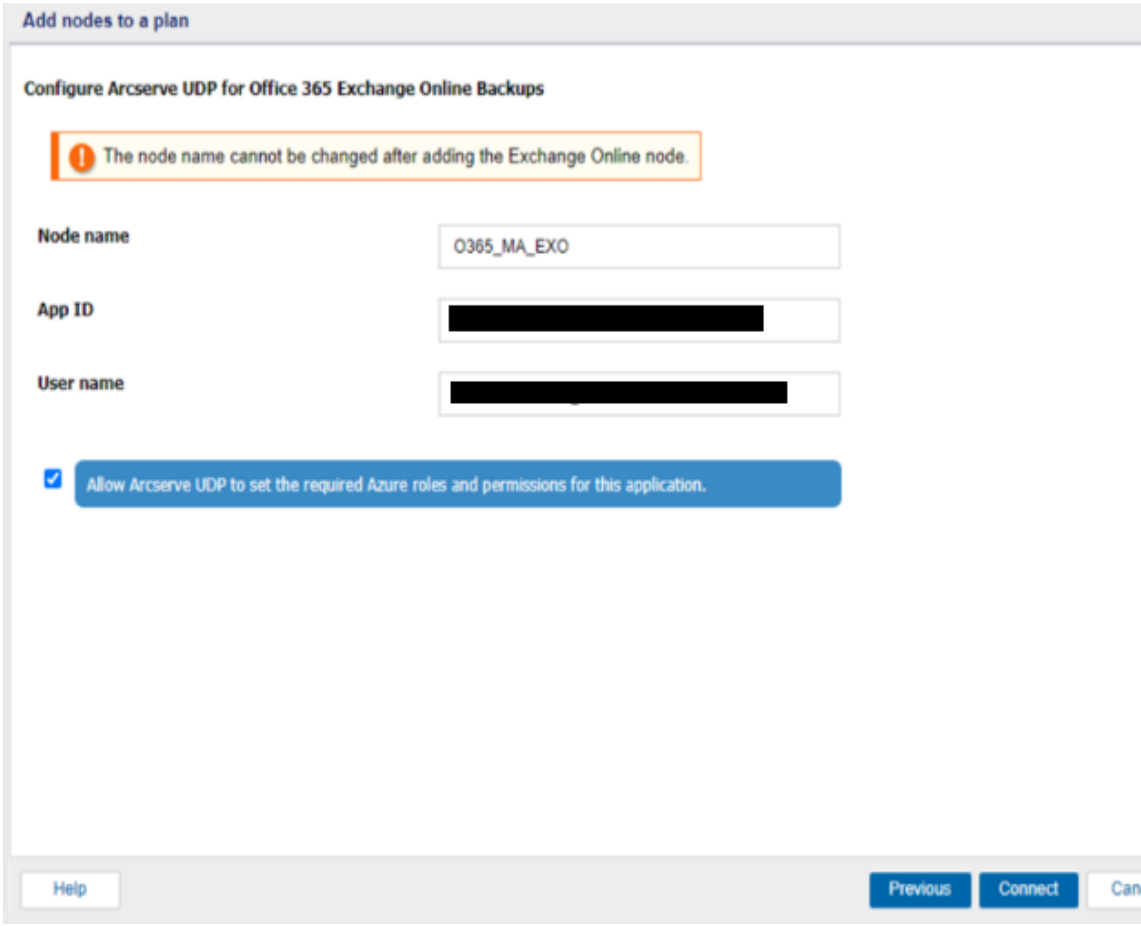
Note: By default, the backup plan selects **Use security certificate** for modern authentication.

Follow these steps:

a. To authenticate the UDP application with security certificates, see How to Use Self-Signed Certificate.

b. Enter the following details:

- Node Name
- App ID
- User Name



c. To apply the required roles and permissions, select the **Allow Arcserve UDP to set the required Azure roles and permissions for this application** check box, and then click **Connect**.

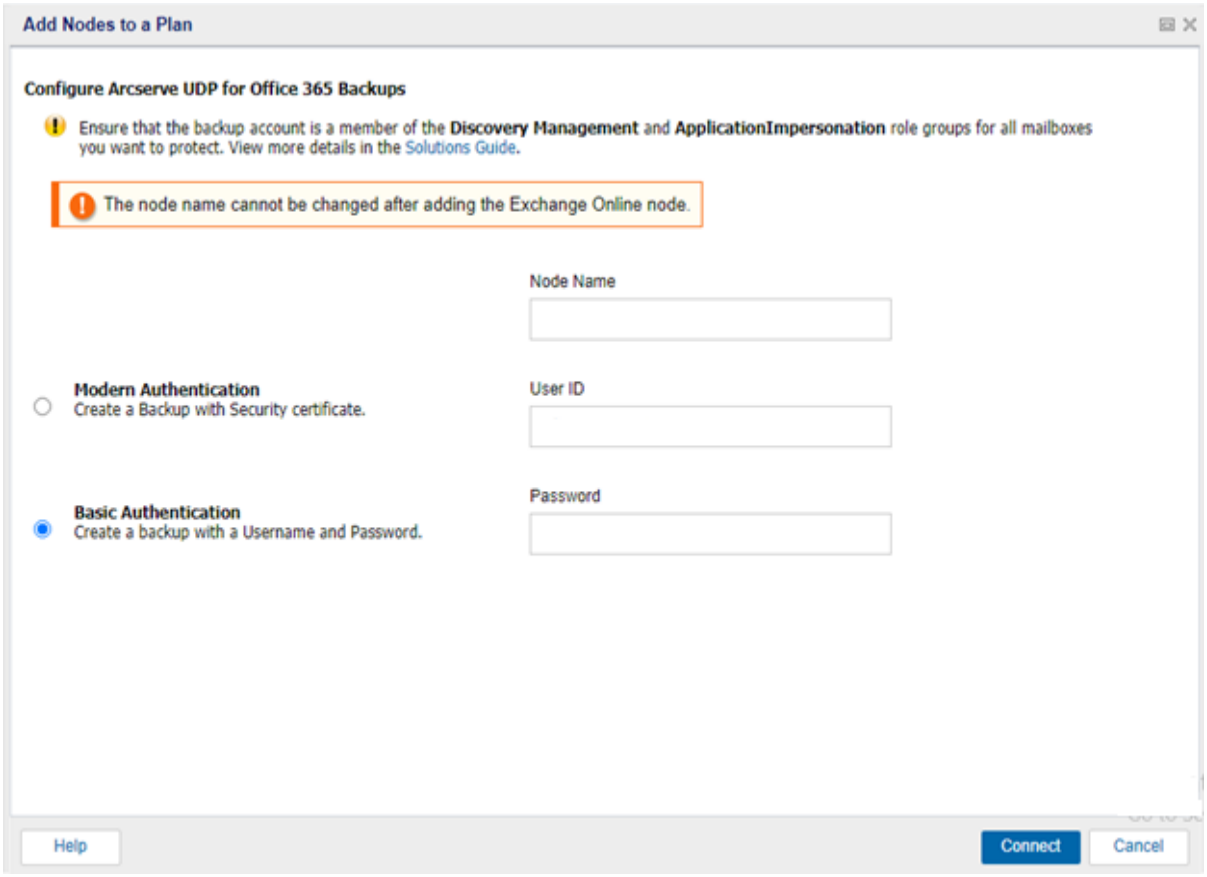
Notes:

- Exchange online node name cannot be changed after adding the node.
- Enter the user name of Exchange Online backup account that meet the prerequisites in **Admin Username**.
- Modern Authentication does not apply to the following:
  - Customers using Microsoft 365 (Office 365) for their Arcserve Cloud Hybrid instances or Arcserve Cloud Backup for Office 365, which was created before 18th Oct, 2020 does not allow modern authentication.
  - Customers who continue to use basic authentication

Basic Authentication

For Basic Authentication, to proceed with the backup, provide the following information, and then click **Connect**:

- Node name
- User ID
- Password



Note: Updating or changing the user account may change the number of protected mailboxes. You need to verify that the new or updated service accounts have impersonation rights for the mailboxes to be protected.

5. To protect Exchange Online mailboxes, do one of the following:

   - To protect all exchange online mailboxes and groups, select the **Protect all of the Office 365 Exchange Source** check box.

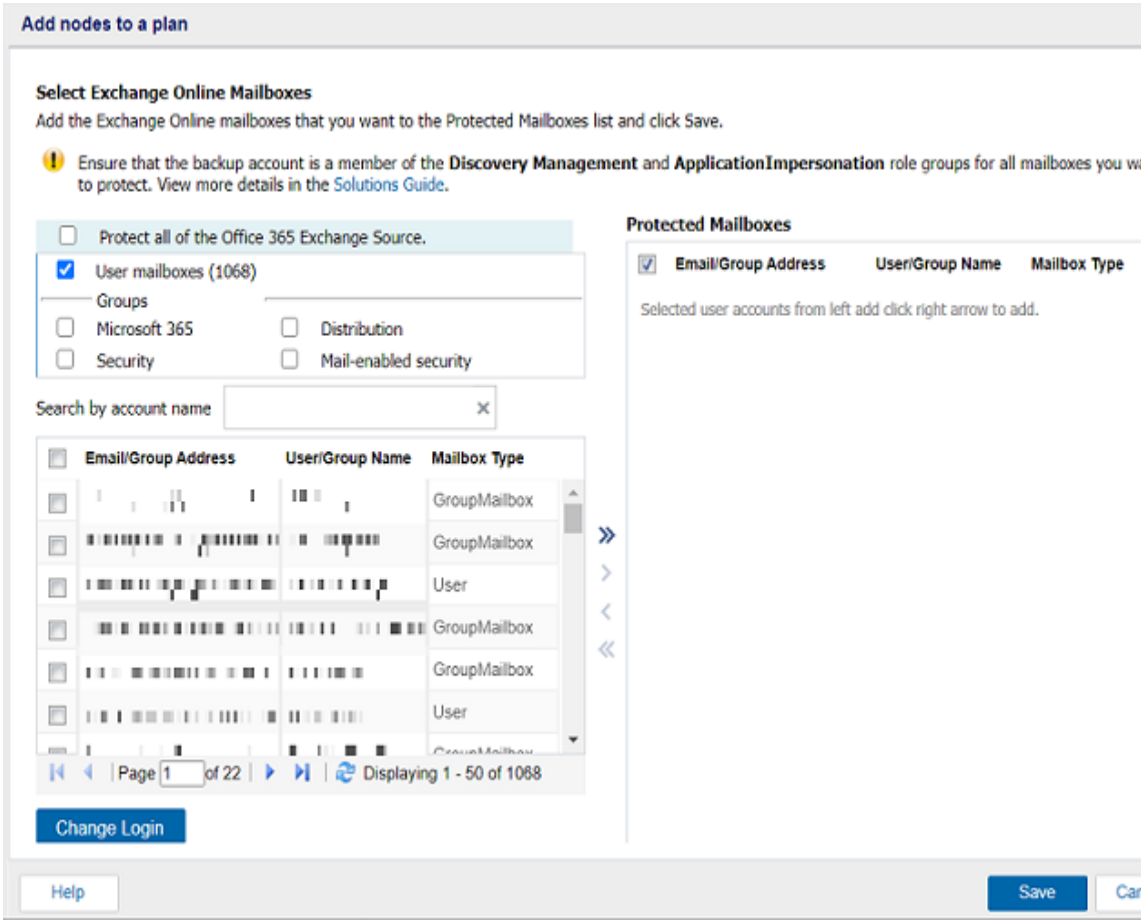
     Note: This option is automatically selected by default when the Select Exchange Online Mailboxes screen appears.

   - For User mailboxes, follow these steps:

a. Select the **User mailboxes** check box.

   A list of user mailboxes appears.

b. (Optional) To search an account from the list, in the Search by account name field, type the account name.



c. From the list of user mailboxes, select the mailboxes that you want to protect.

d. Click the right arrow (>) to move the selected mailboxes to the Protected Mailboxes list.

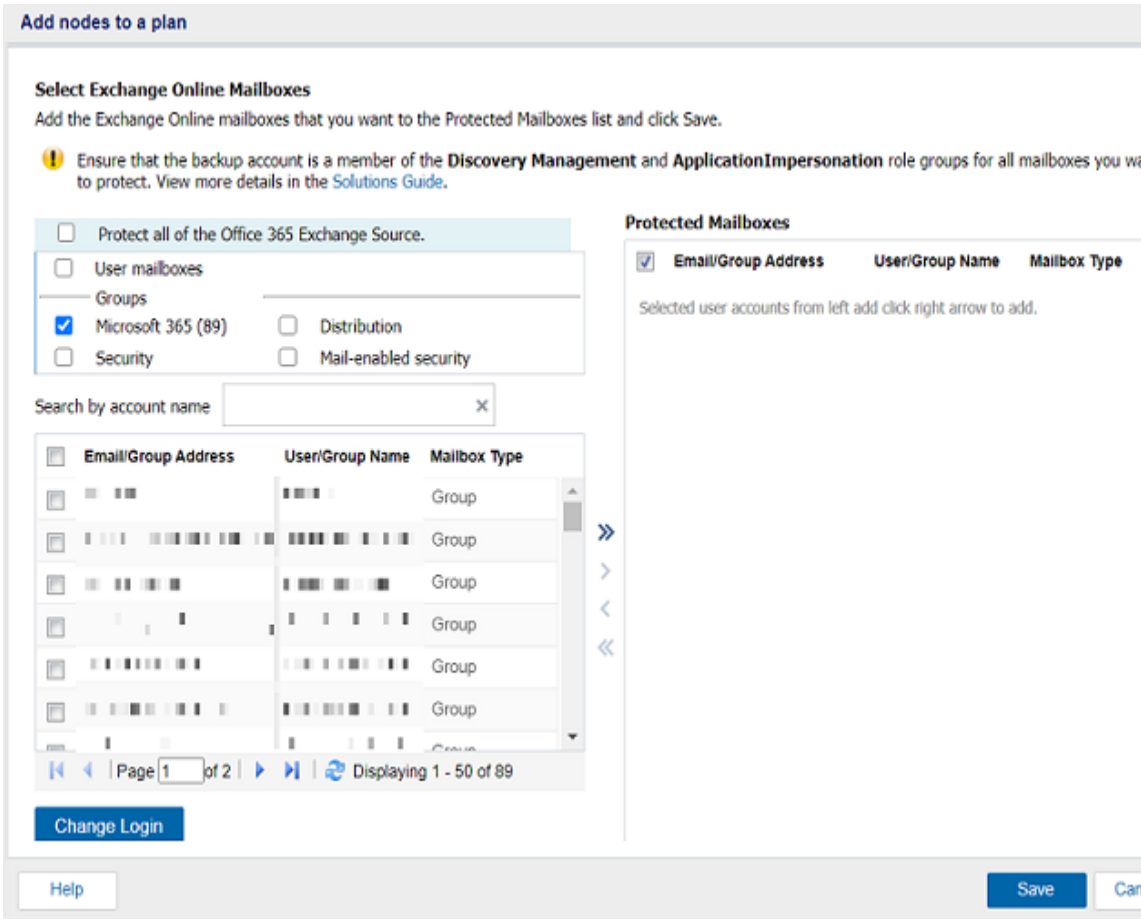
- For Groups, follow these steps:

  a. Select the **Microsoft 365** check box.

     A list of groups appears.

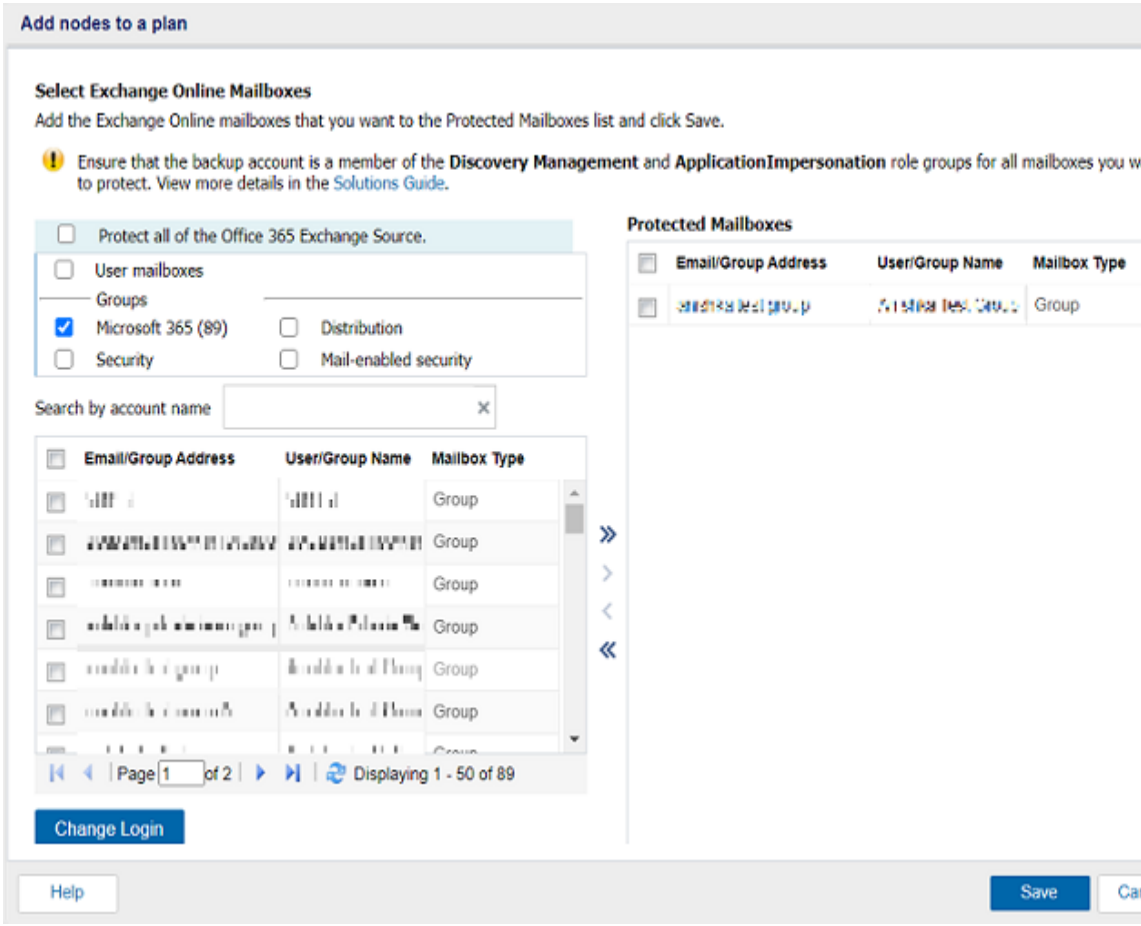
     Note: If you want to select groups along with sub-groups such as **Distribution**, **Security**, or **Mail-enabled security**, select the **Microsoft 365** check box, and then select a sub-group check box as needed.

b. (Optional) To search an account from the list, in the Search by account name field, type the account name.



c. From the list of groups, select the groups that you want to protect.

d. Click the right arrow (>) to move the selected groups to the Protected Mailboxes list.



6. Click **Save**.
7. On the Sources tab, go to **Folders to Exclude from Backup**, and then select the desired check box.
8. From the Advanced Option, select the desired check box.
   - To allow Exchange Online Protection support Archiving Mailbox, select the check box of **Backup up In-Place Archiving**.

     Note: For more information about Archiving Mailbox, refer to the link.
   - Select the **Backup up Recoverable items** check box to enable to protect the mailbox that enables the In-Place Hold or Litigation Hold feature.

     Note: For Archiving In-Place Hold and Litigation Hold for Exchange Online, refer to the link.

   Note: To enable both the features in the mailbox at the same time to back up the recoverable items in Archiving mailbox, select both the options **Backup up In-Place Archiving** and **Backup up Recoverable items**.

Specify the Destination

The destination is a location where you store the backup data. You must at least specify the destination to save the plan.

Follow these steps:

1. Click the **Destination** tab.
2. Select the **Arcserve UDP Recovery Point Server** option. **Arcserve UDP Recovery Point Server** specifies that the backup destination is a recovery point server. If you select this option, the data is stored as recovery points. You cannot store data as recovery sets.
3. Provide the following details:
    a. Select a recovery point server.
    b. Select a non-deduplication or deduplication data store. The list displays all the data stores created on the specified recovery point server.
    c. Provide a session password. The session password is optional when the backup destination is an unencrypted RPS data store.
    d. Confirm the session password.

The destination is specified.

Specify the Schedule

The Schedule page lets you define a schedule for Backup, Merge, and Throttle functions to repeat at specific intervals. After you define a schedule, the jobs run automatically per the schedule. You can add multiple schedules and can provide retention settings.

A Backup Schedule refers to regular schedule that is repeated multiple times a day based on the number of hours or minutes you select. Besides the regular schedule, a backup schedule also provides options to add daily, weekly, and monthly schedules.

Note: For more information on scheduling and retention settings, see Understanding Advanced Scheduling and Retention.

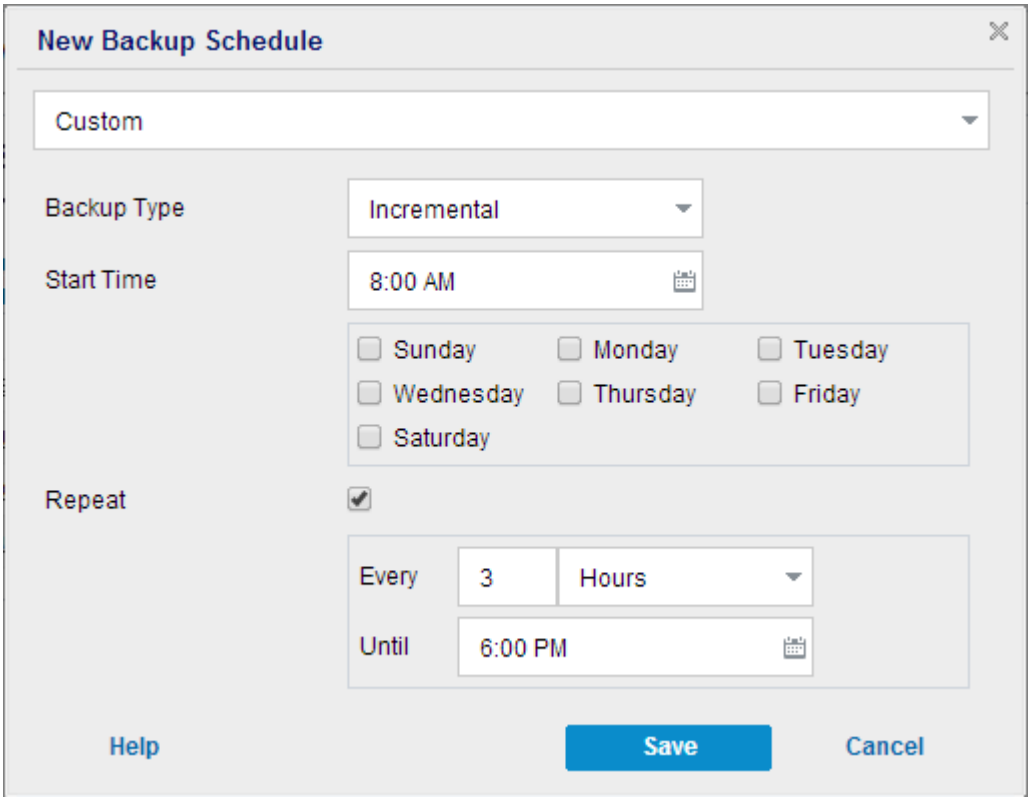
Follow these steps:

1. Add backup, merge, and throttle schedules.

   Add Backup Schedule

   a. Click **Add** and select **Add Backup Schedule**.

      The **New Backup Schedule** dialog opens.



   b. Select one of the following options:

      Custom

Specifies the backup schedule that repeats multiple times a day.

Daily

Specifies the backup schedule that occurs once a day. By default, all the days of the week are selected for Daily backup. If you do not want to run the backup job on a specific day, clear the check box for that day of the week.

Weekly

Specifies the backup schedule that occurs once a week.

Monthly

Specifies the backup schedule that occurs once a month.

c. Select the backup type.

Full

Determines the backup schedule for Full Backups. As scheduled, Arcserve UDP performs a full backup of all used blocks from the source machine. A full backup typically consumes time depending on the backup size.

Incremental

Determines the backup schedule for Incremental Backups.

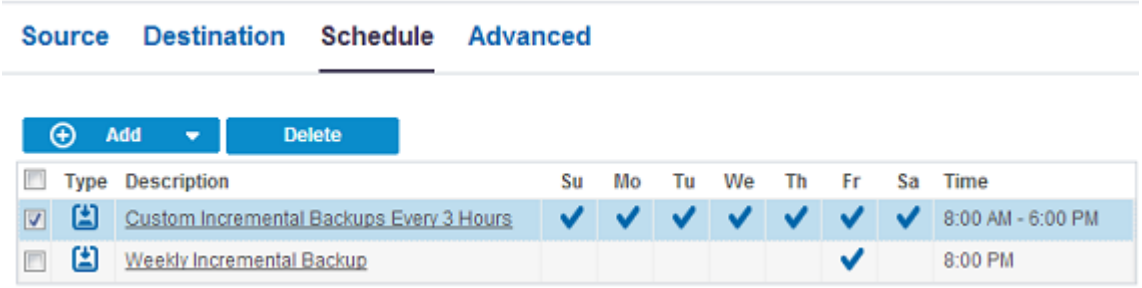
As scheduled, Arcserve UDP incrementally backs up only those blocks that have changed after the last successful backup. The advantages of Incremental Backups are that the backup is a fast backup and it produces a small backup image. This is the most optimal way to perform a backup.

d. Specify the backup start time.

e. (Optional) Select the **Repeat** check box and specify the repeat schedule.

f. Click **Save**.

The Backup Schedule is specified and appears on the **Schedule** page.



Add Merge Schedule

a. Click **Add** and select **Add Merge Schedule**.

   The **Add New Merge Schedule** dialog opens.

b. Specify the start time to start the merge job.

c. Specify **Until** to specify an end time for the merge job.

d. Click **Save**.

The Merge Schedule is specified and appears on the **Schedule** page.

Add Throttle Schedule

a. Click **Add** and select **Add Throttle Schedule**.

   The **Add New Throttle Schedule** dialog opens.

b. Specify the throughput limit in MB per minutes unit.

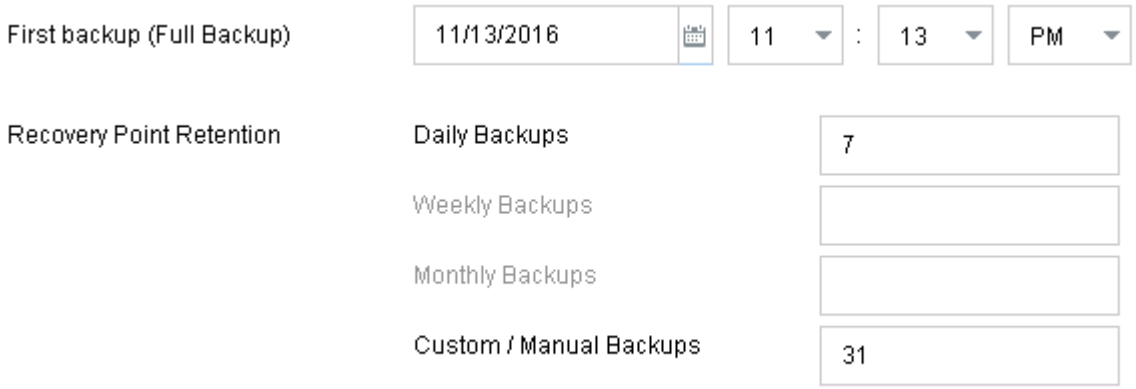
c. Specify the start time to start the backup throughput job.

d. Specify **Until** to specify an end time for the throughput job.

e. Click **Save**.

The Throttle Schedule is specified and displayed on the **Schedule** page.

2. Specify the start time for the scheduled backup.



3. Specify the recovery points retention settings for Custom, Daily, Weekly, and Monthly schedule.

   These options are enabled if you have added the corresponding backup schedule. If you modify the retention settings on this page, the changes are reflected on the **Backup Schedule** dialog.

   The schedule is specified.

Understanding Advanced Scheduling and Retention

The scheduling option lets you specify a Custom schedule or a Daily / Weekly / Monthly schedule, or both the schedules. In the Custom schedule, you can configure the backup schedule for each day of the week and you can add up to four backup schedules each day. You can select a specific day of a week and create a time window to define when to run backup and at what frequency.

Schedule	Supported Job	Comments
Backup	Backup job	Define time windows to run backup jobs.
Backup throttling	Backup job	Define time windows to control the backup speed.
Merge	Merge job	Define when to run merge jobs.
Daily schedule	Backup job	Define when to run daily backup jobs.
Weekly schedule	Backup job	Define when to run weekly backup jobs.
Monthly schedule	Backup job	Define when to run monthly backup jobs.

You can also specify the retention settings for the recovery points.

Note: Set the retention settings within each plan to control how data for the nodes assigned to that plan are retained at the target data store.

Schedules for Daily / Weekly / Monthly backups are independent to the Custom schedule, and each other. You can configure only to run Daily backup, Weekly backup or Monthly backup, without configuring the Custom schedule.

Backup Job Schedule

You can add four time windows per day in your backup schedule. A valid time window is from 12:00 AM until 11:59 PM. You cannot specify a time window such as 6:00 PM to 6:00 AM. In such cases, you must manually specify two different time windows.

For each time window, the start time is inclusive, and the end time is exclusive. For example, you have configured to run Incremental Backup every one hour between 6:00 AM and 9:00 AM and the backup will start at 6:00 AM. This means the backup runs at 6:00 AM, 7:00 AM, 8:00 AM, but NOT at 9:00 AM.

Note: If you want to run the backup job repeatedly until the end of day, set the schedule until 12:00 AM. For example, to run backup job every 15 minutes for the entire day, set the schedule from 12:00 AM to 12:00 AM, every 15 minutes.

Backup Throttle Schedule

Backup throttle schedule lets you control the backup throughput speed that in turn controls the resource usage (disk I/O, CPU, network bandwidth) of the server being backed up. This is useful if you do not want to affect the server

performance during business hours. You can add four time windows per day in your backup throttle schedule. For each time window, you can specify a value, in MB per minute. This value is used to control the backup throughput. Valid values are from 1 MB/minutes to 99999 MB/minutes.

If a backup job extends its specified time, then the throttle limit adjusts according to the specified time window. For example, you have defined the backup throttle limit as 500 MB/minute from 8:00 AM to 8:00 PM, and 2500 MB/minute from 8:00 PM to 10:00 PM. If a backup job starts at 7:00 PM and it runs for three hours, then, from 7:00 PM to 8:00 PM the throttle limit is 500 MB/minute and from 8:00 PM to 10:00 PM the throttle limit is 2500 MB/minute.

If you do not define any backup schedule and backup throughput schedule, the backup runs as fast as it can.

Merge Schedule

Lets you merge recovery points based on the provided schedule.

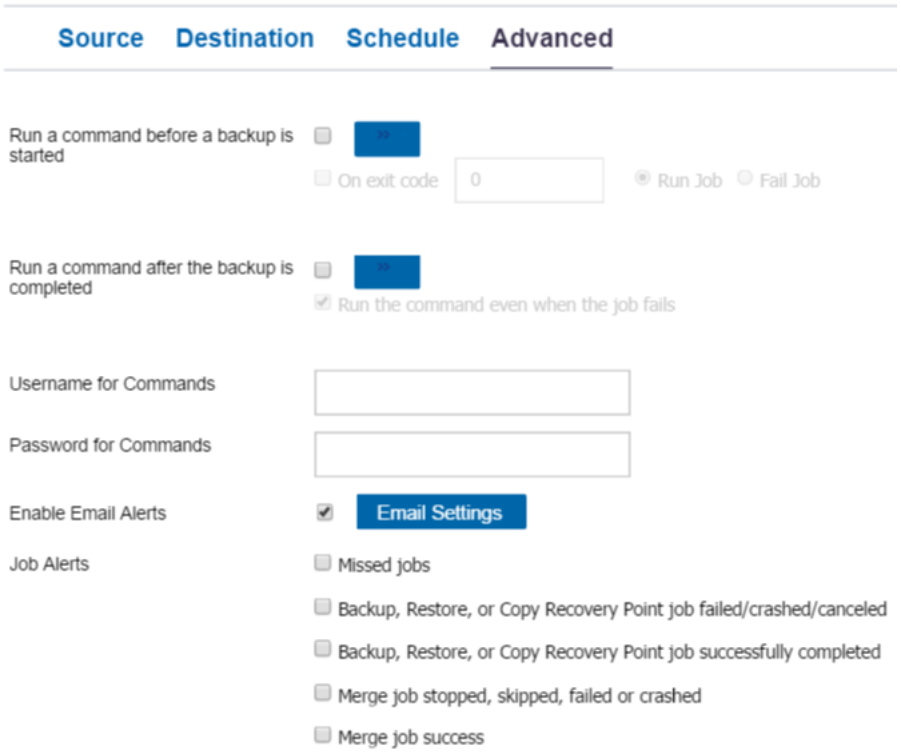
Consider the following points for the merge job:

- At any given time only one merge job can run for a node.
- If a merge job starts, it has to complete before the next merge job can start. This means, if one or more sets of recovery points are merged, new recovery point cannot be added to this merge process, until the merge process of the current set of recovery point completes.
- If a merge job is processing more than one set of recovery points (for example set [1~4], set [5~11], and set [12~14]; they are three sets), recovery point server processes these sets one by one.
- If a merge job is resumed after a pause, the job detects at which point it is paused and resumes the merge from the break-point.

Specify the Advanced Settings

The **Advanced** tab lets you specify some advanced settings for the backup job. The advanced settings include providing the location of any scripts, and email settings.

The following image displays the **Advanced** tab:



Follow these steps:

1. Specify the following details.

   Run a command before a backup is started

   Lets you run a script before the backup job starts. Specify the path where the script is stored inside the proxy node. Click **On exit code** and specify the exit code for **Run Job** or **Fail Job**. **Run Job** indicates that the backup job continues when the script returns the exit code. **Fail Job** indicates that the backup job stops when the script returns the exit code.

   Run a command after a backup is completed

   Lets you run a script after the backup job is completed. Specify the complete path where the script is stored.

   Run a command even when the job fails

If this check box is selected, the script specified in **Run a command after a backup is completed** is executed even when the backup job fails. Otherwise, that script is executed only when backup job completes successfully.

Username for Commands

Lets you specify the username to run the commands.

Password for Commands

Lets you specify the password to run the commands.

Enable Email Alerts

Lets you enable email alerts. You can configure email settings and specify the types of alerts that you want to receive in an email. When you select this option, the following options are enabled for your selection.

Email Settings

Lets you configure the email settings. Click **Email Settings** and configure the email server and proxy server details. For more information about how to configure Email Settings, refer to Email and Alert Configuration.

Job Alerts

Lets you select the types of job alert emails that you want to receive.

2. Click **Save**.

Note: When you select a node as a backup source or backup proxy, Arcserve UDP checks whether the agent is installed on proxy node and if it is the latest version. Arcserve UDP then displays a verification dialog that lists all the nodes that either have an outdated version of the agent or does not have the agent installed. To install/upgrade the agent on these nodes, select the installation method and click **Save**.

The changes are saved and a green checkmark appears next to the task name. The plan page closes.

Note: If you have to add another task, you must select the plan from the **resources** tab and modify the plan. To modify the plan, click the plan from the center pane. The plan opens and you can modify it. You may add the **Copy Recovery Point**, **Copy to Tape**, **Replicate**, and **Replicate from a remote RPS** tasks as follow up tasks.

The plan is automatically deployed to the proxy server node.

The exchange online backup plan for the proxy server is created. The backup runs per the schedule that you have configured on the **Schedule** tab. You can also perform a manual backup at any time.

(Optional) Perform a Manual Backup

Typically, based on the schedule settings you applied, the backups perform automatically. In addition to the scheduled backup, a manual backup provides you the option to back up your nodes on a need basis. For example, if you have a repeat schedule for Full and Incremental backups, and you want to make major changes to your machine, you can perform an immediate manual backup without waiting for the next scheduled backup.

To perform a manual backup of Exchange Online nodes, follow these steps:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and then click **All Nodes**.

   The Exchange Online nodes are displayed in the center pane.
3. Select the Exchange Online nodes (for example, Mailbox@<organizationname.com>) that you want to backup and make sure a plan is assigned to it. The node name is the account that is used when adding the Exchange Online node and connecting it.
4. Right-click the node, and then select **Backup Now**.

   The Run a backup now dialog opens.
5. Select the type of backup that you want to perform, and then click **OK**.

   The Progress Information dialog appears.
6. Wait until the backup job completes, and then click **OK**.
7. To monitor the backup jobs, navigate to **jobs** > **Jobs in Progress**.

   Wait until the backup job is finished.

   The manual backup is successfully performed.

To perform a manual backup of an Exchange Online plan, follow these steps:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Plans**, and then click **All Plans**.

   The Exchange Online backup plans are displayed in the center pane.
3. Select the plan that you want to backup and make sure a plan is assigned to it.
4. Right-click the plan, and then select **Backup Now**.

   The Run a backup now dialog opens.
5. Select the type of backup that you want to perform, and then click **OK**.

   The Progress Information dialog appears.

6. Wait until the backup job completes, and then click **OK**.
7. To monitor the backup jobs, navigate to **jobs** > **Jobs in Progress**.

   Wait until the backup job is finished.

   The manual backup is successfully performed.

Verify the Backup

To verify your backup job, confirm that you have successfully created the backup plan. After you verify that the plan is created successfully, confirm whether the backup job is running as scheduled. You can verify the status of backup jobs from the **Jobs** tab.

Follow these steps to verify plans:

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and then click **All Nodes**.

   A list of all nodes is displayed on the center pane.
3. Verify that plans are mapped with nodes.

Follow one of these methods to verify backup jobs:

Method 1

1. Click the **jobs** tab.
2. From the left pane, click **All Jobs Completed**.

   The status of each job is listed on the center pane.
3. Verify that the backup job is successful.

   The backup job is verified.

Method 2

1. Click the **jobs** tab.
2. From the left pane, click **Jobs Successfully Completed**.

   The list of backup jobs completed appear on the center pane.
3. Click the job you want to verify.

   The right pane refreshes.
4. Under Job Details, click the **View Logs** hyperlink.
5. Verify that the backup job is successful.

   The backup job is verified.

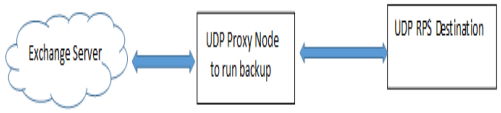
Applying Best Practices

The best practices section for Exchange Online contains the following topics:

- Configuration for optimal performance
- Best Practices for Backup
- Best Practice for Restore
- Frequently Asked Questions

Configuration for optimal performance

Possible topology:



You can have proxy installed either on RPS or separately. However, for optimal performance we recommend to put both on the same node.

Recommended configuration for proxy

- Architecture: 64 bit Windows machine. For more information about supported operating systems for Arcserve UDP, click here.
- Memory: 8 GB or higher
- CPU: 2 cores

Recommended configuration for Recovery Point Server:

Refer Release notes for system requirements.

Best Practices for Backup

- Expect longer backup duration in the first full backup that gets the data from Microsoft Exchange Online server through a WAN link.
- As the backup duration for the first full backup is long, ensure required configuration for hardware, network, and resource (for example, disk space on destination, memory, CPU, and so on) availability during the backup to avoid disruption due to any environmental failures.
- Resume backup from the breakpoint during any of the following scenarios:
  - Ensure that a backup is not running during a planned outage or network downtime. If running, then you need to cancel the backup. When you cancel, the backup retains the recovery point backed up partially and the next backup schedule can resume from the breakpoint.
  - If during backup the machine encounters an unexpected shut down or process termination, the recovery point of the running job is removed. Start the backup again. Backup does not start from breakpoint. All data is backed up again.
- We always recommend to use one node with auto-protection mechanism to protect all mailboxes, including newly created mailboxes after creating the Plan.
- In Arcserve UDP, the supported size of the protected data is limited to 8 Terabytes (Compressed) by default. To configure the size, you can create the following registry value on the proxy node:

  [HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll\<NodeGUID>]

  " VirtualDiskSize "=dword:00000002

  The above example sets the size to 2 TB.

  Note: Increasing the size of the virtual disk converts the next incremental backup job to full and also impacts the speed of the following jobs:

  Copy to tape, Replication, Copy Recovery Point to cloud.
- Select the size of virtual disk based on your current source data size and its growth.

  For example: If the total size of all users is 5 TB and average data growth per day is 1%, that is 50 GB, you need 5 TB + 50GB * 30days = 6.5 TB of disk space to accommodate first full and 30 daily incremental backups for a month.

As a result, the minimum size of the disk required is 7 TB (uncompressed). To retain more recovery points, use the same method to calculate the size.

- Use the throttling schedules properly to ensure backup gets desired bandwidth during off business hours. Run a small backup with just few users to assess the network utilization by the backup job and set the throttle values accordingly.
- Ensure that the backup account meets the mandatory prerequisites. For more information, view link.
- If the network link used to connect the Exchange Online server is unstable and results in frequent backup cancelation, configure Recovery Point Retention policy to retain a large number of non-merged sessions.

  Default value for daily incremental backups: 7

  Default value for manual backups: 31

  Configure the backup schedule(s) to retain at least one daily, weekly, and monthly backup.
- Ensure that setup has sufficient licenses. You can manage license using license manager. For details, Refer Q.4 in the FAQ section.

Best Practice for Restore

Ensure that the account used for restore has required permissions to impersonate the selected target user. For more details, view link.

Exchange Online Frequently Asked Questions

1. Can I use RPS directly to backup Exchange Online users without using any proxy?

   Yes, you can.

2. Do I need administrator/ Group Admin account to backup Exchange Online mailboxes?

   No, the backup account only needs specific set of permissions. For more details, view link.

3. What type of throughput can I expect for Exchange Online backups if my internet bandwidth is good?

   Besides network bandwidth, Throughput is also limited by the rate at which data is read from Exchange server that is controlled by Exchange server design.

   Sometimes, Exchange server refuses the connection to Arcserve UDP to read data. In such cases, Arcserve UDP keeps trying till it gets a successful connection. The longer wait time can lead to the lower throughput numbers.

4. If I have more users configured in the backup plan than the number of licenses, can I run partial backups of the licensed users?

   No, you cannot. if you run backup for more users than the number of licenses available, then the backup fails for all users. For example, if you have configured 100 users to protect in the Arcserve UDP console whereas you have 90 available licenses, the backup will fail. To run the backup successfully, you need to remove 10 or more users from the protected users list.

5. How can I improve the backup throughput performance?

   Consider the following points to ensure optimal throughput performance:

   - Divide a big backup into smaller manageable backup jobs.
   - Exclude the folders that you do not want to back up, from the Plan settings. For example, Clutter, Sync Issues, and so on.
   - Keep a periodic check on the internal and external network infrastructure.

Configuring for Multi-Factor Authentication

When an organization has multi-factor authentication (MFA) enabled for the users, the office 365 backup plan needs to be configured using the App Password for the backup service account.

Perform the following steps to configure Arcserve UDP to Support multi-factor authentication:

1. Enable the backup service account to Set app password
2. Create app password for the backup service account

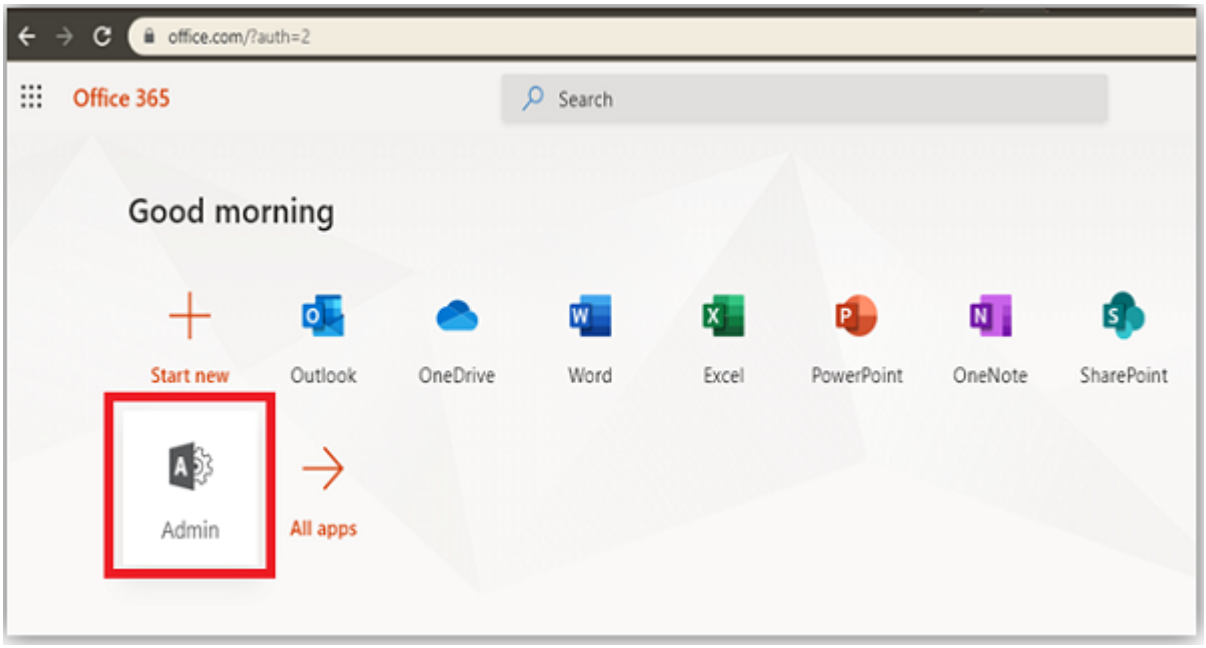
Note: MFA authentication (App password) is currently supported for O365 Exchange Online and SharePoint Online backups only.

Enable the backup service account to Set app password

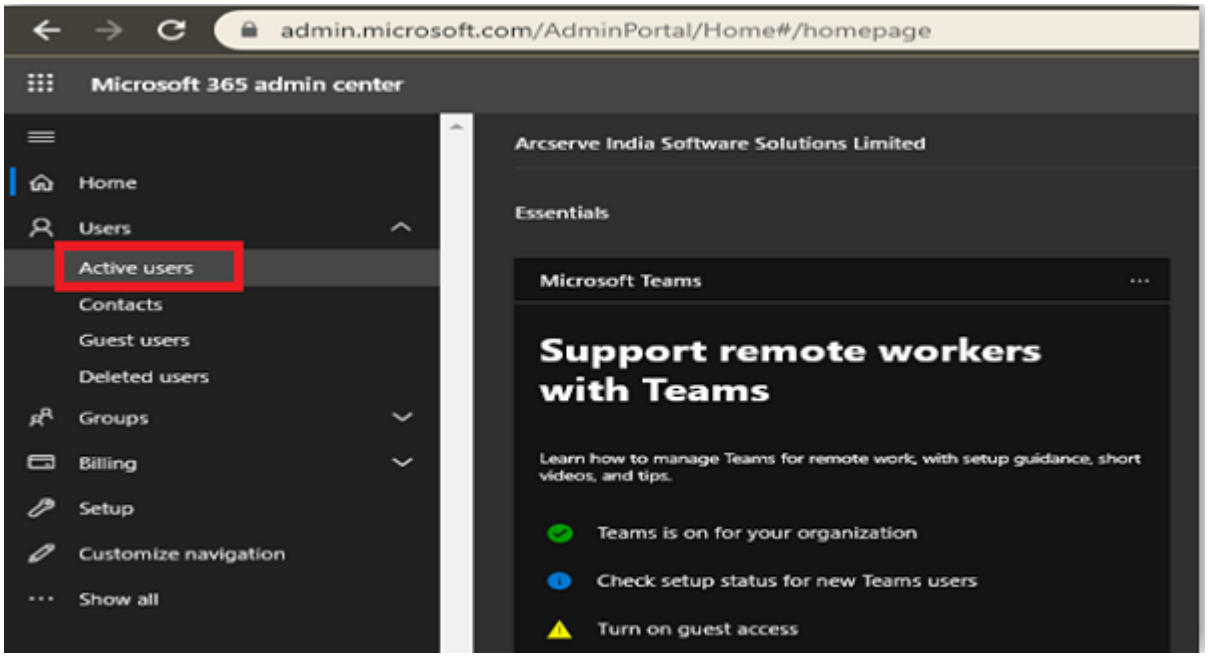
To configure, the first step is to enable the backup service account to Set app password.

Follow these steps:

1. Sign into Microsoft Office 365 using credentials of an administrator account and click the **Admin** icon.

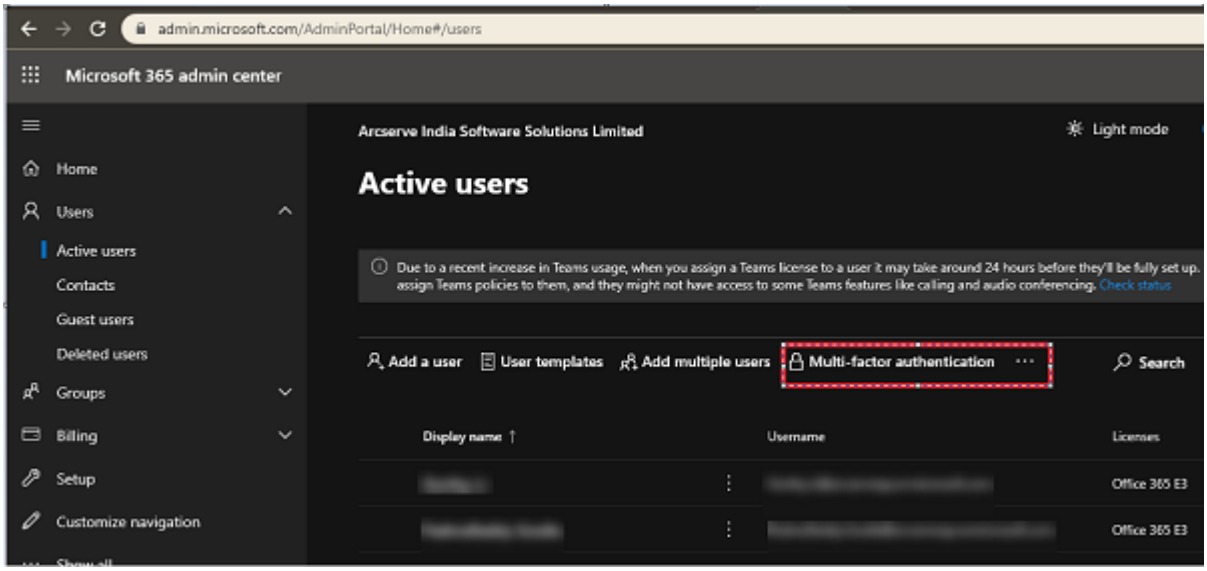


2. From the Microsoft 365 admin center screen, navigate to **Users** > **Active users**.



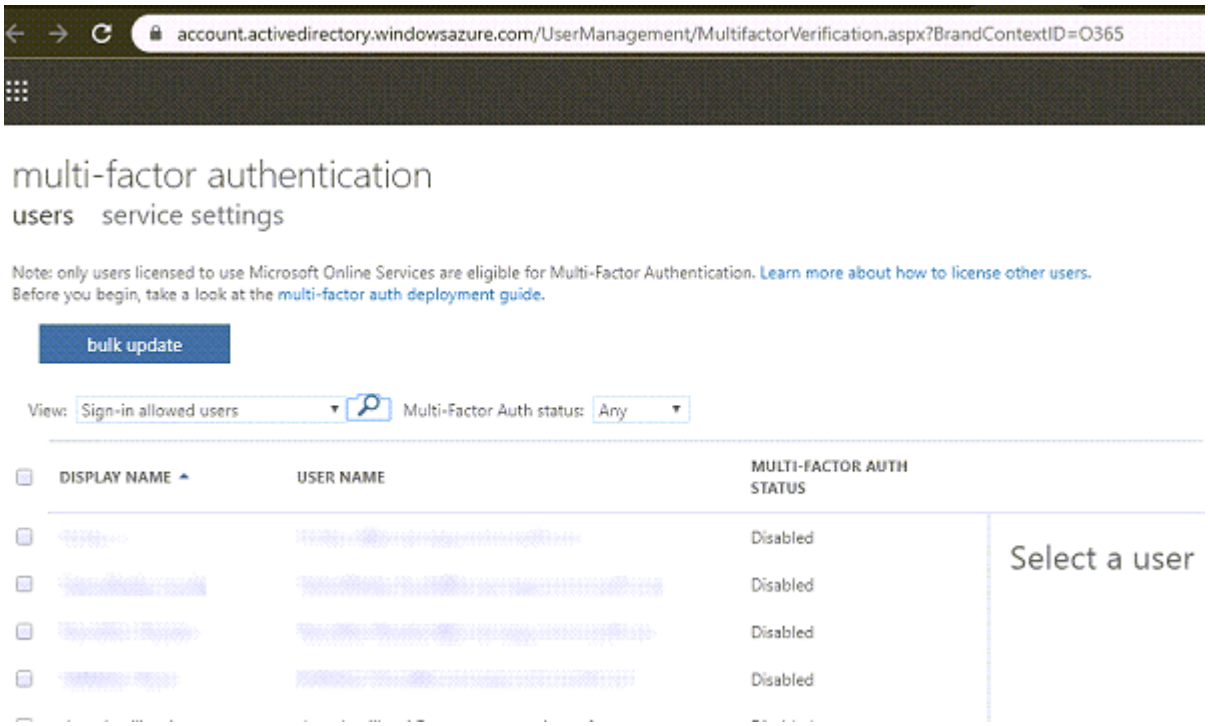
3. From the Active users screen, click the **Multi-factor authentication** option.

   Important! If you do not see the **(...)** option, then you are not a global admin for your subscription.

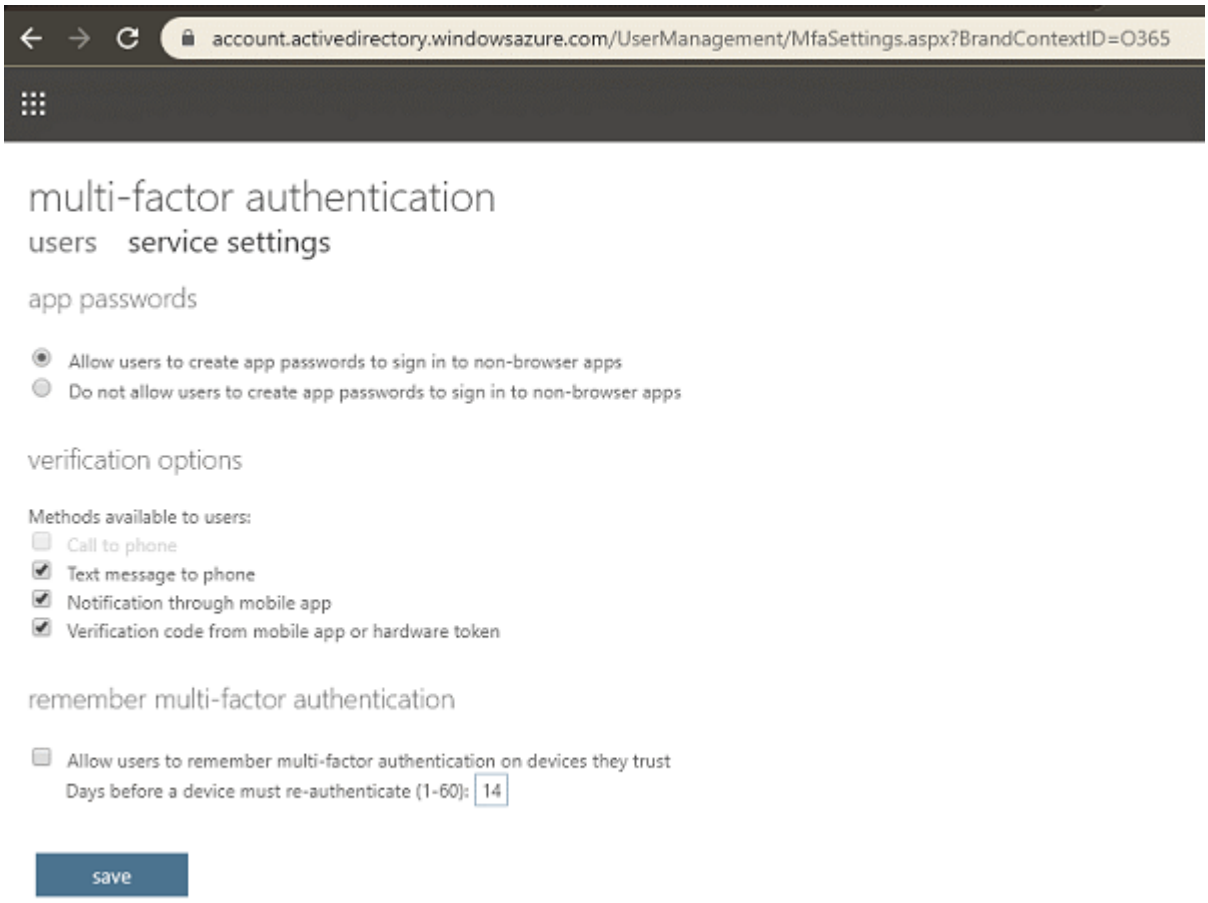


   You are led to the Setup Azure multi-factor authentication.

   Steps 4 and 5 only need to be set once. Then, skip to step 6



4. From the multi-factor authentication screen, click **service settings**.
5. From app passwords, select the checkbox of **Allow users to create app passwords to sign in to non-browser apps**.



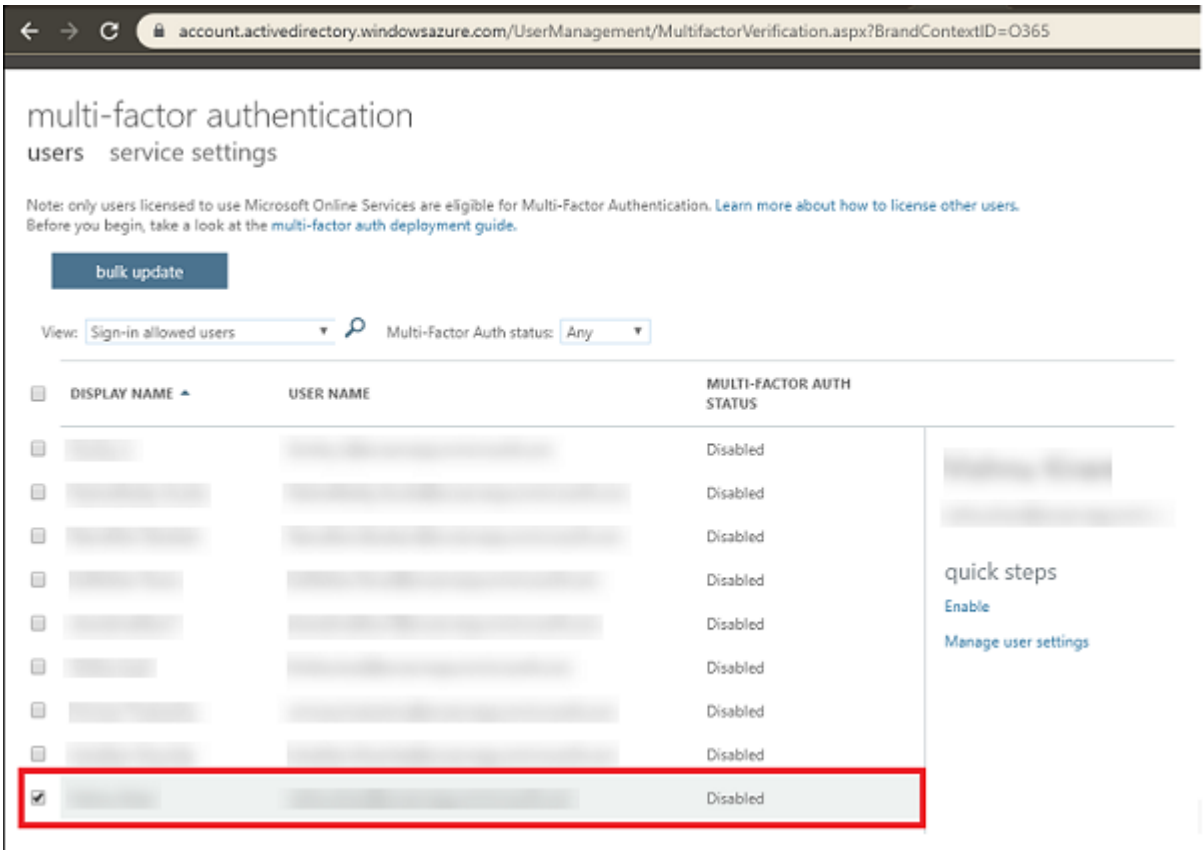
You can then use client Office apps after you create a new password.

6. Click **save** and close the window.

   You return to the users screen.

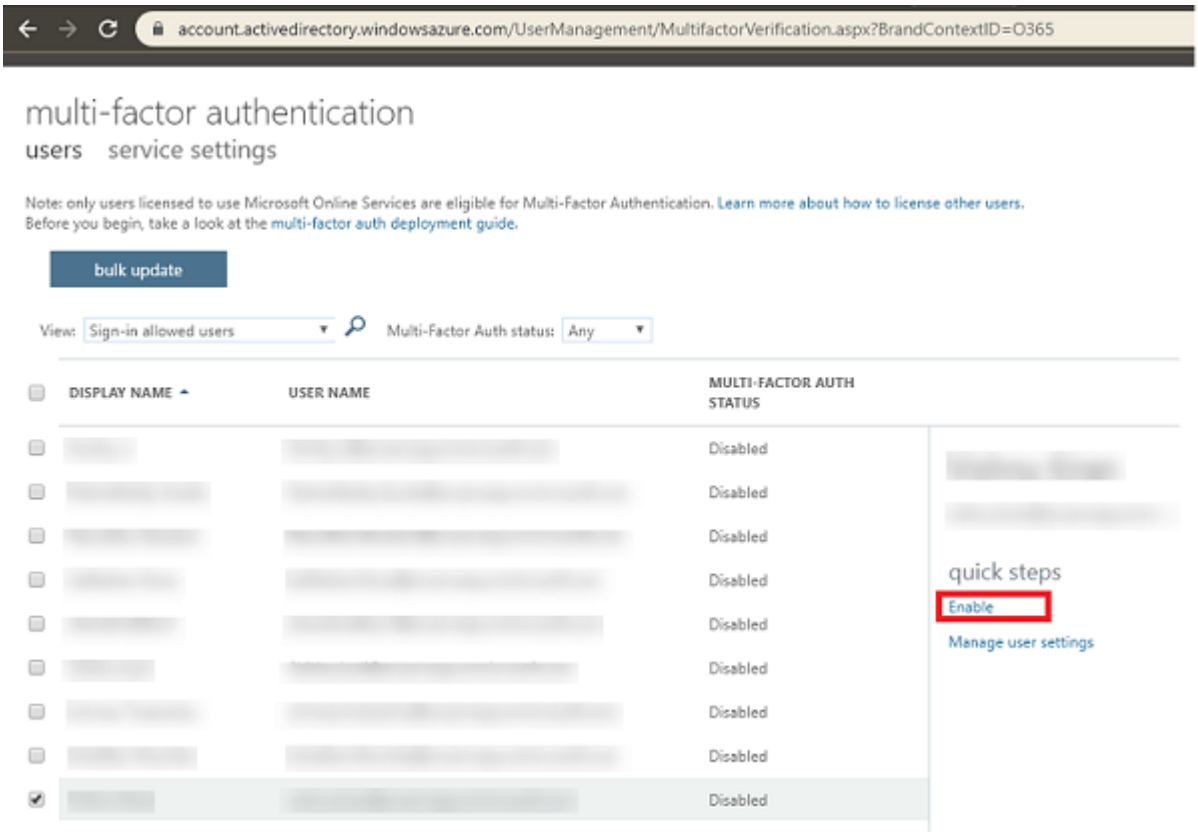
7. From the users screen, perform the following steps:

a. Select the check box for the users to enable MFA.



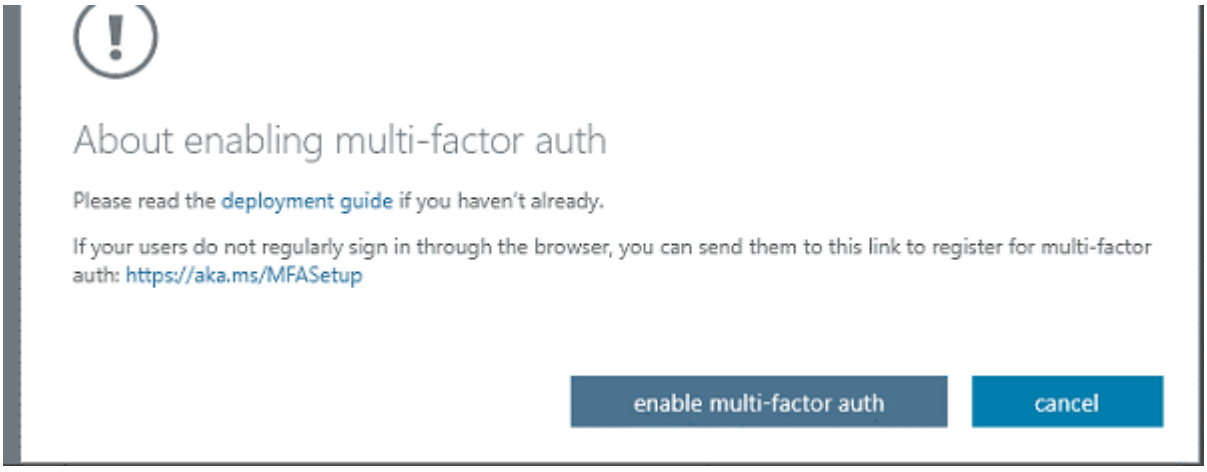
The right pane displays name of the user and under quick steps, you can view Enable and Manage user settings.

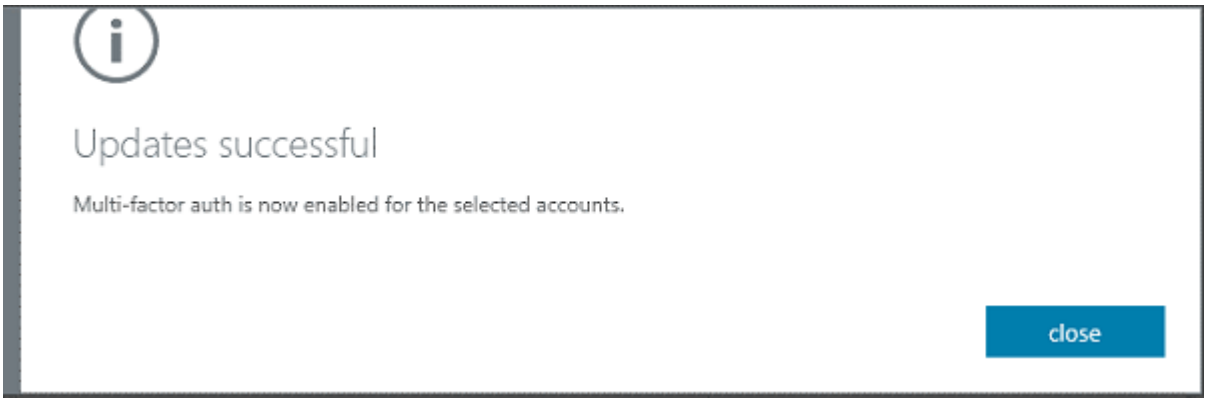
b. Click **Enable**.



A dialog box appears.

c. From the dialog box, click **enable multi-factor auth.**





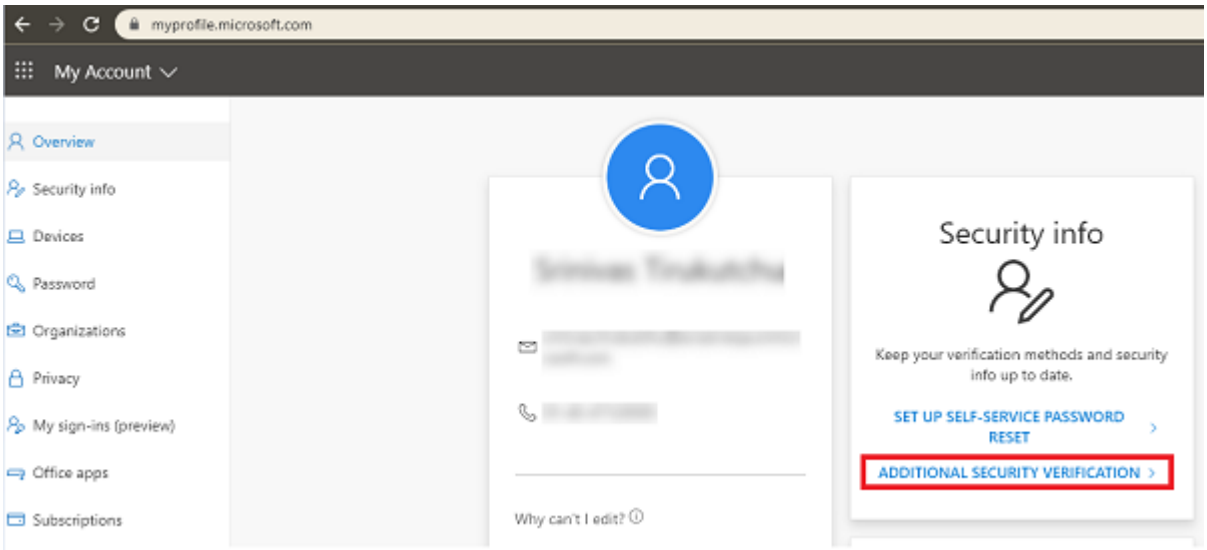
The backup service account to Set app password is enabled.

Creating app password for the backup service account

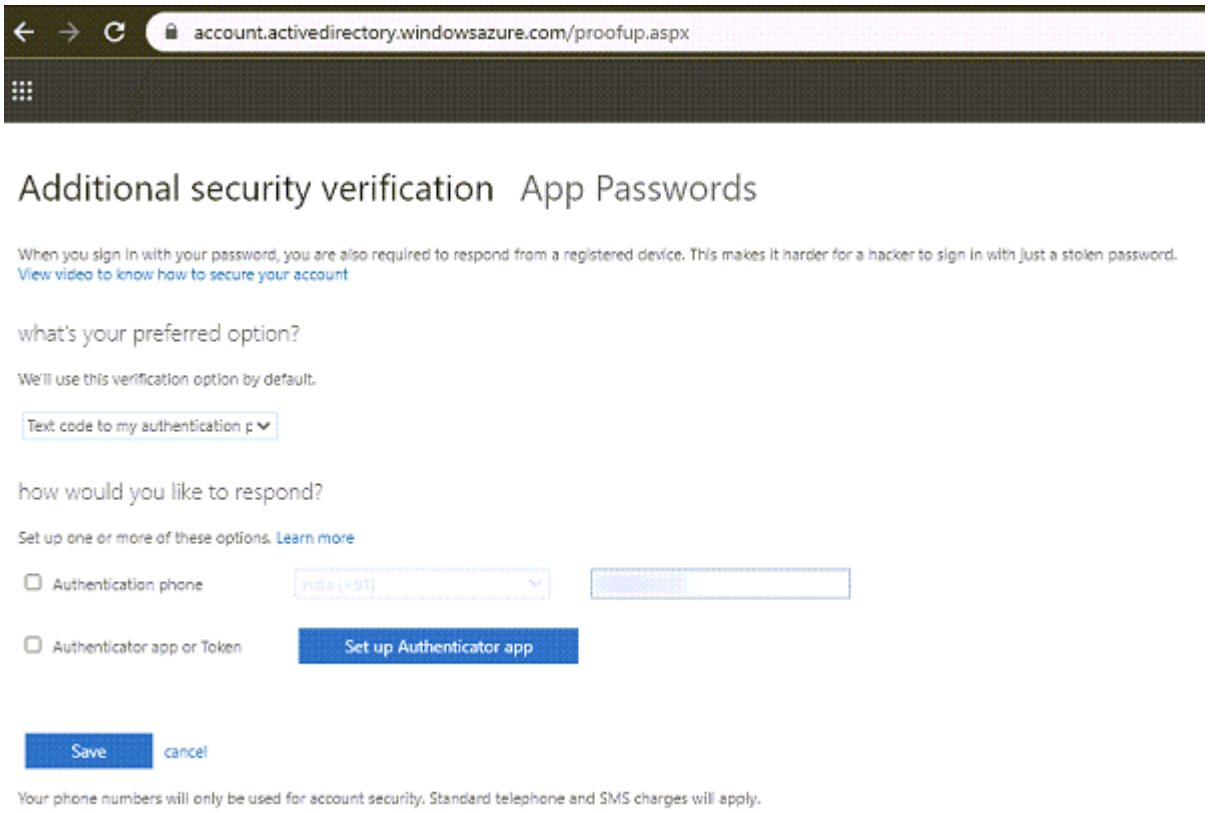
After enabling the backup service account to Set app password, the backup plan needs to use App Password for the backup service account. You need to create app password for the backup service account.

Follow these steps:

1. Sign into Office 365 (https://myprofile.microsoft.com/) with your work or school account and password.
2. Click **Additional Security Verification.**

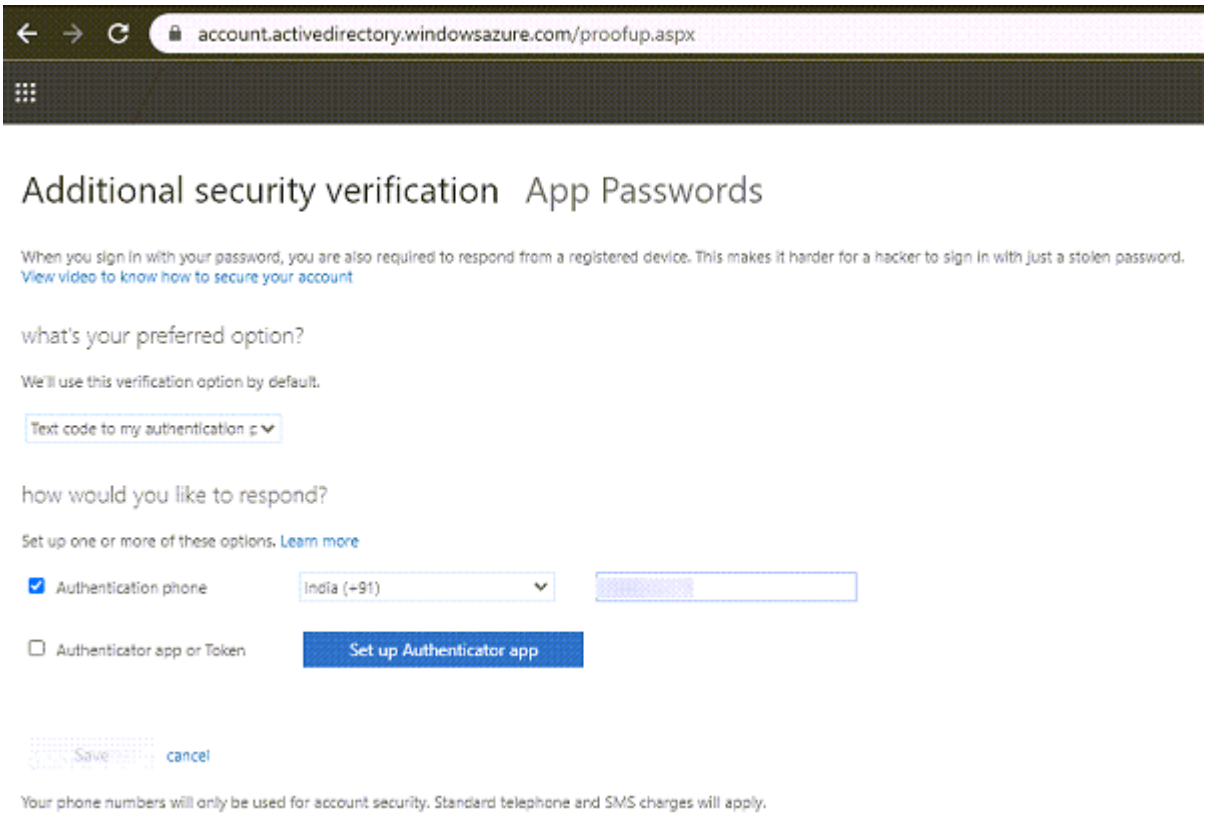


The Additional security verification screen appears.



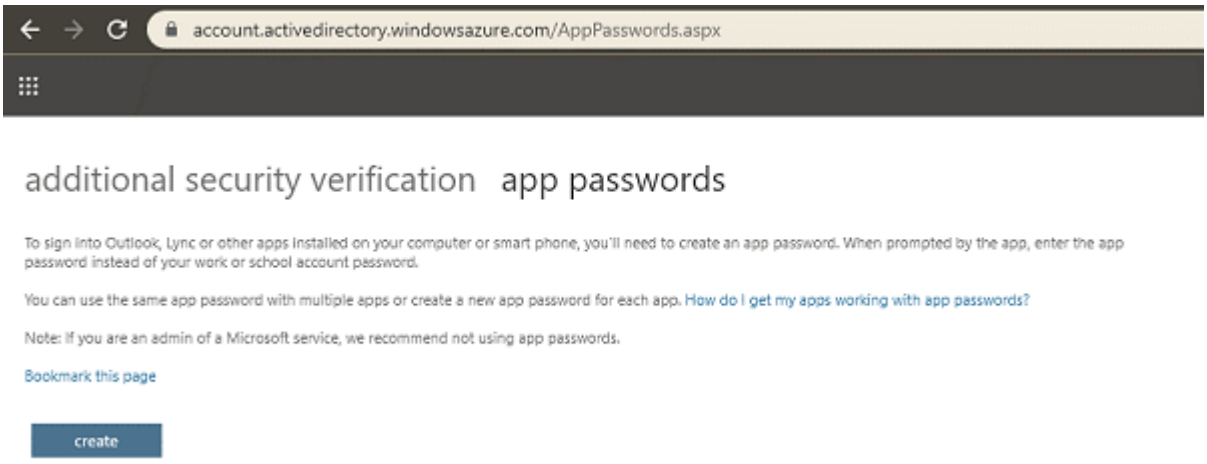
3. Perform the following steps on the Additional security verification screen:
    a. Select your authentication method, and then follow the prompts on the page.

       For example, if you select the authentication phone, you should select your country and enter the phone number. You can also select a method to get the verification code.



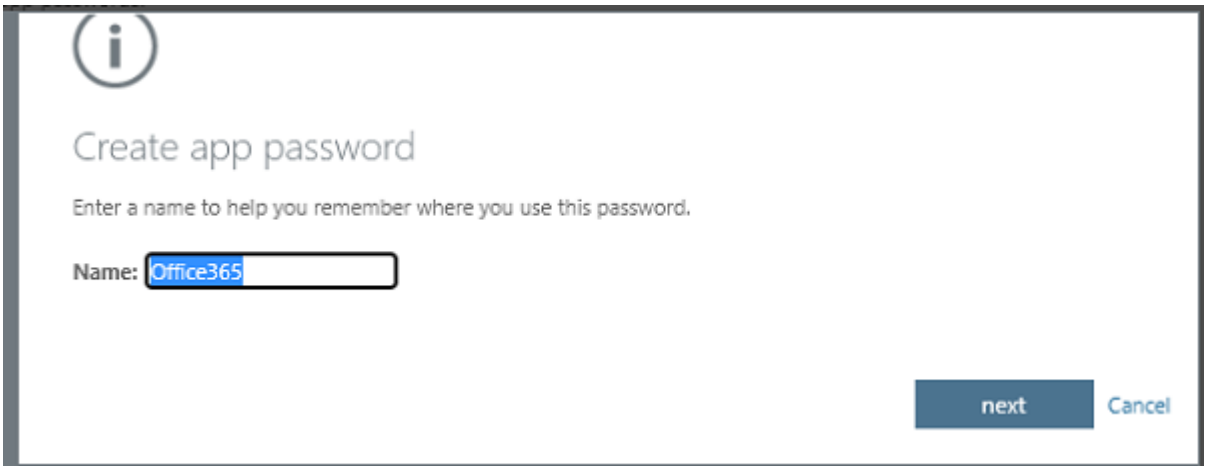
b. Click **App Passwords**.

The app passwords screen appears.



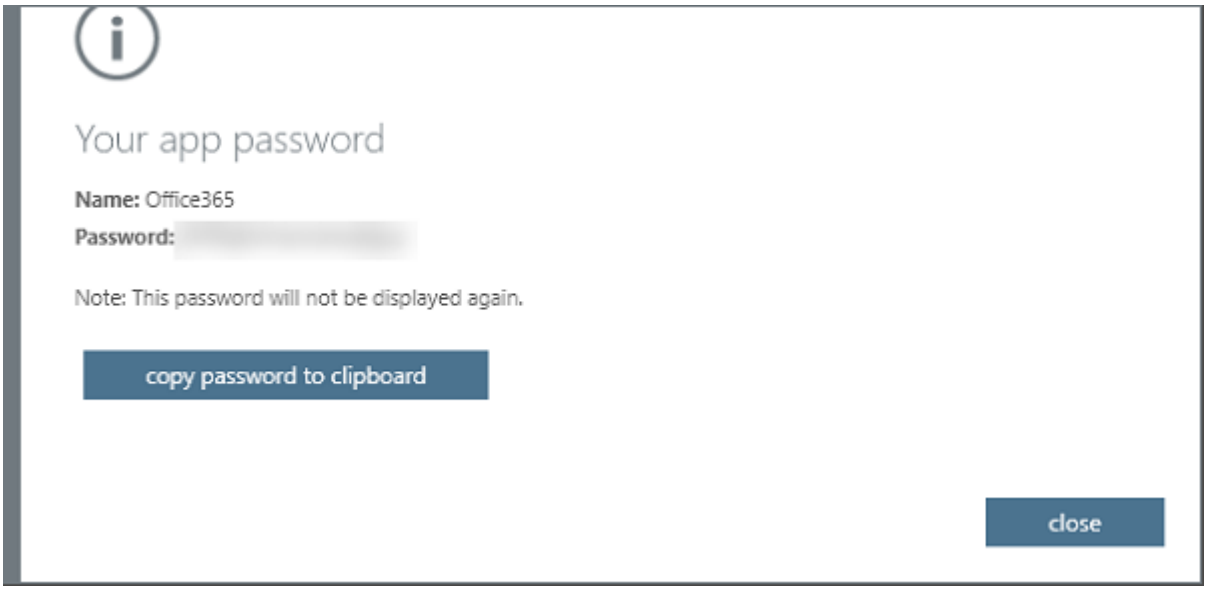
c. Click **create**.

The Create app password screen appears.

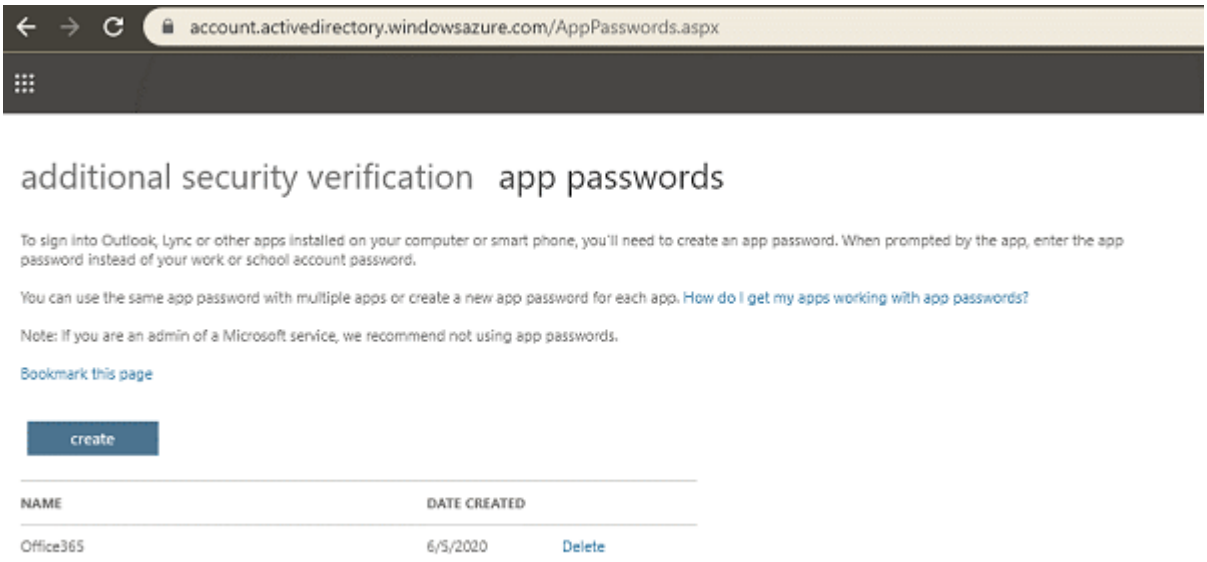


d. From the Create app password screen, enter a name, and then click **next**.

   You will receive an app password that you can use with Outlook, Apple Mail, and other Email options.



e. Select the **copy password to clipboard** option and the password is copied to your clipboard.



App password for the backup service account is created.

How to Create a SharePoint Online Backup Plan

The SharePoint Protection is used to backup and restore Microsoft SharePoint Online site and list item. The SharePoint Online is one of the major products in Microsoft Office 365. To protect your SharePoint content, you need to create a Plan.

What To Do Next?

1. Review the Prerequisites and Considerations
2. Create a SharePoint Online Backup Plan
3. Verify the Backup Plan
4. Configuration for Multi-Factor Authentication

Review the Prerequisites

Verify the following prerequisites before performing a backup and restore:

- You have the SharePoint Site Collection URL to backup.
- The backup account is a member of Site Collection Administrators groups or assigned with *SharePoint admin* Role.

  To add an account to Site Collection Administrators group, refer to the link.
- Review the Prerequisites before you create a plan using Modern Authentication.

Create a Personal Site in SharePoint

This section provides information about how to create a personal site in SharePoint.

Follow these steps:

1. Go to the Microsoft Login page, and then log into the Microsoft Office 365 Admin account.
2. Navigate to the **SharePoint admin center** page, and then on the left pane, click **More features**.
3. Under Classic site collections page, click **Open**.

   The Site Collections page appears.
4. On the Site Collections page, click **New** > **Private Site Collection**.

   The new site collection dialog appears.
5. On the new site collection dialog, provide the following details and then click **OK**.
   - **Title** – Type a name for the site.
   - **Web Site Address** – From the drop-down lists, select a domain name and a URL path as needed, and then type a URL name for the site.
   - **Template Selection** – In the Template Selection section, do the following:
     - From the **Select a language** drop-down list, select a language for the site.
     - Under **Select a template**, select the template for the site.

- **Time Zone** – Select the time zone based on the location of the site.
- **Administrator** – Type the name of the person you want to be the site administrator. You can also use the Check Names or Browse button to search for users and make them as site administrator.
- **Storage Quota** – Type the size (MBs) of the site you want to allocate.
- **Server Resource Quota** – Accept the default resource quota.

The Personal Site in SharePoint is created.

Create a SharePoint Online Backup Plan

A backup plan includes a **Backup: Office 365 SharePoint Online** task that performs a backup of SharePoint Online items (Lists, Libraries, Documents, and so on) and stores data either at a deduplication or non-deduplication data store. Additionally, this backup plan allows UDP to discover and specify SharePoint Online items for which a backup task is performed. Each task consists of parameters that define the source, destination, schedule, and other backup details.

Watch video and view how to create the SharePoint Online Backup plan.

Follow these steps:

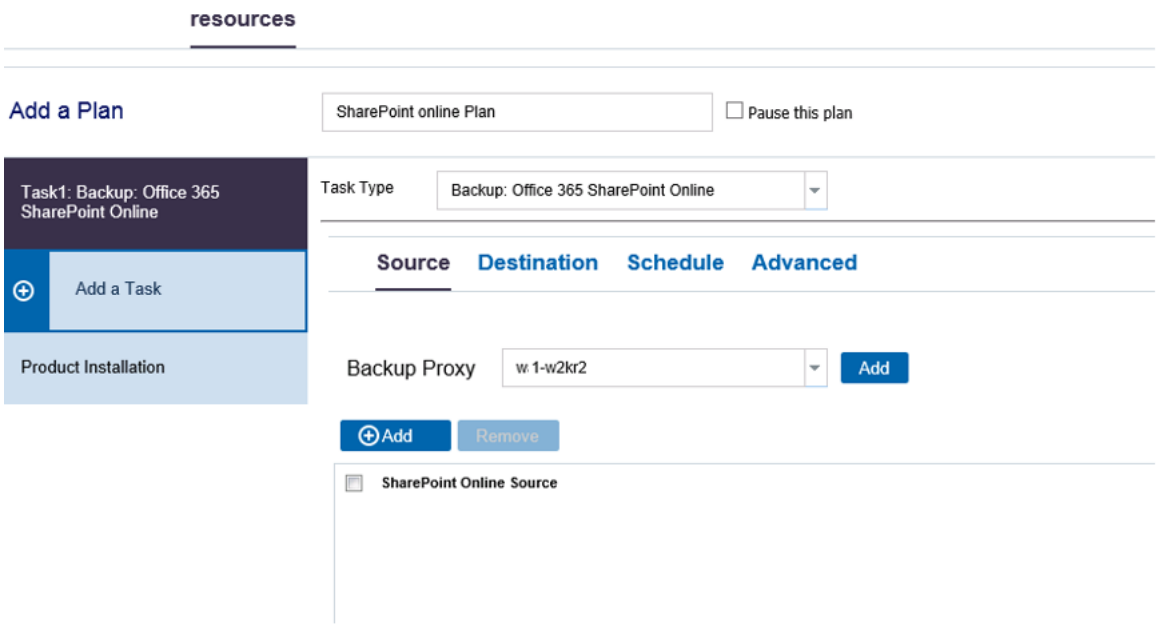
1. Click the **resources** tab on the Console.
2. From the left pane, navigate to **Plans**, and click **All Plans**.

   If you have created plans earlier, those plans are displayed on the center pane.
3. On the center pane, click **Add a Plan**.

   The **Add a Plan** dialog opens.
4. Enter a plan name.
5. (Optional) Select the **Pause this plan** check box.

   The plan does not run until you clear the check box to resume the plan.

   Note: If a plan is paused, then any in-progress job is not paused but all corresponding scheduled jobs associated with that plan are paused. However, you can manually run a job. For example, you can manually run backup job and replication job for a node even if the respective plan is paused. In such case, the following task to the on-demand (manual) job does not run. For example, there is a replication task after an on-demand backup job, the replication job does not run for the on-demand backup job. You need to manually run the replication job. When you resume the plan, the pending jobs do not resume immediately. After you resume the plan, the pending jobs run from the next scheduled time.
6. From the Task Type drop-down list, select **Backup: Office 365 SharePoint Online**.



Now specify the Source, Destination, Schedule, and Advanced details.

Specify the Source

The Source page lets you specify the SharePoint Online source nodes that you want to protect. You can select more than one SharePoint Online source nodes in a plan.

Note: You can save a plan without adding any source nodes, but the plan is not deployed unless you add a node.

Follow these steps:

1. Click the **Source** tab.
2. Select the **Backup Proxy** from the drop-down list.
3. Add SharePoint Online node using one of the following methods:
   - **Method 1:**
     a. Click **Add**, and then click **Select Source to Protect in Arcserve UDP**.

        The Add Nodes to Plan dialog is displayed.

        Note: Use this option only if you have already added SharePoint node.

     b. Select a node, and then click **Connect**.

        Note: To find the SharePoint Online nodes that you want to protect, use **Search**.
   - **Method 2:**

     Click **Add**, and then click **Add SharePoint Online Source**.

     The Add nodes to a plan dialog appears.

     Note:
     - Use this option to discover and specify online items for SharePoint Online Source.
     - Unlike other nodes, you cannot add the SharePoint Online node from the All Nodes page. You can add a SharePoint Online node only in a plan while creating or modifying a plan.
4. To apply authentication, do one of the following:

   Modern Authentication

   For modern authentication, you can run a backup using security certificate.

Use security certificate

Note: By default, the backup plan selects **Use security certificate** for modern authentication.

Follow these steps:

1. To authenticate the UDP application with security certificates, see How to Use Self-Signed Certificate.
2. Enter the following details:
   - Node Name
   - App ID
   - User Name
3. To apply the required roles and permissions, select the **Allow Arcserve UDP to set the required Azure roles and permissions for this application** check box, and then click **Connect**.

Notes:

- SharePoint Online Source node names cannot change after creating the node.
- You can use a single account to protect multiple SharePoint Online nodes.
- Fill in the app password if Multi-Factor Authentication is enabled and the tenant is set to use Basic Authentication.

Basic Authentication

For Basic Authentication, to proceed with the backup, provide the following information, and then click **Connect**:

- Node name
- User ID
- Password

5. To protect SharePoint Online items, do the following, and then click **Save**.
   - From the Site Type drop-down list, select the type of site you want to protect such as SharePoint Sites or Personal Sites.
   - For Protect All Sites/Subsites, perform the following:

     Notes:

     - This option is automatically selected by default when the Select SharePoint Online Content screen appears.

- To protect all items in the selected site type, select the **Protect All Sites/Subsites** check box.
- To protect specific items in the selected site type, clear the **Protect All Sites/Subsites** check box, and then select the specific items that you want to protect.

- While configuring the existing SharePoint Online plan, if you have selected sites manually for protection and now want to select the **Protect All Sites/Subsites** check box, follow these steps:
  a. Modify the existing SharePoint Online plan.
  b. Clear the SharePoint sites that are manually selected for protection.
  c. Select the **Protect All Sites/Subsites** check box, and then save the plan.

- To exclude specific items from being protected, select the **Exclude selected sites** check box, and then select specific items in the selected site type.

The SharePoint Online sources that you want to protect are added to the plan.

Specify the Destination

The destination is a location where you store the backup data. You must at least specify the destination to save the plan.

Follow these steps:

1. Click the **Destination** tab.

   The **Arcserve UDP Recovery Point Server** option. **Arcserve UDP Recovery Point Server** specifies that the backup destination is a recovery point server.

   You cannot store data as recovery sets.

2. Perform the following steps:

   a. Select a recovery point server.

   b. Select a non-deduplication or deduplication data store.

      The list displays all the data stores created on the specified recovery point server.

   c. Provide a session password.

      The session password is optional when the backup destination is an unencrypted RPS data store.

   d. Confirm the session password.

The destination is specified.

Specify the Schedule

The Schedule page lets you define a schedule for Backup, Merge, and Throttle functions to repeat at specific intervals. After you define a schedule, the jobs run automatically per the schedule. You can add multiple schedules and provide retention settings.

A Backup Schedule refers to regular schedule that is repeated multiple times a day based on the number of hours or minutes you select. Besides the regular schedule, a backup schedule also provides options to add daily, weekly, and monthly schedules.

Note: For more information on scheduling and retention settings, see Understanding Advanced Scheduling and Retention.

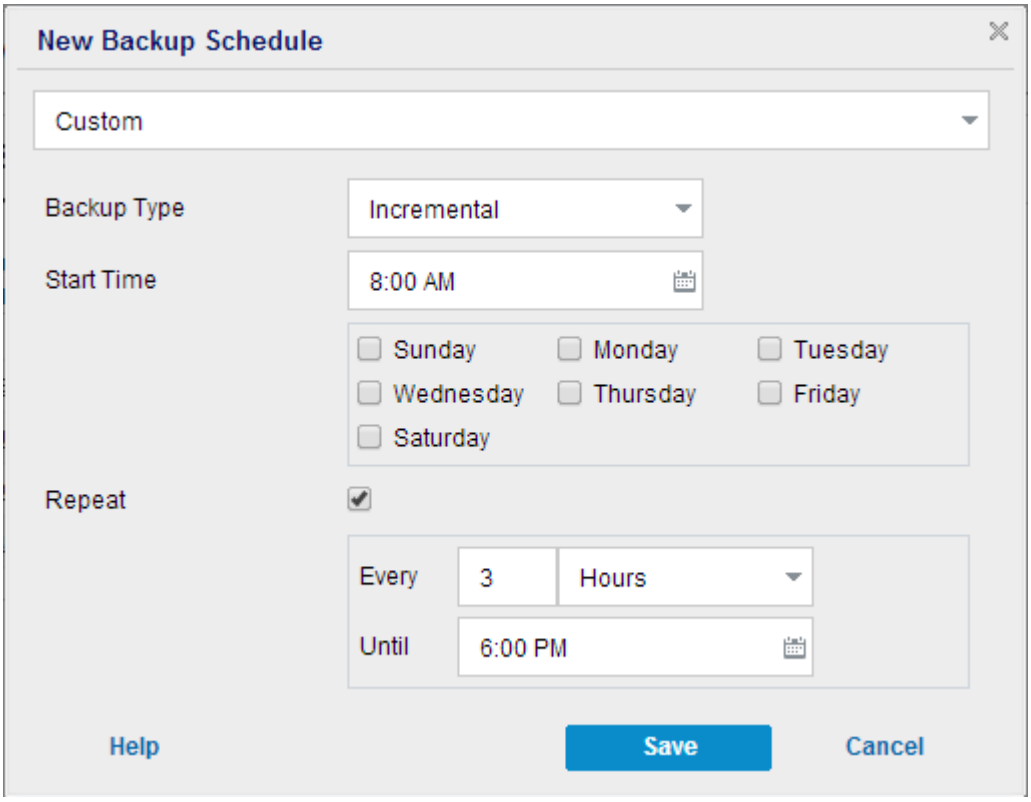
Follow these steps:

1. Add backup, merge, and throttle schedules.

   Add Backup Schedule

   a. Click **Add** and select **Add Backup Schedule**.

      The **New Backup Schedule** dialog opens.



   b. Select one of the following options:

      Custom

Specifies the backup schedule that repeats multiple times a day.

Daily

Specifies the backup schedule that occurs once a day. By default, all the days of the week are selected for Daily backup. If you do not want to run the backup job on a specific day, clear the check box for that day of the week.

Weekly

Specifies the backup schedule that occurs once a week.

Monthly

Specifies the backup schedule that occurs once a month.

c. Select the backup type.

Full

Determines the backup schedule for Full Backups. As scheduled, Arcserve UDP performs a full backup of all used blocks from the source machine. A full backup typically consumes time depending on the backup size.

Incremental

Determines the backup schedule for Incremental Backups.

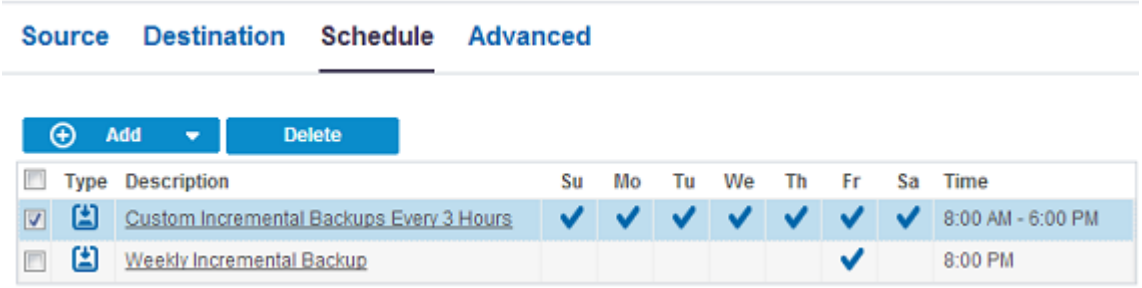
As scheduled, Arcserve UDP incrementally backs up only those blocks that have changed after the last successful backup. The advantages of Incremental Backups are that the backup is a fast backup and it produces a small backup image. This is the most optimal way to perform a backup.

d. Specify the backup start time.

e. (Optional) Select the **Repeat** check box and specify the repeat schedule.

f. Click **Save**.

The Backup Schedule is specified and appears on the **Schedule** page.



Add Merge Schedule

a. Click **Add** and select **Add Merge Schedule**.

   The **Add New Merge Schedule** dialog opens.

b. Specify the start time for the merge job.

c. Specify **Until** to provide an end time for the merge job.

d. Click **Save**.

The Merge Schedule is specified and appears on the **Schedule** page.

Add Throttle Schedule

a. Click **Add** and select **Add Throttle Schedule**.

   The **Add New Throttle Schedule** dialog opens.

b. Specify the throughput limit in MB per minutes unit.

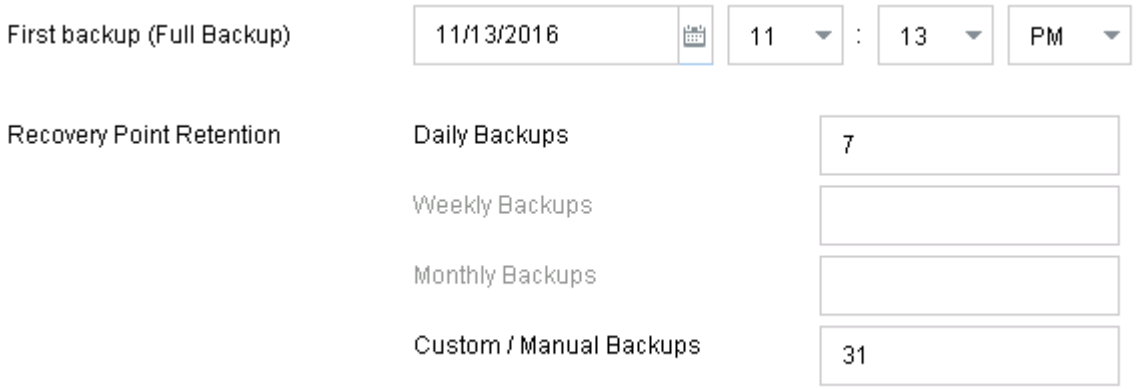
c. Specify the start time for the backup throughput job.

d. Specify **Until** to provide an end time for the throughput job.

e. Click **Save**.

The Throttle Schedule is specified and displayed on the **Schedule** page.

2. Specify the start time for the scheduled backup.



3. Specify the recovery points retention settings for Custom, Daily, Weekly, and Monthly schedule.

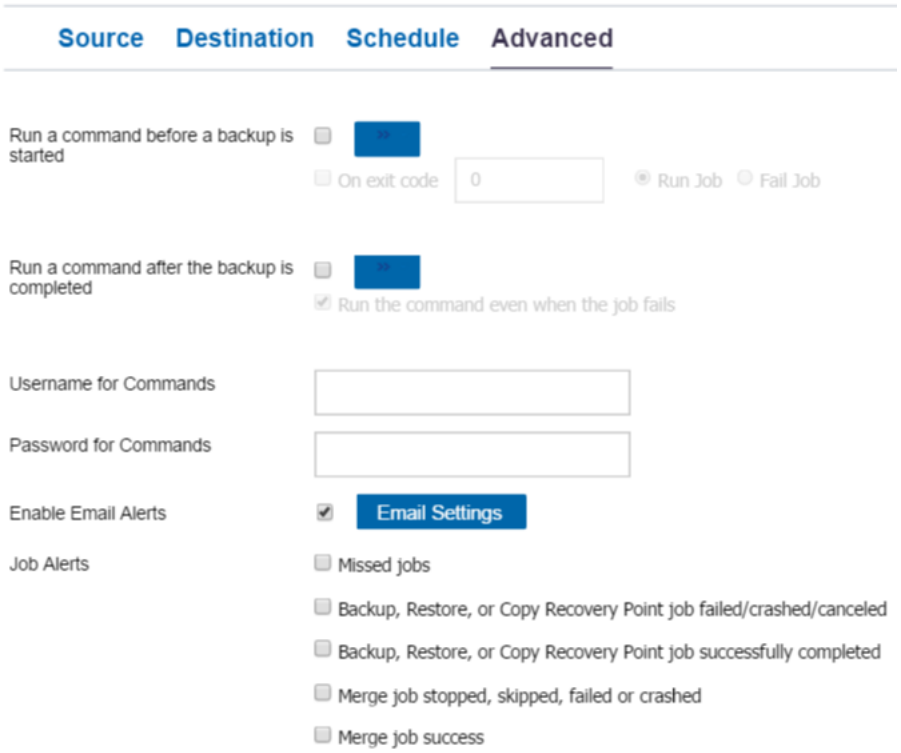
   The options are enabled if you have added the corresponding backup schedule. If you modify the retention settings on this page, the changes are reflected on the **Backup Schedule** dialog.

   The schedule is specified.

Specify the Advanced Settings

The **Advanced** tab lets you specify some advanced settings for the backup job. The advanced settings include providing the location of any scripts, and email settings.

The following image displays the **Advanced** tab:



Follow these steps:

1. Specify the following details.

   Run a command before a backup is started

   Lets you run a script before the backup job starts. Specify the path where the script is stored inside the proxy node. Click **On exit code** and specify the exit code for **Run Job** or **Fail Job**. **Run Job** indicates that the backup job continues when the script returns the exit code. **Fail Job** indicates that the backup job stops when the script returns the exit code.

   Run a command after a backup is completed

   Lets you run a script after the backup job is completed. Specify the complete path where the script is stored.

   Run a command even when the job fails

If this check box is selected, the script specified in **Run a command after a backup is completed** is executed even when the backup job fails. Otherwise, that script is executed only when backup job completes successfully.

Username for Commands

Lets you specify the username to run the commands.

Password for Commands

Lets you specify the password to run the commands.

Enable Email Alerts

Lets you enable email alerts. You can configure email settings and specify the types of alerts that you want to receive in an email. When you select this option, the following options are enabled for your selection.

Email Settings

Lets you configure the email settings. Click **Email Settings** and configure the email server and proxy server details. For more information about how to configure Email Settings, refer to Email and Alert Configuration.

Job Alerts

Lets you select the types of job alert emails that you want to receive.

2. Click **Save**.

Note: When you select a node as a backup source or backup proxy, Arcserve UDP checks whether the agent is installed on proxy node and if it is the latest version. Arcserve UDP then displays a verification dialog that lists all the nodes that either have an outdated version of the agent or does not have the agent installed. To install/upgrade the agent on these nodes, select the installation method and click **Save**.

The changes are saved and a green checkmark appears next to the task name. The plan page closes.

Note: If you have to add another task, you must select the plan from the **resources** tab and modify the plan. To modify the plan, click the plan from the center pane. The plan opens and you can modify it. You may add the **Copy Recovery Point**, **Copy to Tape**, **Replicate**, and **Replicate from a remote RPS** tasks as follow up tasks.

The plan is automatically deployed to the proxy server node.

The SharePoint Online backup plan for the proxy server is created. The backup runs per the schedule that you have configured on the **Schedule** tab. You can also perform a manual backup at any time.

(Optional) Perform a Manual Backup

Typically, based on the schedule settings you applied, the backups perform automatically. In addition to the scheduled backup, a manual backup provides you the option to back up your nodes on a need basis. For example, if you have a repeat schedule for Full and Incremental backups, and you want to make major changes to your machine, you can perform an immediate manual backup without waiting for the next scheduled backup.

To perform a manual backup of SharePoint Online nodes, follow these steps:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and then click **All Nodes**.

   The SharePoint Online nodes are displayed in the center pane.
3. Select the SharePoint Online nodes (for example, Mailbox@<organizationname.com>) that you want to backup and make sure a plan is assigned to it. The node name is the account that is used when adding the SharePoint Online node and connecting it.
4. Right-click the node, and then select **Backup Now**.

   The Run a backup now dialog opens.
5. Select the type of backup that you want to perform, and then click **OK**.

   The Progress Information dialog appears.
6. Wait until the backup job completes, and then click **OK**.
7. To monitor the backup jobs, navigate to **jobs** > **Jobs in Progress**.

   Wait until the backup job is finished.

   The manual backup is successfully performed.

To perform a manual backup of SharePoint Online plan, follow these steps:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Plans**, and then click **All Plans**.

   The SharePoint Online backup plans are displayed in the center pane.
3. Select the plan that you want to backup and make sure a plan is assigned to it.
4. Right-click the plan, and then select **Backup Now**.

   The Run a backup now dialog opens.
5. Select the type of backup that you want to perform, and then click **OK**.

   The Progress Information dialog appears.

6. Wait until the backup job completes, and then click **OK**.
7. To monitor the backup jobs, navigate to **jobs** > **Jobs in Progress**.

   Wait until the backup job is finished.

   The manual backup is successfully performed.

Verify the Backup

To verify your backup job, confirm that you have successfully created the backup plan. After you verify that the plan is created successfully, confirm whether the backup job is running as scheduled. You can verify the status of backup jobs from the **Jobs** tab.

Follow these steps to verify plans:

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and then click **All Nodes**.

   A list of all nodes is displayed on the center pane.
3. Verify that plans are mapped with nodes.

Follow one of these methods to verify backup jobs:

Method 1

1. Click the **jobs** tab.
2. From the left pane, click **All Jobs Completed**.

   The status of each job is listed on the center pane.
3. Verify that the backup job is successful.

   The backup job is verified.

Method 2

1. Click the **jobs** tab.
2. From the left pane, click **Jobs Successfully Completed**.

   The list of backup jobs completed appear on the center pane.
3. Click the job you want to verify.

   The right pane refreshes.
4. Under Job Details, click the **View Logs** hyperlink.
5. Verify that the backup job is successful.

   The backup job is verified.

Configuring for Multi-Factor Authentication

When an organization has multi-factor authentication (MFA) enabled for the users, the office 365 backup plan needs to be configured using the App Password for the backup service account.

Perform the following steps to configure Arcserve UDP to Support multi-factor authentication:

1. Enable the backup service account to Set app password
2. Create app password for the backup service account

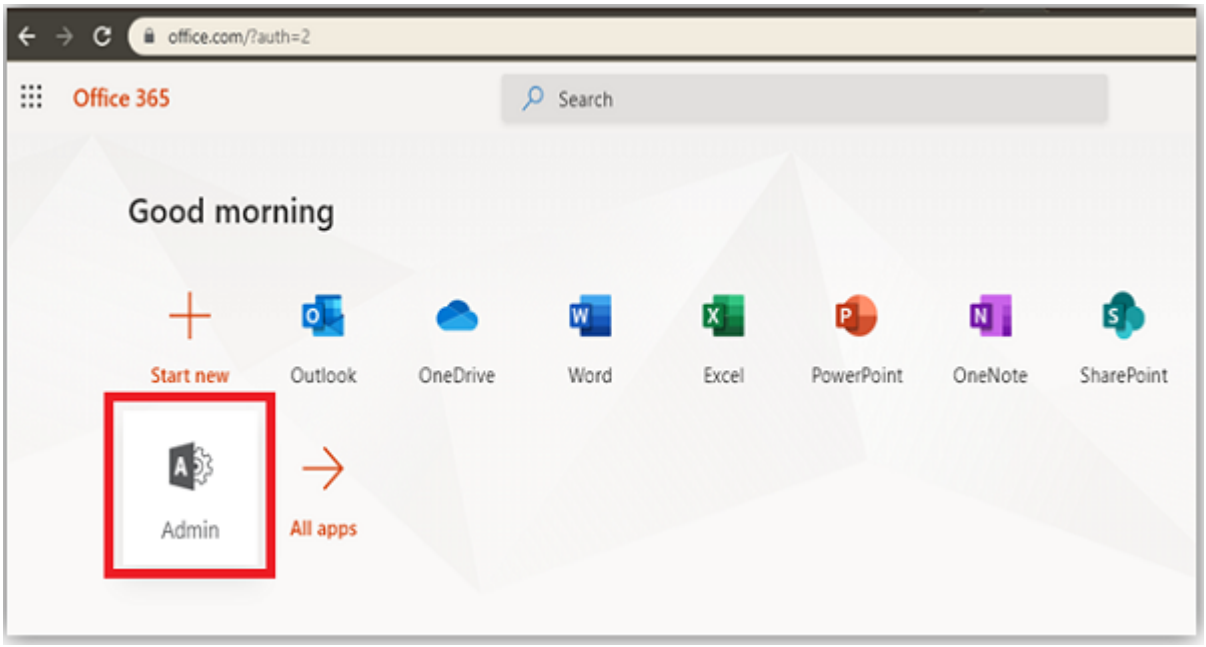
Note: MFA authentication (App password) is currently supported for O365 Exchange Online and SharePoint Online backups only.

Enable the backup service account to Set app password

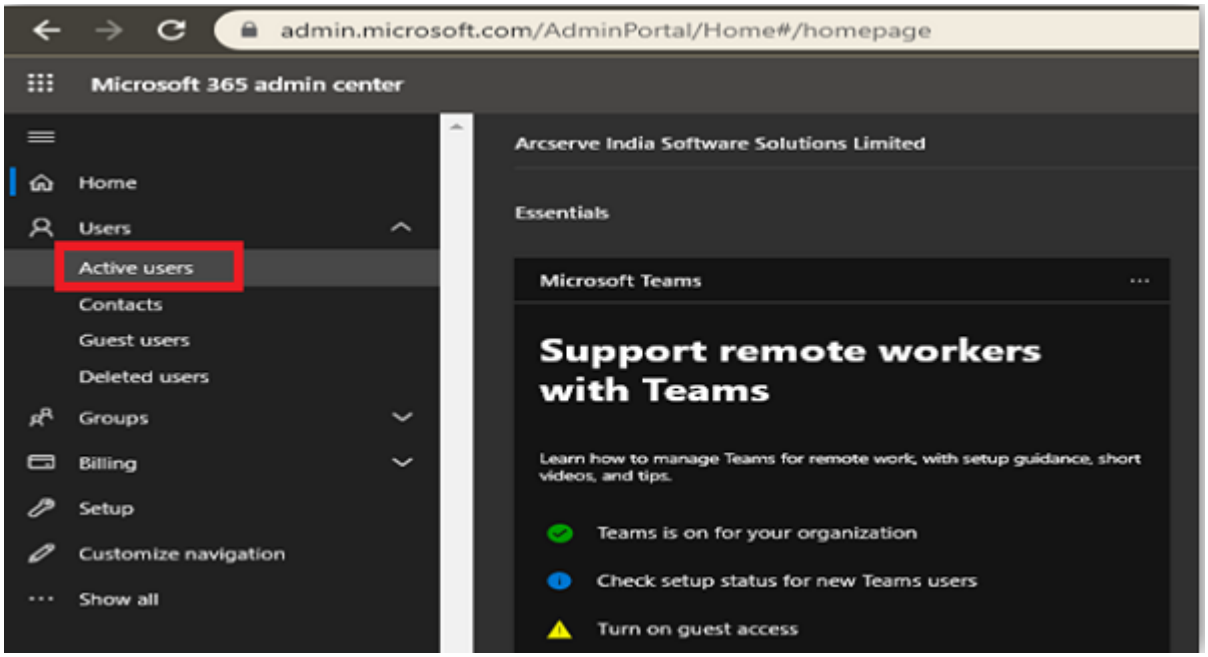
To configure, the first step is to enable the backup service account to Set app password.

Follow these steps:

1. Sign into Microsoft Office 365 using credentials of an administrator account and click the **Admin** icon.

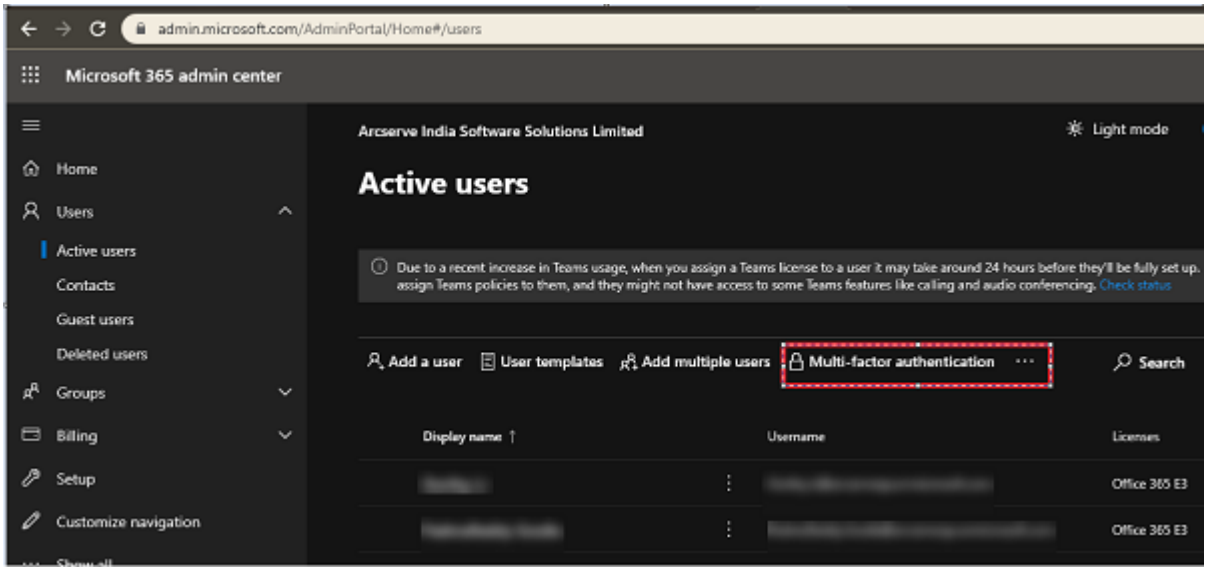


2. From the Microsoft 365 admin center screen, navigate to **Users** > **Active users**.



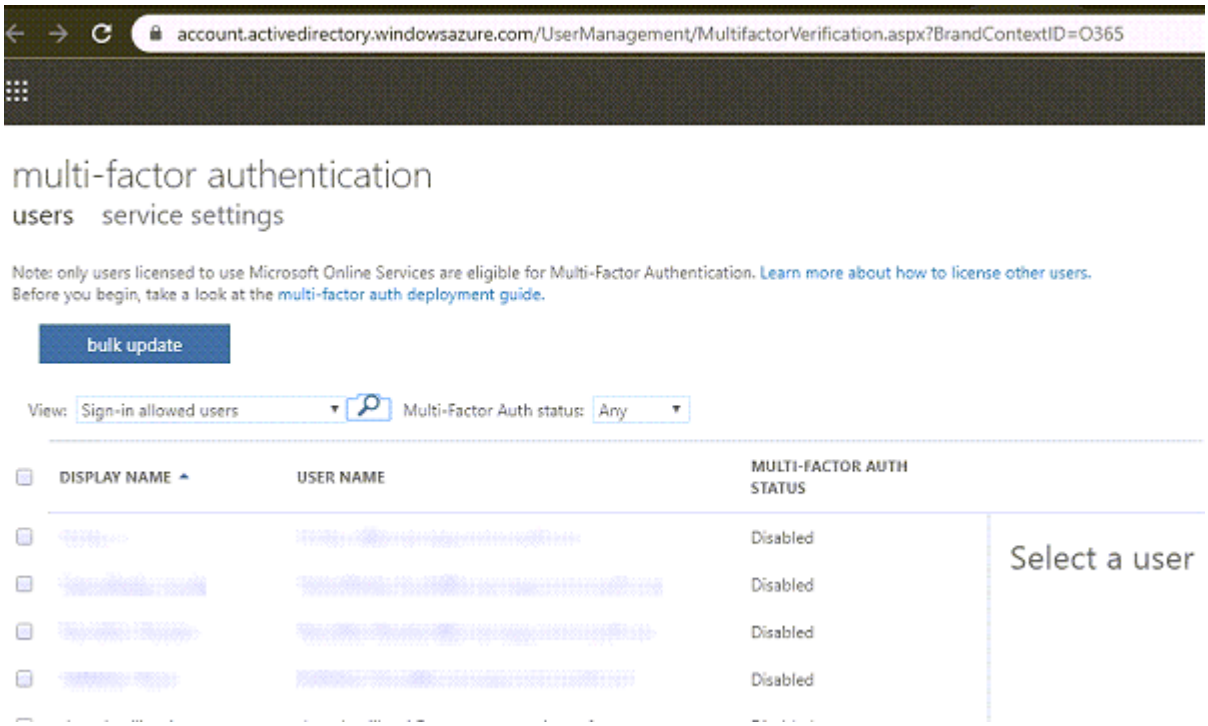
3. From the Active users screen, click the **Multi-factor authentication** option.

   Important! If you do not see the **(...)** option, then you are not a global admin for your subscription.

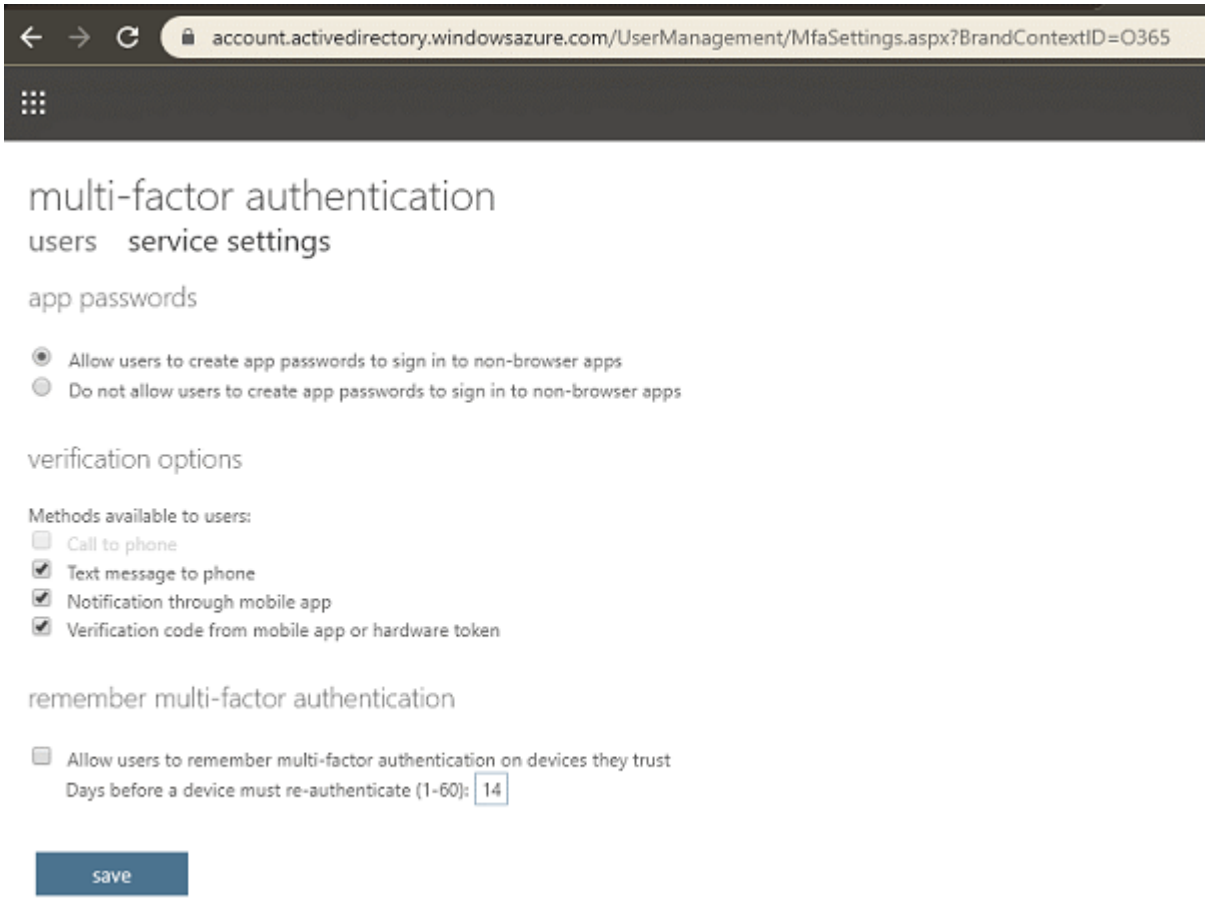


   You are led to the Setup Azure multi-factor authentication.

   Steps 4 and 5 only need to be set once. Then, skip to step 6



4. From the multi-factor authentication screen, click **service settings**.
5. From app passwords, select the checkbox of **Allow users to create app passwords to sign in to non-browser apps**.



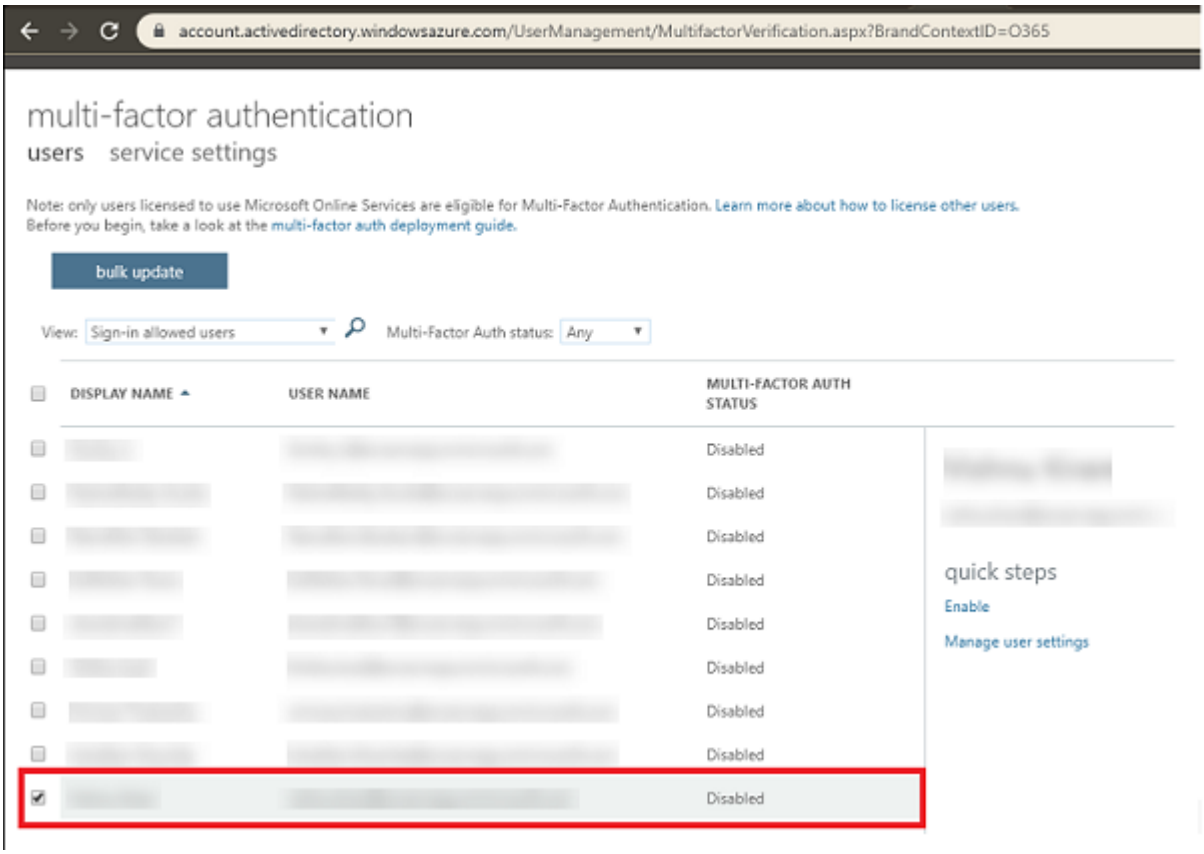
You can then use client Office apps after you create a new password.

6. Click **save** and close the window.

   You return to the users screen.

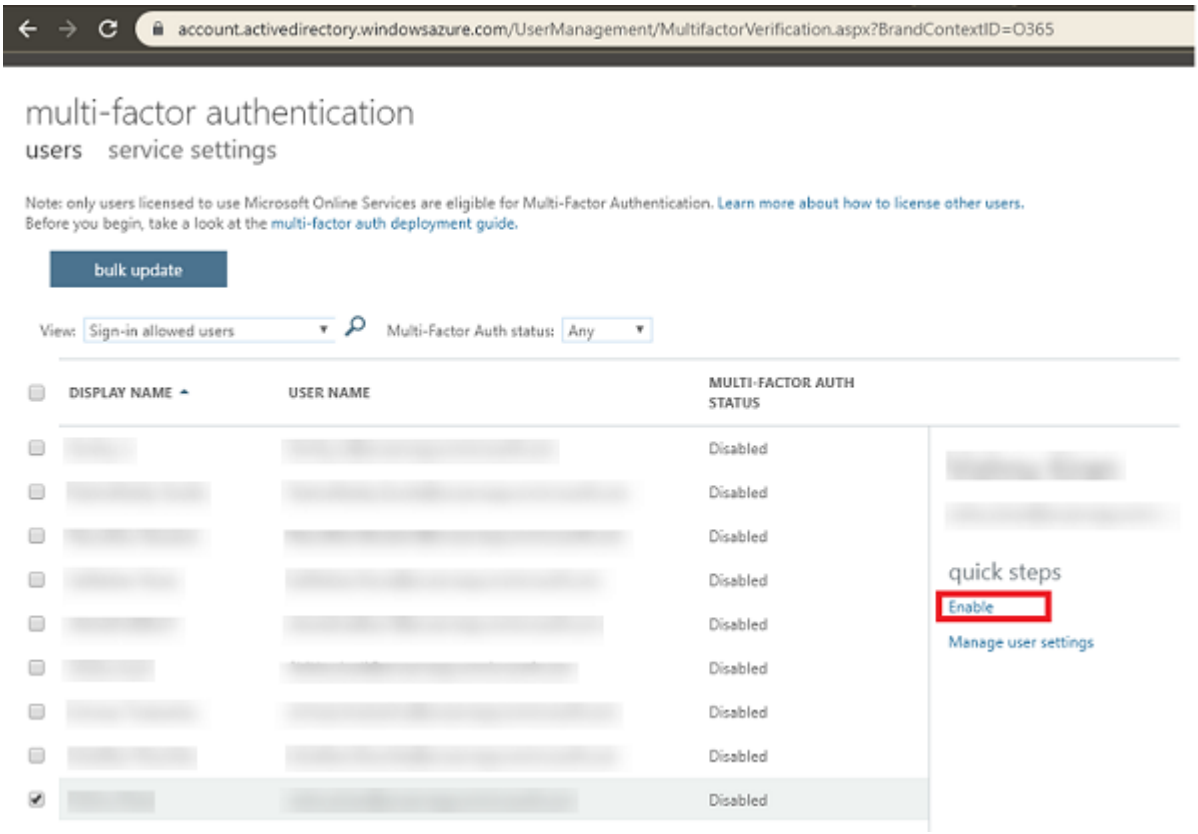
7. From the users screen, perform the following steps:

a. Select the check box for the users to enable MFA.



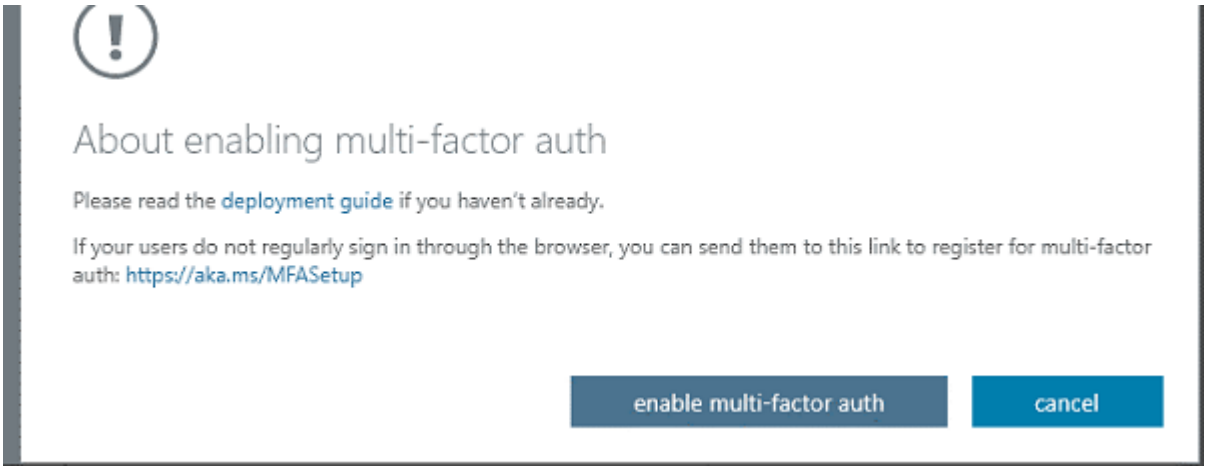
The right pane displays name of the user and under quick steps, you can view Enable and Manage user settings.

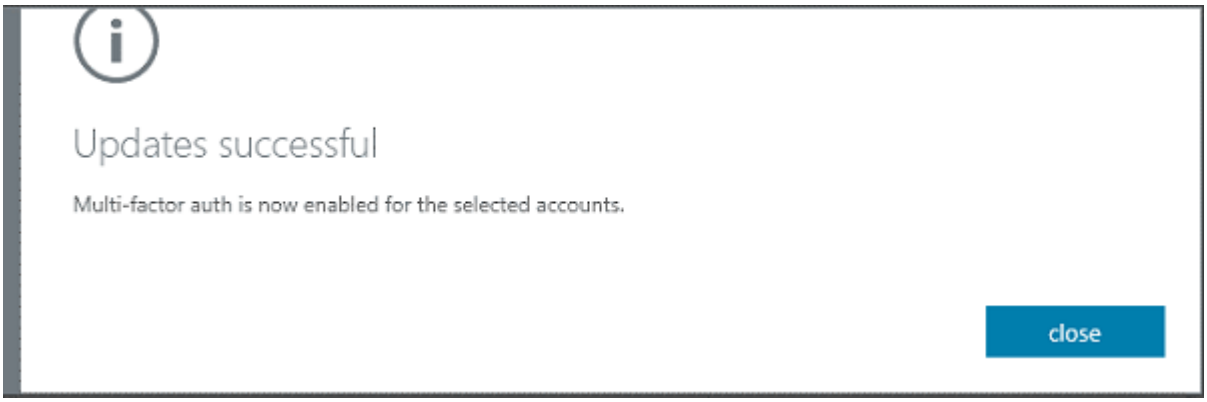
b. Click **Enable**.



A dialog box appears.

c. From the dialog box, click **enable multi-factor auth.**





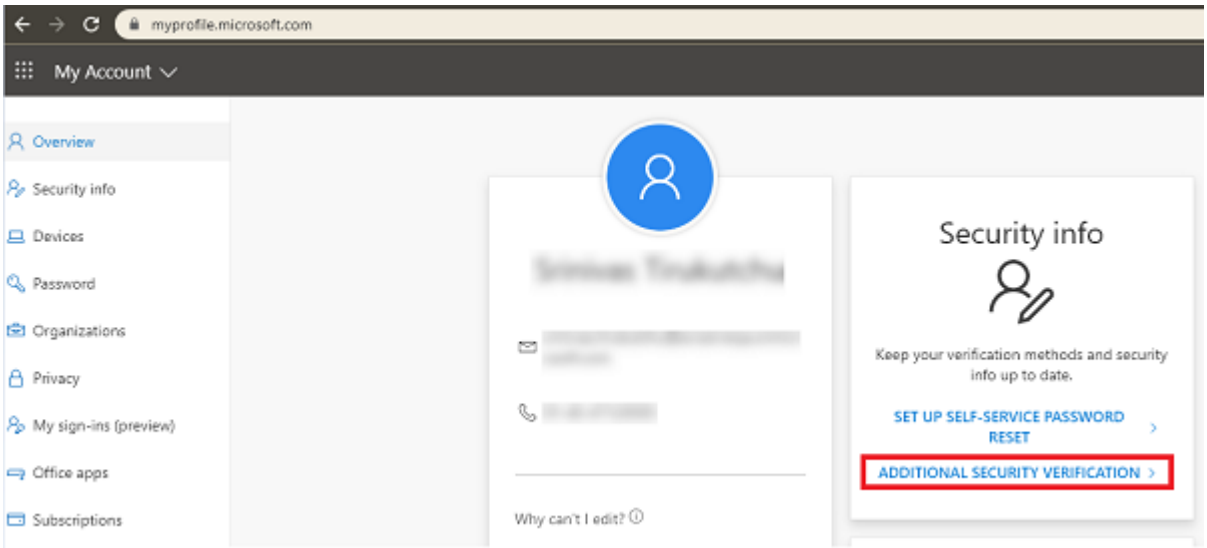
The backup service account to Set app password is enabled.

Creating app password for the backup service account

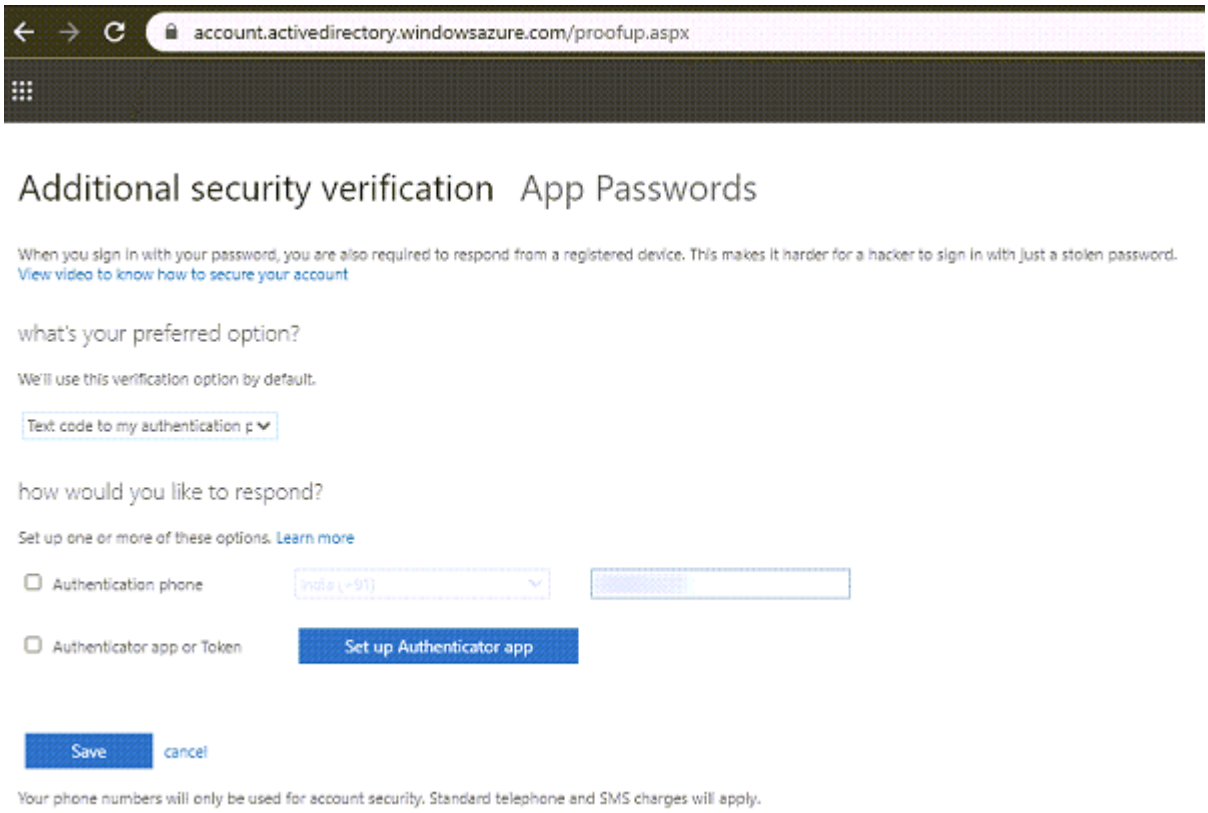
After enabling the backup service account to Set app password, the backup plan needs to use App Password for the backup service account. You need to create app password for the backup service account.

Follow these steps:

1. Sign into Office 365 (https://myprofile.microsoft.com/) with your work or school account and password.
2. Click **Additional Security Verification.**

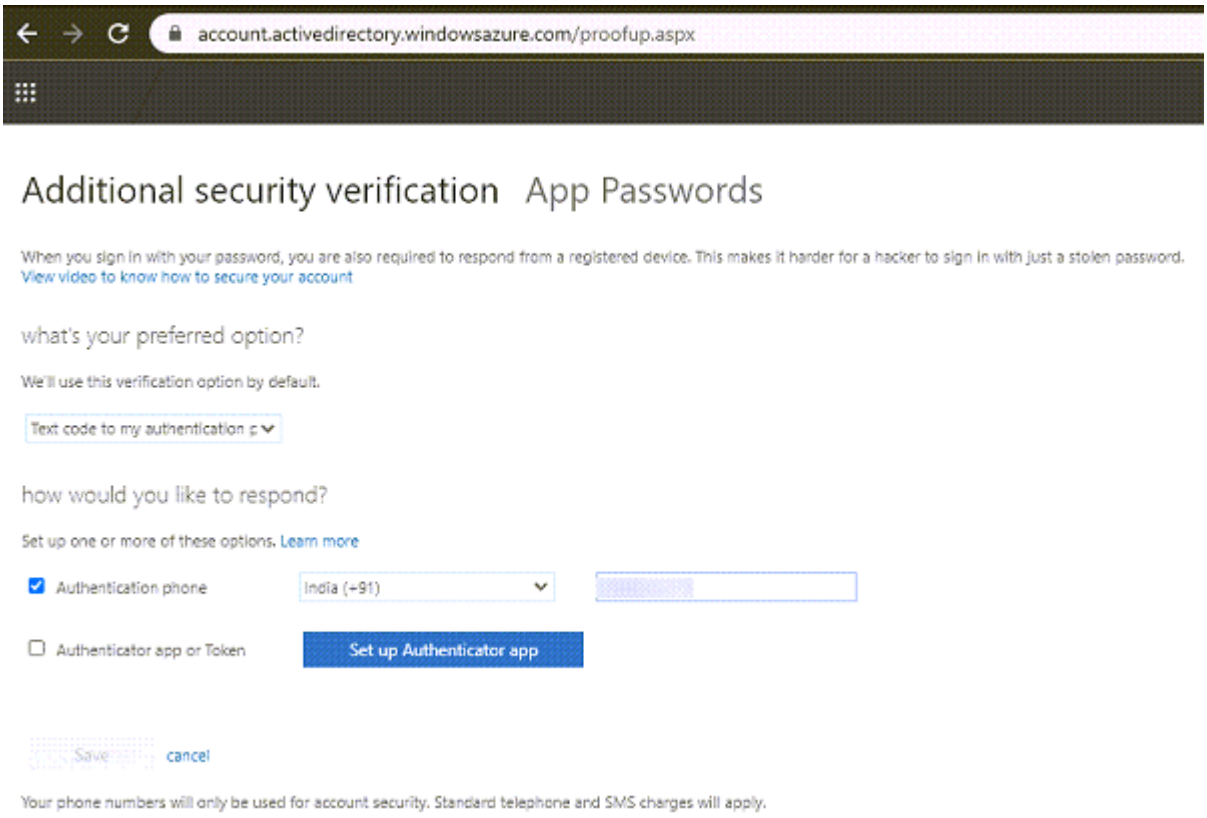


The Additional security verification screen appears.



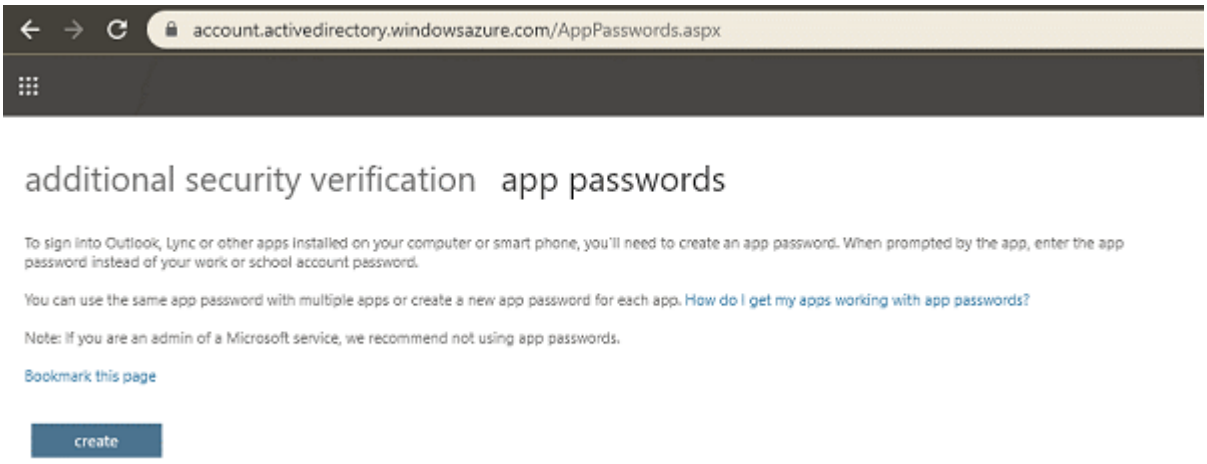
3. Perform the following steps on the Additional security verification screen:
    a. Select your authentication method, and then follow the prompts on the page.

       For example, if you select the authentication phone, you should select your country and enter the phone number. You can also select a method to get the verification code.



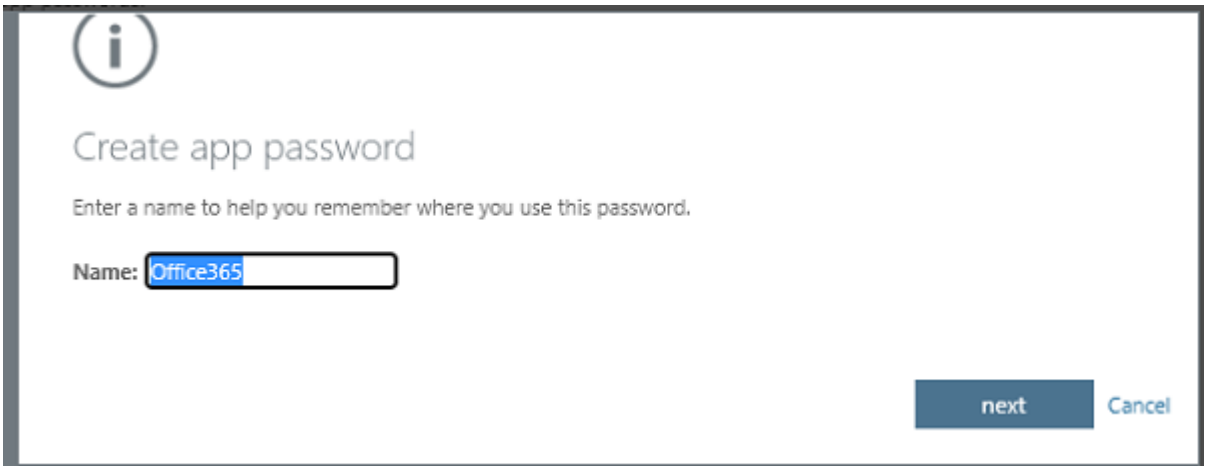
b. Click **App Passwords**.

The app passwords screen appears.



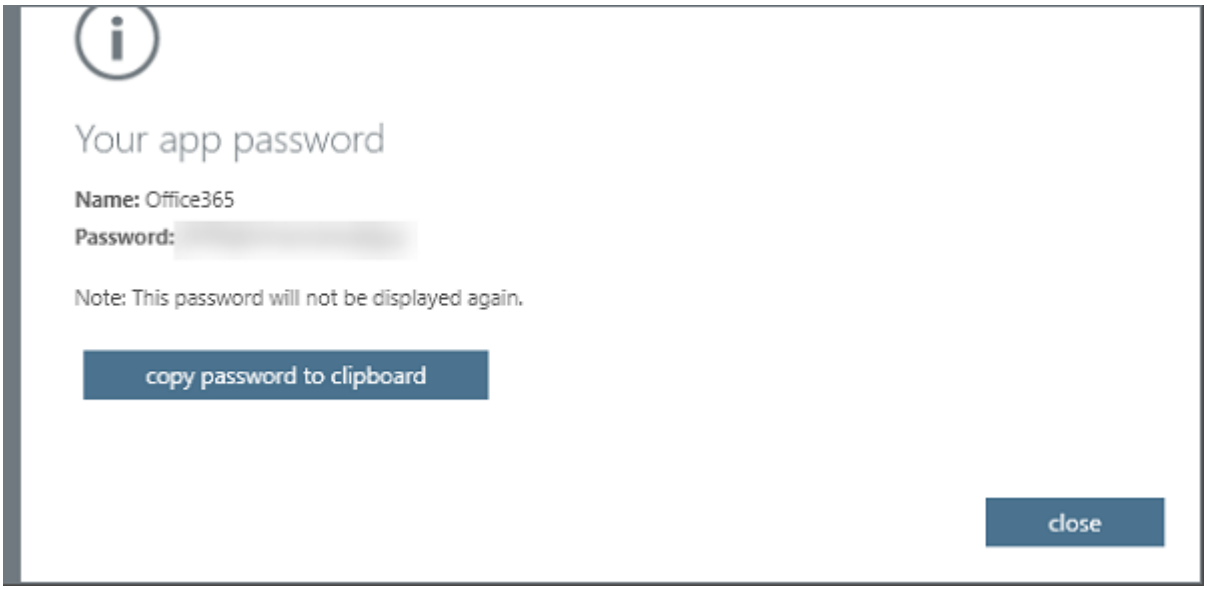
c. Click **create**.

The Create app password screen appears.

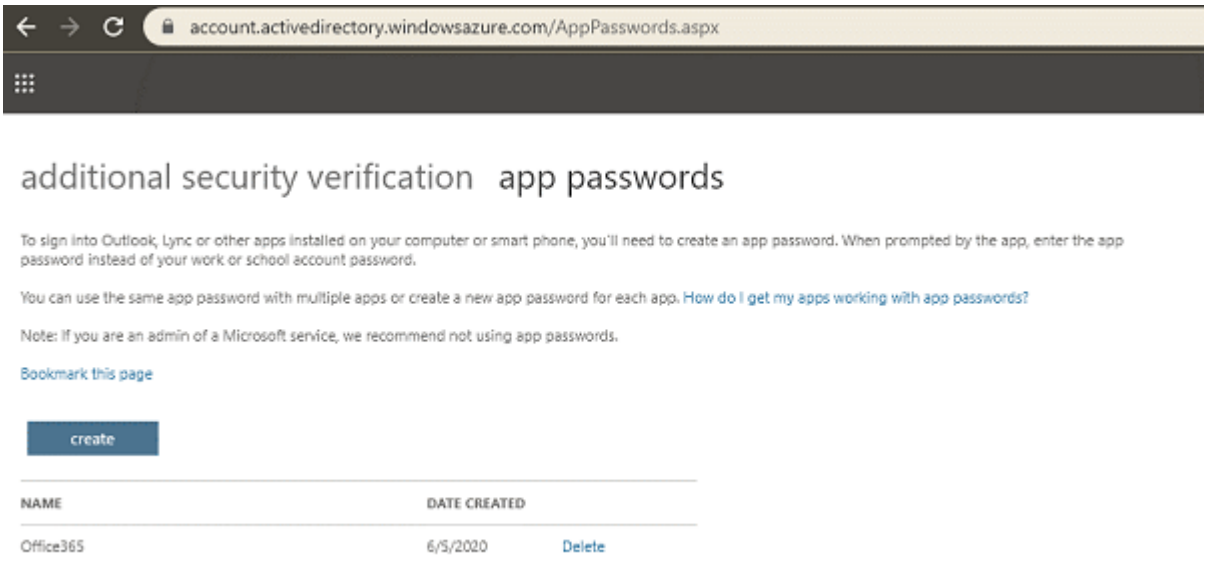


d. From the Create app password screen, enter a name, and then click **next**.

   You will receive an app password that you can use with Outlook, Apple Mail, and other Email options.



e. Select the **copy password to clipboard** option and the password is copied to your clipboard.



App password for the backup service account is created.

How to Create a Microsoft Office 365 OneDrive Backup Plan

OneDrive, part of Microsoft Office 365 Cloud service, facilitates cloud storage and sharing of files. To protect your Onedrive items (Files, Folders, and so on) from Microsoft cloud, you need to create a plan. The plan for OneDrive consists of a backup task. This backup task lets you specify the OneDrive nodes that you want to protect, the backup destination, and the backup schedule.

What To Do Next?

1. Review the Prerequisites and Considerations
2. Create a OneDrive Backup Plan
3. (Optional) Perform a Manual Backup

Review the Prerequisite and Consideration

Prerequisite:

- Running PowerShell script is enabled on the proxy server. If not enabled, then run the command: *Set-ExecutionPolicy RemoteSigned*
- Proxy server must connect to Microsoft Azure.
- Review the following before you create a plan using Basic Authentication:
  - Make sure the PowerShell version of the Proxy server is minimum 5.0.
  - Install the Azure cmdlet in the Proxy server using the following PowerShell command:

    Install-Module AzureAD
- Review the Prerequisites before you create a plan using Modern Authentication.

Consideration:

Microsoft Cloud services are different for national clouds of Germany, China, and US Government. For details, refer to the link. You need to modify the setting when deploying product in these regions.

Follow these steps:

1. Open the config file whose path is $UDP Installation Path.

   $/Engine/Bin/Office365/Arcserve.Office365.Onedrive.json

2. Find the setting **Region**, and set the value to Germany, China or US Government.

   Default value: Normal

Create a OneDrive Backup Plan

A backup plan includes a backup task that performs a backup of OneDrive data items (Files, Folders, and so on) and stores data either at a non-deduplication data store or deduplication data store. Additionally, this backup plan allows UDP to discover and specify OneDrive groups for which a backup task is performed. Each task consists of parameters that define the source, destination, schedule, and other backup details.

Follow these steps:

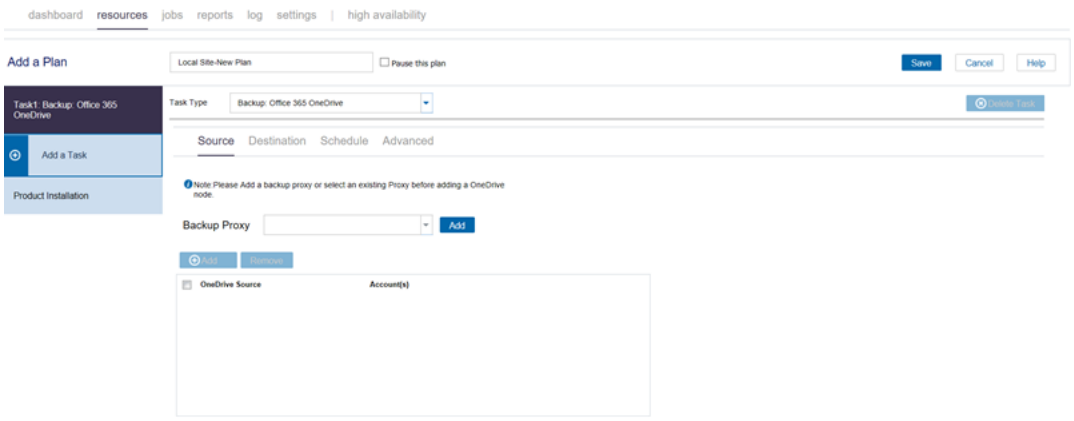
1. Click the **resources** tab on the Console.
2. From the left pane, navigate to **Plans**, and click **All Plans**.

   If you have created plans earlier, those plans are displayed on the center pane.
3. On the center pane, click **Add a Plan**.

   The **Add a Plan** dialog opens.
4. Enter a plan name.
5. (Optional) Select the **Pause this plan** check box.

   The plan does not run until you clear the check box to resume the plan.

   Note: If a plan is paused, then any in-progress job is not paused but all corresponding scheduled jobs associated with that plan are paused. However, you can manually run a job. For example, you can manually run backup job and replication job for a node even if the respective plan is paused. In such case, the following task to the on-demand (manual) job does not run. For example, there is a replication task after an on-demand backup job, the replication job does not run for the on-demand backup job. You need to manually run the replication job. When you resume the plan, the pending jobs do not resume immediately. After you resume the plan, the pending jobs run from the next scheduled time.
6. From the Task Type drop-down list, select **Backup: Office 365 OneDrive**.



Now specify the Source, Destination, Schedule, and Advanced details.

Specify the Source

The Source page lets you specify the OneDrive source nodes that you want to protect. You can select more than one OneDrive source nodes in a plan. If you have not added any nodes to the Console, you can add OneDrive source nodes from the Source page.

Note: You can save a plan without adding any source nodes, but the plan is not deployed unless you add any nodes.

Follow these steps:

1. Click the **Source** tab.
2. Select the **Backup Proxy** from the drop-down list.
3. Add OneDrive node using one of the following methods:

   Method 1

   a. Click **Add** and then click **Select Source to Protect in Arcserve UDP**.

      The Add Nodes to Plan dialog is displayed.

   b. Select a node, and then click **Connect**.

      Note: You can also search for the OneDrive nodes that you want to protect in Search.

   c. Select the **Protect all the OneDrive Accounts** check box to protect all the OneDrive accounts across all the pages. To select few accounts, click the arrow placed on the right side of check box, and then specify the accounts.

      The OneDrive accounts that you selected are added.

   Method 2

a. Click **Add**, and then click **Add OneDrive Source**.

Notes:

- Use this option to discover and specify distribution groups for OneDrive.
- Unlike other nodes, you cannot add the OneDrive node from All Nodes page. You can add a OneDrive node only in a plan or when you modify a plan.

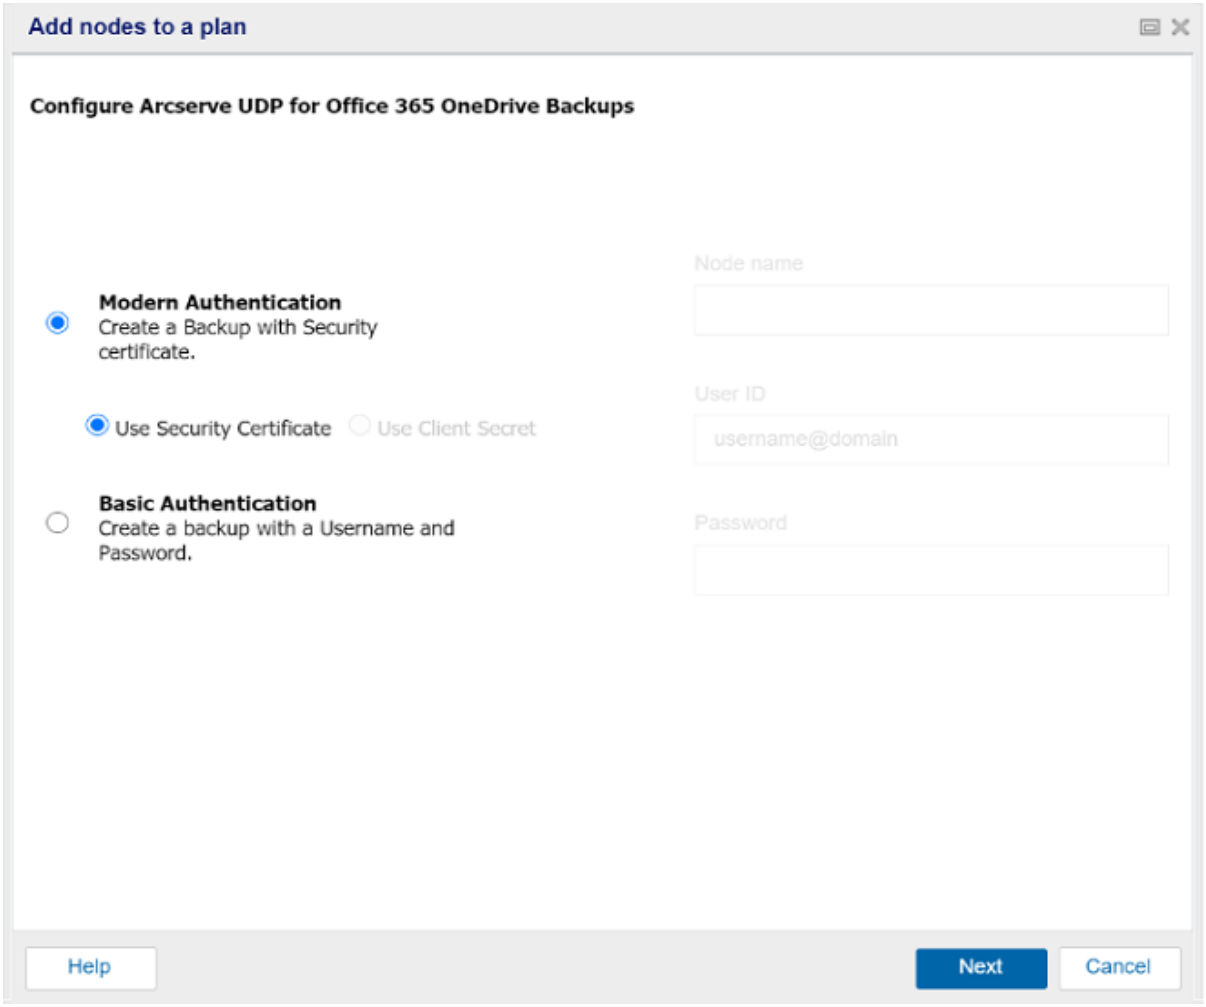
b. To add OneDrive node by plan, specify the node name, user ID, and password. You cannot change the nodename after creating the OneDrive node. Multiple OneDrive nodes can use the same user account (service account) of OneDrive.

Note: We recommend selecting user accounts of Azure Active Directory Administrators.

4. To apply authentication, do one of the following:

Modern Authentication

For modern authentication, you can run a backup using security certificate.

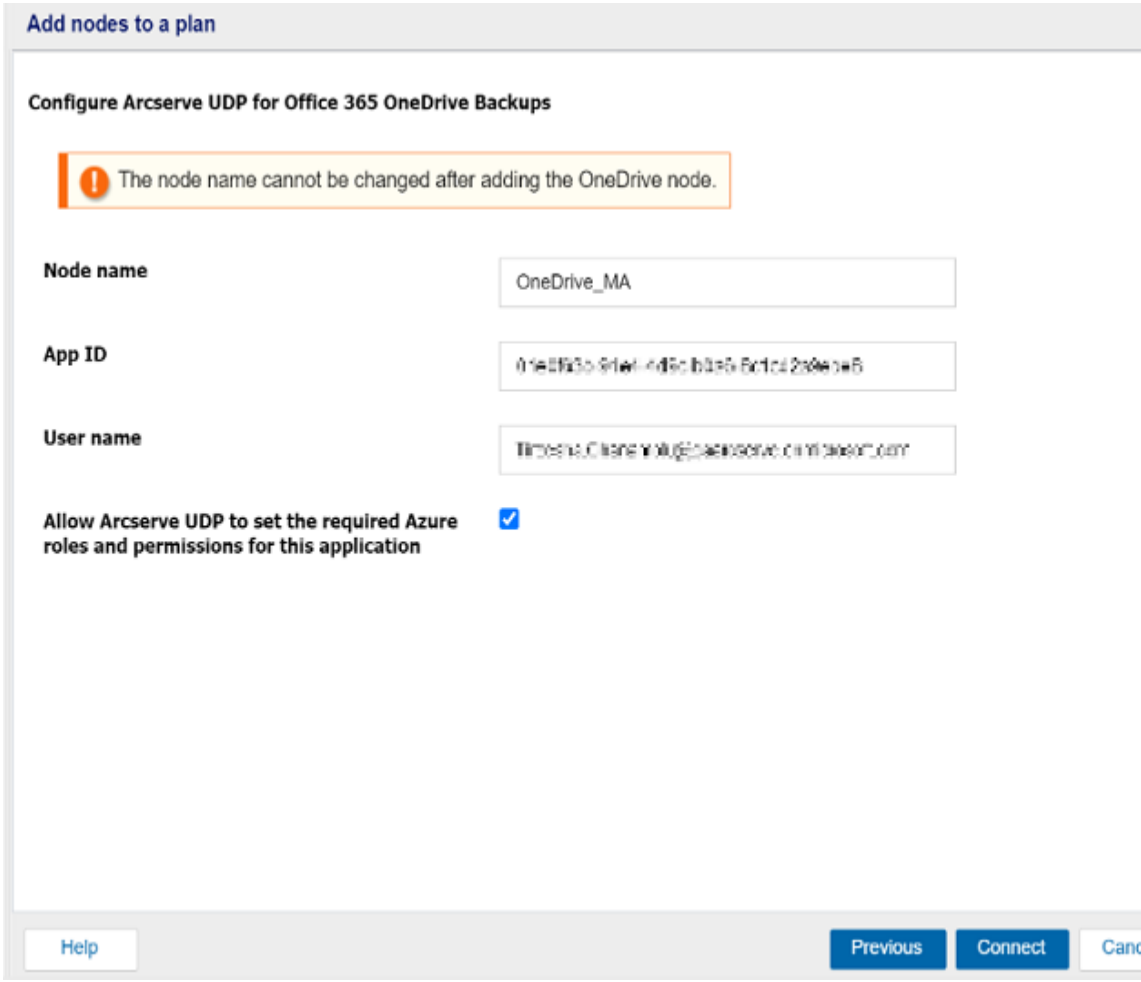


Use security certificate

Note: By default, the backup plan selects **Use security certificate** for modern authentication.

Follow these steps:

1. To authenticate the UDP application with security certificates, see How to Use Self-Signed Certificate.
2. Enter the following details:
   - Node Name
   - App ID
   - User Name



Note: The OneDrive node names cannot change after creating a node.

3. To apply the required roles and permissions, select the **Allow Arcserve UDP to set the required Azure roles and permissions for this application** check box, and then click **Connect**.

Basic Authentication

For Basic Authentication, to proceed with the backup, follow these steps:

a. Select Basic Authentication, type Node name, Username and Password, and then click **Connect**.

   The Arcserve UDP URL opens in a browser and requests permission for all Arcserve UDP read-write data from the OneDrive portal.

b. In the browser, use the Microsoft Azure administrator credentials to sign in.

   The Azure portal opens.

c. From the Microsoft Azure portal, perform the following steps to configure the UDP App:

i. Click **API permissions**.

ii. On the right pane, click the **Grant permissions for** button.

iii. Click **Yes** to agree to Grant permissions.

iv. After granting permission, in Arcserve UDP, navigate to **Add a Plan > Add Nodes to a Plan**, and then click the **Next** button.

Note: If you close the grant permission URL and want to reopen it, click the **Here** button in the UDP Add Node to Plan screen.

Arcserve UDP lists all the account under current tenant.

5. To protect OneDrive accounts, do one of the following:

   - To protect all OneDrive accounts and groups, select the **Protect all the OneDrive Accounts** check box.

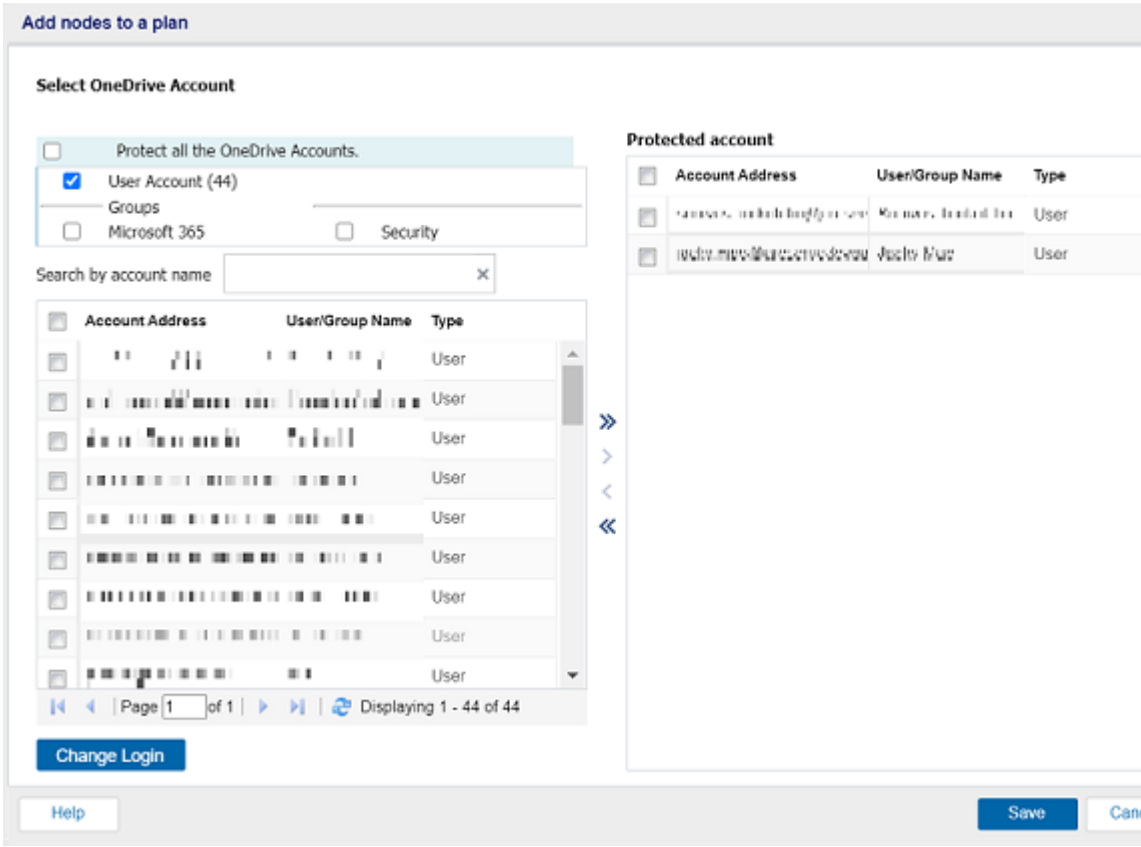
     Note: This option is automatically selected by default when the Select OneDrive Account screen appears.

   - For User Account, follow these steps:

     a. Select the **User Account** check box.

        A list of user accounts appears.

     b. (Optional) To search an account from the list, in the Search by account name field, type the account name.



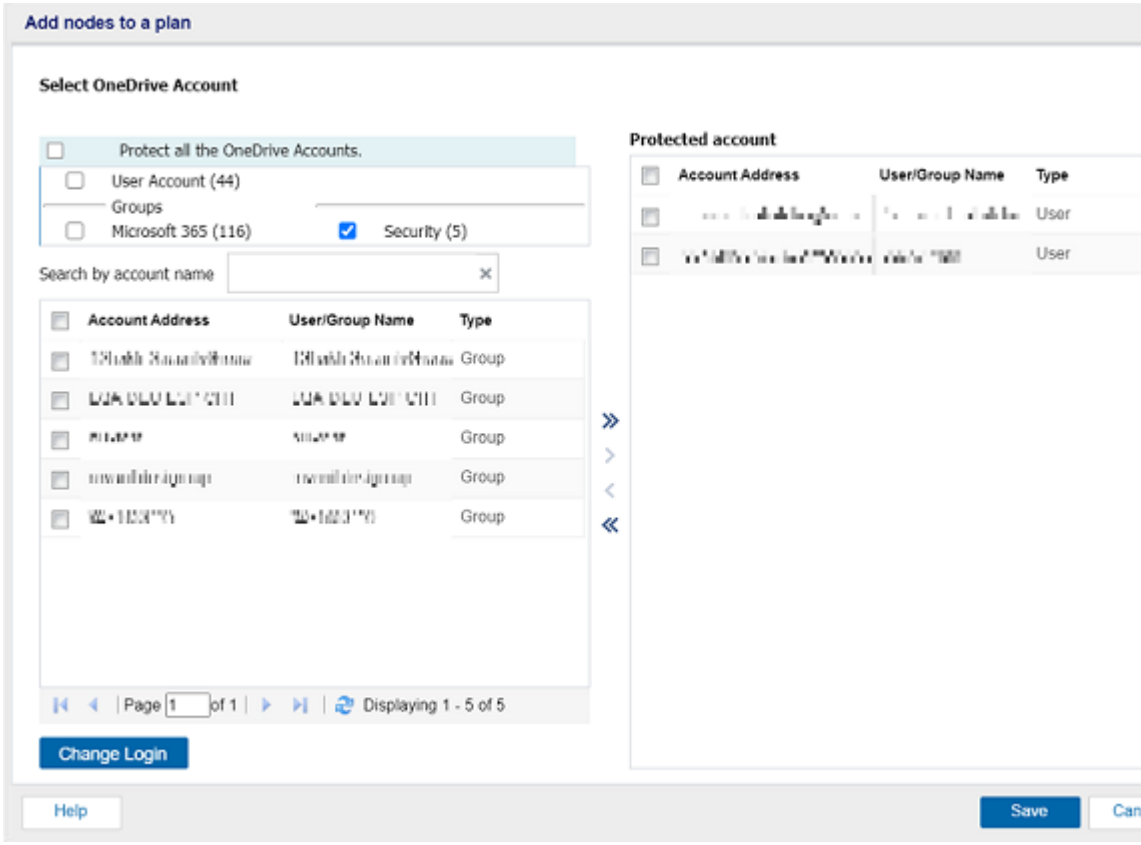
c. From the list of user accounts, select the accounts that you want to protect.

d. Click the right arrow (>) to move the selected accounts to the Protected account list.

- For Groups, follow these steps:

a. Select the **Microsoft 365** or **Security** check box.

Based on selection, a list of groups appears.



b. (Optional) To search an account from the list, in the Search by account name field, type the account name.

c. From the list of groups, select the groups that you want to protect.

d. Click the right arrow (>) to move the selected groups to the Protected account list.

6. Click **Save**.

   The selected OneDrive accounts are added.

Specify the Destination

The destination is a location where you store the backup data. You must at least specify the destination to save the plan.

Follow these steps:

1. Click the **Destination** tab.

   The **Arcserve UDP Recovery Point Server** option is automatically selected. **Arcserve UDP Recovery Point Server** specifies that the backup destination is a recovery

point server. If you select this option, the data is stored as recovery points. You cannot store data as recovery sets.

2. Provide the following details:
   a. Select a recovery point server.
   b. Select a non-deduplication or deduplication data store. The list displays all the data stores created on the specified recovery point server.
   c. Provide a session password. The session password is optional when the backup destination is an unencrypted RPS data store.
   d. Confirm the session password.

The destination is specified. Now, specify the Schedule and Advanced details.

Specify the Schedule

The Schedule page lets you define a schedule for Backup, Merge, and Throttle functions to repeat at specific intervals. After you define a schedule, the jobs run automatically per the schedule. You can add multiple schedules and can provide retention settings.

A Backup Schedule refers to regular schedule that is repeated multiple times a day based on the number of hours or minutes you select. Besides the regular schedule, a backup schedule also provides options to add daily, weekly, and monthly schedules.

Note: For more information on scheduling and retention settings, see Understanding Advanced Scheduling and Retention.

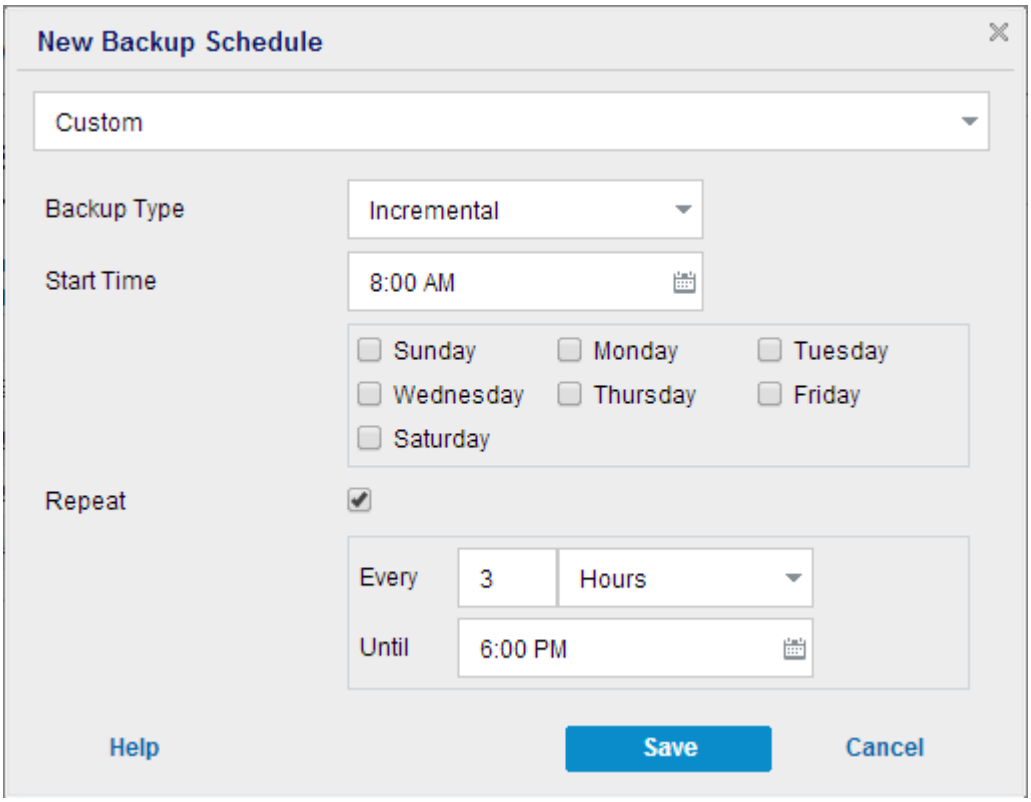
Follow these steps:

1. Add backup, merge, and throttle schedules.

   Add Backup Schedule

   a. Click **Add** and select **Add Backup Schedule**.

      The **New Backup Schedule** dialog opens.



b. Select one of the following options:

Custom

Specifies the backup schedule that repeats multiple times a day.

Daily

Specifies the backup schedule that occurs once a day. By default, all the days of the week are selected for Daily backup. If you do not want to run the backup job on a specific day, clear the check box for that day of the week.

Weekly

Specifies the backup schedule that occurs once a week.

Monthly

Specifies the backup schedule that occurs once a month.

c. Select the backup type.

Full

Determines the backup schedule for Full Backups. As scheduled, Arcserve UDP performs a full backup of all used blocks from the source machine. A full backup typically consumes time depending on the backup size.

Incremental

Determines the backup schedule for Incremental Backups.

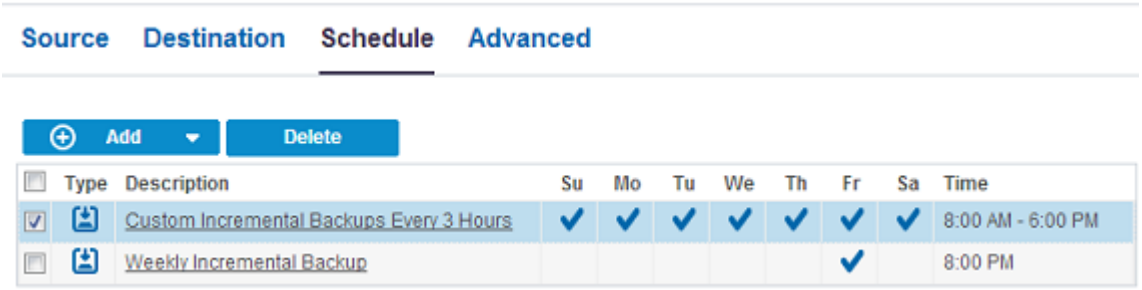
As scheduled, Arcserve UDP incrementally backs up only those blocks that have changed after the last successful backup. The advantages of Incremental Backups are that the backup is a fast backup and it produces a small backup image. This option is the most optimal way to perform a backup.

d. Specify the backup start time.

e. (Optional) Select the **Repeat** check box and specify the repeat schedule.

f. Click **Save**.

The Backup Schedule is specified and appears on the **Schedule** page.



Add Merge Schedule

a. Click **Add** and select **Add Merge Schedule**.

The **Add New Merge Schedule** dialog opens.

b. Specify the start time to start the merge job.

c. Specify **Until** to specify an end time for the merge job.

d. Click **Save**.

The Merge Schedule is specified and appears on the **Schedule** page.

Add Throttle Schedule

a. Click **Add** and select **Add Throttle Schedule**.

The **Add New Throttle Schedule** dialog opens.

b. Specify the throughput limit in MB per minutes unit.

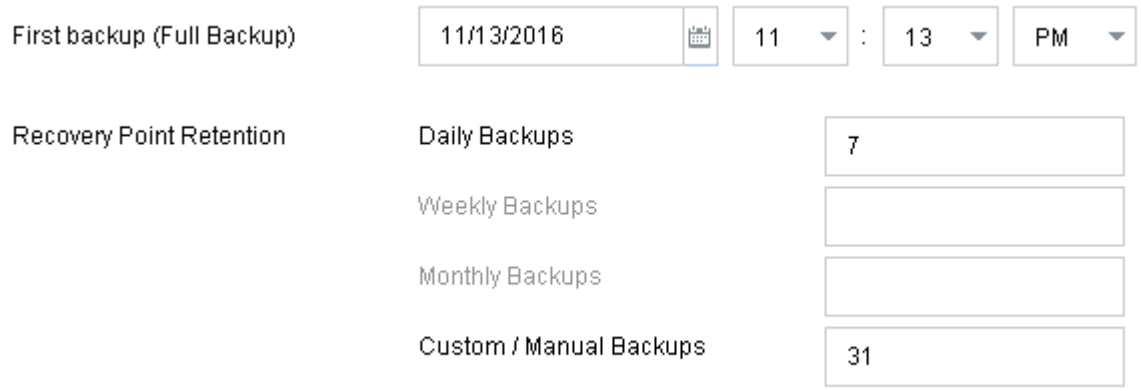
c. Specify the start time to start the backup throughput job.

d. Specify **Until** to specify an end time for the throughput job.

e. Click **Save**.

The Throttle Schedule is specified and displayed on the **Schedule** page.

2. Specify the start time for the scheduled backup.



3. Specify the recovery points retention settings for Custom, Daily, Weekly, and Monthly schedule.

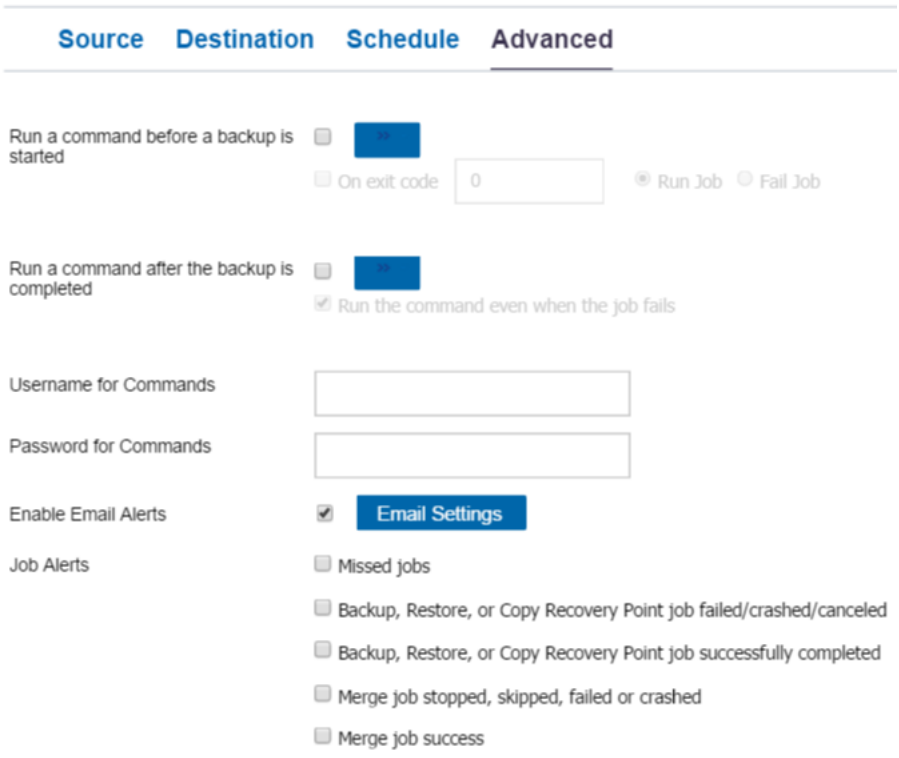
   These options are enabled if you have added the corresponding backup schedule. If you modify the retention settings on this page, the changes are reflected on the **Backup Schedule** dialog.

   The schedule is specified. Now specify Advanced details.

Specify the Advanced Settings

The **Advanced** tab lets you specify some advanced settings for the backup job. The advanced settings include providing the location of any scripts, and email settings.

The following image displays the **Advanced** tab:



Follow these steps:

1. Specify the following details.

   Run a command before a backup is started

   Lets you run a script before the backup job starts. Specify the path where the script is stored inside the proxy node. Click **On exit code** and specify the exit code for **Run Job** or **Fail Job**. **Run Job** indicates that the backup job continues when the script returns the exit code. **Fail Job** indicates that the backup job stops when the script returns the exit code.

   Run a command after a backup is completed

   Lets you run a script after the backup job is completed. Specify the complete path where the script is stored.

   Run a command even when the job fails

   If this check box is selected, the script specified in **Run a command after a backup is completed** is executed even when the backup job fails. Otherwise, that script is executed only when backup job completes successfully.

   Username for Commands

   Lets you specify the username to run the commands.

   Password for Commands

   Lets you specify the password to run the commands.

   Enable Email Alerts

   Lets you enable email alerts. You can configure email settings and specify the types of alerts that you want to receive in an email. When you select this option, the following options are enabled for your selection.

   Email Settings

   Lets you configure the email settings. Click **Email Settings** and configure the email server and proxy server details. For more information about how to configure Email Settings, refer to Email and Alert Configuration.

   Job Alerts

   Lets you select the types of job alert emails that you want to receive.

2. Click **Save**.

Note: When you select a node as a backup source or backup proxy, Arcserve UDP checks whether the agent is installed on proxy node and if it is the latest version. Arcserve UDP then displays a verification dialog that lists all the nodes that either have an outdated version of the agent or does not have the agent installed. To

install/upgrade the agent on these nodes, select the installation method and click **Save**.

The changes are saved and a green checkmark appears next to the task name. The plan page closes.

Note: If you have to add another task, you must select the plan from the **resources** tab and modify the plan. To modify the plan, click the plan from the center pane. The plan opens and you can modify it. You may add the **Copy Recovery Point**, **Copy to Tape**, **Replicate**, and **Replicate from a remote RPS** tasks as follow up tasks.

The plan is automatically deployed to the proxy server node.

The exchange online backup plan for the proxy server is created. The backup runs per the schedule that you have configured on the **Schedule** tab. You can also perform a manual backup at any time.

(Optional) Perform a Manual Backup

Typically, based on the schedule settings you applied, the backups perform automatically. In addition to the scheduled backup, a manual backup provides you the option to back up your nodes on a need basis. For example, if you have a repeat schedule for Full and Incremental backups, and you want to make major changes to your machine, you can perform an immediate manual backup without waiting for the next scheduled backup. You can submit the backup job from both, Console and proxy user interface. Using the Job Monitor, you can view the job status and cancel the ongoing job.

To perform a manual backup of OneDrive nodes, Follow these steps:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and click **All Nodes**.

   The OneDrive nodes are displayed in the center pane.
3. Select the OneDrive nodes that you want to backup and that has a plan assigned to it.
4. On the center pane, click **Actions**, **Backup Now**.

   The **Run a backup now** dialog opens.
5. Select a backup type and optionally provide a name for the backup job.
6. Click **OK**.

   The backup job runs.

Follow these steps: to perform a manual backup of a OneDrive plan

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Plans**, and click **All Plans**.

   The OneDrive backup plans are displayed in the center pane.
3. Select the plan that you want to backup and make sure a plan is assigned to it.
4. On the center pane, click **Actions**, **Backup Now**.

   The **Run a backup now** dialog opens.
5. Select a backup type and optionally provide a name for the backup job.
6. Click **OK**.

   The backup job runs.

   The manual backup is successfully performed. Now, you can verify the Backup.

Verify the Backup

To verify your backup job, confirm that you have successfully created the backup plan. After you verify that the plan is created successfully, confirm whether the backup job is running as scheduled. You can verify the status of backup jobs from the **Jobs** tab.

Follow these steps to verify plans:

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and then click **All Nodes**.

   A list of all nodes is displayed on the center pane.
3. Verify that plans are mapped with nodes.

Follow one of these methods to verify backup jobs:

Method 1

1. Click the **jobs** tab.
2. From the left pane, click **All Jobs Completed**.

   The status of each job is listed on the center pane.
3. Verify that the backup job is successful.

   The backup job is verified.

Method 2

1. Click the **jobs** tab.
2. From the left pane, click **Jobs Successfully Completed**.

   The list of backup jobs completed appear on the center pane.

3. Click the job you want to verify.

   The right pane refreshes.
4. Under Job Details, click the **View Logs** hyperlink.
5. Verify that the backup job is successful.

   The backup job is verified.

Managing Microsoft Teams Backup and Restore

This section provides information about how to backup and restore Microsoft Teams data using Exchange Online and SharePoint Online. Arcserve UDP provides the backup of messages (chats) and files shared across the private, public, and group channels.

Whenever you create a private, public, and group channel in Microsoft Teams, a SharePoint Online site gets created for it automatically. This SharePoint Online site has a default document library folder, in which all the files shared across all the conversations get saved automatically. You can also customize the permissions and security options for sensitive information.

This section contains the following topics:

How to Backup and Restore Microsoft Teams Messages (Chats) using Exchange Online

You can backup individual chats using the user mailbox in UDP Exchange Online node backup. All the messages are stored in the Exchange Online Server in the TeamsMessagesData folder, which is under a root folder. This section provides information about how to backup and restore Microsoft Teams messages (chats) using Exchange Online.

This section contains the following topics:

Backup Microsoft Teams Messages (Chats) using Exchange Online

This section provides information about how to backup Microsoft Teams messages (chats) using Exchange Online.

This section contains the following topics:

How to Create an Exchange Online Backup Plan

Exchange Online is an email application hosted on Microsoft cloud. To protect your Exchange Online mail items (Mails, Calendar items, Contacts, and so on) from Microsoft cloud, you need to create a plan. The plan for Exchange Online consists of a backup task. This backup task lets you specify the Exchange Online nodes you want to protect, the backup destination, and the backup schedule.

This section contains the following topics:

Review the Prerequisites and Considerations

Prerequisites:

For Backup account:

- Use a backup service account with Global admin permissions.
- Add impersonation permission for backup user to the Exchange Online backup account to connect the Exchange Online organization, perform backups and restore.

  Note: If you do not add the backup account to the Discovery Management role group and do not assign the Application Impersonation permission, backup fails.

- Associate the Backup user account with one exchange online mailbox.

For Backup proxy:

- Log in to the Console.
- Install Arcserve UDP Agent on the proxy machine where you want to run backup/restore.
- Review the Prerequisites before you create a plan using Modern Authentication.

Considerations:

- While backing up a large set of users, you can use exchange groups for better backup and optimization.
- Default setting for exchange online backup uses 4 backup threads (one thread per user) at a time. You can also modify the threads in configuration file of *Engine\BIN\Office365\Arcserve.Office365.Exchange.config* as follows:

  <!--#region for multi thread-->

  <!--MultiThreadEnable default value:0. if enable, set 1.-->

  <add key="MultiThreadEnable" value ="1"/>

  <!--set how many thread will be used to backup mailbox. default value is 4-->

  <add key="MaxDegreeOfParallelismForMailbox" value="4"/>

  <!--#endregion-->

  We recommend setting the value from 1-5. The maximum possible value is 10. But, we do not recommend setting the value from 6-10.

Note: UDP 8.0 supports the following:

- User mailbox
- Shared mailbox
- Mail-enabled Public folders
- Room mailboxes
- Equipment mailboxes
- Group mailboxes

Add the Required Role and Group to the Exchange Online Backup Account to Perform Backup and Restore

Add the backup account to the Discovery Management role group and assign Application Impersonation permission to it.

Follow these steps:

1. Add the required role and group using any one of the following ways:

   Using Office 365 portal

a. Log into Office 365 portal as an Administrator or with an account that has Global Admin permissions.

   The Exchange admin center page appears.

b. Go to **permissions** and from the Add drop-down list, double-click **Discovery Management**.

   The Discovery Management dialog box appears.

   Note: Member of the Discovery Management role group can search mailboxes in the Exchange organization for data that meets specific criteria.

c. Under **Roles**, click **+** to add the ApplicationImpersonation role.

   The Select a Role dialog box appears.

d. From the **Display Name** drop-down, select **ApplicationImpersonation**.

   Note: The ApplicationImpersonation role enables applications to impersonate users in an organization to perform tasks on behalf of the user.

e. Click **Add** to add the ApplicationImpersonation role, and then click **OK**.

f. On the Discovery Management dialog box, under **Members**, click **+** to add the backup account as a member.

   The Select Members dialog box appears.

g. From the Name drop-down, select the backup account, click **Add** to add the member, and then click **OK**.

   The selected backup account is displayed under **Members** on the **Discovery Management** dialog.

h. Click **Save**.

Using Remote Powershell

a. To connect to the Exchange Online tenant using remote PowerShell, refer to the link.

b. Once connected, to add the backup account as a member of Discovery Management role group, use the:

`"Add-RoleGroupMember" cmdlet`

For example: Add-RoleGroupMember "discovery management" -member userName@domain.onmicrosoft.com.

c. To assign application impersonation role to the backup account, use the:

`"New-ManagementRoleAssignment" cmdlet`

For example:

New-ManagementRoleAssignment -Name: impersonationAssignmentName -Role:ApplicationImpersonation - User: "username@domain.onmicrosoft.com"

The ApplicationImpersonation role and Members group are added to the Exchange Online backup account.

Create an Exchange Online Backup Plan

A backup plan includes a backup task that performs a backup of Exchange Online mail data items (Mails, Calendar items, Contacts, so on) and stores data either at a non-deduplication data store or deduplication data store. Additionally, this backup plan allows UDP to discover and specify Exchange Online groups for which a backup task is performed. Each task consists of parameters that define the source, destination, schedule, and other backup details.

Follow these steps:

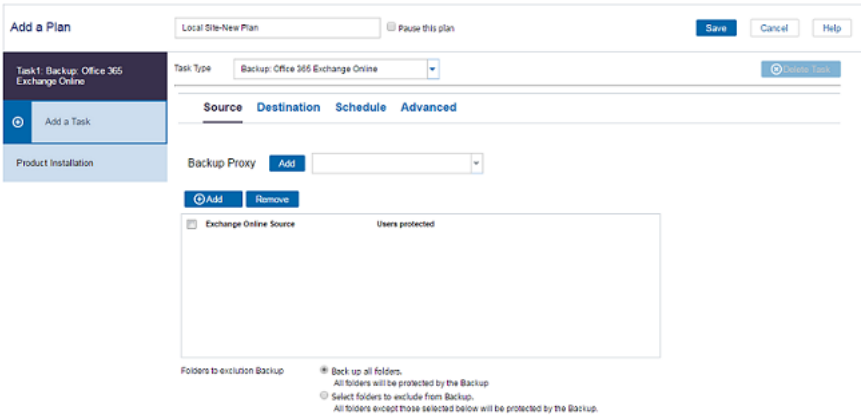
1. Click the **resources** tab on the Console.
2. From the left pane, navigate to **Plans**, and click **All Plans**.

   If you have created plans earlier, those plans are displayed on the center pane.
3. On the center pane, click **Add a Plan**.

   The **Add a Plan** dialog opens.
4. Enter a plan name.
5. (Optional) Select the **Pause this plan** check box.

   The plan does not run until you clear the check box to resume the plan.

   Note: If a plan is paused, then any in-progress job is not paused but all corresponding scheduled jobs associated with that plan are paused. However, you can manually run a job. For example, you can manually run backup job and replication job for a node even if the respective plan is paused. In such case, the following task to the on-demand (manual) job does not run. For example, there is a replication task after an on-demand backup job, the replication job does not run for the on-demand backup job. You need to manually run the replication job. When you resume the plan, the pending jobs do not resume immediately. After you resume the plan, the pending jobs run from the next scheduled time.
6. From the Task Type drop-down list, select **Backup: Office 365 Exchange Online**.



Now specify the Source, Destination, Schedule, and Advanced details.

Specify the Source

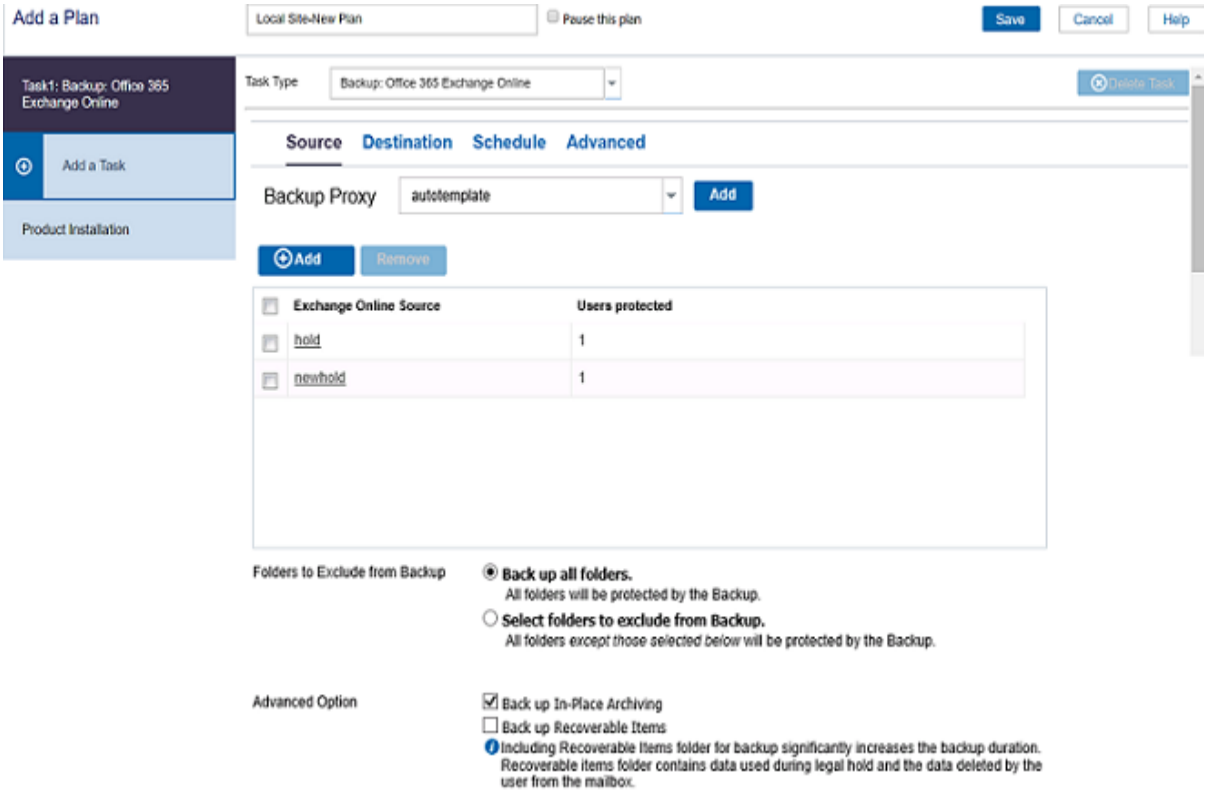
The Source page lets you specify the Exchange Online source nodes that you want to protect. You can select more than one Exchange Online source nodes in a plan. If you have not added any nodes to the Console, you can add Exchange Online source nodes from the Source page.

Note: You can save a plan without adding any source nodes but the plan is not deployed unless you add any nodes.

You can also manage Exchange Online nodes using Public folder Mailbox Support for Exchange Online Protection.

Follow these steps:

1. Click the **Source** tab.



2. For Backup Proxy, do one of the following:
   - From the drop-down list, select an existing backup proxy.
   - To add a new backup proxy, next to the drop-down list, click **Add**.
3. Add Exchange Online node using one of the following options:

   Method 1

a. Click **Add** and then click **Select Source to Protect in Arcserve UDP**.

   The Add Nodes to Plan dialog is displayed.

b. Select a node, and then click **Connect**.

   Note: You can also search for the Exchange Online nodes that you want to protect in **Search**.

c. Select the **Protect all of the Office 365 Exchange Sources** check box to protect all the Exchange Online accounts across all the pages.

   Note: To add all the Exchange Online accounts to the protected list, you may click the right (>) arrow.

   The Exchange Online accounts that you selected are added.

Method 2

Click **Add** and then click **Add Exchange online Source**.

The Add nodes to a plan dialog appears.

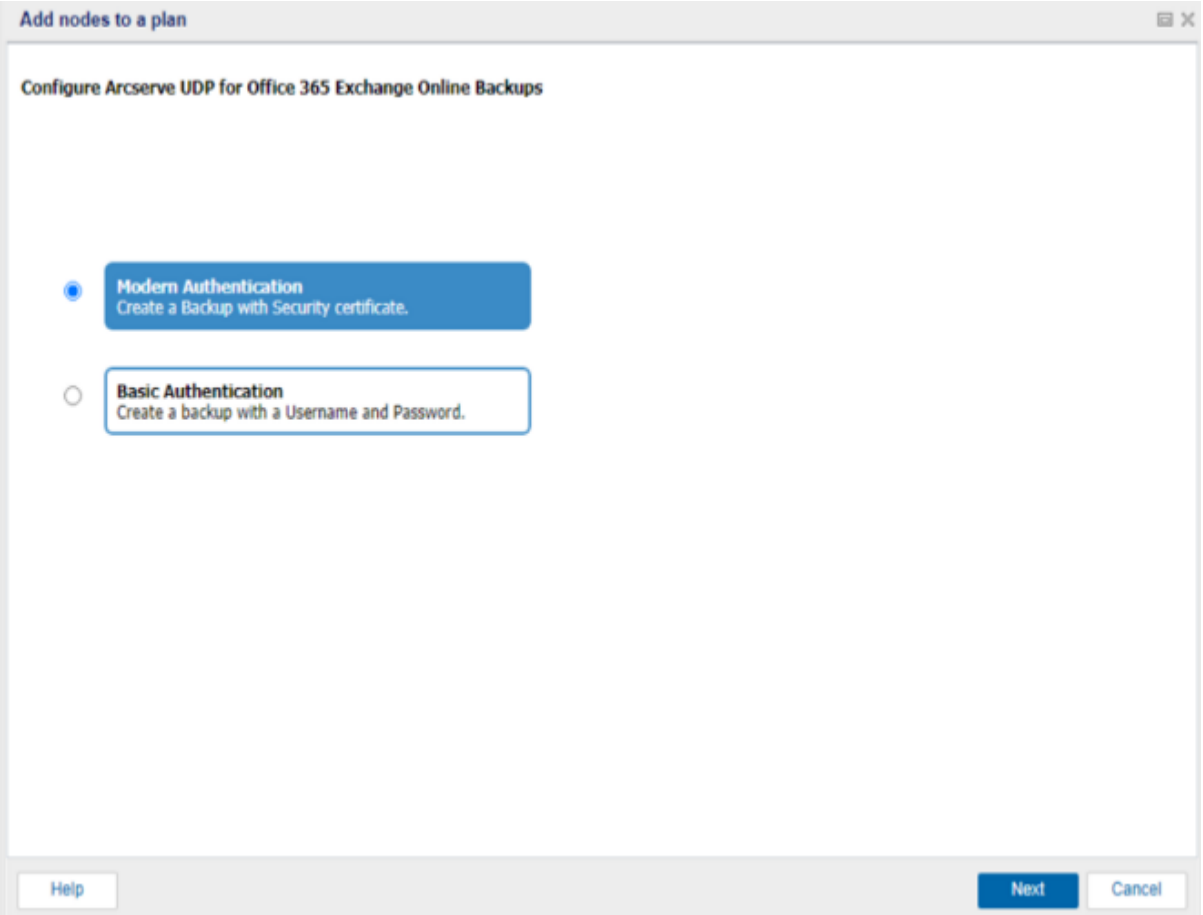
Note: Unlike other nodes, you cannot add the Exchange Online node from **All Node**s page. You can add an Exchange Online node only in a plan or when you modify a plan.

Multiple Exchange online nodes can use the same user account (service account) of Exchange online.

4. To apply authentication, do one of the following:

Modern Authentication

For modern authentication, you can run a backup using security certificate.



Use security certificate

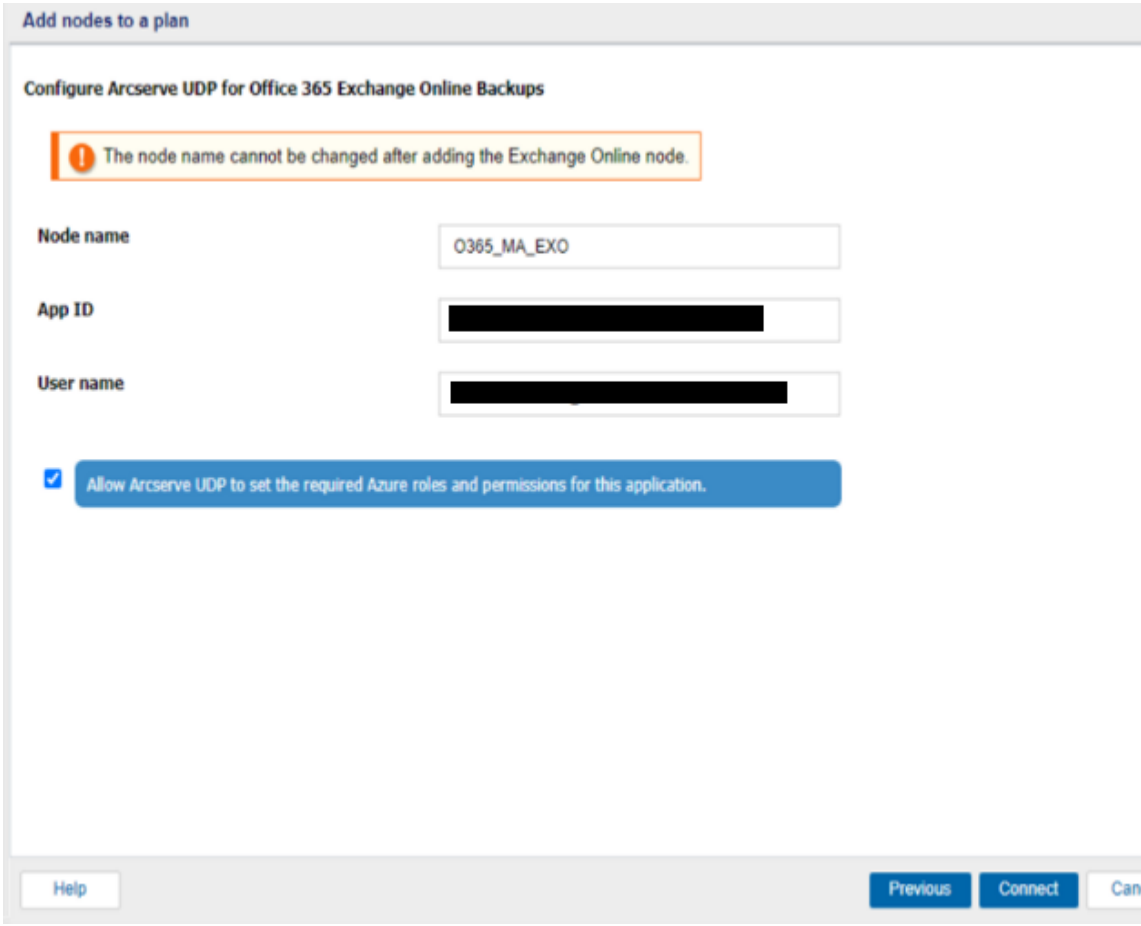
Note: By default, the backup plan selects **Use security certificate** for modern authentication.

Follow these steps:

a. To authenticate the UDP application with security certificates, see How to Use Self-Signed Certificate.

b. Enter the following details:

- Node Name
- App ID
- User Name



c. To apply the required roles and permissions, select the **Allow Arcserve UDP to set the required Azure roles and permissions for this application** check box, and then click **Connect**.

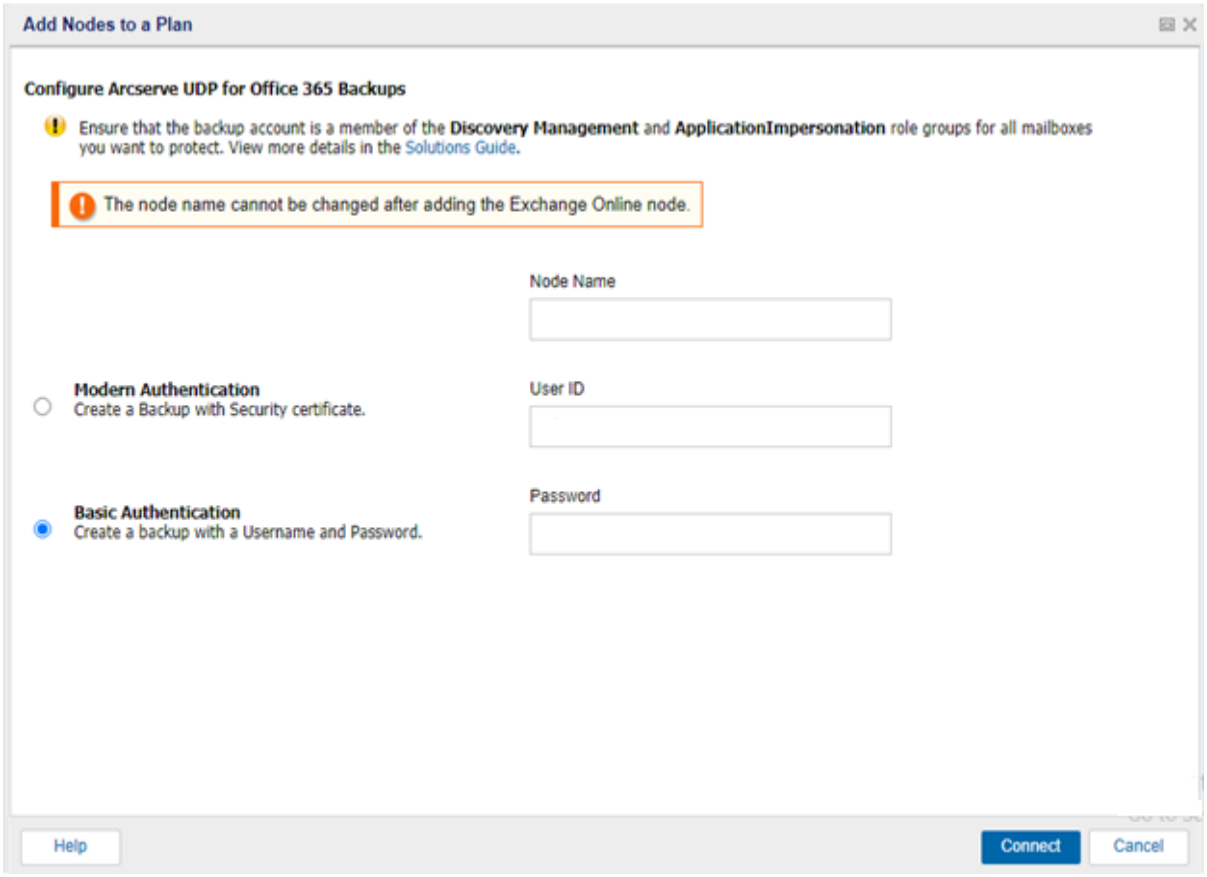
Notes:

- Exchange online node name cannot be changed after adding the node.
- Enter the user name of Exchange Online backup account that meet the prerequisites in **Admin Username**.
- Modern Authentication does not apply to the following:
  - Customers using Microsoft 365 (Office 365) for their Arcserve Cloud Hybrid instances or Arcserve Cloud Backup for Office 365, which was created before 18th October, 2020 does not allow modern authentication.
  - Customers who continue to use basic authentication.

Basic Authentication

For Basic Authentication, to proceed with the backup, provide the following information, and then click **Connect**:

- Node name
- User ID
- Password



Note: Updating or changing the user account may change the number of protected mailboxes. You need to verify that the new or updated service accounts have impersonation rights for the mailboxes to be protected.

5. To protect Exchange Online mailboxes, do one of the following:

   - To protect all exchange online mailboxes and groups, select the **Protect all of the Office 365 Exchange Source** check box.

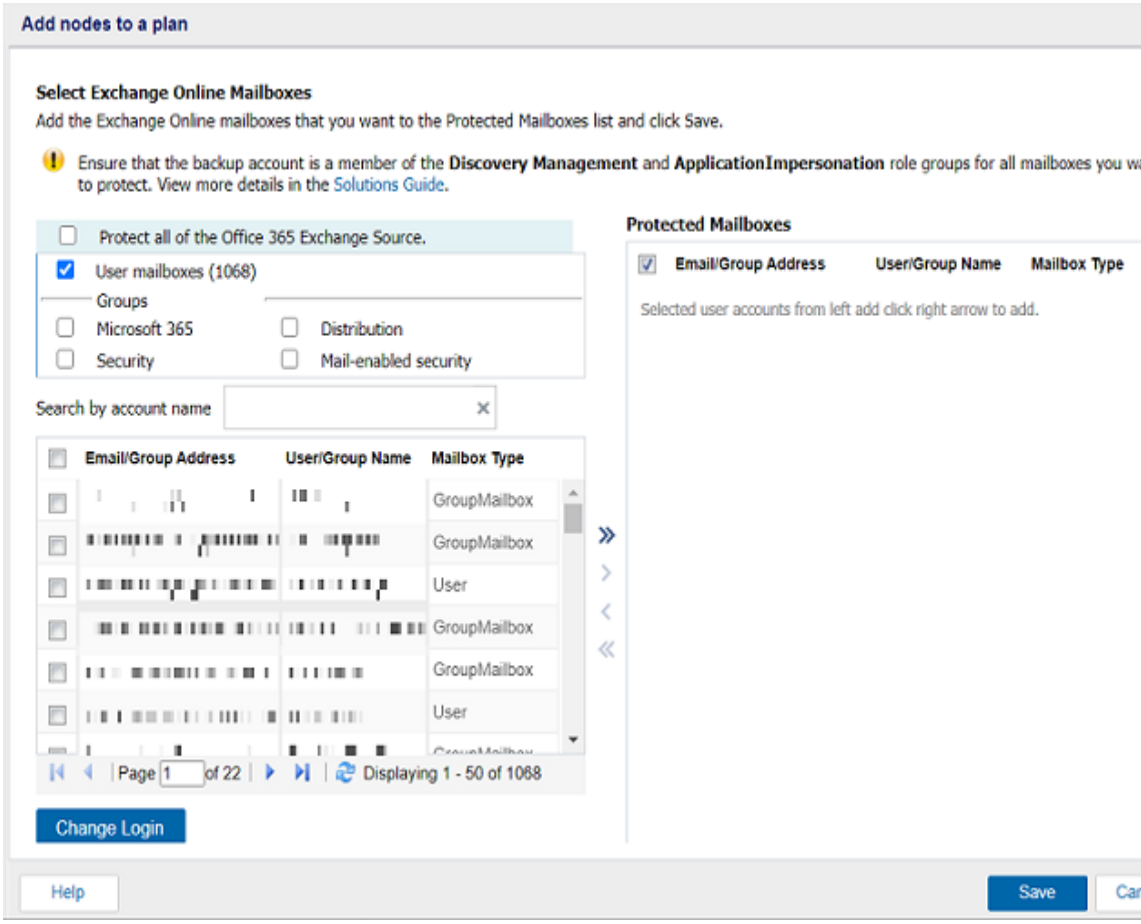
     Note: This option is automatically selected by default when the Select Exchange Online Mailboxes screen appears.

   - For User mailboxes, follow these steps:

a. Select the **User mailboxes** check box.

   A list of user mailboxes appears.

b. (Optional) To search an account from the list, in the Search by account name field, type the account name.



c. From the list of user mailboxes, select the mailboxes that you want to protect.

d. Click the right arrow (>) to move the selected mailboxes to the Protected Mailboxes list.

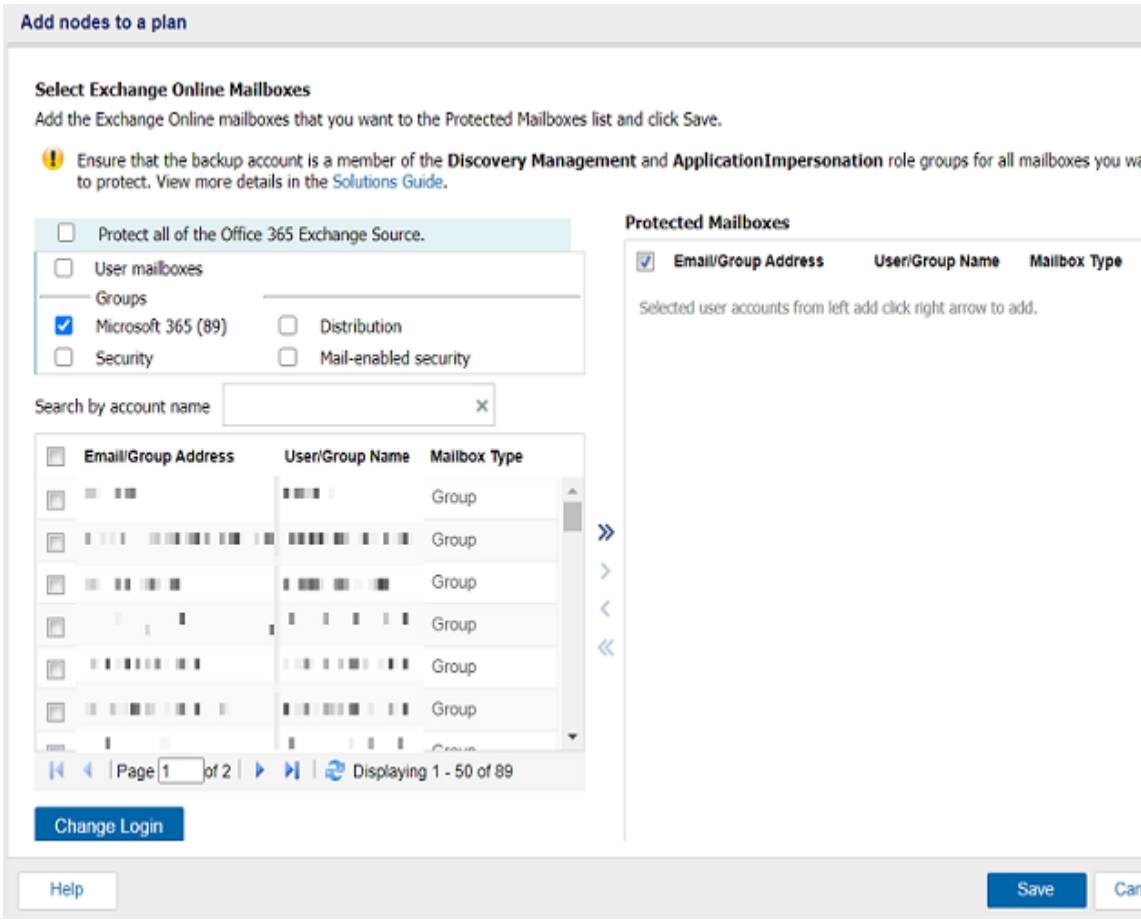
- For Groups, follow these steps:

  a. Select the **Microsoft 365** check box.

     A list of groups appears.

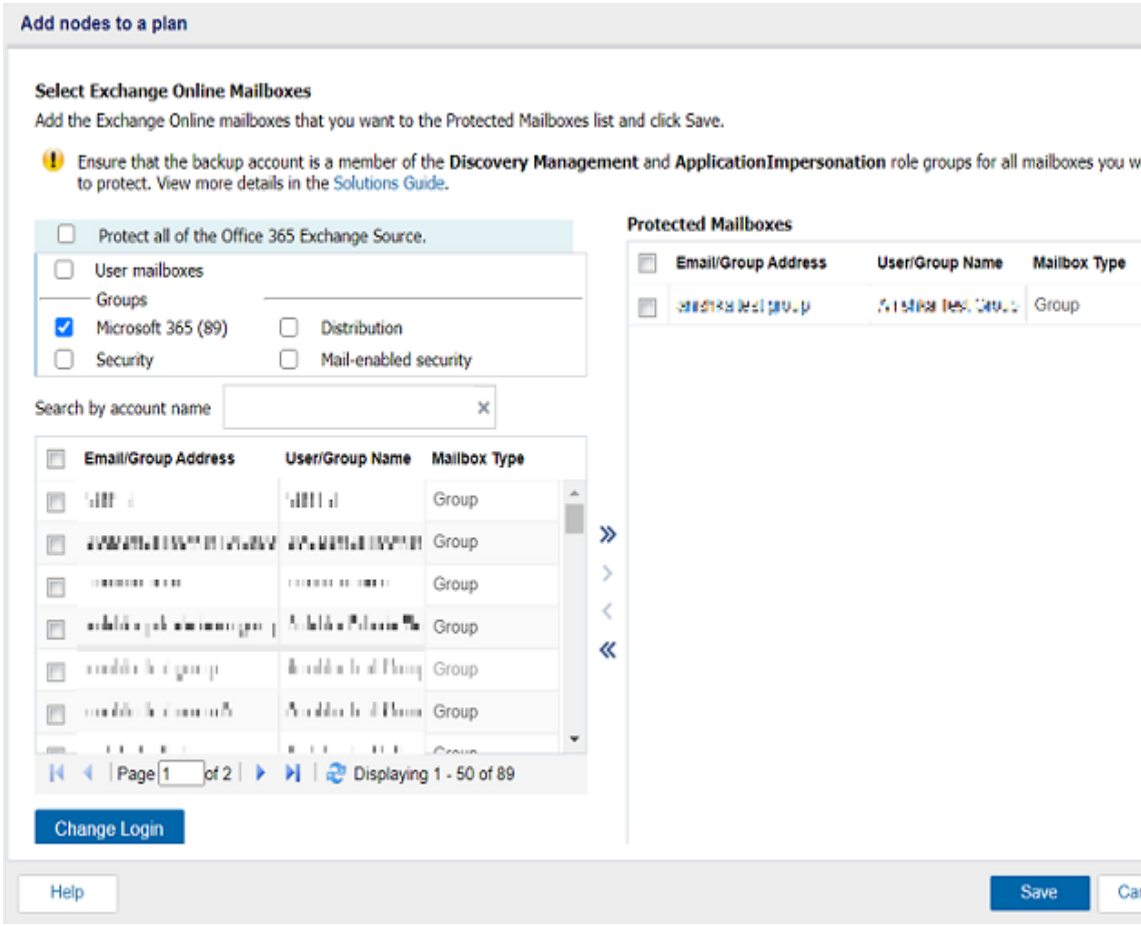
     Note: If you want to select groups along with sub-groups such as **Distribution**, **Security**, or **Mail-enabled security**, select the **Microsoft 365** check box, and then select a sub-group check box as needed.

b. (Optional) To search an account from the list, in the Search by account name field, type the account name.



c. From the list of groups, select the groups that you want to protect.

d. Click the right arrow (>) to move the selected groups to the Protected Mailboxes list.



6. Click **Save**.
7. On the Sources tab, go to **Folders to Exclude from Backup**, and then select the desired check box.
8. From the Advanced Option, select the desired check box.
    - To allow Exchange Online Protection support Archiving Mailbox, select the **Backup up In-Place Archiving** check box.

      Note: For more information about Archiving Mailbox, refer to the link.
    - Select the **Backup up Recoverable items** check box to enable to protect the mailbox that enables the In-Place Hold or Litigation Hold feature.

      Note: For Archiving In-Place Hold and Litigation Hold for Exchange Online, refer to the link.

Note: To enable both the features in the mailbox at the same time to back up the recoverable items in Archiving mailbox, select both the options **Backup up In-Place Archiving** and **Backup up Recoverable items**.

Specify the Destination

The destination is a location where you store the backup data. You must at least specify the destination to save the plan.

Follow these steps:

1. Click the **Destination** tab.
2. Select the **Arcserve UDP Recovery Point Server** option. **Arcserve UDP Recovery Point Server** specifies that the backup destination is a recovery point server. If you select this option, the data is stored as recovery points. You cannot store data as recovery sets.
3. Provide the following details:
    a. Select a recovery point server.
    b. Select a non-deduplication or deduplication data store. The list displays all the data stores created on the specified recovery point server.
    c. Provide a session password. The session password is optional when the backup destination is an unencrypted RPS data store.
    d. Confirm the session password.

The destination is specified.

Specify the Schedule

The Schedule page lets you define a schedule for Backup, Merge, and Throttle functions to repeat at specific intervals. After you define a schedule, the jobs run automatically per the schedule. You can add multiple schedules and can provide retention settings.

A Backup Schedule refers to regular schedule that is repeated multiple times a day based on the number of hours or minutes you select. Besides the regular schedule, a backup schedule also provides options to add daily, weekly, and monthly schedules.

Note: For more information on scheduling and retention settings, see Understanding Advanced Scheduling and Retention.

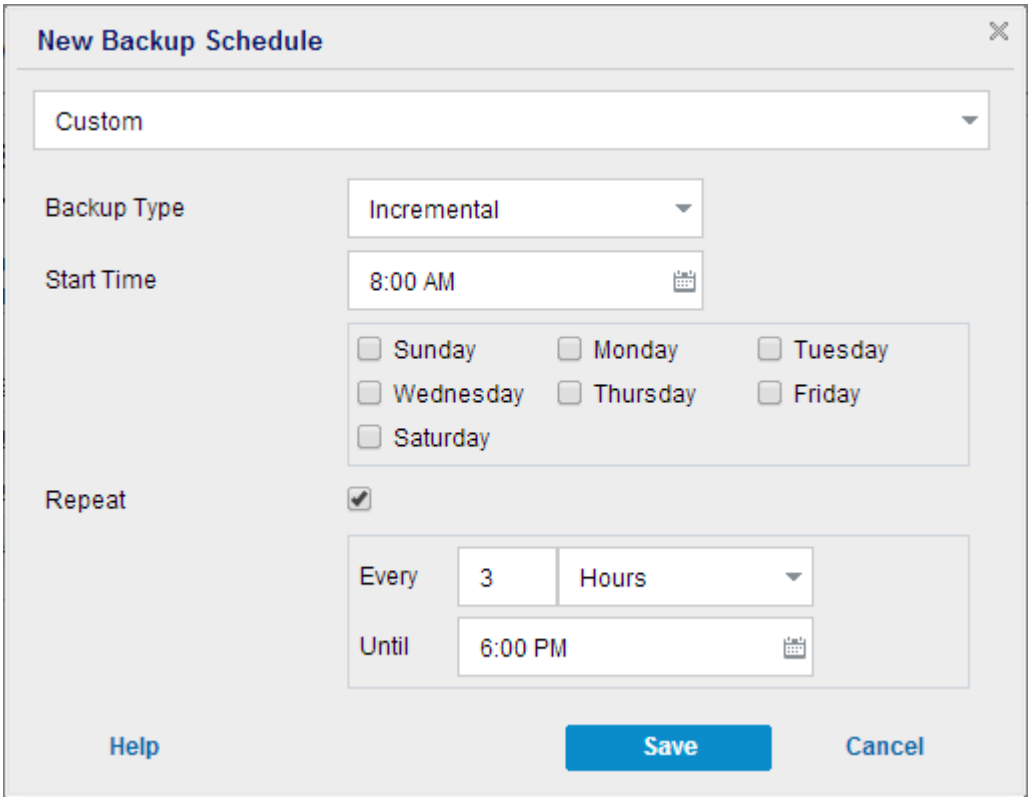
Follow these steps:

1. Add backup, merge, and throttle schedules.

   Add Backup Schedule

   a. Click **Add** and select **Add Backup Schedule**.

      The **New Backup Schedule** dialog opens.



   b. Select one of the following options:

      Custom

Specifies the backup schedule that repeats multiple times a day.

Daily

Specifies the backup schedule that occurs once a day. By default, all the days of the week are selected for Daily backup. If you do not want to run the backup job on a specific day, clear the check box for that day of the week.

Weekly

Specifies the backup schedule that occurs once a week.

Monthly

Specifies the backup schedule that occurs once a month.

c. Select the backup type.

Full

Determines the backup schedule for Full Backups. As scheduled, Arcserve UDP performs a full backup of all used blocks from the source machine. A full backup typically consumes time depending on the backup size.

Incremental

Determines the backup schedule for Incremental Backups.

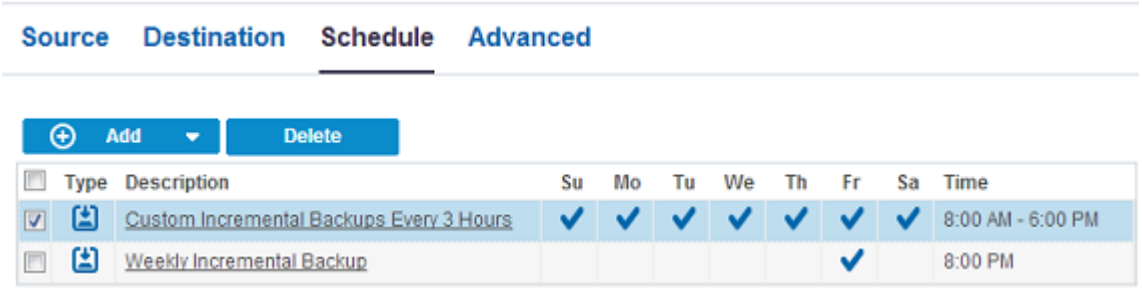
As scheduled, Arcserve UDP incrementally backs up only those blocks that have changed after the last successful backup. The advantages of Incremental Backups are that the backup is fast and produces a small backup image. This is the most optimal way to perform a backup.

d. Specify the backup start time.

e. (Optional) Select the **Repeat** check box and specify the repeat schedule.

f. Click **Save**.

The Backup Schedule is specified and appears on the **Schedule** page.



Add Merge Schedule

a. Click **Add** and select **Add Merge Schedule**.

   The **Add New Merge Schedule** dialog opens.

b. Specify the start time to start the merge job.

c. Specify **Until** to specify an end time for the merge job.

d. Click **Save**.

The Merge Schedule is specified and appears on the **Schedule** page.

Add Throttle Schedule

a. Click **Add** and select **Add Throttle Schedule**.

   The **Add New Throttle Schedule** dialog opens.

b. Specify the throughput limit in MB per minutes unit.

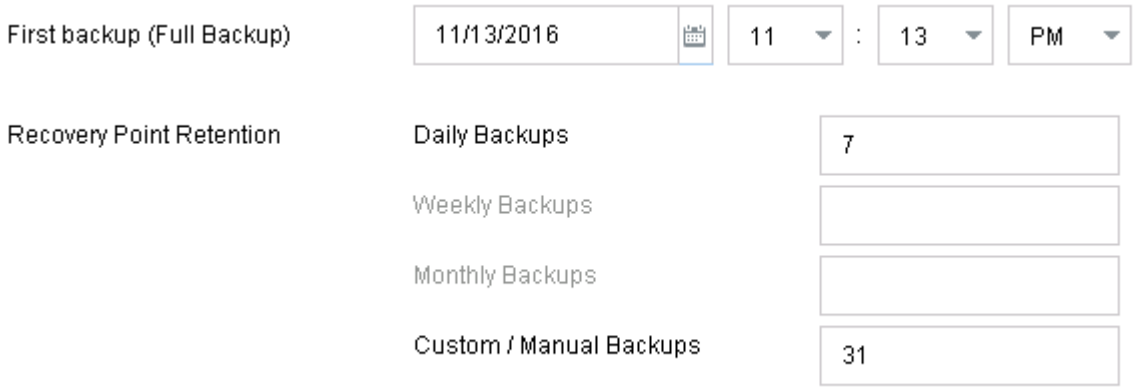
c. Specify the start time to start the backup throughput job.

d. Specify **Until** to specify an end time for the throughput job.

e. Click **Save**.

The Throttle Schedule is specified and displayed on the **Schedule** page.

2. Specify the start time for the scheduled backup.



3. Specify the recovery points retention settings for Custom, Daily, Weekly, and Monthly schedule.

   These options are enabled if you have added the corresponding backup schedule. If you modify the retention settings on this page, the changes are reflected on the **Backup Schedule** dialog.

   The schedule is specified.

Understanding Advanced Scheduling and Retention

The scheduling option lets you specify a Custom schedule or a Daily / Weekly / Monthly schedule, or both the schedules. In the Custom schedule, you can configure the backup schedule for each day of the week and you can add up to four backup schedules each day. You can select a specific day of a week and create a time window to define when to run backup and at what frequency.

Schedule	Supported Job	Comments
Backup	Backup job	Define time windows to run backup jobs.
Backup throttling	Backup job	Define time windows to control the backup speed.
Merge	Merge job	Define when to run merge jobs.
Daily schedule	Backup job	Define when to run daily backup jobs.
Weekly schedule	Backup job	Define when to run weekly backup jobs.
Monthly schedule	Backup job	Define when to run monthly backup jobs.

You can also specify the retention settings for the recovery points.

Note: Set the retention settings within each plan to control how data for the nodes assigned to that plan are retained at the target data store.

Schedules for Daily / Weekly / Monthly backups are independent to the Custom schedule, and each other. You can configure only to run Daily backup, Weekly backup or Monthly backup, without configuring the Custom schedule.

Backup Job Schedule

You can add four time windows per day in your backup schedule. A valid time window is from 12:00 AM until 11:59 PM. You cannot specify a time window such as 6:00 PM to 6:00 AM. In such cases, you must manually specify two different time windows.

For each time window, the start time is inclusive, and the end time is exclusive. For example, you have configured to run Incremental Backup every one hour between 6:00 AM and 9:00 AM and the backup will start at 6:00 AM. This means the backup runs at 6:00 AM, 7:00 AM, 8:00 AM, but NOT at 9:00 AM.

Note: If you want to run the backup job repeatedly until the end of day, set the schedule until 12:00 AM. For example, to run backup job every 15 minutes for the entire day, set the schedule from 12:00 AM to 12:00 AM, every 15 minutes.

Backup Throttle Schedule

Backup throttle schedule lets you control the backup throughput speed that in turn controls the resource usage (disk I/O, CPU, network bandwidth) of the server being backed up. This is useful if you do not want to affect the server

performance during business hours. You can add four time windows per day in your backup throttle schedule. For each time window, you can specify a value, in MB per minute. This value is used to control the backup throughput. Valid values are from 1 MB/minutes to 99999 MB/minutes.

If a backup job extends its specified time, then the throttle limit adjusts according to the specified time window. For example, you have defined the backup throttle limit as 500 MB/minute from 8:00 AM to 8:00 PM, and 2500 MB/minute from 8:00 PM to 10:00 PM. If a backup job starts at 7:00 PM and it runs for three hours, then, from 7:00 PM to 8:00 PM the throttle limit is 500 MB/minute and from 8:00 PM to 10:00 PM the throttle limit is 2500 MB/minute.

If you do not define any backup schedule and backup throughput schedule, the backup runs as fast as it can.

Merge Schedule

Lets you merge recovery points based on the provided schedule.

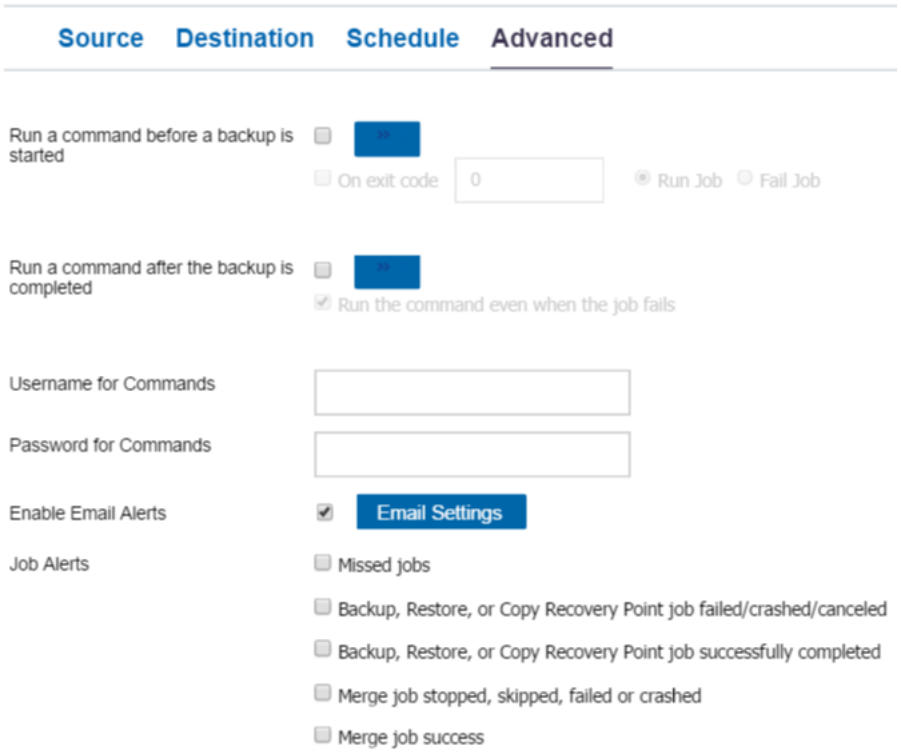
Consider the following points for the merge job:

- At any given time only one merge job can run for a node.
- If a merge job starts, it has to complete before the next merge job can start. This means, if one or more sets of recovery points are merged, new recovery point cannot be added to this merge process, until the merge process of the current set of recovery point completes.
- If a merge job is processing more than one set of recovery points (for example set [1~4], set [5~11], and set [12~14]; they are three sets), recovery point server processes these sets one by one.
- If a merge job is resumed after a pause, the job detects at which point it is paused and resumes the merge from the break-point.

Specify the Advanced Settings

The **Advanced** tab lets you specify some advanced settings for the backup job. The advanced settings include providing the location of any scripts, and email settings.

The following image displays the **Advanced** tab:



Follow these steps:

1. Specify the following details.

   Run a command before a backup is started

   Lets you run a script before the backup job starts. Specify the path where the script is stored inside the proxy node. Click **On exit code** and specify the exit code for **Run Job** or **Fail Job**. **Run Job** indicates that the backup job continues when the script returns the exit code. **Fail Job** indicates that the backup job stops when the script returns the exit code.

   Run a command after a backup is completed

   Lets you run a script after the backup job is completed. Specify the complete path where the script is stored.

   Run a command even when the job fails

If this check box is selected, the script specified in **Run a command after a backup is completed** is executed even when the backup job fails. Otherwise, that script is executed only when backup job completes successfully.

Username for Commands

Lets you specify the username to run the commands.

Password for Commands

Lets you specify the password to run the commands.

Enable Email Alerts

Lets you enable email alerts. You can configure email settings and specify the types of alerts that you want to receive in an email. When you select this option, the following options are enabled for your selection.

Email Settings

Lets you configure the email settings. Click **Email Settings** and configure the email server and proxy server details. For more information about how to configure Email Settings, see Email and Alert Configuration.

Job Alerts

Lets you select the types of job alert emails that you want to receive.

2. Click **Save**.

Note: When you select a node as a backup source or backup proxy, Arcserve UDP checks whether the agent is installed on proxy node and if it is the latest version. Arcserve UDP then displays a verification dialog that lists all the nodes that either have an outdated version of the agent or does not have the agent installed. To install/upgrade the agent on these nodes, select the installation method and click **Save**.

The changes are saved, and a green checkmark appears next to the task name. The plan page closes.

Note: If you have to add another task, you must select the plan from the **resources** tab and modify the plan. To modify the plan, click the plan from the center pane. The plan opens and you can modify it. You may add the **Copy Recovery Point**, **Copy to Tape**, **Replicate**, and **Replicate from a remote RPS** tasks as follow up tasks.

The plan is automatically deployed to the proxy server node.

The exchange online backup plan for the proxy server is created. The backup runs per the schedule that you have configured on the **Schedule** tab. You can also perform a manual backup at any time.

(Optional) Perform a Manual Backup

Typically, backups are performed automatically and are controlled by the schedule settings. In addition to the scheduled backup, a manual backup provides you the option to back up your nodes on a need basis. For example, if you have a repeat schedule for Full, Incremental backup, and you want to make major changes to your machine, you should perform an immediate manual backup without waiting for the next scheduled backup to occur.

To perform a manual backup of Exchange Online nodes, follow these steps:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and click **All Nodes**.

   The Exchange Online nodes are displayed in the center pane.
3. Select the Exchange Online nodes (for example, Mail-box@<organizationname.com>) that you want to backup and that has a plan assigned to it. The node name is the account that is used when adding the Exchange Online node and connecting it.
4. Right-click the node, and then select **Backup Now**.

   The Run a backup now dialog opens.
5. Select the type of backup that you want to perform, and then click **OK**.

   The Progress Information dialog appears.
6. Wait until the backup job completes, and then click **OK**.
7. To monitor the backup jobs, navigate to **jobs** > **Jobs in Progress**.

   Wait until the backup job is finished.

   The manual backup is successfully performed.

To perform a manual backup of an Exchange Online plan, follow these steps:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Plans**, and click **All Plans**.

   The Exchange Online backup plans are displayed in the center pane.
3. Select the plan that you want to backup and that has a plan assigned to it.
4. Right-click the plan, and then select **Backup Now**.

   The Run a backup now dialog opens.
5. Select the type of backup that you want to perform, and then click **OK**.

   The Progress Information dialog appears.

6. Wait until the backup job completes, and then click **OK**.
7. To monitor the backup jobs, navigate to **jobs** > **Jobs in Progress**.

   Wait until the backup job is finished.

   The manual backup is successfully performed.

Verify the Backup

To verify your backup, confirm that you have successfully created the backup plan. After you verify that the plan is created successfully, confirm whether the backup job is running as scheduled. You can verify the status of backup jobs from the **Jobs** tab.

Follow these steps to verify plans:

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and then click **All Nodes**.

   A list of all nodes is displayed on the center pane.
3. Verify that plans are mapped with nodes.

Follow one of these methods to verify backup jobs:

Method 1

1. Click the **jobs** tab.
2. From the left pane, click **All Jobs**.

   The status of each job is listed on the center pane.
3. Verify that the backup job is successful.

   The backup job is verified.

Method 2

1. Click the **jobs** tab.
2. From the left pane, click **Jobs Successfully Completed**.

   The list of backup jobs completed appear on the center pane.
3. Click the job you want to verify.

   The right pane refreshes.
4. Under Job Details, click the **View Logs** hyperlink.
5. Verify that the backup job is successful.

   The backup job is verified.

Restore Microsoft Teams Messages (Chats) using Exchange Online

You can restore Microsoft Teams messages (chats) using Exchange Online from Microsoft cloud using any computer. You can restore data from recovery point to the original or an alternate location.

Perform the following tasks to restore Exchange Online mailbox data:

1. Select the Exchange Online Mailbox Items to Restore
2. Define the Restore Options
3. Restore the Recoverable Items
4. Restore the Recovery Point Content
5. Verify that Content is Restored

Select the Exchange Online Mail Items to Restore

You can restore Exchange Online mail data from a recovery point. When you select a recovery date, and then specify the time, all the associated recovery points for that duration are displayed. You can then browse and select the backup content (including applications) to restore.

Follow these steps:

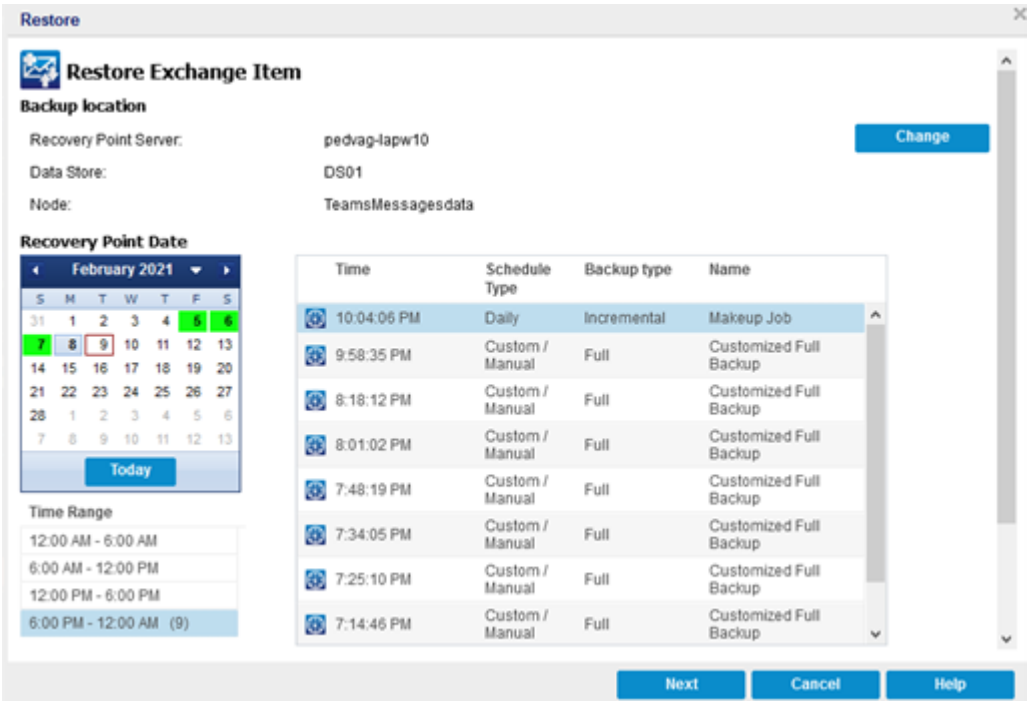
1. Log into Arcserve UDP.
2. Click the **resources** tab.
3. Select **All Nodes** in the left pane.

   All the added nodes are displayed in the center pane.
4. In the center pane, select the Exchange Online node and then click **Actions**.
5. Click **Restore** from the Actions drop-down list.

   The Restore Exchange Item dialog opens.

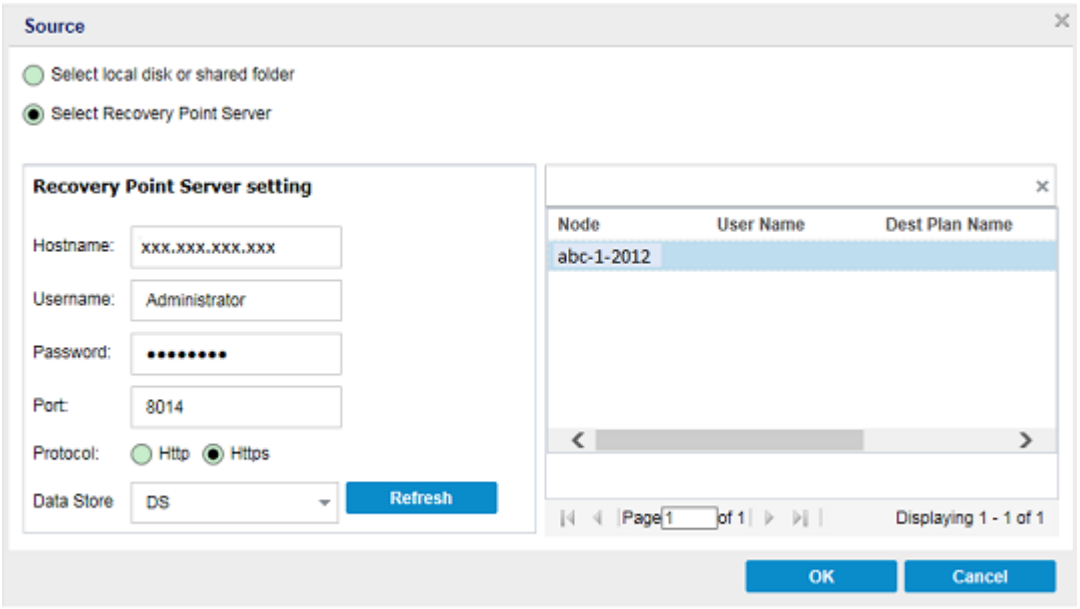
   Note: You are automatically logged in to the agent node and the **Restore Exchange Item** dialog opens.

   You can see the Recovery Point Server details in the Backup Location.



6. (Optional) To modify the backup location, click **Change** .

   The Source dialog opens. You can select the backup location in this dialog.



7. To specify Source, select one of the following options, and then click **OK**:

   Select local disk or shared folder

   Note: In Arcserve UDP, we do not recommend selecting the **Select local disk or shared folder** option.

   Select Recovery Point Server

   a. Specify the Recovery Point Server setting details, and then click **Refresh**.

      All the agents are listed in the Data Protection Agent column in the Source dialog.

   b. Select the agent from the displayed list, and then click **OK**.

      The recovery points are listed in the Restore Exchange Item dialog.

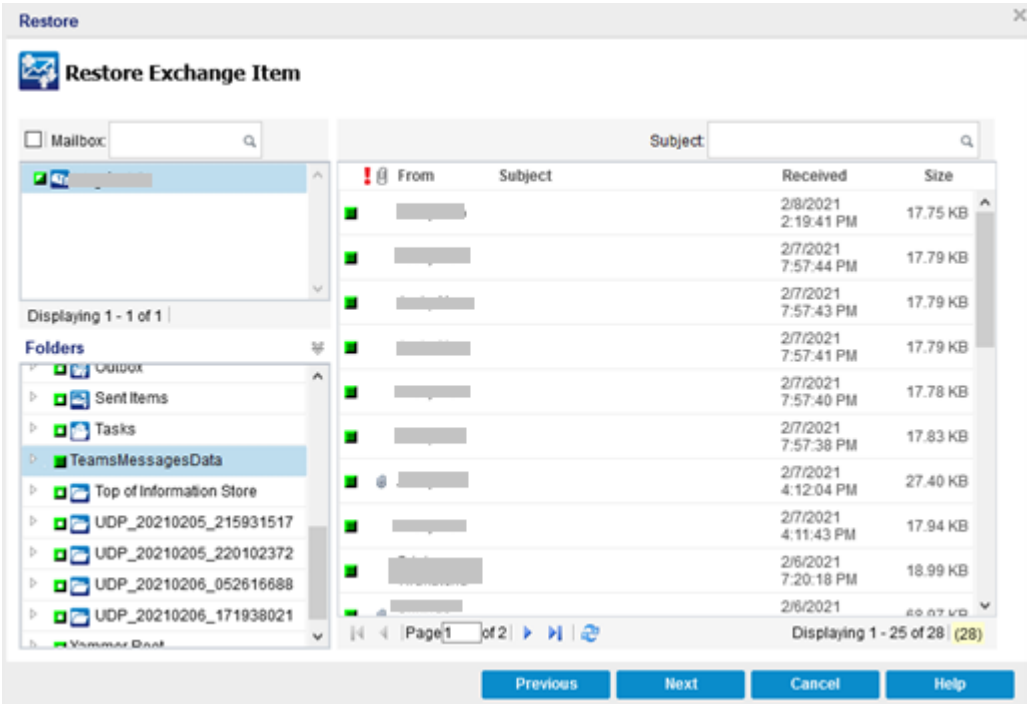
      Note: From the recovery point **Folder**, you may see many folders with the same name along with the exchange online nodes. This happens because the node GUID changes and a new recovery point folder is created when you delete a node and add again in the test plan.

8. Select the calendar date for the backup image to restore, and then click **Next**.

   All the dates containing recovery points for the specified backup source are high-lighted in green.

   The corresponding recovery points for that date are displayed, with the time of the backup, the type of backup that was performed (Full, Incremental, or Verify), and the name of the backup.

9. From the **Mailbox** pane, click the mailbox that you want to restore.



All mail items related to the mailbox are displayed in the Folders pane.

10. Select the **TeamsMessagesData** folder, and then click **Next**.

    Notes:

    - You can select the entire content or partial content of the Exchange object to restore. To select partial content, expand the object, and click the check box of that content.
    - You can select multiple Exchange objects to restore.

    The Restore Options dialog opens.

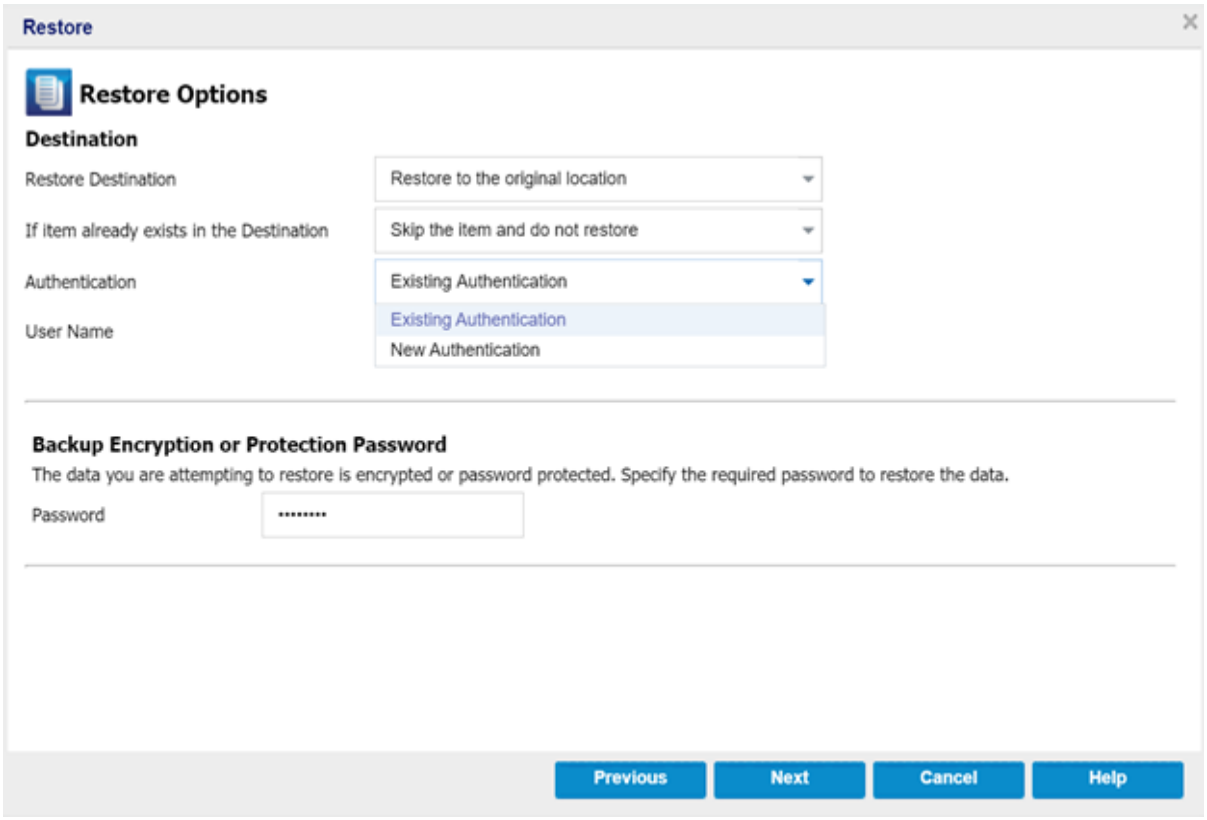
    The Exchange Online mail items to restore are selected. Now, you can define the restore options.

Define the Restore Options

After you specify the Exchange Online information to restore, define the restore options for the selected content.

Follow these steps:

1. On the Restore Options dialog, select the restore destination.



The available destination options are:

Restore to original location

Restores the mail data to the same location from where you took a backup.

Notes:

- If you restore to original location, the data does not reflect to the Teams messages. This is not a complete restore to original location as you cannot perform the restore of root folder.
- If you restore a mail item to original location using the overwrite option twice, then after restoring second time, the first restore mail

item is not overwritten. As a result, two similar mail items appear in original location.

- If you restore a mail item to original location using the skip option, and in the original folder an item similar to the restored item already exists, then the backup job displays incomplete result.

Restore to an alternate location

Restores the mail data to another mailbox or another folder in the original mailbox. When you select this option, you can browse and select the destination.

2. Specify one of the following options from the **If Item already exists in the Destination** drop-down:

Skip the item and do not restore

Skips over the items and does not restore.

Default: The Skip the item and do not restore.

Overwrite the item in the destination

Overwrites the item in the destination.

Note: The If Item already exists in the Destination drop-down list is available if you select the **Restore to original location** option from the Restore Destination drop-down list.

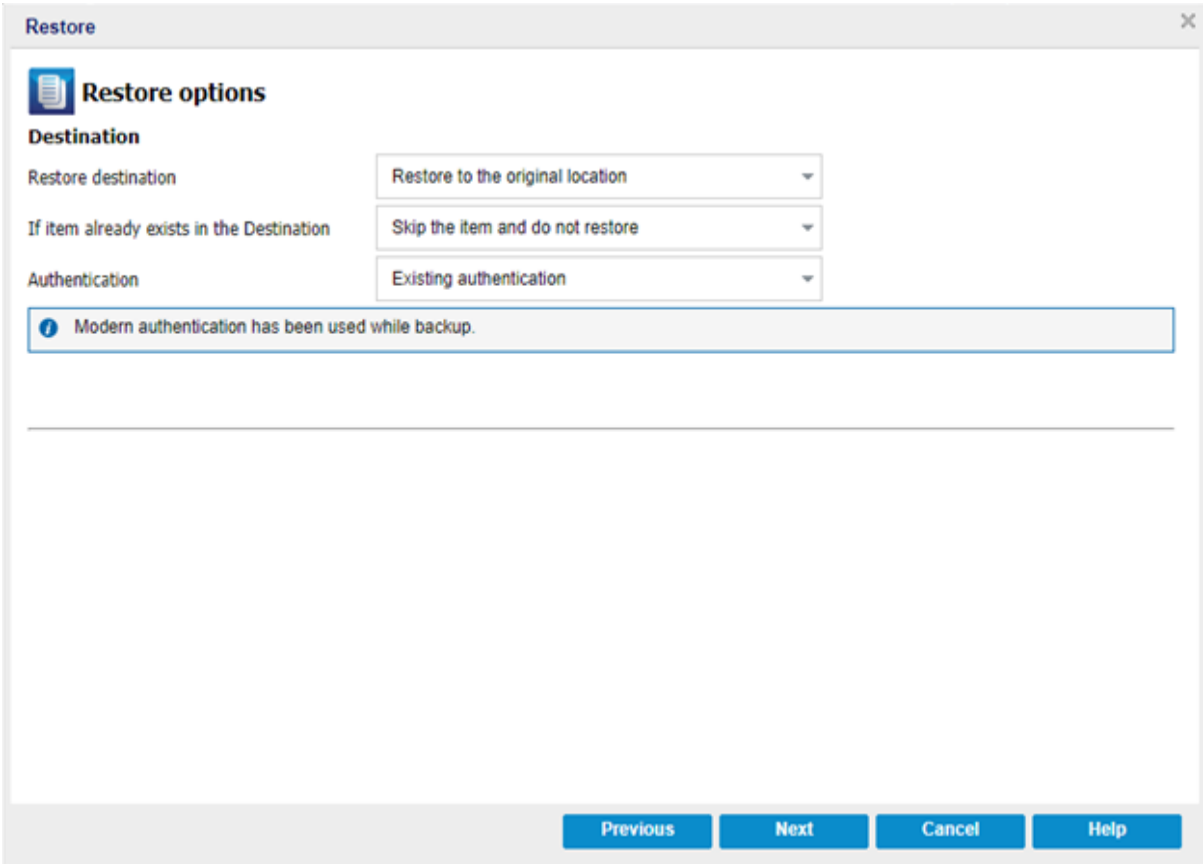
Authentication

Validates and identifies the user account. Select one of the following options:

- **Existing Authentication**

  When you select Existing Authentication, the authentication type that was selected during Node creation is used.

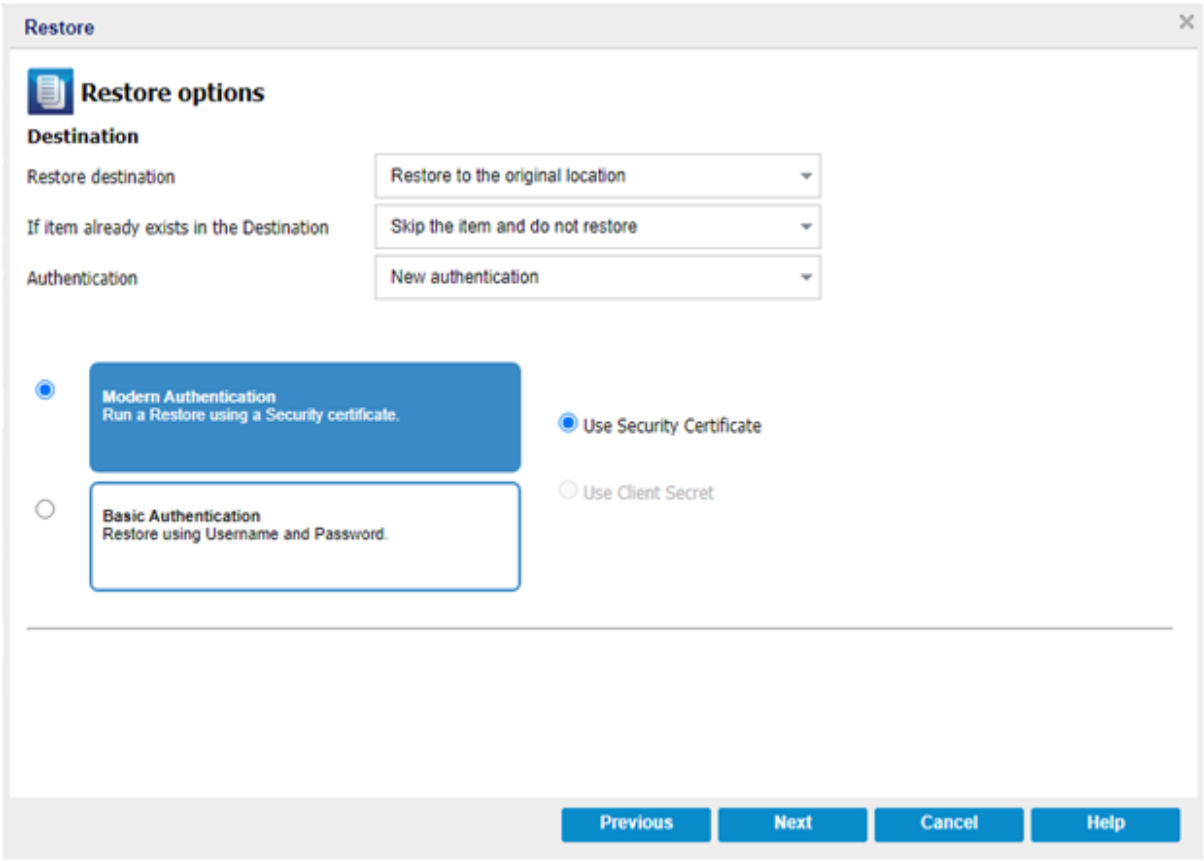
  - If the Node Authentication is basic, parameters such as Username, Password and Backup Encryption or Protection Password [If Encrypted] are loaded automatically to proceed with the restore.
  - If the Node Authentication is modern, parameters such as Username and Backup Encryption or Protection Password [If Encrypted] are loaded automatically to proceed with the restore.



- **New Authentication**

  During restore, you can select one of the following to change the authentication type:

  - Modern Authentication
  - Basic Authentication



3. To apply authentication, do one of the following:

   Modern Authentication

   For modern authentication, you can run a restore using security certificate.

   Use security certificate

   Note: By default, the backup plan selects **Use security certificate** for modern authentication.

   Follow these steps:

   a. To authenticate using a self-signed certificate, see How to Use Self-Signed Certificate.

   b. Enter the Application ID and Username.

   c. To apply the required roles and permissions, select the **Allow Arcserve UDP to set the required Azure roles and permissions for this application** check box, and then click **Next**.

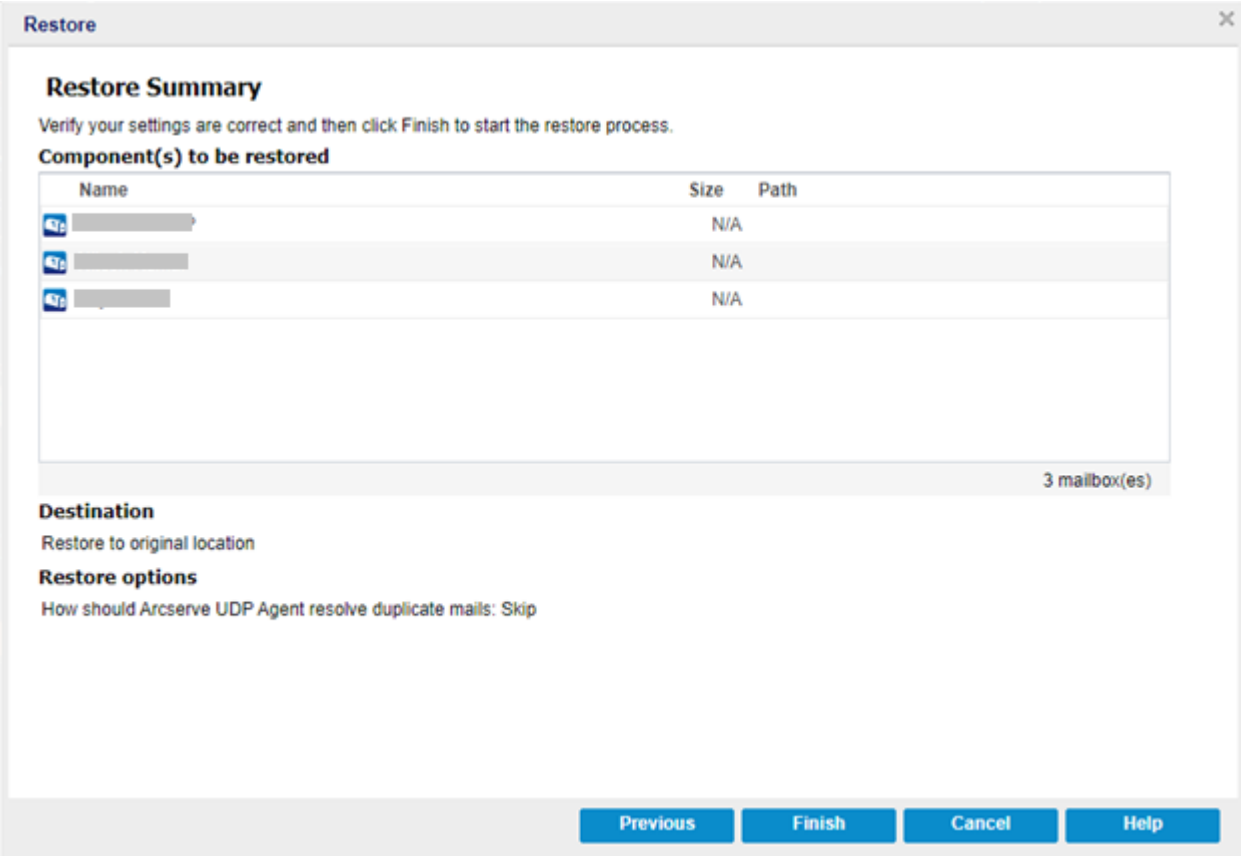
   Basic Authentication

   For Basic Authentication, to proceed with the restore, enter User Name and Password.

Note: This option fails if you have Modern Authentication enabled on the Tenant or Account.

4. Click **Next**.

   The Restore Summary dialog opens.



5. On the Restore Summary window, verify the settings, and then click **Finish**.

The restore options are defined to restore the Exchange Online information.

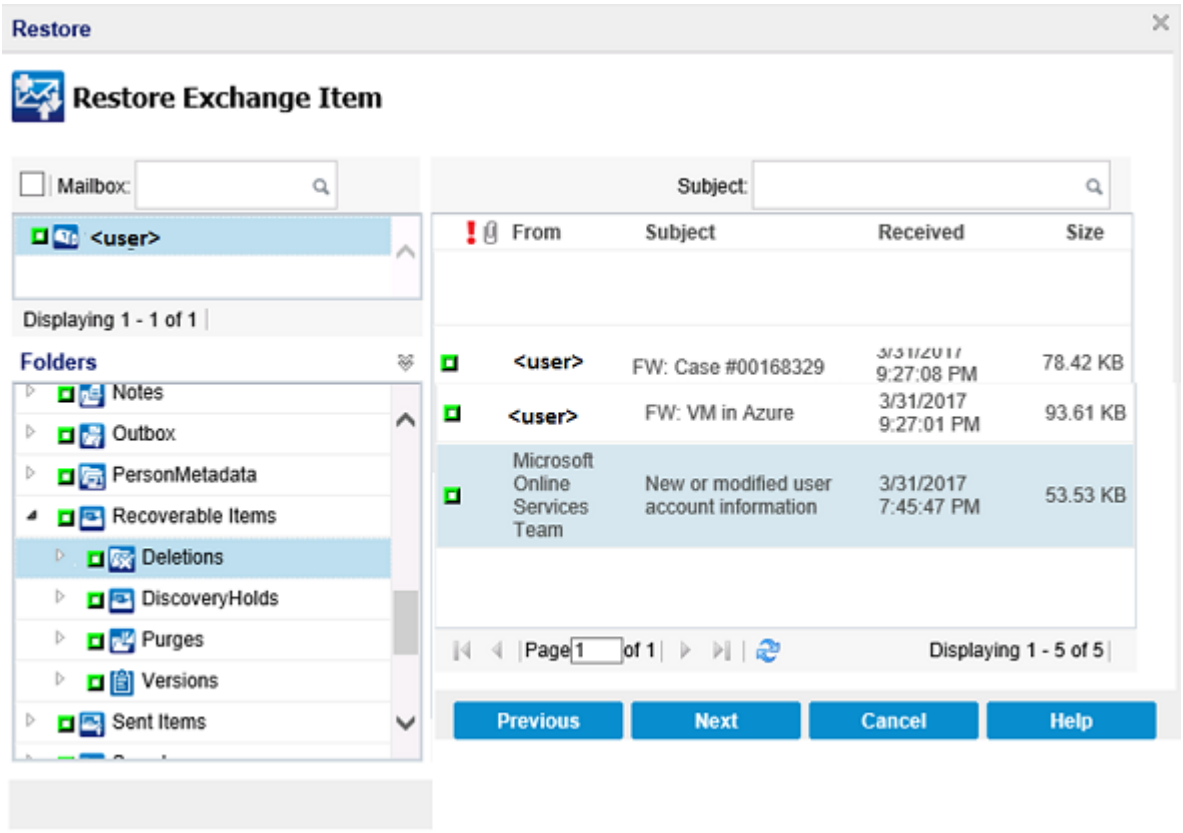
Restore the Recoverable Items

You can restore the recoverable items from the mailbox that enables the In-Place Hold or Litigation Hold feature, from Exchange online node UI to browse backup session. To enable this feature, see Specify the Source.

After you enable the feature, using the Restore wizard you can restore the recoverable items.

Follow these steps:

1. On the Restore Exchange Item dialog, select required folders under **Recoverable Items**, and then click **Next**.



The Restore Option screen is displayed.

2. From the Restore Option screen, perform the following details, and then click **Next**:
   - Select Restore Destination.
   - Skip or overwrite if the item already exists in the destination.

     You can perform either Original or Alternate restore.

     Original restore

The items under Recoverable Items are restored at the *UDP_RecoverableItems_yyyyMMdd_HHmmssfff* folder that is created on the target mailbox. Skip and Overwrite options for restore are not applicable to these items.

Alternate restore

As with user mailbox, restores the selected items to the destination path in a unique time stamped folder, such as *titled /restore_yyyyMMdd-mmssff*.

- Provide user id and password of the destination where you want to restore.

The selected recoverable items are restored.

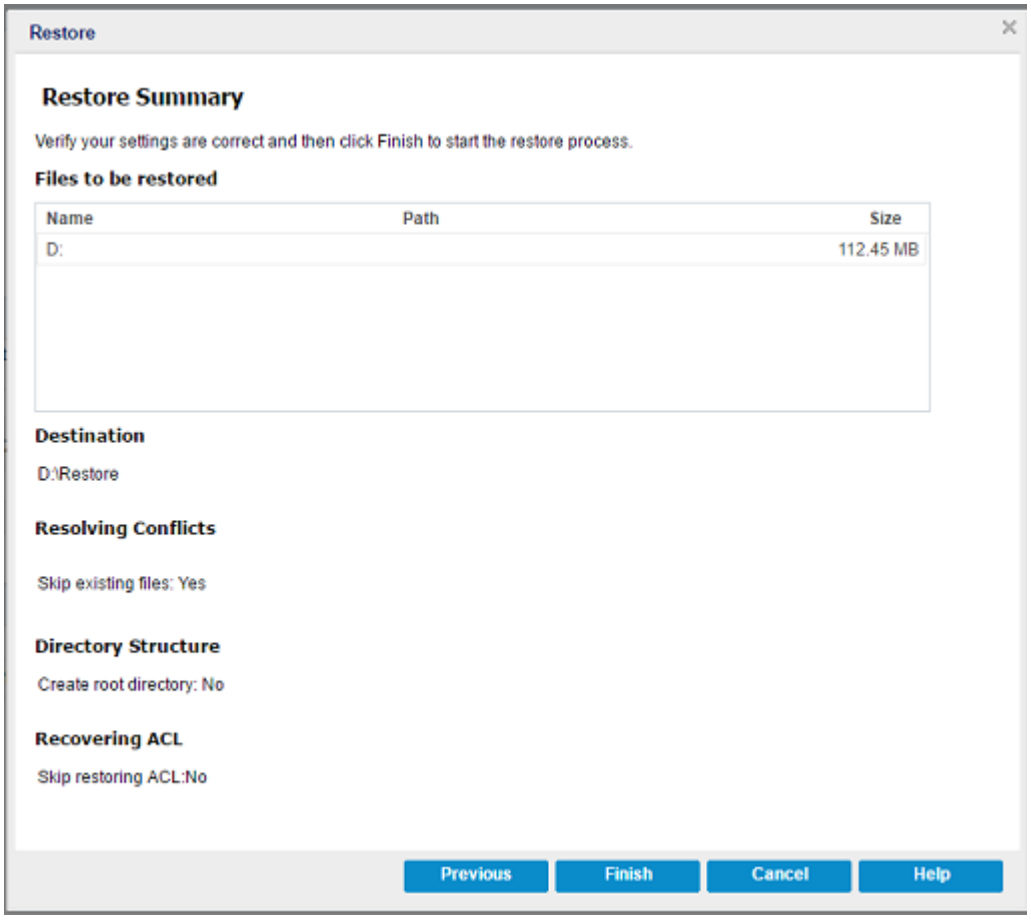
The recovery point content is stored.

Restore the Recovery Point Content

After you define the restore options, verify that your settings are correct and confirm the restore process. **Restore Summary** helps you to review all the restore options that you defined and modify them if necessary.

Follow these steps:

1. On the Restore Summary dialog, review the displayed information to verify that all the restore options and settings are correct.



- If the summary information is incorrect, click **Previous** and go back to the applicable dialog to change the incorrect setting.
- If the summary information is correct, click **Finish** to launch the restore job.

The recovery point content is stored.

Verify that Content is Restored

After the completion of the restore process, verify that content was restored to the specified destination.

Follow these steps:

1. Log into the destination mailbox.
2. Check the mailbox item that you restored.
3. Verify the restored content.

   The restored content is successfully verified.

How to Backup and Restore Microsoft Teams Files using SharePoint Online

Whenever you create a private, public, and group channel in Microsoft Teams, a SharePoint Online site gets created for it automatically. This SharePoint Online site has a default document library folder, in which all the files shared across all the conversations get saved automatically. You can also customize the permissions and security options for sensitive information.

When you create any group in Microsoft Teams, a SharePoint Online site with the **Team – Group_name** name is created automatically. All the files shared in the public channels are saved in this site.

When you create any private channel under the Microsoft Teams group, another SharePoint Online site with the **Teams_Group_name – Private_Channel_name** name is created for that private channel apart from the site created for the Team group.

This section provides information about how to backup and restore Microsoft Teams files using SharePoint Online.

This section contains the following topics:

Backup Microsoft Teams Files using SharePoint Online

This section provides information about how to backup Microsoft Teams files using SharePoint Online.

This section contains the following topics:

How to Create a SharePoint Online Backup Plan

The SharePoint Protection is used to backup and restore Microsoft SharePoint Online site and list item. The SharePoint Online is one of the major products in Microsoft Office 365. To protect your SharePoint content, you need to create a Plan.

This section contains the following topics:

Review the Prerequisites

Verify the following prerequisites before performing a backup and restore:

- You have the SharePoint Site Collection URL to backup.
- The backup account is a member of Site Collection Administrators groups or assigned with *SharePoint admin* Role.

  To add an account to Site Collection Administrators group, refer to the link.
- Review the Prerequisites before you create a plan using Modern Authentication.

Create a SharePoint Online Backup Plan

A backup plan includes a **Backup: Office 365 SharePoint Online** task that performs a backup of SharePoint Online items (Lists, Libraries, Documents, and so on) and stores data at a deduplication or non-deduplication data store. Additionally, this backup plan allows UDP to discover and specify SharePoint Online items for which a backup task is performed. Each task consists of parameters that define the source, destination, schedule, and other backup details.

Watch video and view how to create the SharePoint Online Backup plan.

Follow these steps:

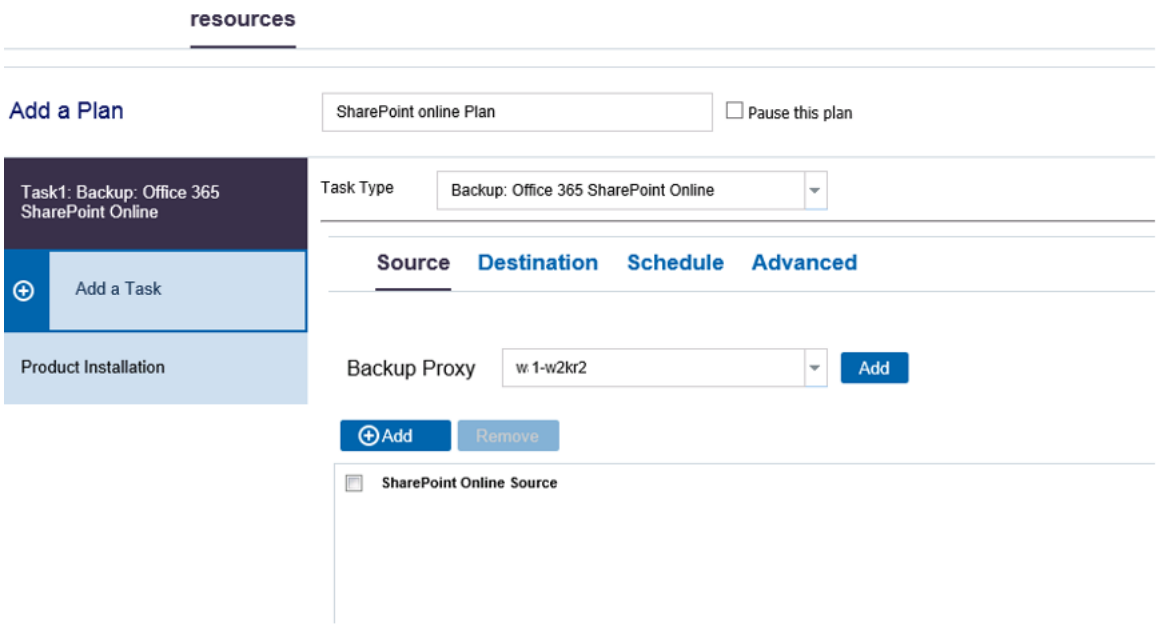
1. Click the **resources** tab on the Console.
2. From the left pane, navigate to **Plans**, and then click **All Plans**.

   If you have created plans earlier, those plans are displayed on the center pane.
3. On the center pane, click **Add a Plan**.

   The Add a Plan dialog opens.
4. Enter a plan name.
5. (Optional) Select the **Pause this plan** check box.

   The plan does not run until you clear the check box to resume the plan.

   Note: If a plan is paused, then any in-progress job is not paused but all corresponding scheduled jobs associated with that plan are paused. However, you can manually run a job. For example, you can manually run backup job and replication job for a node even if the respective plan is paused. In such case, the following task to the on-demand (manual) job does not run. For example, there is a replication task after an on-demand backup job, the replication job does not run for the on-demand backup job. You need to manually run the replication job. When you resume the plan, the pending jobs do not resume immediately. After you resume the plan, the pending jobs run from the next scheduled time.
6. From the Task Type drop-down list, select **Backup: Office 365 SharePoint Online**.



Now specify the Source, Destination, Schedule, and Advanced details.

Specify the Source

The Source page lets you specify the SharePoint Online source nodes that you want to protect. You can select more than one SharePoint Online source nodes in a plan.

Note: You can save a plan without adding any source nodes, but the plan is not deployed unless you add a node.

Follow these steps:

1. Click the **Source** tab.
2. Select the **Backup Proxy** from the drop-down list.
3. Add SharePoint Online node using one of the following methods:
   - **Method 1:**
     a. Click **Add**, and then click **Select Source to Protect in Arcserve UDP**.

        The Add Nodes to Plan dialog is displayed.

        Note: Use this option only if you have already added SharePoint node.

     b. Select a node, and then click **Connect**.

        Note: To find the SharePoint Online nodes that you want to protect, use **Search**.
   - **Method 2:**

     Click **Add**, and then click **Add SharePoint Online Source**.

     The Add nodes to a plan dialog appears.

     Note:

     - Use this option to discover and specify online items for SharePoint Online Source.
     - Unlike other nodes, you cannot add the SharePoint Online node from the All Nodes page. You can add a SharePoint Online node only in a plan while creating or modifying a plan.
4. To apply authentication, do one of the following:

   Modern Authentication

   For modern authentication, you can run a backup using security certificate.

Use security certificate

Note: By default, the backup plan selects **Use security certificate** for modern authentication.

Follow these steps:

1. To authenticate the UDP application with security certificates, see How to Use Self-Signed Certificate.
2. Enter the following details:
   - Node Name
   - App ID
   - User Name
3. To apply the required roles and permissions, select the **Allow Arcserve UDP to set the required Azure roles and permissions for this application** check box, and then click **Connect**.

Notes:

- SharePoint Online Source node names cannot change after creating the node.
- You can use a single account to protect multiple SharePoint Online nodes.
- Fill in the app password if Multi-Factor Authentication is enabled and the tenant is set to use Basic Authentication.

Basic Authentication

For Basic Authentication, to proceed with the backup, provide the following information, and then click **Connect**:

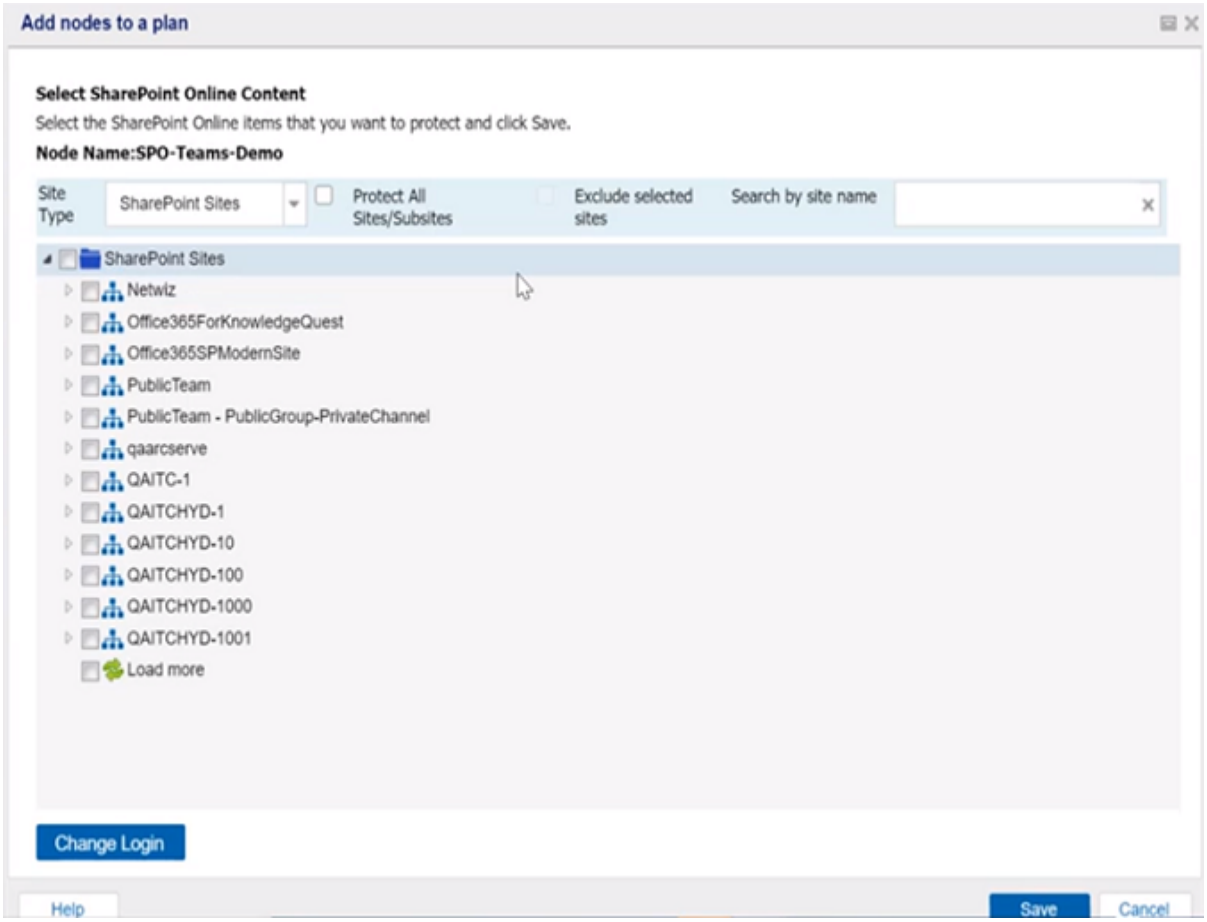
- Node name
- Site Collection URL
- User ID
- Password

5. To protect SharePoint Online items, do the following:
   - From the Site Type drop-down list, select the type of site you want to protect such as SharePoint Sites or Personal Sites.
   - For Protect All Sites/Subsites, perform the following:

     Notes:

- This option is automatically selected by default when the Select SharePoint Online Content screen appears.
  - To protect all items in the selected site type, select the **Protect All Sites/Subsites** check box.
  - To protect specific items in the selected site type, clear the **Protect All Sites/Subsites** check box, and then select the specific items that you want to protect.
- While configuring the existing SharePoint Online plan, if you have selected sites manually for protection and now want to select the **Protect All Sites/Subsites** check box, follow these steps:
  a. Modify the existing SharePoint Online plan.
  b. Clear the SharePoint sites that are manually selected for protection.
  c. Select the **Protect All Sites/Subsites** check box, and then save the plan.

- To exclude specific items from being protected, select the **Exclude selected sites** check box, and then select specific items in the selected site type.



6. Click **Save**.

   The SharePoint Online sources that you want to protect are added to the plan.

Specify the Destination

The destination is a location where you store the backup data. You must at least specify the destination to save the plan.

Follow these steps:

1. Click the **Destination** tab.

   The Arcserve UDP Recovery Point Server option specifies that the backup destination is a recovery point server.

   You cannot store data as recovery sets.

2. Perform the following steps:

   a. Select a recovery point server.

   b. Select a non-deduplication or deduplication data store.

      The list displays all the data stores created on the specified recovery point server.

   c. Provide a session password.

      The session password is optional when the backup destination is an unencrypted RPS data store.

   d. Confirm the session password.

The destination is specified.

Specify the Schedule

The Schedule page lets you define a schedule for Backup, Merge, and Throttle functions to repeat at specific intervals. After you define a schedule, the jobs run automatically per the schedule. You can add multiple schedules and provide retention settings.

A Backup Schedule refers to regular schedule that is repeated multiple times a day based on the number of hours or minutes you select. Besides the regular schedule, a backup schedule also provides options to add daily, weekly, and monthly schedules.

Note: For more information on scheduling and retention settings, see Understanding Advanced Scheduling and Retention.

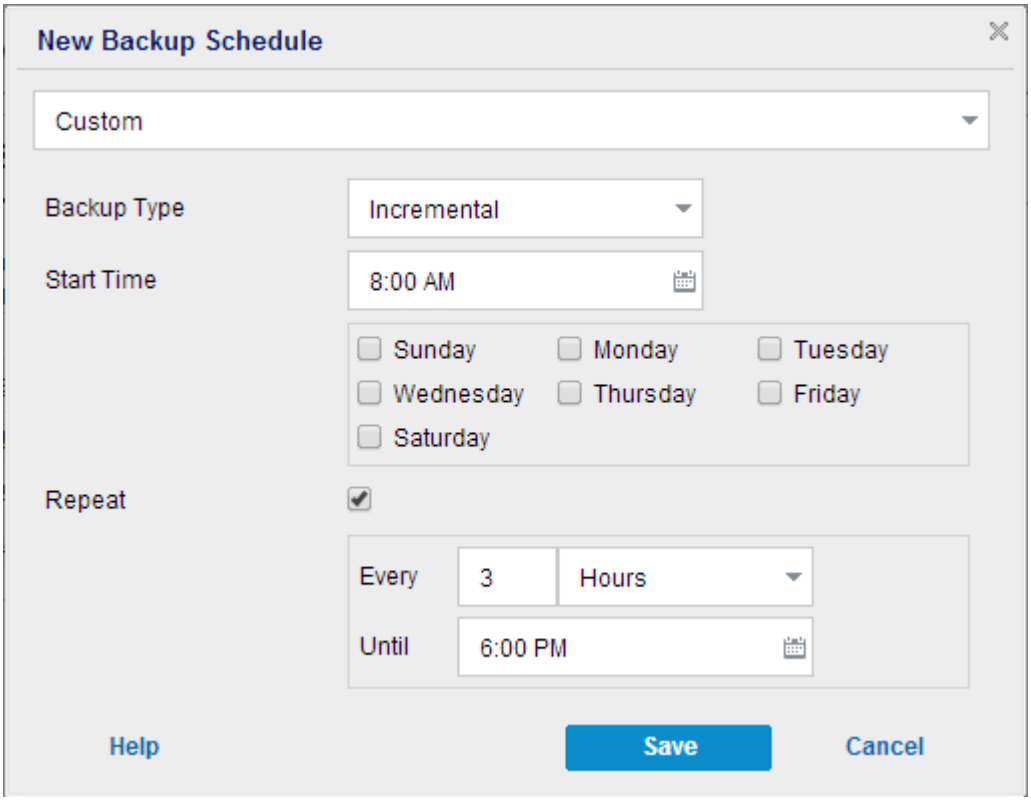
Follow these steps:

1. Add backup, merge, and throttle schedules.

   Add Backup Schedule

   a. Click **Add** and select **Add Backup Schedule**.

      The **New Backup Schedule** dialog opens.



   b. Select one of the following options:

      Custom

Specifies the backup schedule that repeats multiple times a day.

Daily

Specifies the backup schedule that occurs once a day. By default, all the days of the week are selected for Daily backup. If you do not want to run the backup job on a specific day, clear the check box for that day of the week.

Weekly

Specifies the backup schedule that occurs once a week.

Monthly

Specifies the backup schedule that occurs once a month.

c. Select the backup type.

Full

Determines the backup schedule for Full Backups. As scheduled, Arcserve UDP performs a full backup of all used blocks from the source machine. A full backup typically consumes time depending on the backup size.

Incremental

Determines the backup schedule for Incremental Backups.

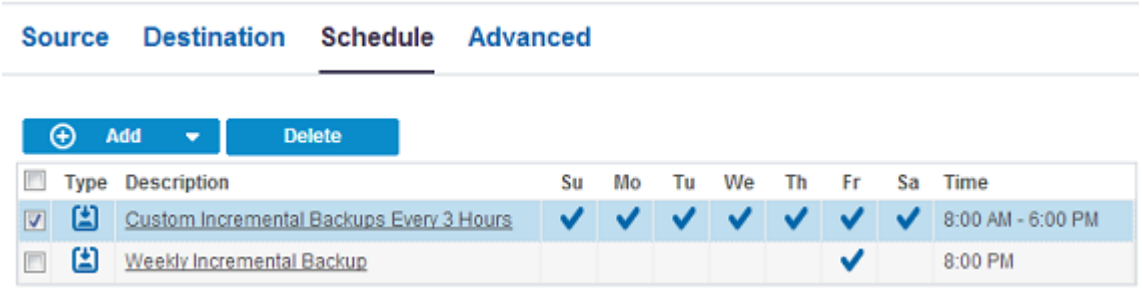
As scheduled, Arcserve UDP incrementally backs up only those blocks that have changed after the last successful backup. The advantages of Incremental Backups are that the backup is fast and produces a small backup image. This is the most optimal way to perform a backup.

d. Specify the backup start time.

e. (Optional) Select the **Repeat** check box and specify the repeat schedule.

f. Click **Save**.

The Backup Schedule is specified and appears on the **Schedule** page.



Add Merge Schedule

a. Click **Add** and select **Add Merge Schedule**.

   The **Add New Merge Schedule** dialog opens.

b. Specify the start time for the merge job.

c. Specify **Until** to provide an end time for the merge job.

d. Click **Save**.

The Merge Schedule is specified and appears on the **Schedule** page.

Add Throttle Schedule

a. Click **Add** and select **Add Throttle Schedule**.

   The **Add New Throttle Schedule** dialog opens.

b. Specify the throughput limit in MB per minutes unit.

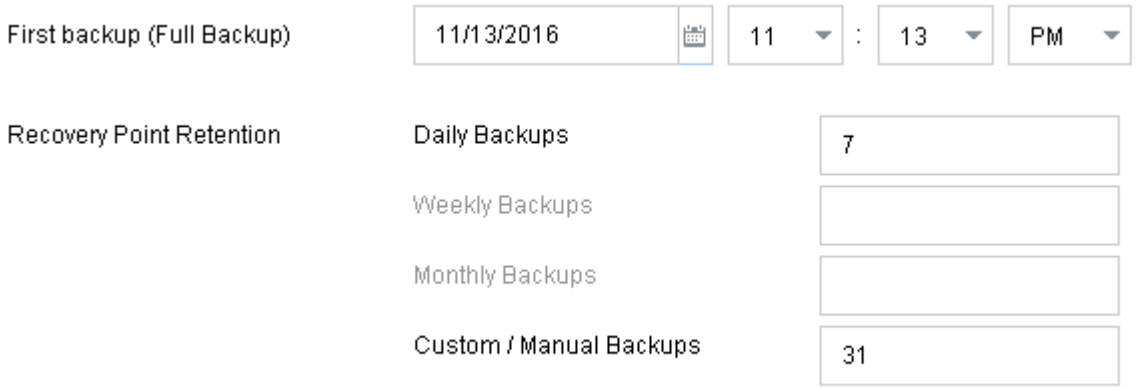
c. Specify the start time for the backup throughput job.

d. Specify **Until** to provide an end time for the throughput job.

e. Click **Save**.

The Throttle Schedule is specified and displayed on the **Schedule** page.

2. Specify the start time for the scheduled backup.



3. Specify the recovery points retention settings for Custom, Daily, Weekly, and Monthly schedule.

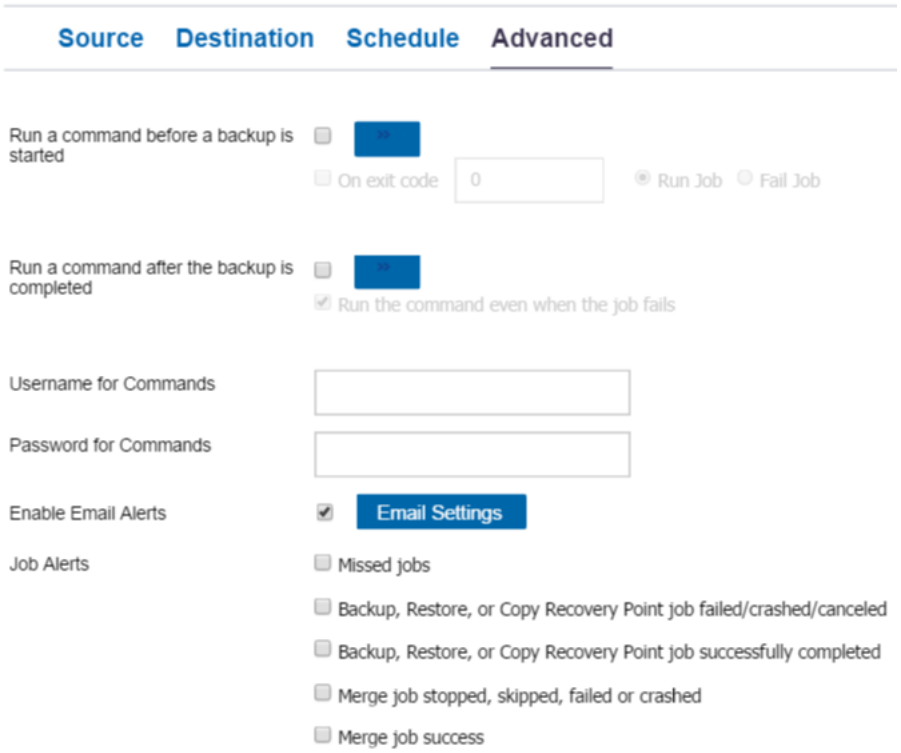
   The options are enabled if you have added the corresponding backup schedule. If you modify the retention settings on this page, the changes are reflected on the **Backup Schedule** dialog.

   The schedule is specified.

Specify the Advanced Settings

The **Advanced** tab lets you specify some advanced settings for the backup job. The advanced settings include providing the location of any scripts, and email settings.

The following image displays the **Advanced** tab:



Follow these steps:

1. Specify the following details.

   Run a command before a backup is started

   Lets you run a script before the backup job starts. Specify the path where the script is stored inside the proxy node. Click **On exit code** and specify the exit code for **Run Job** or **Fail Job**. **Run Job** indicates that the backup job continues when the script returns the exit code. **Fail Job** indicates that the backup job stops when the script returns the exit code.

   Run a command after a backup is completed

   Lets you run a script after the backup job is completed. Specify the complete path where the script is stored.

   Run a command even when the job fails

If this check box is selected, the script specified in **Run a command after a backup is completed** is executed even when the backup job fails. Otherwise, that script is executed only when backup job completes successfully.

Username for Commands

Lets you specify the username to run the commands.

Password for Commands

Lets you specify the password to run the commands.

Enable Email Alerts

Lets you enable email alerts. You can configure email settings and specify the types of alerts that you want to receive in an email. When you select this option, the following options are enabled for your selection.

Email Settings

Lets you configure the email settings. Click **Email Settings** and configure the email server and proxy server details. For more information about how to configure Email Settings, see Email and Alert Configuration.

Job Alerts

Lets you select the types of job alert emails that you want to receive.

2. Click **Save**.

Note: When you select a node as a backup source or backup proxy, Arcserve UDP checks whether the agent is installed on proxy node and if it is the latest version. Arcserve UDP then displays a verification dialog that lists all the nodes that either have an outdated version of the agent or does not have the agent installed. To install/upgrade the agent on these nodes, select the installation method and click **Save**.

The changes are saved, and a green checkmark appears next to the task name. The plan page closes.

Note: If you have to add another task, you must select the plan from the **resources** tab and modify the plan. To modify the plan, click the plan from the center pane. The plan opens and you can modify it. You may add the **Copy Recovery Point**, **Copy to Tape**, **Replicate**, and **Replicate from a remote RPS** tasks as follow up tasks.

The plan is automatically deployed to the proxy server node.

The SharePoint Online backup plan for the proxy server is created. The backup runs per the schedule that you have configured on the **Schedule** tab. You can also perform a manual backup at any time.

(Optional) Perform a Manual Backup

Typically, based on the schedule settings you applied, the backups perform automatically. In addition to the scheduled backup, a manual backup provides you the option to back up your nodes on a need basis. For example, if you have a repeat schedule for Full and Incremental backups, and you want to make major changes to your machine, you can perform an immediate manual backup without waiting for the next scheduled backup.

To perform a manual backup of SharePoint Online nodes, follow these steps:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and then click **All Nodes**.

   The SharePoint Online nodes are displayed in the center pane.
3. Select the SharePoint Online nodes (for example, Mailbox@<organizationname.com>) that you want to backup and make sure a plan is assigned to it. The node name is the account that is used when adding the SharePoint Online node and connecting it.
4. Right-click the node, and then select **Backup Now**.

   The Run a backup now dialog opens.
5. Select the type of backup that you want to perform, and then click **OK**.

   The Progress Information dialog appears.
6. Wait until the backup job completes, and then click **OK**.
7. To monitor the backup jobs, navigate to **jobs** > **Jobs in Progress**.

   Wait until the backup job is finished.

   The manual backup is successfully performed.

To perform a manual backup of SharePoint Online plan, follow these steps:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Plans**, and then click **All Plans**.

   The SharePoint Online backup plans are displayed in the center pane.
3. Select the plan that you want to backup and make sure a plan is assigned to it.
4. Right-click the plan, and then select **Backup Now**.

   The Run a backup now dialog opens.
5. Select the type of backup that you want to perform, and then click **OK**.

   The Progress Information dialog appears.

6. Wait until the backup job completes, and then click **OK**.
7. To monitor the backup jobs, navigate to **jobs** > **Jobs in Progress**.

   Wait until the backup job is finished.

   The manual backup is successfully performed.

Verify the Backup

To verify your backup job, confirm that you have successfully created the backup plan. After you verify that the plan is created successfully, confirm whether the backup job is running as scheduled. You can verify the status of backup jobs from the **Jobs** tab.

Follow these steps to verify plans:

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and then click **All Nodes**.

   A list of all nodes is displayed on the center pane.
3. Verify that plans are mapped with nodes.

Follow one of these methods to verify backup jobs:

Method 1

1. Click the **jobs** tab.
2. From the left pane, click **All Jobs Completed**.

   The status of each job is listed on the center pane.
3. Verify that the backup job is successful.

   The backup job is verified.

Method 2

1. Click the **jobs** tab.
2. From the left pane, click **Jobs Successfully Completed**.

   The list of backup jobs completed appear on the center pane.
3. Click the job you want to verify.

   The right pane refreshes.
4. Under Job Details, click the **View Logs** hyperlink.
5. Verify that the backup job is successful.

   The backup job is verified.

Restore SharePoint Online Site Collection Data

You can restore SharePoint Online List/Library or List item in Site. The Site Collection and Site are not supported yet in Arcserve UDP 8.0. You can restore the data to original site with new name, restore the data to original location and export to disk from the recovery points.

Perform the following tasks to restore SharePoint Online List item:

1. Select the SharePoint Online site list Items to Restore
2. Define the Restore Options
3. Verify that Content is Restored

Select the SharePoint Online Site List Items to Restore

You can restore SharePoint Online list items data from a recovery point. When you select a recovery date, and then specify the time, all the associated recovery points for that duration are displayed. You can then browse and select the backup content (including applications) to restore.

Follow these steps:

1. Log into Arcserve UDP.
2. Click the **resources** tab.
3. Select **All Nodes** in the left pane.

   All the added nodes are displayed in the center pane.

   Or

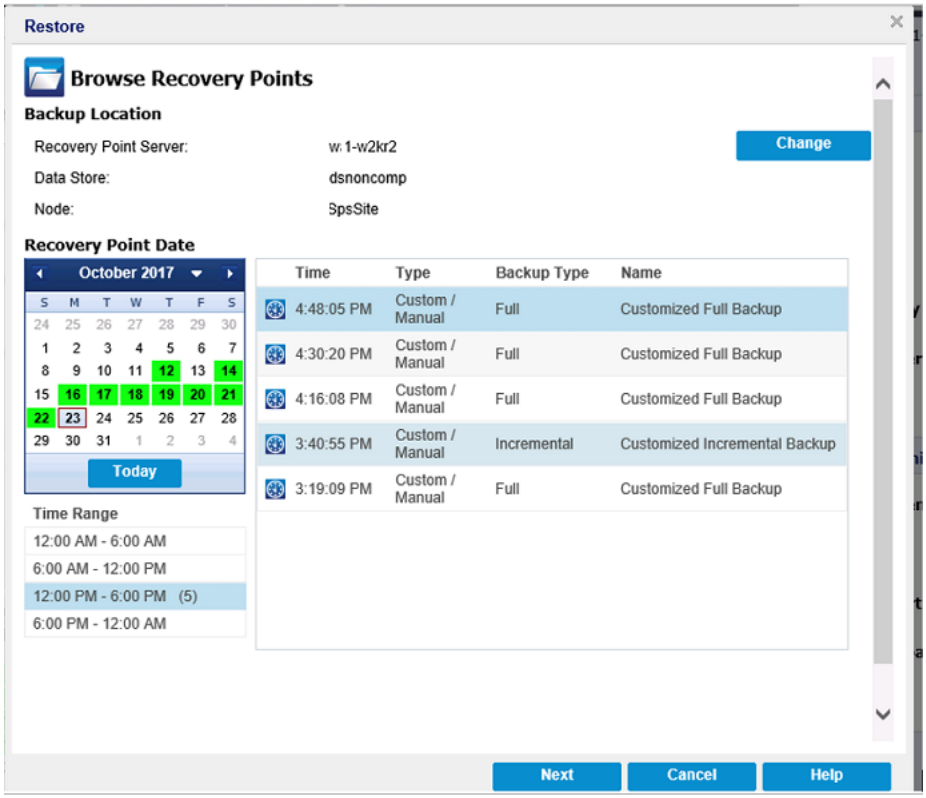
   Select **SharePoint Online Nodes** group.

   All the added SharePoint nodes are displayed in the center pane.
4. In the center pane, select the SharePoint Online node, and then click **Actions**.
5. Click **Restore** from the Actions drop-down menu.

   The Restore SharePoint Item dialog opens.

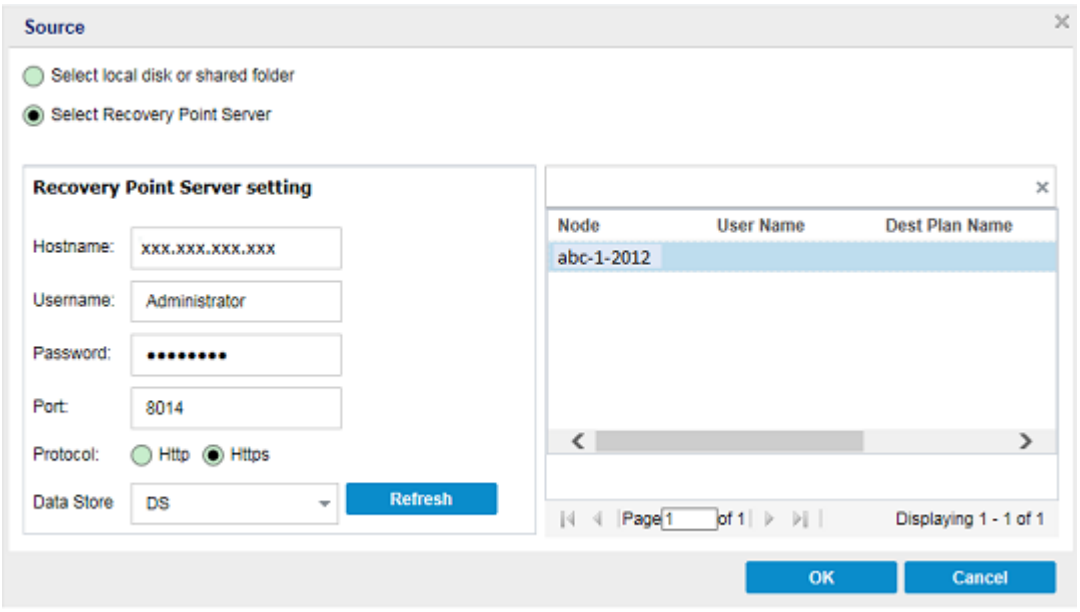
   Note: You are automatically logged into the agent node and the Restore SharePoint Item dialog opens.

   The Backup Location displays the Recovery Point Server details.



6. (Optional) Click **Change** to modify the backup location.

   The Source dialog opens. You can select the backup location in this dialog.



7. To specify Source, select one of the following options, and then click **OK**:

   Select local disk or shared folder

Note: In Arcserve UDP, we do not recommend selecting the **Select local disk or shared folder** option.

Select Recovery Point Server

a. Specify the Recovery Point Server setting details, and then click **Refresh**.

   All the agents are listed in the Data Protection Agent column in the Source dialog.

b. Select the agent from the displayed list/library, and then click **OK**.

   The recovery points are listed in the Restore SharePoint Item dialog.

8. Select the calendar date for the backup image to restore, and then click **Next**.

   All the dates containing recovery points for the specified backup source are high-lighted in green.

   The corresponding recovery points for that date are displayed, with the time of the backup, the type of backup that was performed (Full or Incremental), and the name of the backup.

9. From the Restore SharePoint Online items dialog, expand the site collection.

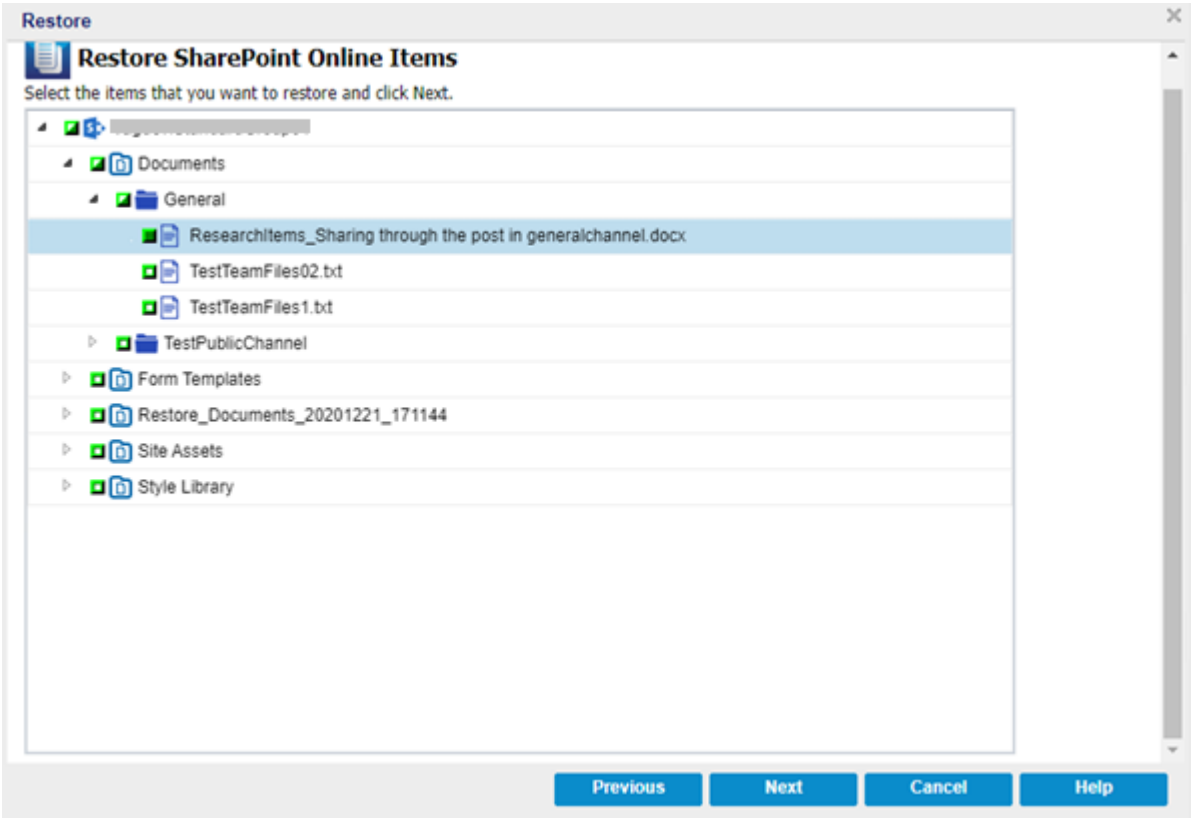
   All lists/Libraries and list are displayed.

10. Select the site or specific channel available in the site collection that you want to restore from SharePoint site collection, and then click **Next**.

    Notes:

    - You can select the entire content or partial content of the SharePoint object to restore. To select partial content, expand the object, and then click the check box of that content.
    - You can select multiple SharePoint lists/Libraries or list items to restore.

    The Restore Options dialog opens.



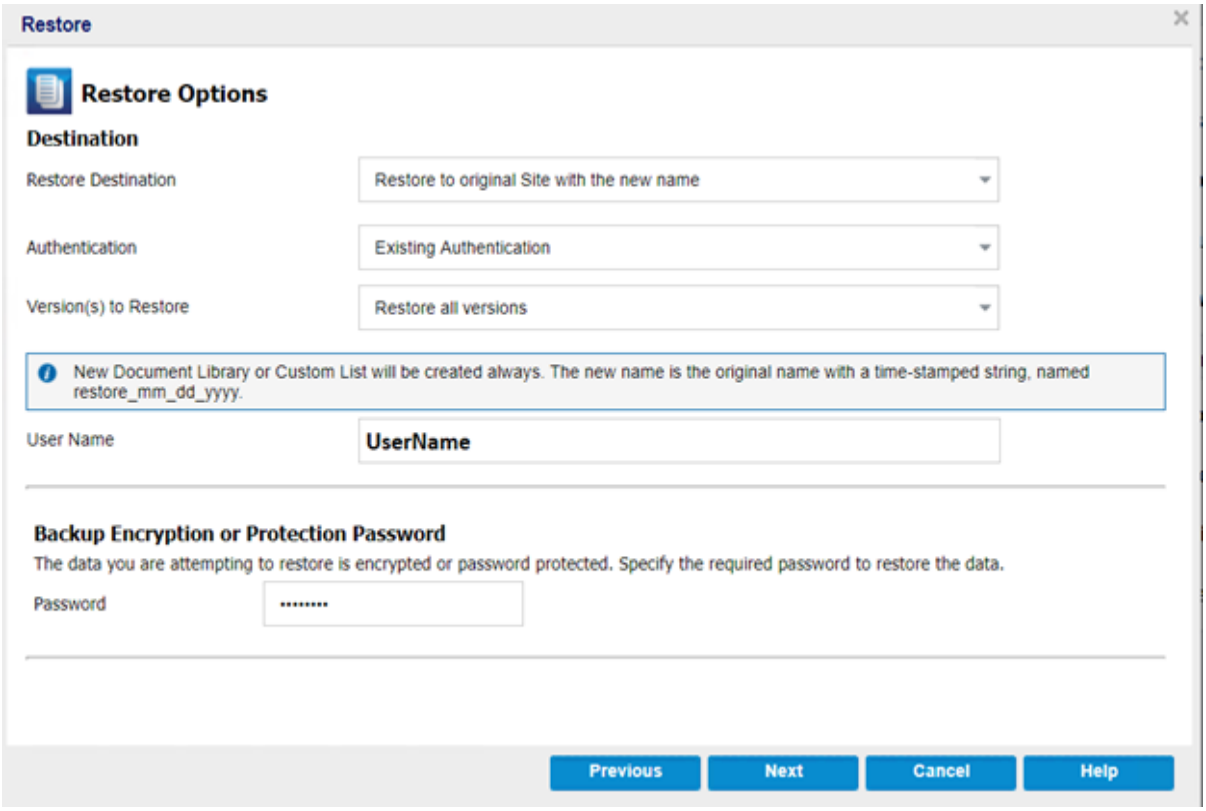
The SharePoint Online lists/Libraries or list items to restore are selected. Now, you can define the restore options.

Define the Restore Options

After you specify the SharePoint Online information to restore, define the restore options for the selected content.

Follow these steps:

1. On the Restore Options dialog, select the restore destination.



The available destination options are:

Restore to original Site with the new name

Restores the list/Library and list items data to the same site with new list name from where you took a backup.

Restore to the original location

Restores the list/Library and list items data to the same location from where you took a backup.

2. Specify one of the following options from the If Item already exists in the Destination drop-down list:

Append as new version if versioning is enable

This option works when your version setting is enabled from library setting in SharePoint Site. Once you choose this option, new version is appended to the current versions of the list items if the list items are existing.

Skip the item and do not restore

Skips over the items and does not restore.

Overwrite the item in the destination

Overwrites the item in the destination.

Export to disk

Restores the lists/Libraries or list items in the site collection to folder or share folder in disk.

Note: For lists restore when selecting restore option of Export to disk, only export attachment(s) of list to disk.

Authentication

Validates and identifies the user account. Select one of the following options:

- **Existing Authentication**

  When you select Existing Authentication, the authentication type that was selected during Node creation is used.

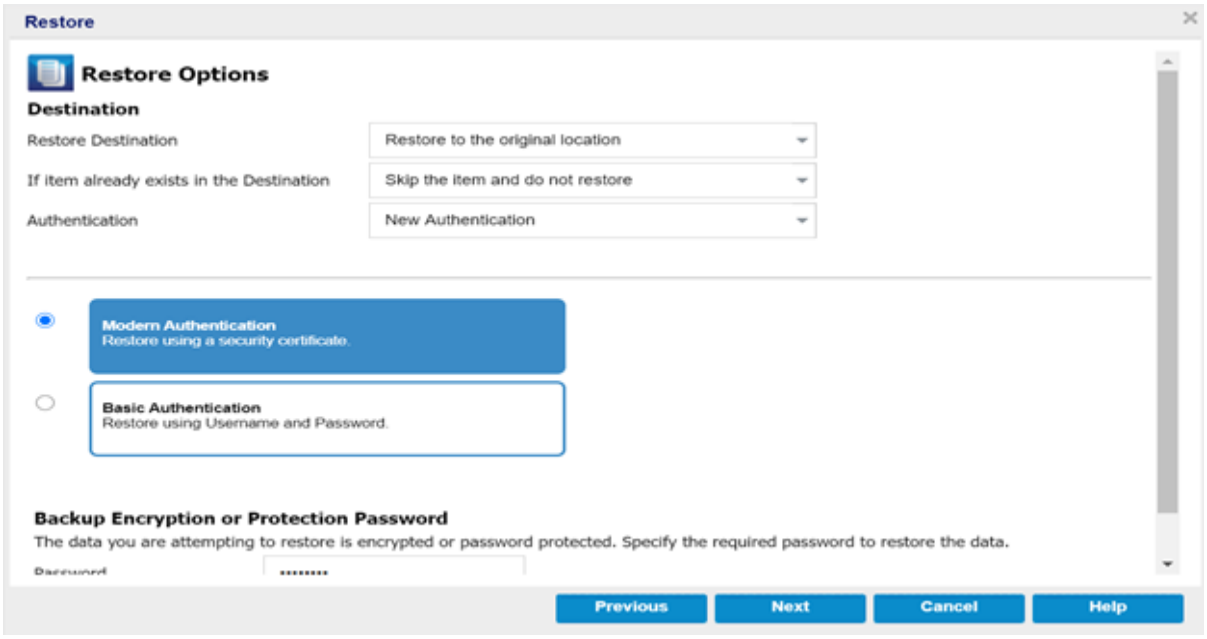
  - If the Node Authentication is basic, parameters such as Username, Password and Backup Encryption or Protection Password [If Encrypted] are loaded automatically to proceed with the restore.
  - If the Node Authentication is modern, parameters such as Username and Backup Encryption or Protection Password [If Encrypted] are loaded automatically to proceed with the restore.

- **New Authentication**

  During restore, you can select one of the following to change the authentication type:

  - Modern Authentication

- Basic Authentication



3. To apply authentication, do one of the following:

   Modern Authentication

   For modern authentication, you can run a restore using security certificate.

   Use security certificate

   Note: By default, the backup plan selects **Use security certificate** for modern authentication.

   Follow these steps:

   a. To authenticate using a self-signed certificate, see How to Use Self-Signed Certificate.

   b. Enter the Application ID and Username.

   c. To apply the required roles and permissions, select the **Allow Arcserve UDP to set the required Azure roles and permissions for this application** check box, and then click **Next**.

   Basic Authentication

   For Basic Authentication, to proceed with the restore, enter UserName and Password.

   Note: This option fails if you have Modern Authentication enabled on the Tenant or Account.

4. Specify the list item Versions that you want to restore when your version setting is enabled from library setting in SharePoint Site.

   Restore all versions

   Restores all versions on backup.

   Restore only the latest version

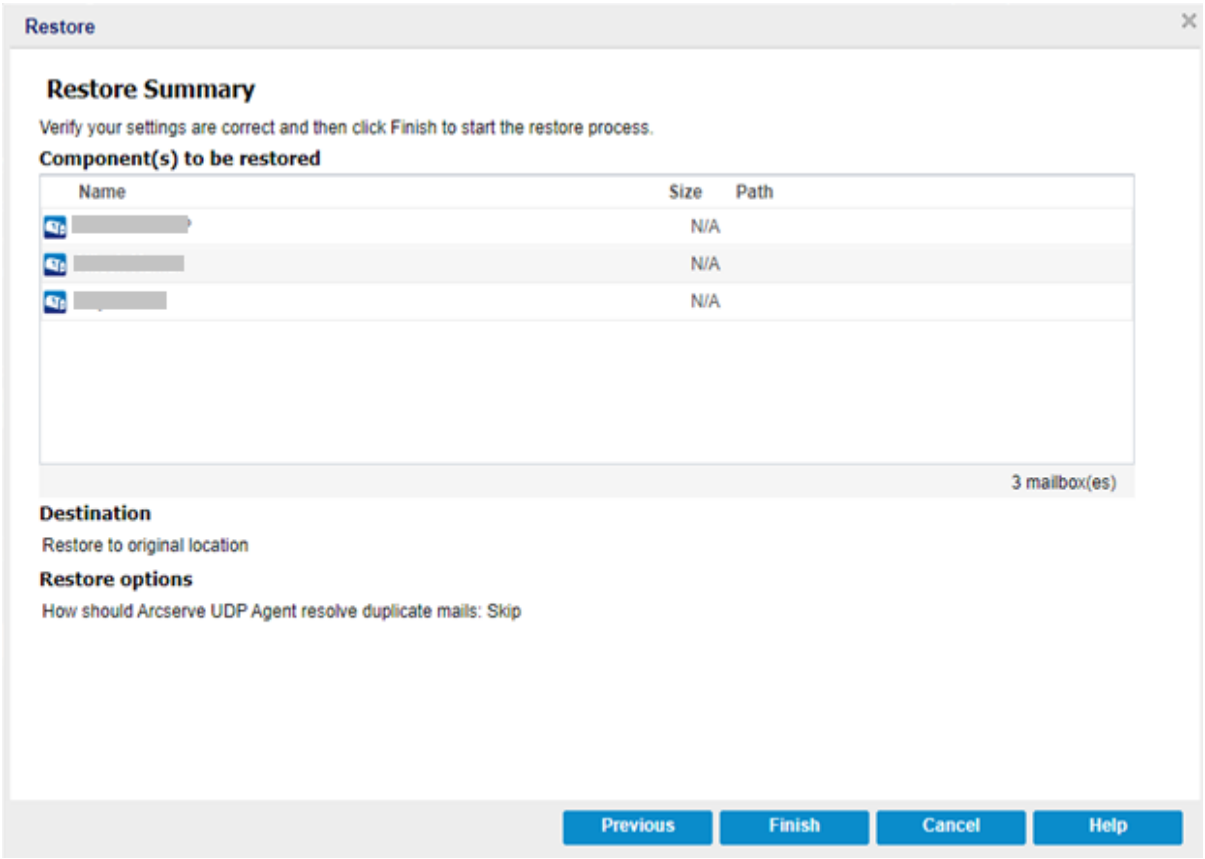
   Restores only the latest version on backup.

   Restore only the latest major version

   Restores only the latest major version on backup.

5. Specify the Session password if have.
6. Click **Next**.

   The Restore Summary dialog opens.



   You can verify the restore information form Restore Summary dialog.

7. Click **Finish** to submit the restore job.

   The restore options are defined to restore the SharePoint Online information.

Verify that Content is Restored

After the completion of the restore process, verify that content was restored to the specified destination.

Follow these steps:

1. Log into the SharePoint site collection.
2. Verify the list/Library and list items.
3. Verify the restored content.

The restored content is successfully verified.

Data Protection Map for Microsoft Teams

Important! Teams data for an entire organization gets protected automatically if Exchange, SharePoint, and OneDrive jobs are set to use the *protect all* option for accounts and sites.

The following table describes about the Data Protection Map for Microsoft Teams:

Teams Item	What to backup in UDP	Where to restore from
Recordings of scheduled meetings	OneDrive account of the user who does the recordings	Agent UI for OneDrive account of the user
1-1 Private Chats/Posts/Conversations	Exchange mailbox of the user	Agent UI for Exchange user
Files sent over 1-1 chat	Exchange mailbox/OneDrive account of the user	Agent UI for Exchange user and OneDrive account of the user
Items in Teams Standard channel		
Chats/Posts/Conversations	Group mailbox of Standard Channel	Agent UI for Group mailbox
Teams calendar for the channel	Group mailbox of Standard Channel	Agent UI for Group mailbox
Files sent over chat conversations	Select the Sharepoint site of the channel/*Protect all* sites or Group mailbox of the channel	Sharepoint agent UI > Site for Standard Channel > Document library OR Exchange agent UI > Group mailbox
Files/Folders uploaded using Files tab	Select the Sharepoint site of the channel/*Protect all* sites	Sharepoint agent UI > Site for standard channel > Document library

Items in Teams Private channel		
Chats/Posts/Conversations	Exchange mailbox of the user	Agent UI for user mailbox
Files sent over chat conversations	Select the Sharepoint site of the channel/*Protect all* sites or User mailbox	Sharepoint agent UI > Site for Standard Channel > Document library OR Exchange agent UI > User mailbox
Files/Folders uploaded using Files tab	Select the Sharepoint site of the channel/*Protect all* sites	Sharepoint agent UI > Site for Standard Channel > Document library

How to Create and Configure Concurrent Instant Virtual Machine(s) on Multiple Hypervisors

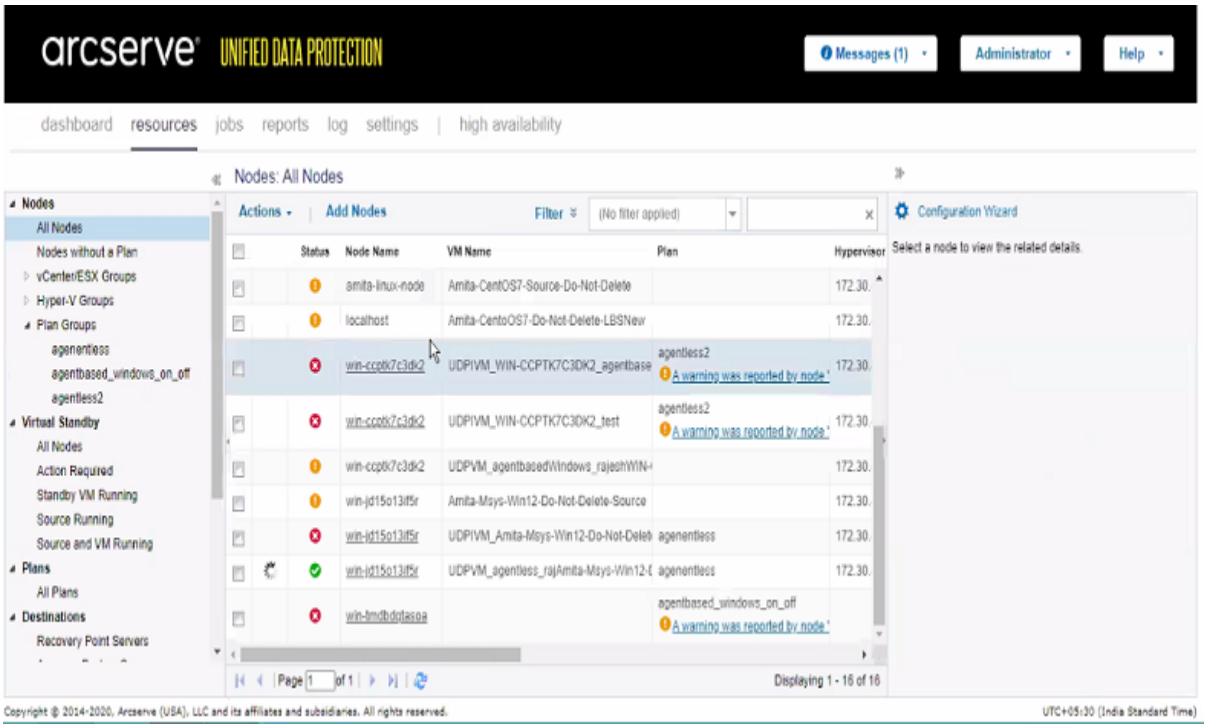
This section provides information about how to create and configure concurrent instant VMs on multiple hypervisors such as Hyper-V, VMware, Nutanix AHV, Amazon EC2, and Microsoft Azure.

You can create concurrent instant VMs from Arcserve UDP Console using the Create an Instant VM wizard.

To open the Create an Instant VM wizard, follow these steps:

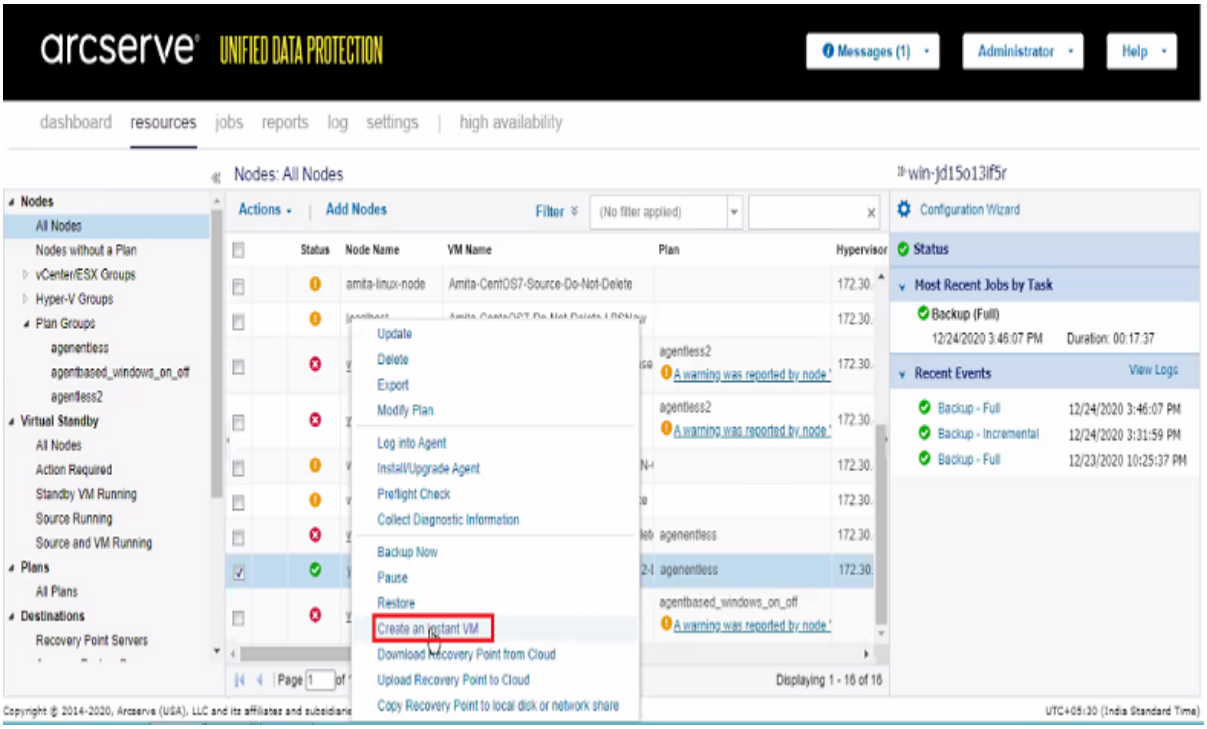
1. Log into the UDP console account.
2. On the UDP console page, click the **resources** tab.
3. From the left pane, navigate to **Nodes** > **All Nodes**.

   The list of existing nodes appears on the center pane.



4. Select one or more nodes that are already associated with a plan.

   Note: If a node is not associated with any plan, the Create an Instant VM option does not appear.



5. Right-click the selected nodes, and then select the **Create an Instant VM** option.

The Create an Instant VM wizard appears. From the wizard, you can configure concurrent instant VMs.

To configure the concurrent instant VMs, follow these steps:

1. Select a source recovery point server(s)
2. Select VM(s) Location
3. Select Recovery server(s)
4. Configure Virtual Machine(s) Settings
5. Verify Summary

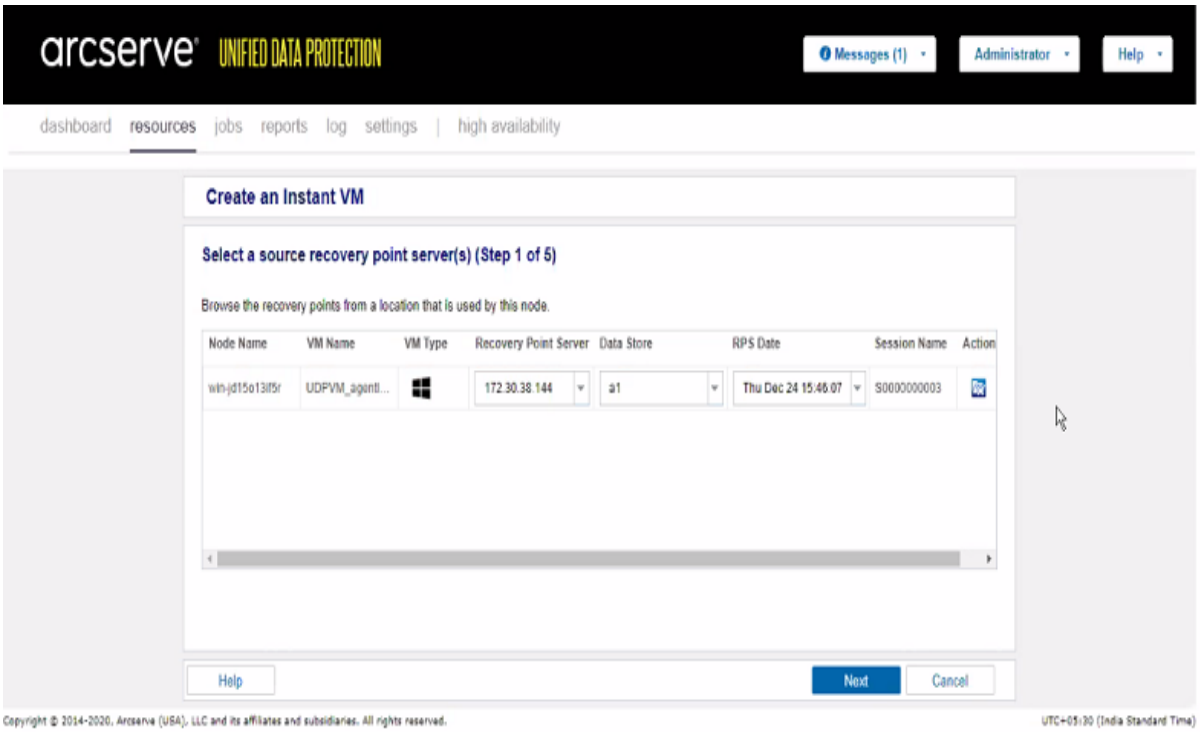
Select a source recovery point server(s)

This section provides information about how to perform the first step of selecting a source recovery point.

The Select a source recovery point server(s) screen displays the list of selected nodes with details such as Node Name, VM Name, VM Type, Recovery Point Server, Data Store, RPS Date, Session Name, and Action. For each node in the list, the console automatically selects the Recovery Point Server, Data Store, and RPS Date details. You can also change these details if there are multiple recovery point servers available on each node.

Follow these steps:

1. From the Select a source recovery point server(s) screen, on each node, perform the following:



- For Recovery Point Server, from the drop-down list, select the required recovery point server.
- For Data Store, from the drop-down list, select the required data store.
- For RPS Date, from the drop-down list, select the required RPS date.

2. Click **Next**.
3. (Optional) If the backup session is encrypted for any of the selected nodes, type the password.

   The VM(s) Location screen appears.

Complete the next step: Select VM(s) Location

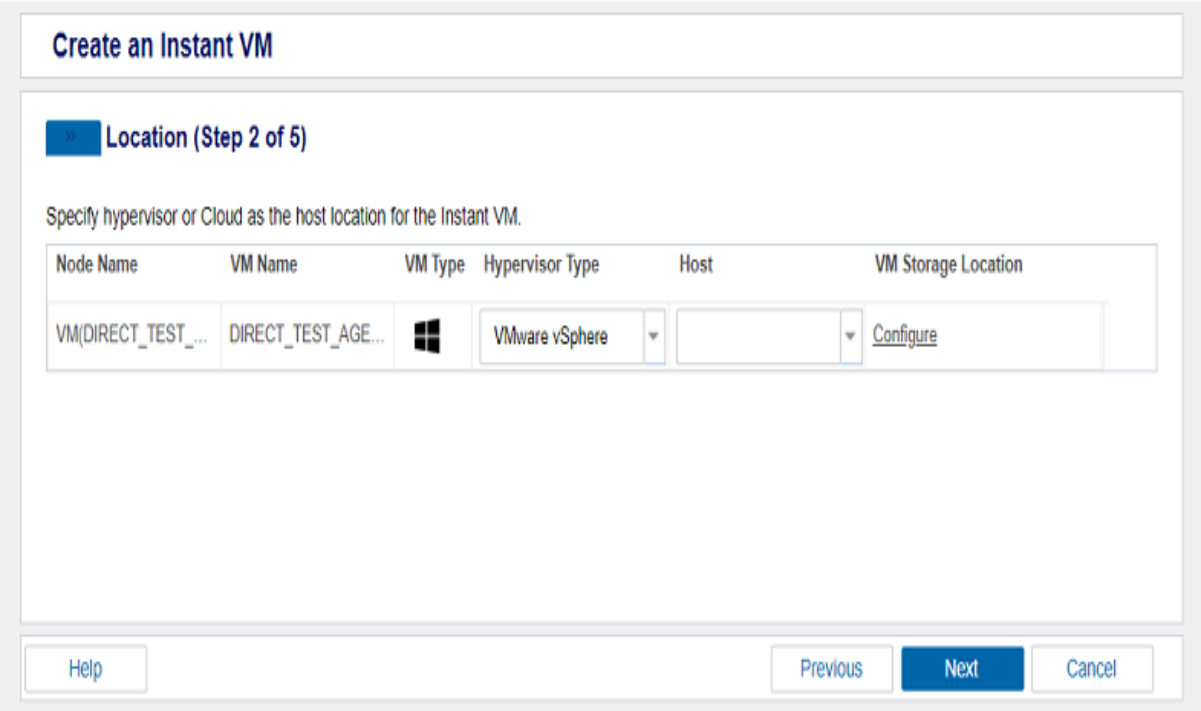
Select VM(s) Location

This section provides information about how to perform the second step of selecting a virtual machine location to host the instant VM.

The VM(s) Location screen displays the list of selected nodes with details such as VM Name, VM Type, Hypervisor Type, and Host. For each node in the list, you can select the hypervisor type and location of the virtual machine to host the instant VM.

Follow these steps:

1. From the VM(s) Location screen, on each node, perform the following:



Windows node

- For Hypervisor Type, from the drop-down list, select the type of the hypervisor.

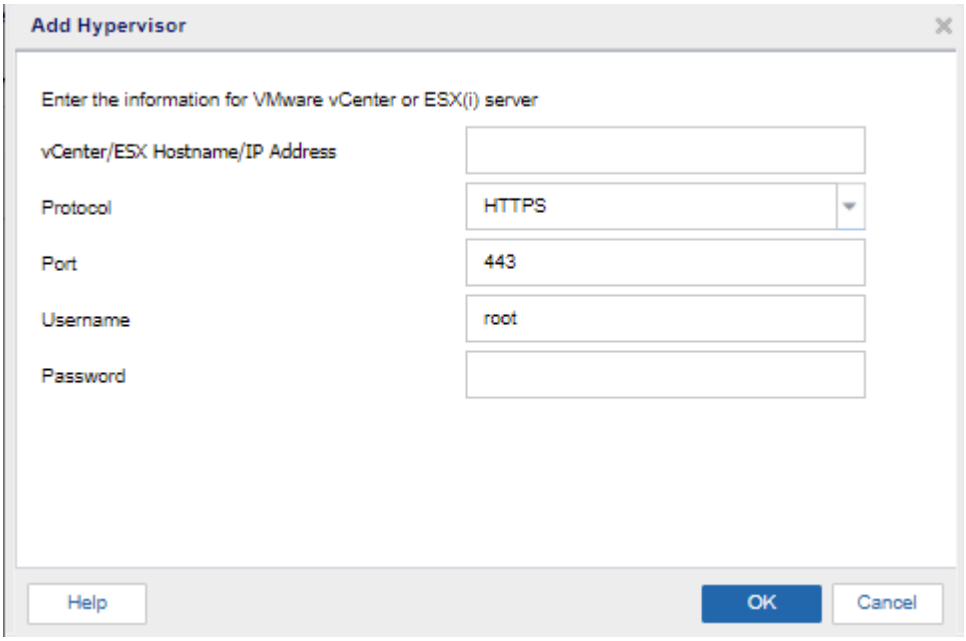
  Note: Arcserve UDP supports VMware and Hyper-V to host the Instant VM for Windows nodes.

- For Host, from the drop-down list, select the existing hypervisor or add a new hypervisor.

  To add a new hypervisor, follow these steps:

  a. From the Host drop-down list, click **+Add New Hypervisor**.

     The Add Hypervisor dialog appears.



b. On the Add Hypervisor dialog, provide the following details as needed:

- **vCenter/ESX Hostname/IP Address:** Type the host name or vCenter IP address.
- **Protocol:** Select the protocol such as HTTP or HTTPS.
- **Port:** Type the port number.
- **Username:** Type the username.
- **Password:** Type the password.

c. Click **OK**.

The new hypervisor is added.

- For VM Storage Location, to configure the selection of VMware storage location, click **Configured**.

Linux node

- For Hypervisor Type, from the drop-down list, select the type of the hypervisor.

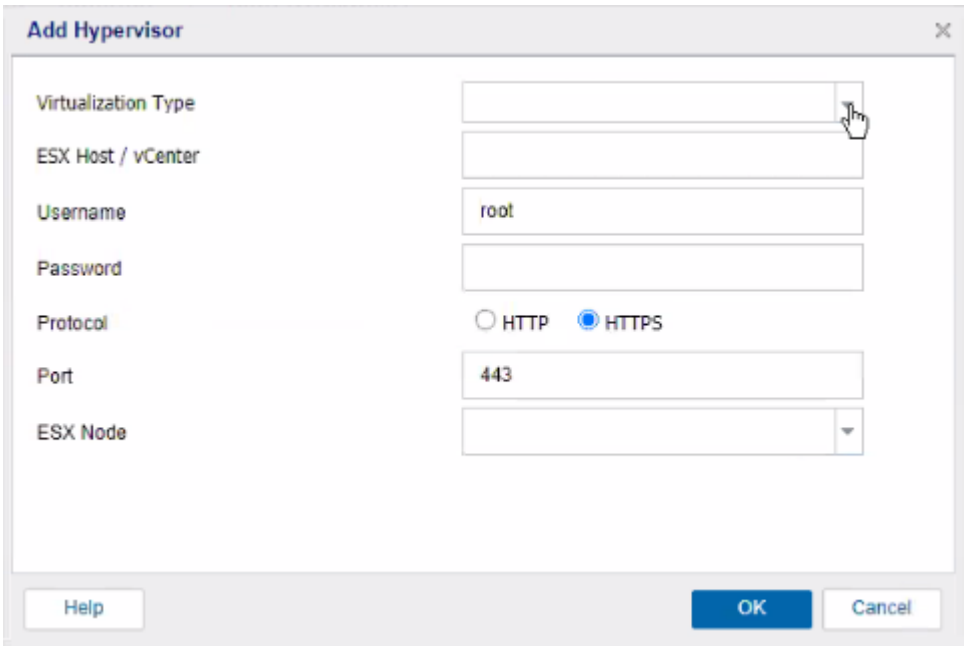
  Note: Arcserve UDP supports VMware, Hyper-V, Nutanix AHV, Amazon EC2, and Microsoft Azure to host the Instant VM for Linux nodes.

- For Host, from the drop-down list, select the existing hypervisor or add a new hypervisor.

  To add a new hypervisor, follow these steps:

a. From the Host drop-down list, click **+Add New Hypervisor**.

The Add Hypervisor dialog appears.



b. On the Add Hypervisor dialog, from the Virtualization Type drop-down list, select one of the following:

VMware

Provide the following details as needed:

- **ESX Host / vCenter:** Type the host name or vCenter IP address.
- **Username:** Type the username.
- **Password:** Type the password.
- **Protocol:** Select the protocol such as HTTP or HTTPS.
- **Port:** Type the port number.
- **ESX node:** From the drop-down, select the one of the available nodes.

Hyper-V

Provide the following details as needed:

- **Hyper-V Host name:** Type the host name for the Hyper-V hypervisor.

- **Username:** Type the username.
- **Password:** Type the password.

Nutanix AHV

Provide the following details as needed:

- **Nutanix AHV Cluster Virtual IP address:** Type the IP address.
- **Username:** Type the username.
- **Password:** Type the password.
- **Port:** Type the port number.

EC2

Provide the following details as needed:

- **Account Name:** Type the name of the cloud account. If the cloud account is not added, click Add. For more information, see Add a Cloud Account.
- **Region:** From the drop-down, select the available region.

Azure

Provide the following details as needed:

- **Account Name:** Type the name of the cloud account. If the cloud account is not added, click Add. For more information, see Add a Cloud Account.
- **Resource group:** From the drop-down, select the available resource group.
- **Region:** From the drop-down, select the available region.

c. Click **OK**.

The new hypervisor is added.

- For VM Storage Location, to configure the selection of VMware storage location, click **Configured**.

2. Click **Next**.

The Recovery server(s) screen appears.

Complete the next step: Select Recover server(s)

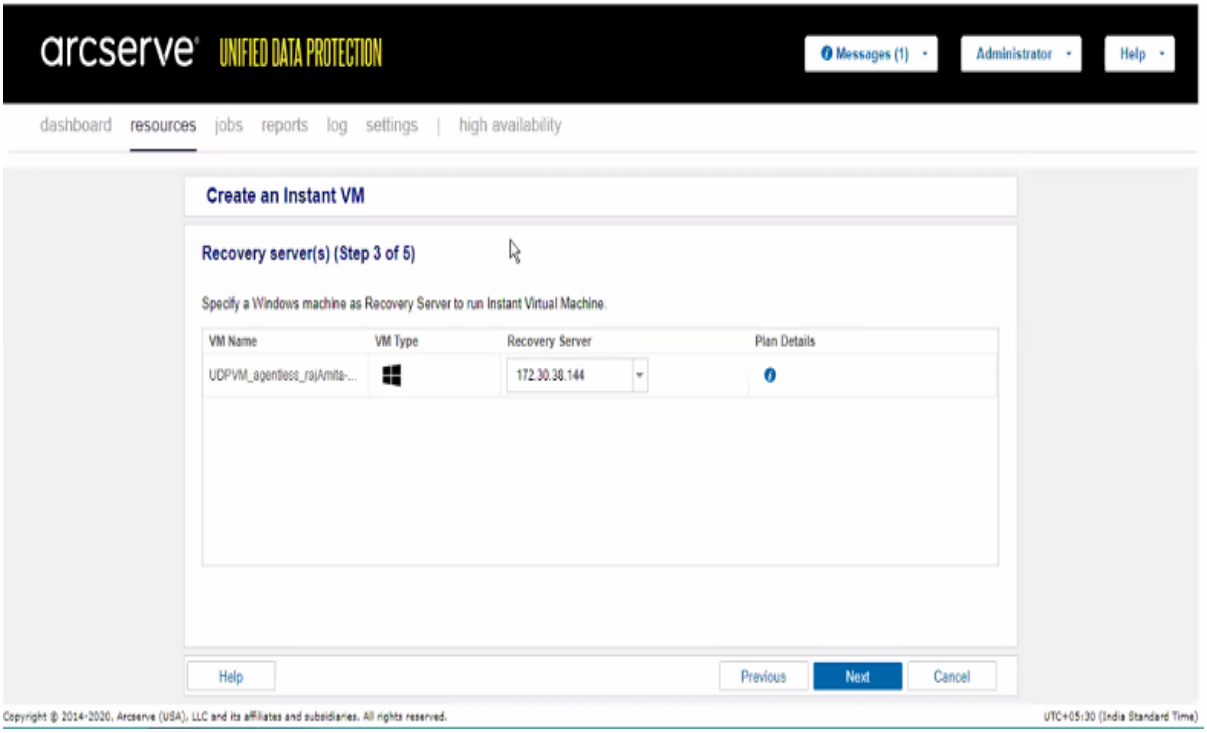
Select Recovery server(s)

This section provides information about how to perform the third step of selecting a recovery server to run the instant VM.

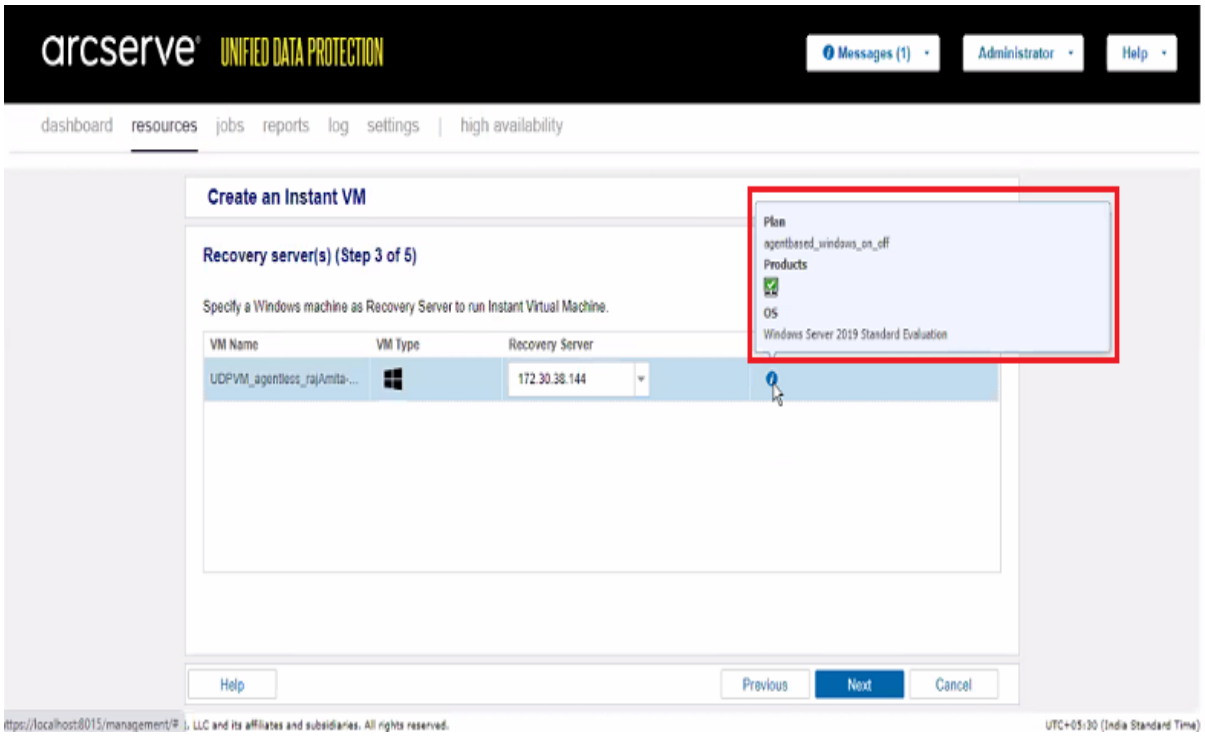
The Recovery server(s) screen displays the list of selected nodes with details such as VM Name, VM Type, Recovery Server, and Plan Details. For each node in the list, you can select the recovery server to run the instant VM.

Follow these steps:

1. From the Recovery server(s) screen, on each node, perform the following:



   - For Recover Server, from the drop-down list, select the required recovery server.
   - For Plan Details, hover over the information icon to verify the details of the plan.



Notes:

- For Linux backup sessions, the recovery server is a Linux Backup Server.
- When the hypervisor is VMware vSphere, the Windows Network File System (NFS) role must be installed in the Recovery Point Server. It gets installed automatically during the Instant VM process.. To manually install the Network File System, see How to manually install Network File System on a Windows Server.

2. Click **Next**.

   The Virtual Machine(s) Settings screen appears.

Complete the next step: Configure Virtual Machine(s) Settings

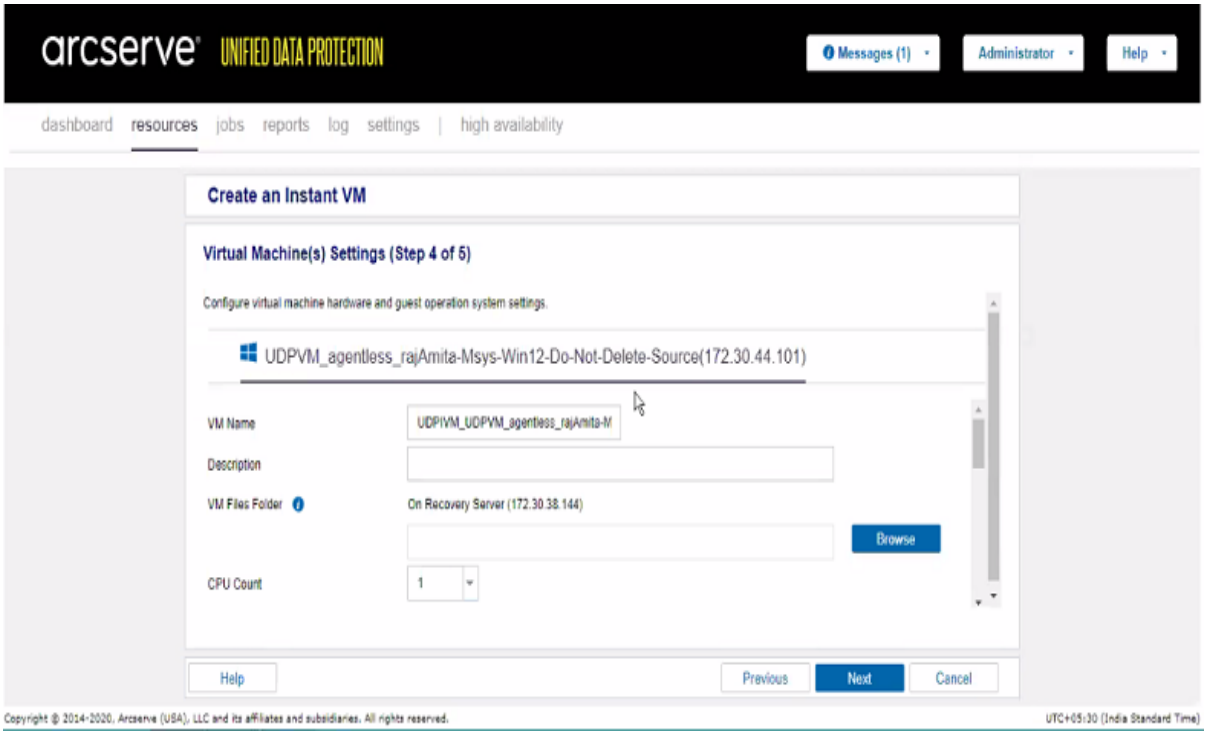
Configure Virtual Machine(s) Settings

This section provides information about how to configure virtual machine settings for creating instant VM.

The Virtual Machine(s) Settings screen displays individual sections for each selected node with the configuration details. On each section, you can provide the details as needed to create the instant VM.

Follow these steps:

1. From the Virtual Machine(s) Settings screen, click the VM node you want to configure.



2. On the selected VM node, perform the following:
   - For VM Name, type the name of the Instant VM. The name of the source node with a prefix is the default name of the Instant VM.
   - (Optional) For Description, type the description as needed for the Instant VM.
   - For VM Files Folder, click the **Browse** button to select the location of the folder on the recovery server for the Instant VM. You can also create a new folder from the VM Files Folder dialog.

     Note: The VM Files Folder option is only available for Windows VM nodes. For Linux VM node, you can select the available datastore from the VM Files datastore drop-down list.
   - For CPU Count, select the number of CPUs for the Instant VM.
   - For Memory Size, type the memory size for the Instant VM.

     Note: You can also use the slider to provide the memory size for the Instant VM.
   - For Network Adapters, do the following:

Add an Adapter

You can add multiple network adapters. After adding a network adapter, you can edit and delete the network adapter from the Actions column.

Note: For a Linux Instant VM, you can use at least one virtual NIC to connect to the Linux Backup Server.

To add a network adapter, follow these steps:

a. On the Network Adapter section, click **Add an Adapter**.

   The Add Network Adapter dialog appears.

b. On the Add Network Adapter dialog, provide the following details, and then click **OK**.

   - **Virtual Network:** From the Virtual Network drop-down, select the required virtual network.
   - **Adapter Type:** From the Adapter Type drop-down, select the type of the Adapter.
   - **TCP/IP Settings:** From the TCP/IP Settings drop-down, select **Source (Automatic)** or **Custom IP/Gateway/DNS/WINS Settings** as needed.

     Note: If you select the Custom IP/Gateway/DNS/WINS Settings option, click Add an Address to provide the TCP/IP details manually.

The Network Adapter is added.

Update DNS

If you already have provided an IP address and DNS address in the Network Adapter and source machine of a domain, you can update the DNS using the following steps:

Note: This feature is only available for Windows Instant VM.

a. On the Network Adapter section, click **Update DNS**.

b. To add a DNS update record, click **Add a DNS Address**.

   Notes:

   - You can remove the DNS update record. To remove, click **Remove**.

- To adjust the sequence of the records, click the Up and Down button.

c. Select a DNS Address and an IP Address from the drop-down list, and then click **OK**.

d. Specify the Time to Live (TTL).

e. Specify the DNS Authentication.

For Microsoft DNS Server, enter the username and password. For Bind Server, you need to specify the full path, including the file name of the key file in the Recovery Server.

- For Advance settings, do the following:
  - Verify the free disk space of the VM Files Folder capacity.

    Note: The **Monitor free disk space of VM Files Folder capacity** check box is selected by default. The capacity bar appears in yellow in the Instant Virtual Machines page if the free space of the VM Files Folder capacity is lesser than the threshold value. The default threshold value is 3%. You may change the value as needed.

  - Select the **Specify disk controller type for the virtual machine** check box, and then from the drop-down, select the type of the disk controller for the virtual machine.

    Note: The **Specify disk controller type for the virtual machine** check box is available if you have selected **VMware vSphere** as the VM location. The instant virtual machine is created applying the specified disk controller in the VMware.

  - To redirect the virtual disk updates to the VMware Datastore, follow these steps:

    a. Select the **Redirect virtual disk updates to VMware Datastore** check box.

    Note: The **Redirect virtual disk updates to VMware Datastore** check box is only available when you select VMware vSphere as VM location.

    b. From the drop-down, select the required VMware Datastore.

The updates of the virtual disk are redirected to the selected VMware datastore.

- To change the Instant Virtual Machine hostname, follow these steps:

  a. Click the **Change Host Name** check box.

     The screen refreshes with the new fields.

  b. In the New Host Name field, type the new host name as needed.

  c. If the source machine is in a domain, provide the following details as needed:

     - **User Account:** Type the user account.
     - **Password:** Type the password.
     - **Confirm Password:** Type the password again to confirm.

     Note: If the source machine is in a domain, you must have the account permissions to change the hostname in domain.

- (Optional) For Linux Instant VM, select the **Recover data automatically after Instant VM is started** check box to enable the automatic recovery of data when the instant VM target is started.

  The default behaviour of Linux Instant VM is to recover the necessary data first and start the VM. If the option is not selected, then when the VM starts the remaining data is not recovered even if used a normal VM. If the option is enabled, the remaining data is recovered at the backend when you are using the VM. You can also preserve the Instant VM target permanently when the data recovery is finished.

  Note: When status of the Linux Instant VM target is Power Off, Instant VM job fails. If that recovery point is merged, the Linux Instant VM fails to Power ON.

3. Click **Next**.

   The Summary screen appears.

Complete the final step: Verify Summary

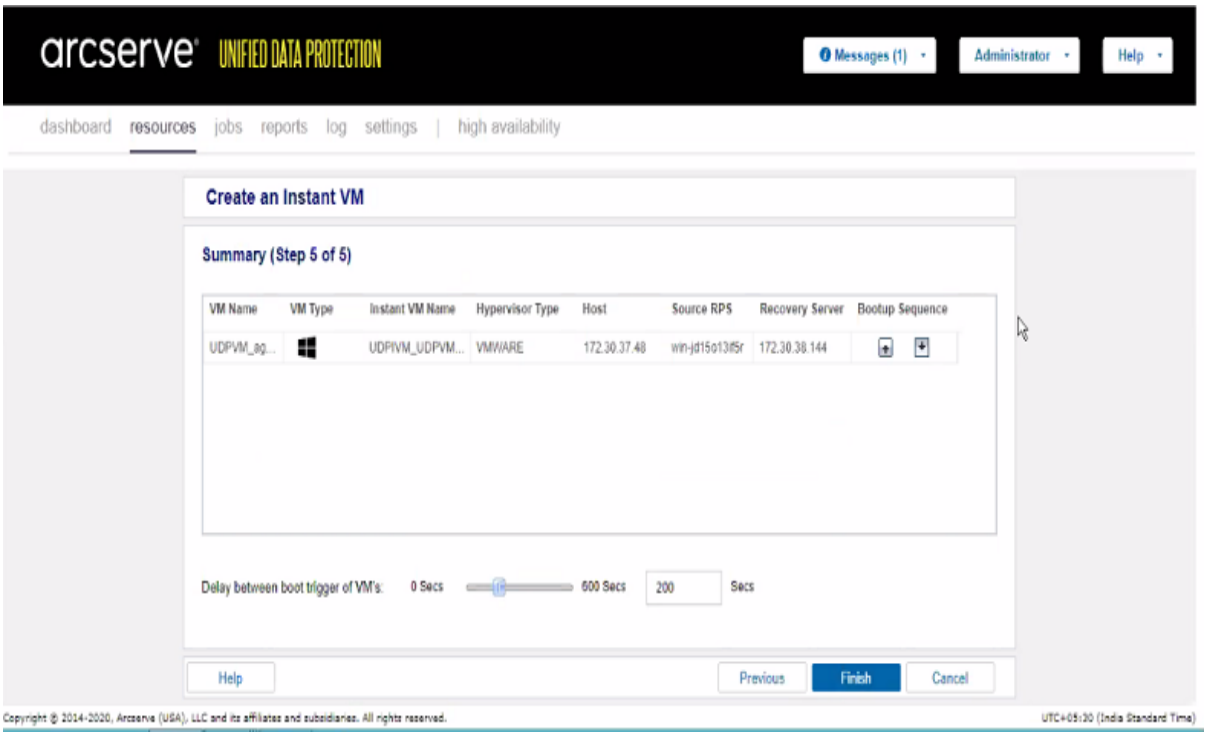
Verify Summary

This is the final step for creating concurrent instant VMs. This section provides information about how to verify summary, adjust bootup sequence, and provide delay between boot trigger of VMs.

The Summary screen displays the list of selected nodes with details such as VM Name, VM Type, Instant VM Name, Hypervisor Type, Host, Source RPS, Recovery Server, and Bootup Sequence. You can verify the details and prioritize the VMs in the order you want to boot.

Follow these steps:

1. From the Summary screen, perform the following:



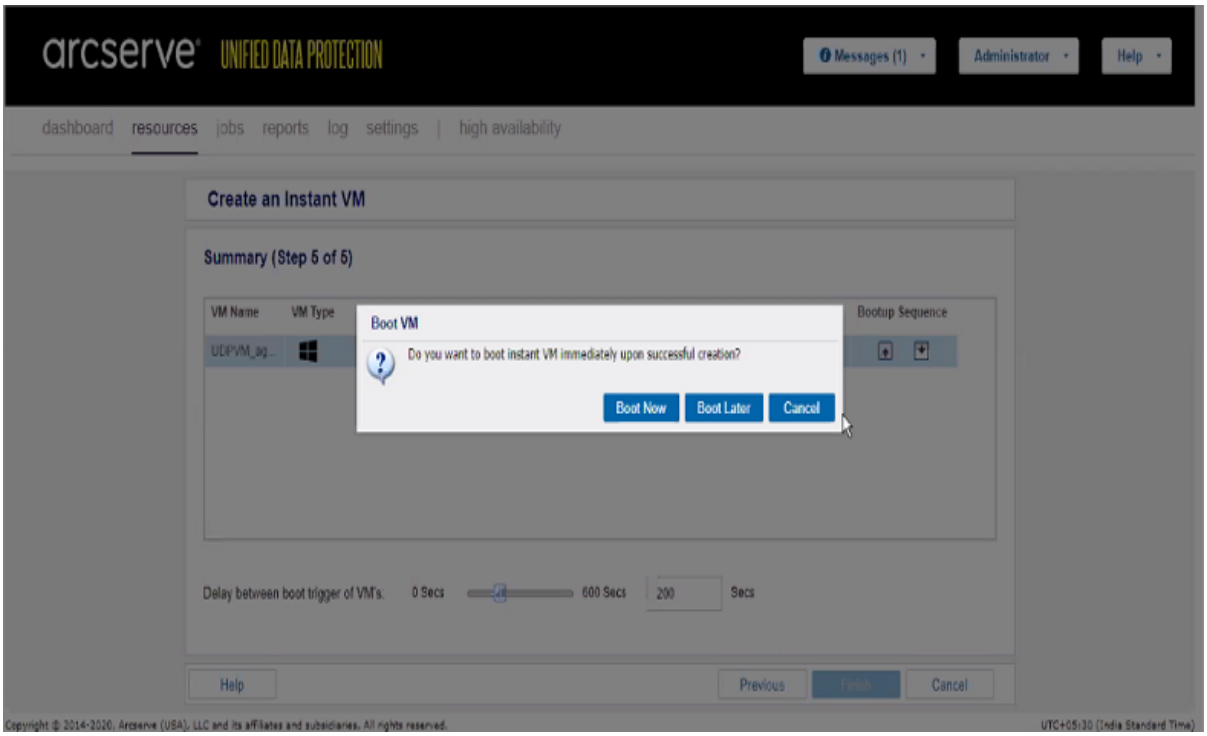
   - From the list of nodes, on Bootup Sequence of each node, click the up or down button to prioritize the VM's bootup sequence.
   - For Delay between boot trigger of VMs, type the delay time in the range of 0 to 600 secs.

     Note: You can also use the slider to set the delay time.

2. Click **Finish**.

   The Boot VM dialog appears.

3. On the Boot VM dialog, select one of the following:



- **Boot Now:** Creates concurrent Instant VMs. After the concurrent Instant VMs are created, the console automatically starts the VMs.
- **Boot Later:** Creates concurrent Instant VMs. After the concurrent Instant VMs are created, you can manually start the VMs.
- **Cancel:** You are returned to the Create a New Instant VM(s) page without creating any concurrent Instant VMs.

Concurrent Instant VM(s) are created.

How to Configure Concurrent Virtual Standby Machines

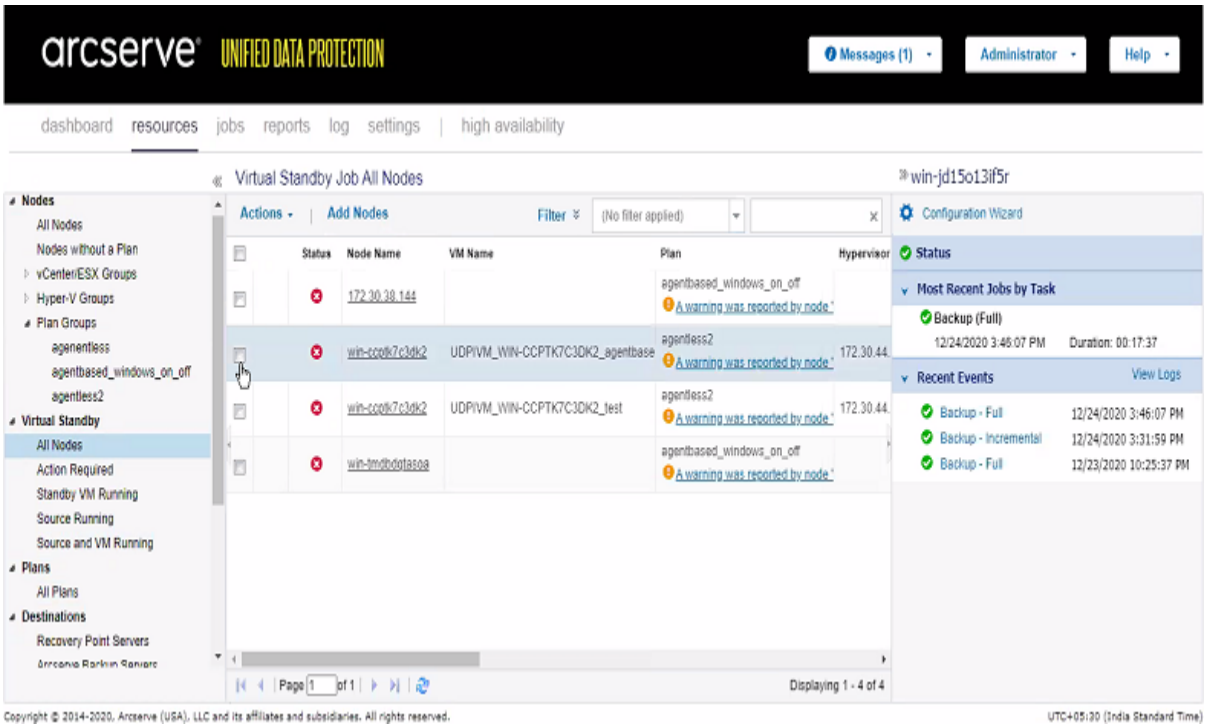
This section provides information about how to configure concurrent Virtual Standby Machine(s).

You can configure concurrent Virtual Standby Machine(s) from Arcserve UDP Console using the Standby VM dialog.

Follow these steps:

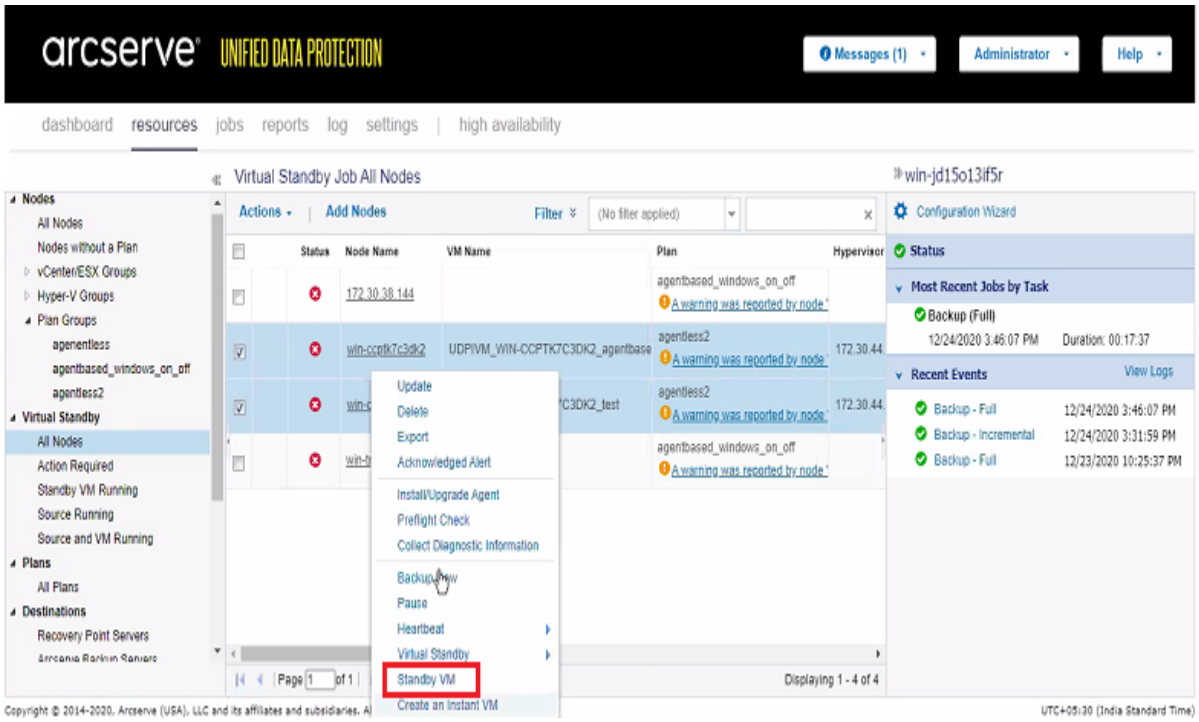
1. Log into Arcserve UDP Console.
2. On the UDP console page, click the **resources** tab.
3. From the left pane, navigate to **Nodes** > **All Nodes**.

   The list of existing nodes appears on the center pane.



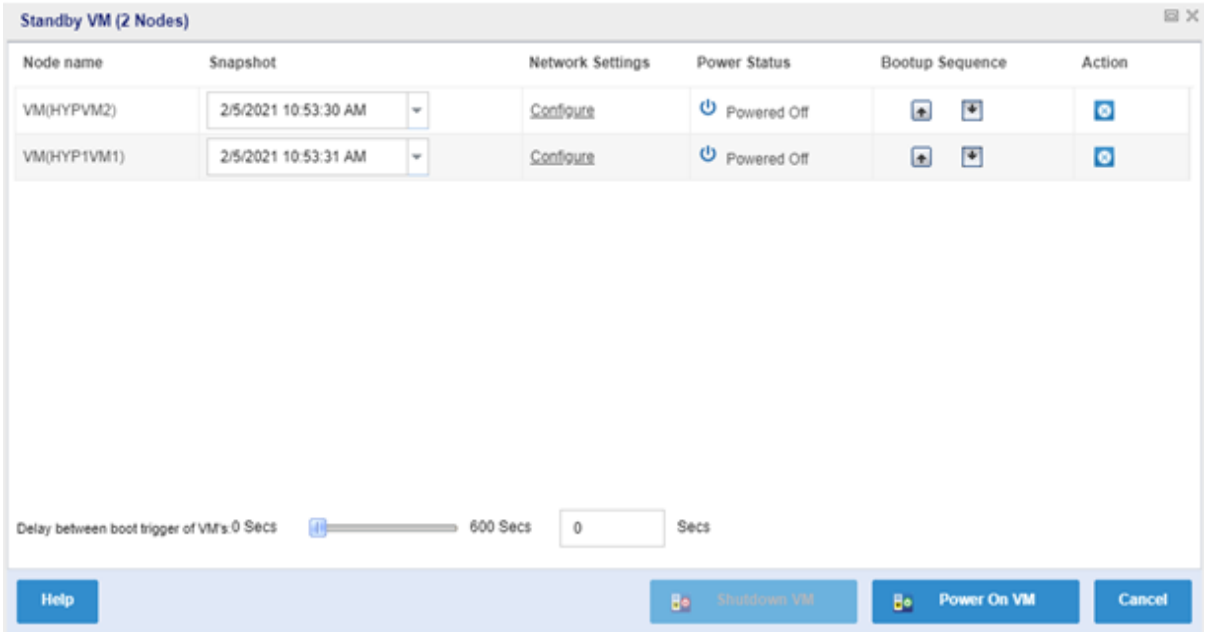
4. Select one or more nodes that are already associated with a plan.

   Note: If a node is not associated with any plan, the Standby VM option does not appear.



5. Right-click the selected nodes, and then select the **Standby VM** option.

   The Standby VM dialog appears with the list of selected nodes.



6. On the Standby VM dialog, do the following:

   - **Power On VM**

     Note: This option is enabled only for the powered off virtual standby machines.

To power on virtual machines concurrently, follow these steps:

a. From the list of selected nodes, for each node, perform the following:

- For Snapshot, from the drop-down list, select the date and time of the recovery point snapshot.

  Note: The Snapshot details appear only when the VSB job completes successfully.

- For Network Settings, select the Configure link. For more information, see Configure the Standby VM Network.
- For Bootup Sequence, click the up or down button to prioritize the Standby VM's bootup sequence.
- Under Action, click the blue-colored cross icon to remove the node.

b. (Optional) For Delay between boot trigger of VMs, type the delay time in the range of 0 to 600 secs.

Note: You can also use the slider to set the delay time.

c. Click **Power On VM**.

The Standby VM(s) are powered on concurrently using the data contained in the recovery point snapshot.

- **Shutdown VM**

  To shut down the virtual standby machines concurrently, click **Shutdown VM**.

  Note: This option is enabled only for the powered on virtual standby machines.

Chapter 12: Using Hardware Snapshot for Backup

This section contains the following topics:

How to use Hardware Snapshot for Backup

Arcserve UDP has the capability to utilize hardware storage snapshots for backup. You can specify whether you want to use the hardware snapshot while creating a backup task. If you select hardware snapshot, then Arcserve UDP first tries to create a hardware snapshot. If hardware snapshot fails, Arcserve UDP automatically reverts to the software snapshot without failing the backup job.

You can use the hardware snapshot for an agent-based backup (Windows physical machines) and host-based agentless backup (VMware and Hyper-V).

Supported storage arrays:

- **NetApp Storage Array:** Arcserve UDP supports hardware snapshot for agent-based backup (Windows physical machines) and host-based agentless backup (VMware and Hyper-V).
- **Nimble Storage Array:** Arcserve UDP supports hardware snapshot for agent based backup (Windows physical machines) and host-based agentless backup (VMware and Hyper-V).
- **HPE 3PAR storeserve array:** Arcserve UDP supports hardware snapshot for agent-based backup (Windows physical machines) and host-based agentless backup (VMware and Hyper-V).
- **Dell EMC Unity VSA Storage Array:** Arcserve UDP supports hardware snapshot for agent-based backup (Windows physical machines) and host-based agentless backup (Hyper-V).

What To Do Next

- Use Hardware Snapshot for VMware Agentless Backup
- Use Hardware Snapshot for Hyper-V Agentless Backup
- Use Hardware Snapshot for Agent-based Backup
- Verify the Backup has Used Hardware Snapshot

Use Hardware Snapshot for VMware Agentless Backup

Follow these steps:

1. Verify if the following prerequisites are met:

   For NetApp Snapshot

   - Arcserve UDP supports NetApp iSCSI//FC LUNs and NetApp NFS Volume exports that are configured as Data stores. iSCSI, FC, and NFS must meet certain conditions to use the hardware snapshot.

     Considerations for NetApp iSCSI/FC Support for VMware

     Considerations for NFS Support for VMware

   - To create a hardware snapshot for VMware, add the storage array to the Console. For more information on adding a storage array, see Add a Storage Array.
   - To use a hardware snapshot, Flexclone license is recommended for NetApp storage arrays running with data ONTAP, operating in the 7-Mode and Cluster mode.

   Note: For more information on configuring NetApp, refer the NetApp documentation or contact the NetApp Support team.

   HPE 3PAR hardware snapshot:

   - To support the VMware VM hardware snapshot, HP RMC must manage the HPE 3PAR storage array.
   - The storage array needs to have Virtual copy license.

     Supported protocols for VMware hardware snapshot: FC and iSCSI
   - To create a hardware snapshot for VMware, add the storage array to the Console. For more information on adding a storage array, see Add a Storage Array.

   Nimble hardware snapshot:

   - Supported protocols for VMware hardware snapshot: FC and iSCSI
   - To create a hardware snapshot for VMware, add the storage array to the Console. For more information on adding a storage array, see Add a Storage Array

     Considerations for Nimble Storage When CHAP Authentication is enabled

2. Log on to the Console and create a plan for backup.

   Note: For more information on creating an agentless backup plan, see How to Create a Host-based Agentless Backup Plan.

3. Verify that you have selected the **Use hardware snapshot wherever possible** option in the **Advanced** tab.

   Note: To run backup only from hardware snapshot (not software snapshot), Arcserve UDP provides the following registry key: *FallbackToSWSnapshot =0(dword)*

   You can apply this registry key for a specific node as well as all nodes.

   - To apply for a specific node create the registry under HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll\(Nodeid)

     Note: The (Nodeid) is created only after one backup.

   - To apply for all nodes create the registry under HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll

   Important! If "AFBackupDll" sub-key is not available, then you need to manually add sub-key *AFBackupDll* under engine as Engine\AFBackupDll\ and then add *DWORD FallbackToSWSnapshot* under *AFBackupDll*.

4. **Save** the plan and **submit** the backup job.

   The backup job runs using the hardware snapshot.

Considerations for NetApp iSCSI/FC Support for VMware

Arcserve UDP provides hardware snapshot support for VMware virtual machine only when the underlying storage array is a NetApp storage array.

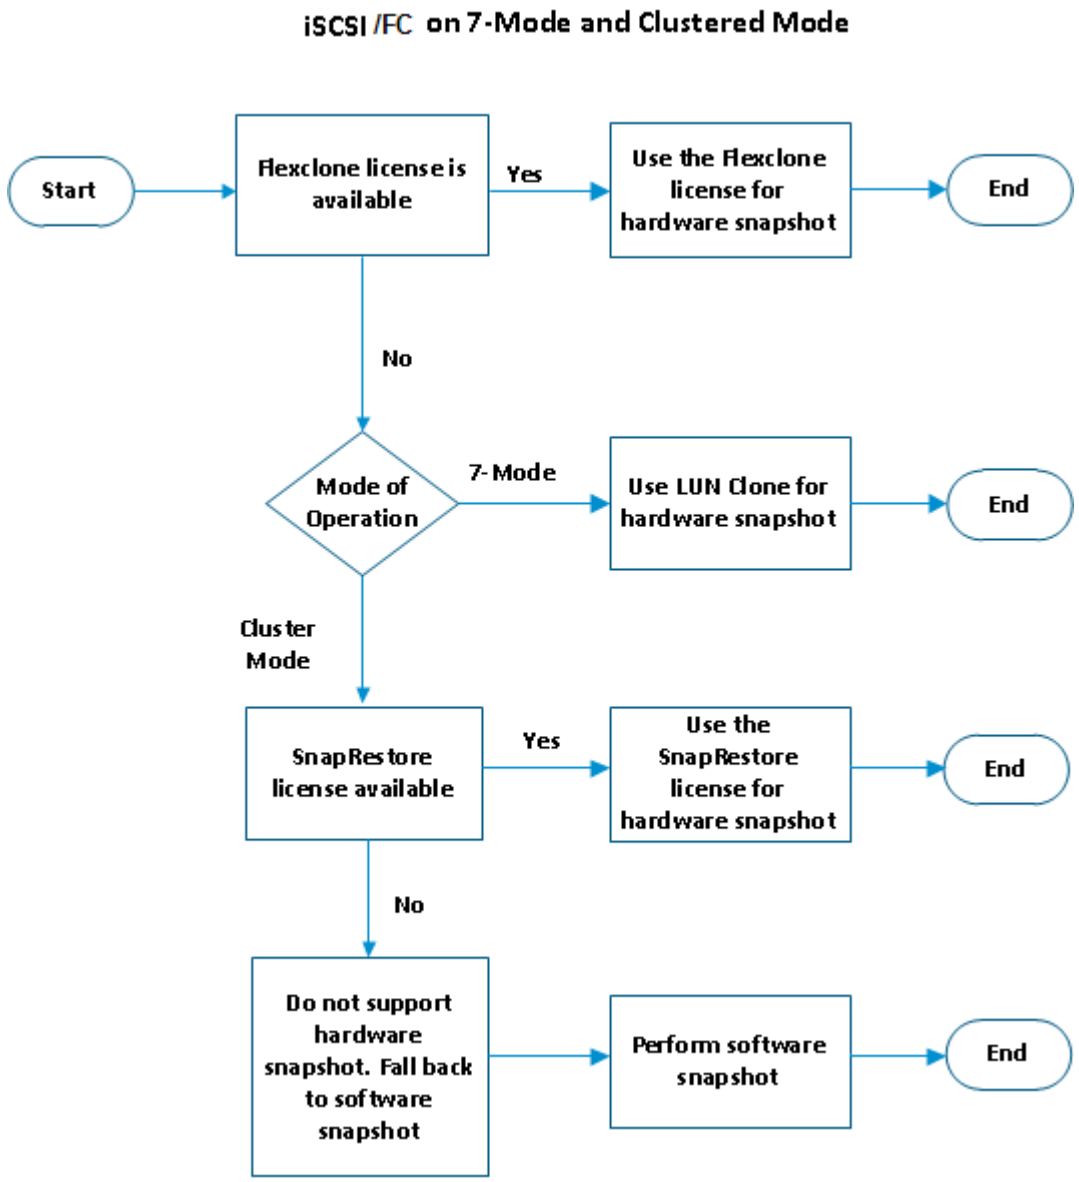
NetApp in 7-mode

If the NetApp storage system operates in the 7-mode, it is not mandatory to install any additional licenses to use hardware snapshots. However, we recommend that you have the FlexClone license installed.

NetApp in Cluster-mode

If the NetApp storage system operates in the Cluster mode (C-mode), you must have a FlexClone or SnapRestore license installed to use hardware snapshots.

The following flowchart explains the conditions applied to hardware snapshot for NetApp iSCSI/FC VMware:



Limitation of a LUN Clone

In a traditional LUN clone, sometimes the backup snapshot cannot be deleted during cleanup operations. Typically, a LUN exists as a file in the filesystem. So, a snapshot captures the file. When you create a LUN clone, another file gets created in the filesystem. So, the next snapshot captures the original file and the duplicate file. When multiple snapshots are captured, the LUN clone becomes a part of the snapshot chain. Now, if you delete a snapshot, the snapshot does not get deleted because it refers to a LUN clone which is, in turn, backed by another snapshot. In such cases, you cannot delete a snapshots until you delete the LUN clone and all the snapshots that reference that LUN clone. As a result, your retention policy for scheduled snapshots may be disrupted.

To avoid this situation, install the FlexClone license on your NetApp storage system and Arcserve UDP will use the FlexClone technology for LUN cloning.

Registry Keys for SnapRestore

When you use the SnapRestore license, restoring LUNs take a longer time depending on the size of the LUN and the environment. This is because deleting a snapshot takes a longer time as the snapshot is busy restoring LUN, which takes a longer time. Arcserve UDP does not have a mechanism to monitor the progress of restoring LUN. So, Arcserve UDP uses the retry mechanism for deleting the snapshot.

Arcserve UDP provides two registry keys (**DeleteRetryTimeoutInMins** and **DeleteRetryCount**) that you can use to improve the performance of snapshot deletion depending on the size of the LUN and the environment. The registry keys are at the following location:

SOFTWARE\Arcserve\Unified Data Protection\Engine

DeleteRetryTimeoutInMins

Specifies the timeout duration (in minutes) to delete a snapshot. If you only have SnapRestore and you do not have the FlexClone license, it is most likely that the snapshot deletion takes time. You can use the registry key to specify the custom values. However, it is advisable to use the FlexClone license.

For example, if the timeout duration is two minutes, UDP agent waits for two minutes for NetApp to delete the snapshot before sending the delete command to the NetApp storage array in its next retry. This registry key is used in combination with the **DeleteRetryCount** key.

Default value: 1 (in minutes)

Type: REG_SZ

DeleteRetryCount

Specifies the retry count to delete a snapshot.

For example, if the retry count is five, UDP agent will try for five times to send the snapshot delete command to the NetApp storage array. After retrying five times to send the snapshot delete command, if snapshot still exists, then you have to manually delete the snapshot and also, increase the retry count to the appropriate value so that next backup will not have this problem. This registry key is used in combination with the **DeleteRetryTimeoutInMins** key.

Default value: 30

Type: REG_SZ

Registry Keys for Disabling Lun Space Reservation

When Lun clone is performed during backup using hardware snapshot, by default space reservation is inherited from source Lun. Arcserve UDP provides a registry

key that you can use to disable the space reservation. The registry key is at the following location:

HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine

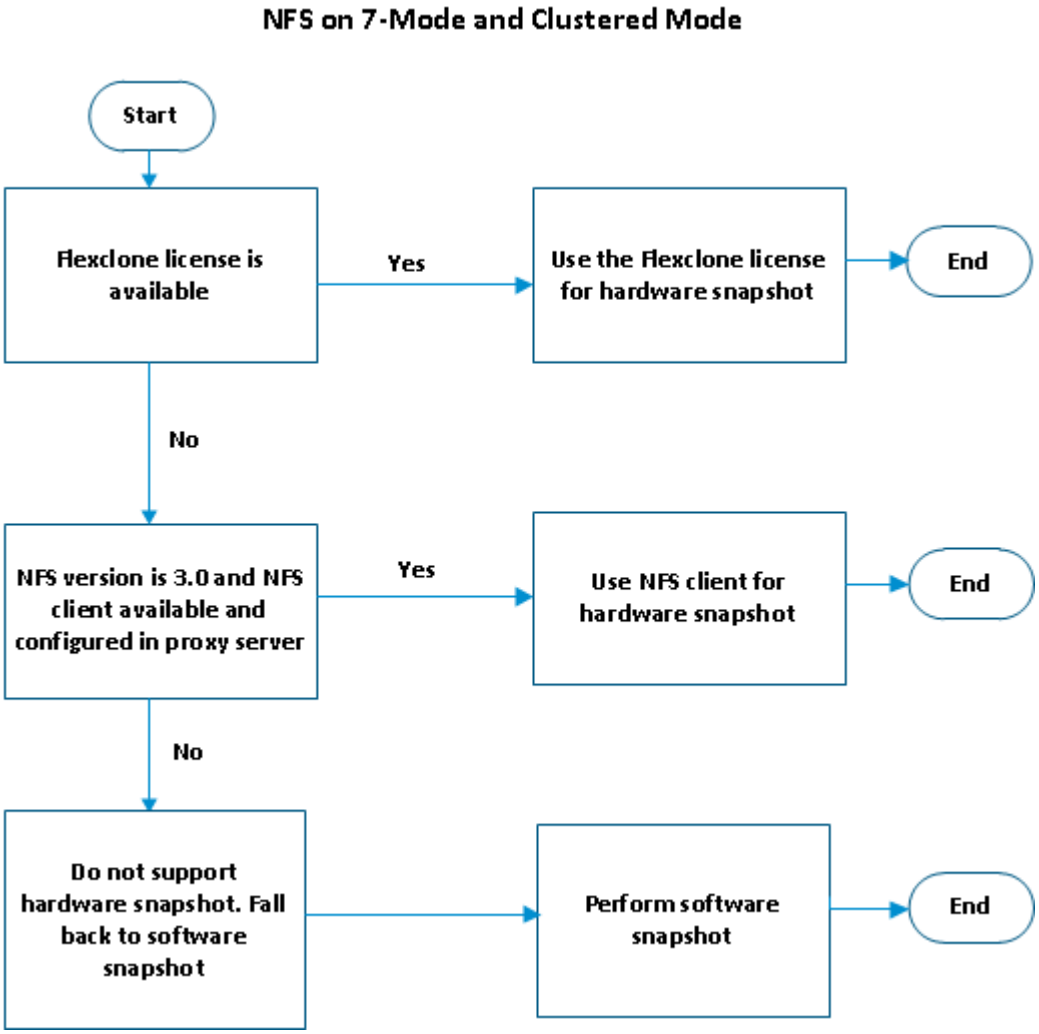
DisableLunSpaceReservation =1

Specifies that the Lun space reservations are disabled.

Note: This is applicable only when FlexClone license is applied.

Considerations for NFS Support for VMware

The following flowchart explains the conditions applied to hardware snapshot for NetApp NFS VMware:



Arcserve UDP supports hardware snapshot for NFS Version 3.0 version Data Stores. To use the hardware snapshot, the backup proxy must have the Microsoft NFS client installed and configured.

The following table displays the NFS versions that VMware VM and Arcserve UDP support. Ensure that you have the correct NFS version with the corresponding VMware version.

VMware Versions	NFS Versions that VMware Support	NFS Versions that Arcserve UDP Support for Hardware Snapshot for FlexClone	NFS Versions that Arcserve UDP Support for Hardware Snapshot for Windows NFS Client
VMware 6.0	Supports NFS	Supports NFS 3.0 and 4.1	Supports NFS 3.0 only and

and higher	3.0 and 4.1		should meet the prerequisites
VMware versions older than 6.0	Supports NFS 3.0 only	Supports NFS 3.0	Supports NFS 3.0 only and should meet the prerequisites

Prerequisite for NFS 3.0

- If the FlexClone license is not present, then, to support hardware snapshot of VMDK files hosted on an NFS 3.0 data store, the NetApp appliance must have the following versions of OnTAP installed:
  - All Data ONTAP 7-Mode systems are supported.
  - Clustered Data ONTAP 8.2 releases starting with the release 8.2.3 are supported. Also, Clustered Data ONTAP 8.3 releases starting with the release 8.3.1 are supported
  - By default, the Windows NFS v3 client support is disabled. To enable it on Storage Virtual Machines (SVMs), use the following command:

  vserver nfs modify -vserver svm_name -v3-ms-dos-client enabled

- The backup proxy, which backs up the VMware VM, should have the Microsoft NFS client installed and configured. The backup proxy should have access to the NFS share. To manually install NFS client on the server, see How to Manually Install the Microsoft NFS Client on a Windows Server.
- You need to restart the Microsoft NFS client service on the proxy server after deploying the agentless backup plan. You have to restart the service only once for the very first time because the plan deployment trigger few changes regarding the NFS client on the proxy server.

How to Manually Install the Microsoft NFS Client on a Windows Server

Follow these steps to manually install the Microsoft NFS client on a Windows server.

1. Log in to the Windows server.
2. Open the Server Manager and click Manage.
3. Click Add Roles and Features.
4. On the Before you begin dialog, click Next.
5. On the Select installation type dialog, select Role-based or feature-based installation, and then click Next.
6. On the Select destination server dialog, if you are installing to the local server, click Next. Otherwise, select a server from the Server Pool list.
7. On the Select server roles dialog, click Next.
8. On the Select features dialog, scroll down the list of available features and select the Client for NFS check box.
9. Click Next.
10. On the Confirm installation selections dialog, review your selections and then click Install.
11. After the installation completes, review the results and then click Close.

Considerations for Nimble Storage When CHAP Authentication is enabled

You can backup hardware snapshot when Nimble storage has CHAP authentication enabled.

Note: CHAP Authentication is not supported on ESXi 5.0 because that version of ESXi does not have SQLite.

You need to complete the following prerequisites:

Prerequisites

Important: The prerequisites apply only if you want the Host-based agentless backup job to use SAN as the transport mode.

- Enable SSH on ESXi where Nimble storage is configured with CHAP authentication.
- Verify if the Proxy machine has the iscsicli command line interface required to configure iSCSI target devices.
- Set the execution policy to RemoteSigned on Proxy machine, to run the Powershell scripts. Use the following command:

  Set-ExecutionPolicy RemoteSigned
- Save ESXi credentials through a Powershell script for which the Nimble storage is configured with CHAP authentication.

  Follow these steps to save the ESXi credentials:

  - Execute "StoreESXCredentials.ps1" Powershell script from Powershell on Proxy.

    You are prompted to enter details about IP, username, and Password.
  - Enter all the three parameters.

    The credentials are stored in a CSV file "ESXCredentials.csv" and the password is encrypted.

  To locate the Powershell script, follow these steps:

  - Log into the proxy machine and open the Powershell.
  - Navigate to the following directory:

    Note: The path changes based on the installation directory.

    C:\Program Files\Arcserve\Unified Data Protection\Engine\BIN
  - Run the Powershell script as given below:

PS C:\Program Files\Arcserve\Unified Data Protection\Engine\BIN>
.\StoreESXCredentials.ps1

Use Hardware Snapshot for Hyper-V Agentless Backup

Important! Hyper-V Agentless Backup uses Hardware snapshot only when all the disks of VM are on Hardware snapshot compatible volumes. Otherwise, backup uses software snapshot.

Follow these steps:

1. Verify if the following prerequisites are met.
   - Install a VSS hardware provider (for example, NetApp) on the Hyper-V servers. To support the transportable snapshot, install the VSS hardware provider on the backup proxy server. A typical configuration of a VSS hardware provider includes:
     - Specifying a server that controls the LUN
     - Specifying the disk array credentials to access the disk array

     Note: For more information on configuring the VSS hardware provider, contact your hardware provider vendor.
   - The Hyper-V server and the proxy server must have a similar operating system version.
   - If the Hyper-V server belongs to a cluster, the proxy server should not be a part of the Hyper-V cluster.
2. Log in to the Console and create a plan for backup.

   Note: For more information on creating an agentless backup plan, see How to Create a Host-based Agentless Backup Plan.
3. Verify that you have selected the **Use hardware snapshot wherever possible** option in the **Advanced** tab.

   Note: To run backup only from hardware snapshot (not software snapshot), Arcserve UDP provides the following registry key: *FallbackToSWSnapshot =0(dword)*

   You can apply this registry key for a specific node as well as all nodes.
   - To apply for a specific node create the registry under HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll\(Nodeid)

     Note: The (Nodeid) is created only after one backup.
   - To apply for all nodes create the registry under HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll

Important! If "AFBackupDll" sub-key is not available, then you need to manually add sub-key *AFBackupDll* under engine as Engine\AFBackupDll\ and then add *DWORD FallbackToSWSnapshot* under *AFBackupDll*.

4. **Save** the plan and **Submit** the backup job.

   The backup job runs using the storage snapshot.

Use Hardware Snapshot for Agent-based Backup

Follow these steps:

1. Verify if the following prerequisites are met:

   - Install a VSS hardware provider that supports hardware snapshot on the Arcserve UDP Agents. A typical configuration of a VSS hardware provider includes:

     - Specifying a server that controls the LUN.
     - Specifying the disk array credentials to access the disk array.

     Note: For more information on configuring the VSS hardware provider, contact your hardware provider vendor.

2. Log in to the Console and create a plan for backup.

   Note: For more information on creating an agent-based plan for Windows, see How to Create a Windows Backup Plan.

3. Verify that you have selected the **Use hardware snapshot wherever possible** option in the **Advanced** tab.

4. Save the plan and submit the backup job.

   The backup job runs using the storage snapshot.

Verify that the Backup has Used Hardware Snapshot

If the prerequisites for a hardware snapshot are not met, Arcserve UDP automatically switches to a software snapshot without failing the backup job. If hardware snapshot fails, the event is recorded in the activity logs.

Review the log messages to ensure that the backup has used the hardware snapshot.

Follow these steps:

1. Navigate to the following path:

   <Installation folder>\Arcserve\Unified Data Protection\Engine\Logs

2. Open the corresponding file for the respective Job ID.

   For example, if the Job ID is JW002, navigate to the **Logs** folder and open the JW002 file.

3. Review the messages in the file to confirm whether the backup has used the storage snapshot.

   You have successfully used the storage snapshot for a backup.

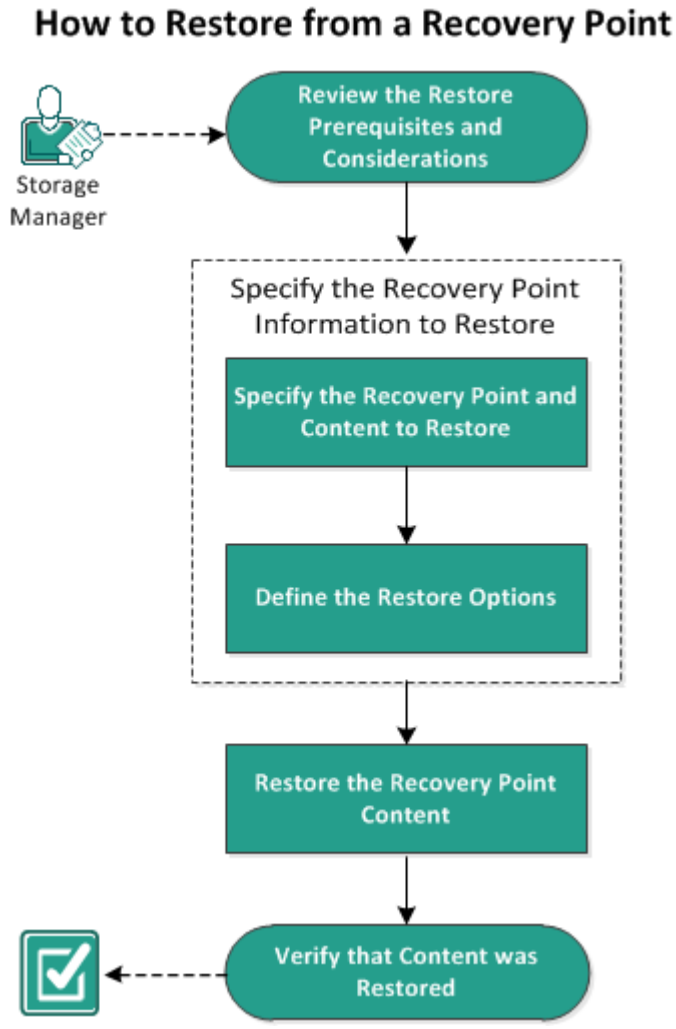
Chapter 13: Restoring Protected Data

This section contains the following topics:

How to Restore From a Recovery Point

Each time Arcserve UDP performs a successful backup, a point-in-time snapshot image of your backup is created (recovery point). This collection of recovery points allows you to locate and specify exactly which backup image you want to restore. If at some later time, you suspect any of the backed up information is missing, corrupted, or not reliable, you can then locate and restore from a previous known good version.

The following diagram illustrates the process to restore from a recovery point:



Perform the following tasks to restore from a recovery point:

1. Review the Restore Prerequisites and Considerations
2. Specify the Recovery Point Information to Restore
    a. Specify the Recovery Point and Content to Restore
    b. Define the Restore Options

3. Restore the Recovery Point Content
4. Verify that Content was Restored

Review the Restore Prerequisites and Considerations

Verify that the following prerequisites exist before performing a restore:

- You have at least one recovery point available to restore.
- You have a valid and accessible recovery point destination to restore the recovery point content from.
- You have a valid and accessible target location to restore the recovery point content to.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Review the following restore considerations:

- If the restore is to a remote destination and if all the drive letters (A - Z) are occupied, the restore to a remote path will not succeed. Arcserve UDP Agent (Windows) needs to use one drive letter to mount the remote destination path.
- (Optional) Understand how the restore process works. For more information, see How File Level Restores Work.
- (Optional) Review the files skipped during restore. For more information, see Files Skipped During Restore.
- When you attempt to restore an optimized backup session to a non-empty volume (unoptimized restore), the restore job may take more time than the estimated time displayed in the job monitor. The amount of data that is processed and the elapsed time may increase based on the data that is optimized on the volume.

Example:

The backup volume size is 100 GB and after optimization the volume size is reduced to 50 GB.

When you perform an unoptimized restore of this volume the restore job monitor displays 100% after restoring 50 GB, but it will take more time to restore the entire 100 GB.

- The following Activity log message will be displayed when restoring the system files:

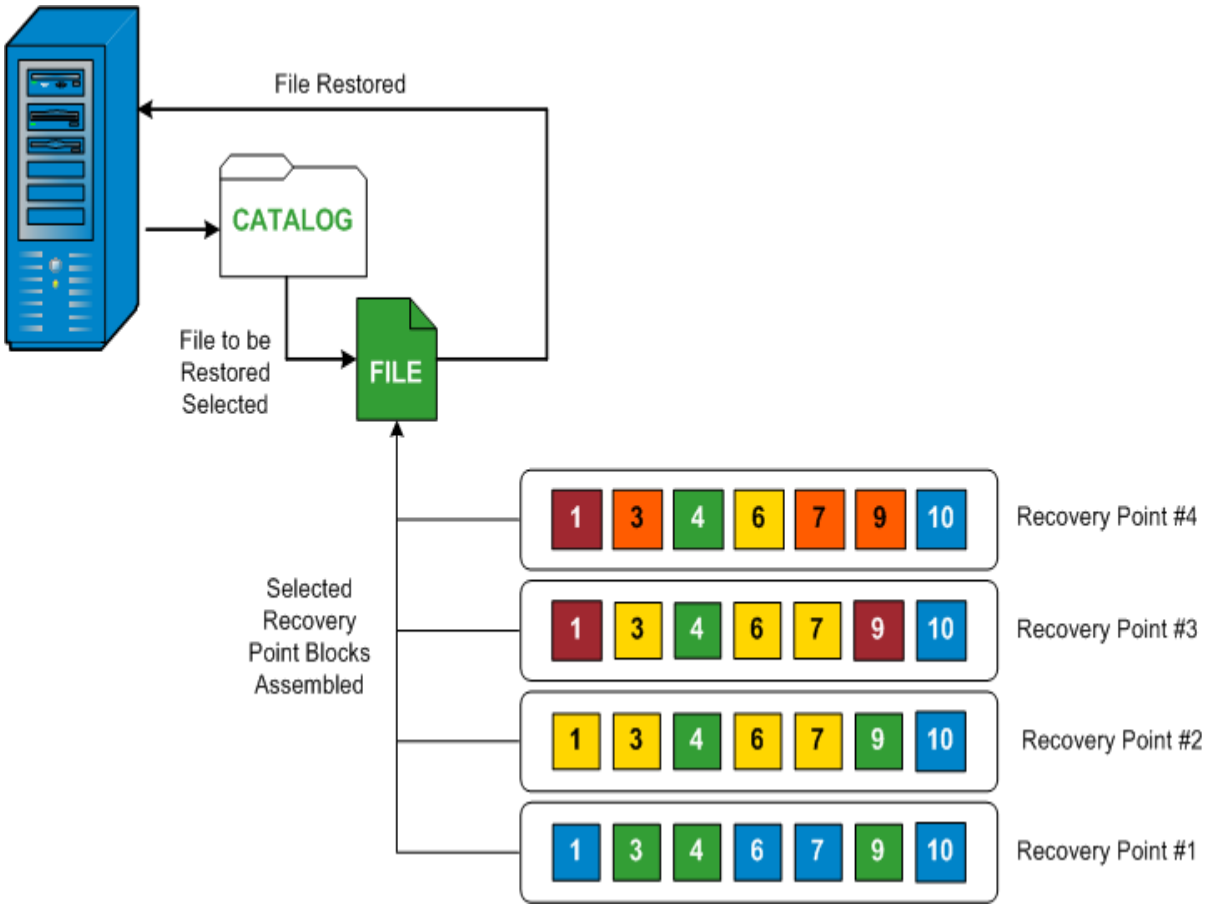
  "System files were skipped. If necessary, you can use the Bare Metal Recovery (BMR) option to restore them."

How File Level Restores Work

During a block-level backup, each backed up file is made up of a collection of blocks that define that particular file. A catalog file is created containing a list of the backed up files, along with the individual blocks that were used for each file and the available recovery points for these files. When you need to restore a particular file, you can search your backup and select the file you want to restore and the recovery point you want to restore from. Then, Arcserve UDP collects the version of the blocks that were used for the recovery point of the specified file, and reassembles and restores the file.

Note: You can also perform a restore without a catalog file from a catalog-less backup recovery point.

The following flow diagram shows the process of how Arcserve UDP restores a specific file:



Files Skipped During Restore

While performing a restore by Arcserve UDP Agent (Windows) some files may be skipped intentionally.

The files and folders in the following table are skipped during a restore if the following two conditions exist:

- Files are skipped when such files exist before the restore and the conflict option is "skip existing files".
- Files and folders listed in the following table are skipped because they are an important component for Windows or Arcserve UDP Agent (Windows).

OS	Folder or Location	File or Folder Name	Remarks
All	Root folder of each volume	CAVolTrc.dat	Used by the Arcserve UDP tracking Driver.
		cavoltrcsnapshot.dat	
		System Volume Information*	Used to save files/folders by a Windows system, for example, volume shadow copy files.
		RECYCLER*	Used only on NTFS partitions. It contains a Recycle Bin for each user that logs on to the computer, sorted by their security identifier (SID).
		$Recycle.Bin*	When you delete a file in Windows NT Explorer or My Computer, the file is stored in the Recycle Bin until you empty the Recycle Bin or restore the file.
	Any folder contain picture files	Thumbs.db	Stores thumbnail images for Windows Explorer thumbnail view.
	Root folder of volume	PageFile.Sys	Windows virtual memory swap file.
		Hiberfil.sys	Hibernate file, used to save the system data when a computer goes into hibernate mode.

The following files and folders are skipped only when you restore to the original or alternate location:

OS	Folder or Location	File or Folder Name	Remark
All	Folder specified in value record under: HKLM\Software\Microsoft\Windows NT\CurrentVersion\WinLogon\SfcDllCache	All files/folders (recursively)	Folder contains a cached dll file which is used for System File Checker

			(SFC) and contents of the system dll cache directory are rebuilt by using SFC.
	%SystemRoot%\SYSTEM32\dllCache		
	Root folder of quorum_device	MSCS*	Used for Microsoft Cluster Server.
	%SystemRoot%\SYSTEM32\	perf?00?.dat	Per-formance data used by the Win-dows per-formance counter.
		perf?00?.bak	Used for Windows File Pro-tection (WFP) records digital sig-natures of the oper-ating sys-tem installs (such as DLL, EXE, SYS, OCX, and so on) to protect them from deletion or from replace-ment by older ver-sions.
		CATROOT*	
	%SystemRoot%\inetsrv\	metabase.bin	Metabase binary file of earlier IIS versions before 6.0.

	File or folder specified in value except "SIS Common Store" under HKLM\SYSTEM\CurrentControlSet\Control\BackupRestore\FilesNotToBackup	All files/folders (recursively)	Files and folders should not be backed up and restored. For more information, see link.
XP W2003	System volume	NTLDR	The main boot loader.
		BOOT.INI	Contains boot configuration (if missing, NTLDR will default to \Windows on the first partition of the first hard drive).
		NTDETECT.COM	Required for booting an NT-based OS. Detects basic hardware information needed for a successful boot.
Vista and later	Root folder of system volume	boot*	Boot folder for Windows.
		bootmgr	Windows boot manager file.

		EFI\Microsoft\Boot*	Used for EFI boot.
	%SystemRoot%\SYSTEM32\	LogFiles\WMI\RTBackup*	Stores ETW trace files (extension .etl) for real time event trace sessions.
		config\RegBack*	Backup of current registry table.
Win-8 and later	System volume	swapfile.sys	System controller file, normally around 256 MB. It is used by Metro style applications that do not fit the traditional paging characteristics (such as usage pattern, growth, space reservation) of pagefile.sys.
		BOOTNXT	Used to boot from OS, other than Windows 8. Created when enabling the startup

			options, and updated by Windows.

The Activity log provides the following information:

- Date Time Information: jobxxxx System Files skipped. You can use Bare-Metal Recovery Option (BMR) to restore them.
- Date Time Information: jobxxxx Files or Directories skipped. Skipped files or directories are available at: C:\Program Files\Arcserve\Unified Data Protection\Engine\Logs\Restore-<YYYYMMDD>-<hhmmss>-<Process ID>-<Job ID>.log.

Specify the Recovery Point Information to Restore

Arcserve UDP provides you with an option to restore data from a recovery point. The goal of running a successful restore job is to identify quickly the data you need and to retrieve it from the appropriate backup location. Each restore job requires a source and destination.

The process involved in restoring from a recovery point is as follows:

1. Specify the Recovery Point and Content to Restore
2. Define the Restore Options

Specify the Recovery Point and Content to Restore

This section provides information about how to specify the recovery point and content to restore.

Follow these steps:

1. Access the restore method selection dialog in one of the following ways:

   From Arcserve UDP Console:

   a. Log into Arcserve UDP Console.

   b. Navigate to **resources** > **Node** > **All Nodes**.

      All the added nodes are displayed in the center pane.

   c. Right-click the node, and then click **Restore**.

      You are automatically logged into the agent node console, and the Restore dialog is opened.

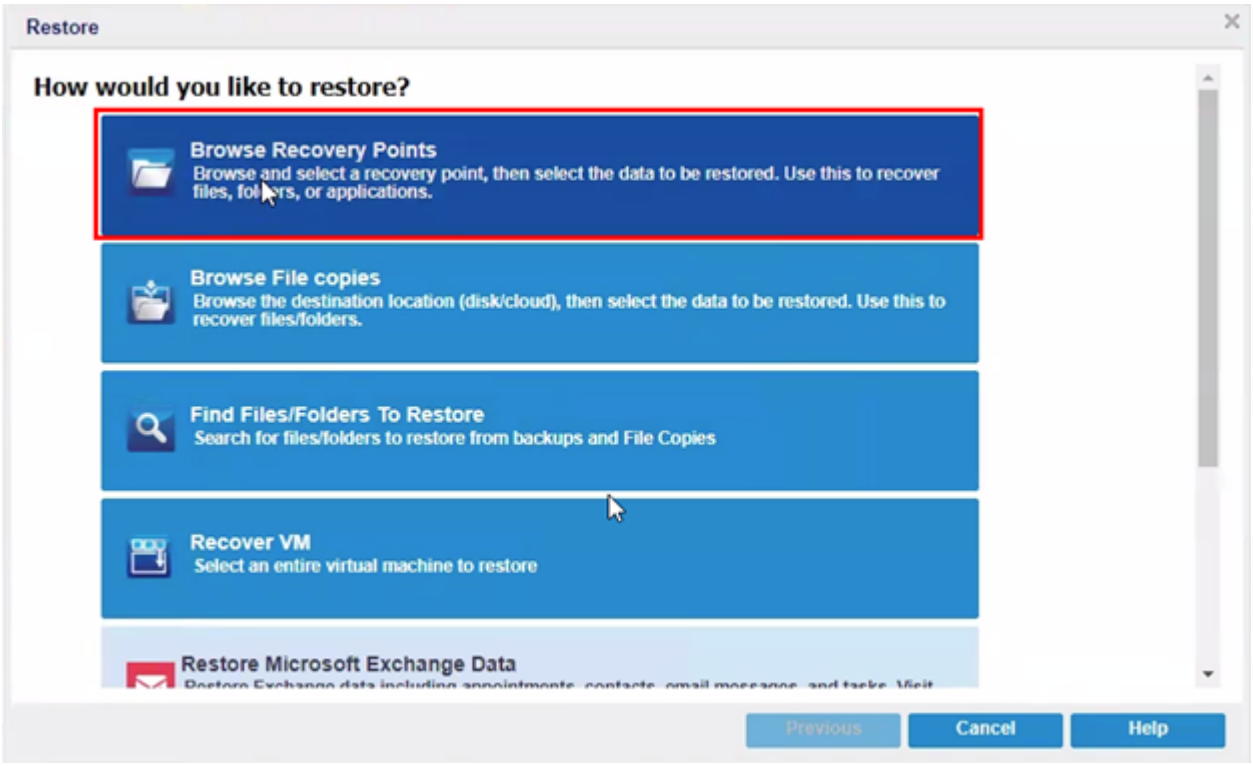
   From Arcserve UDP Agent (Windows):

   a. Log into Arcserve UDP Agent (Windows).

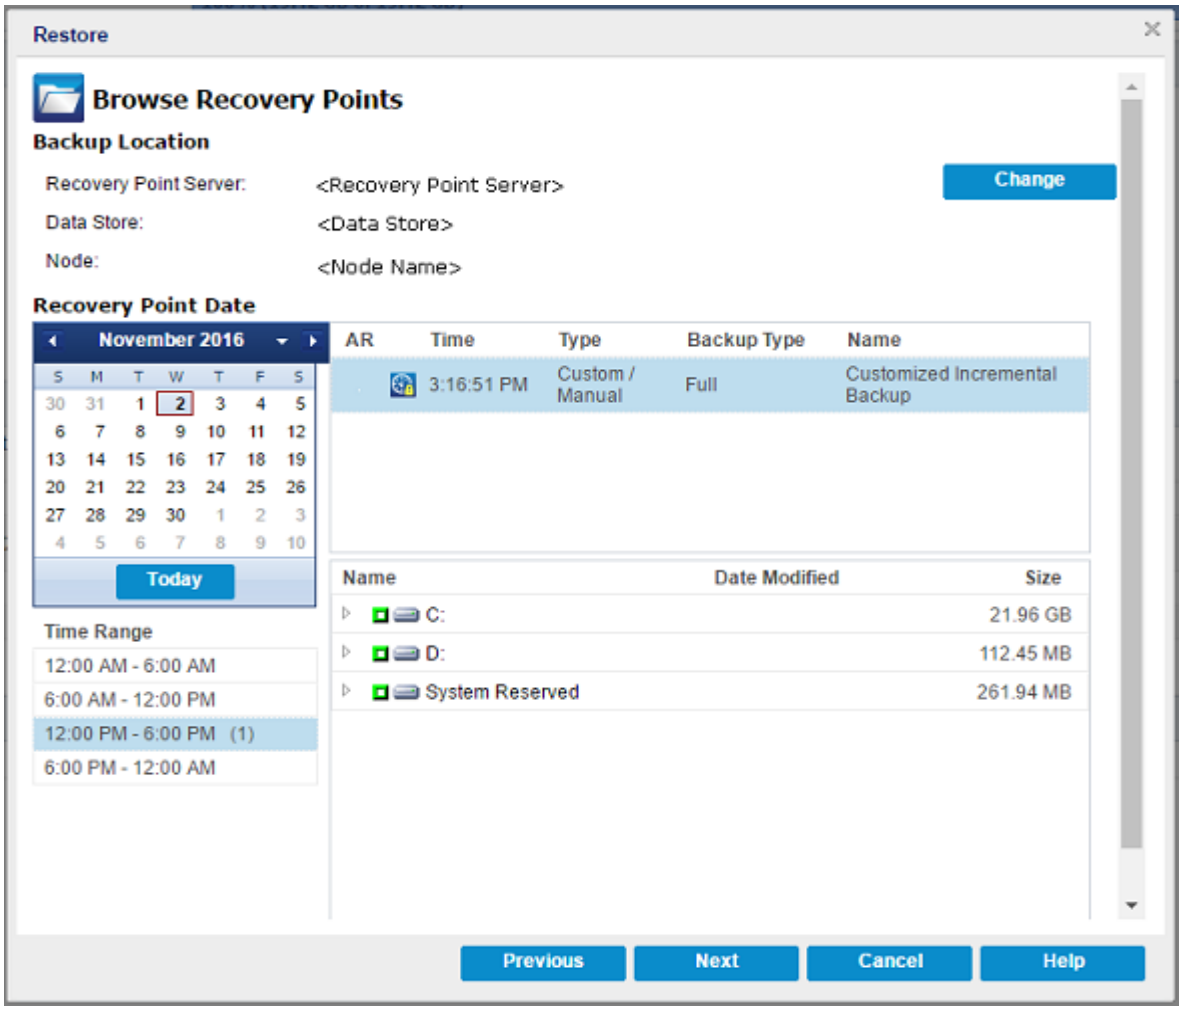
   b. From the home page, select **Restore**.

      The Restore dialog opens.

2. On the Restore dialog, click the **Browse Recovery Points** option.



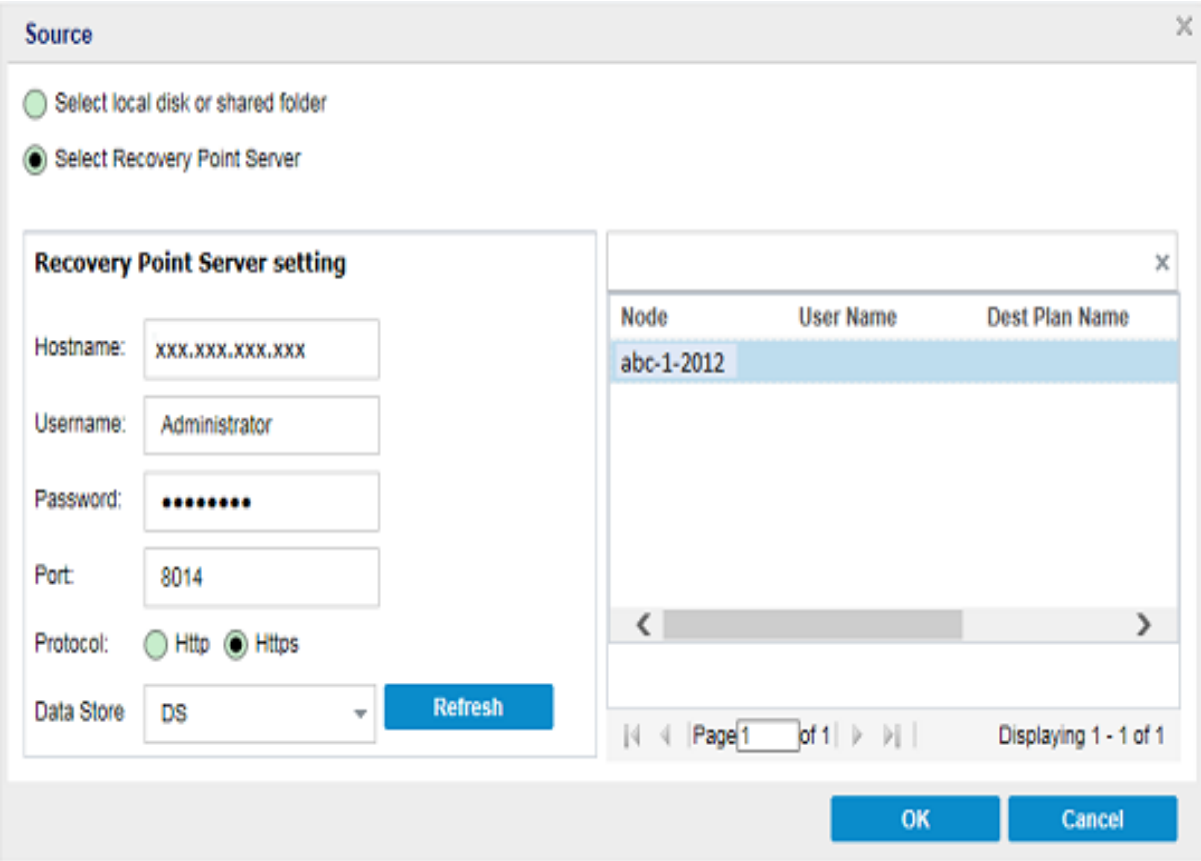
3. On the Browse Recovery Points page, do the following:



a. To update the backup location, click **Change**.

   The Source dialog opens.

b. On the Source dialog, select one of the following backup locations, and then click **OK**:



Select local disk or shared folder

1. Specify or browse to the location where your backup images are stored and select the appropriate backup source.
2. To verify the connection to the specified location, click the green arrow button. If required, enter the Username and Password to access to that source location.

   The Select backup location dialog opens.
3. Select the folder where the recovery points are stored, and then click **OK**.

   The Select backup location dialog closes and you can see the backup location in the Source dialog.

   The recovery points are listed in the Browse Recovery Points dialog.

Select Recovery Point Server

1. Specify the Recovery Point Server setting details, and then click **Refresh**.

   All the agents are listed in the Data Protection Agent column in the Source dialog.

2. Select the agent from the displayed list.

   The recovery points are listed in the Browse Recovery Points dialog.

c. Select the calendar date for the backup image to restore.

   All the dates containing recovery points for the specified backup source are highlighted in green.

   The corresponding recovery points for that date are displayed, with the time of the backup, the type of backup that was performed (Full, Incremental, or Verify), and the name of the backup.

d. Select a recovery point to restore.

   The backup content (including any applications) for the selected recovery point displays.

   Note: If a clock icon appears with a lock symbol, it indicates that the recovery point contains encrypted information and may require a password for restore.

e. Select the content to restore.

   - For a volume-level restore, you can specify to restore the entire volume or selected files/folders within the volume.
   - For an application-level restore, you can specify to restore the entire application or selected components, databases, instances, and so on, within the application.

4. Click **Next**.

   The Restore Options dialog opens.

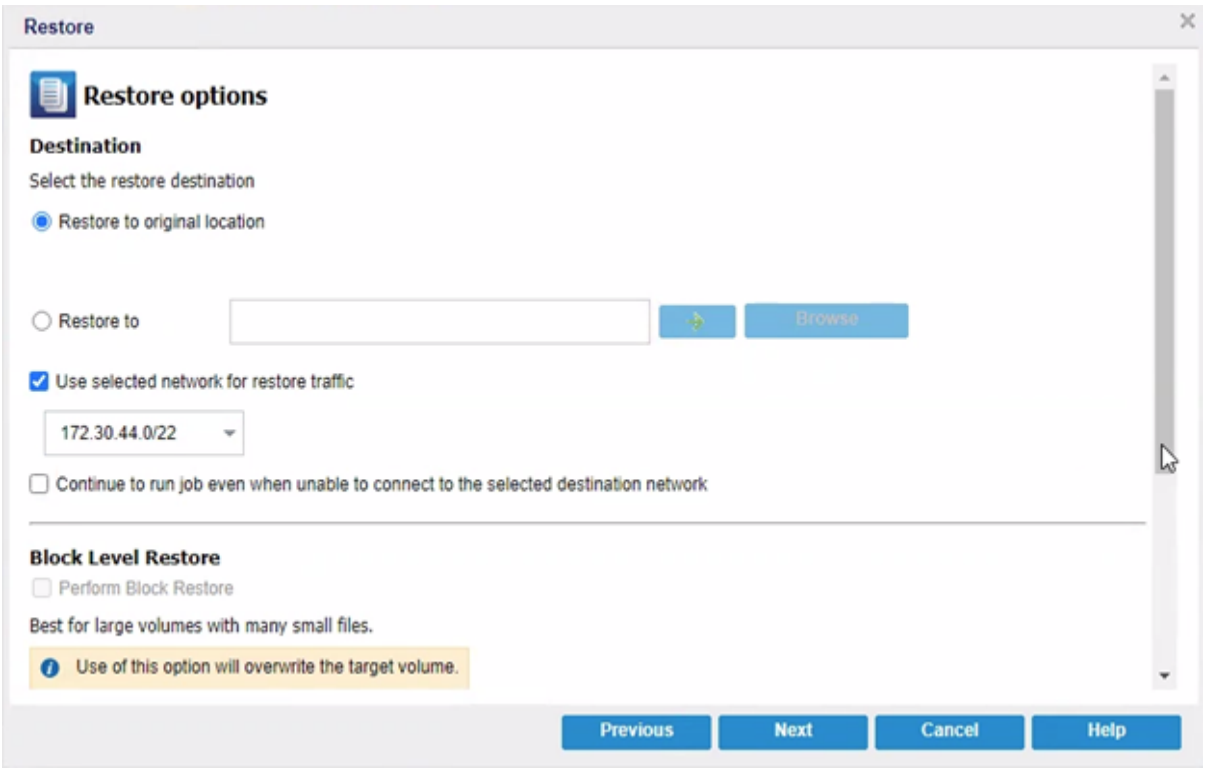
   The recovery point and content to restore is specified.

Define the Restore Options

After you specify a recovery point and content to restore, define the copy options for the selected recovery point. This section explains how to define the restore options for the selected recovery point.

Follow these steps:

1. On the Restore Options dialog, select the restore destination.



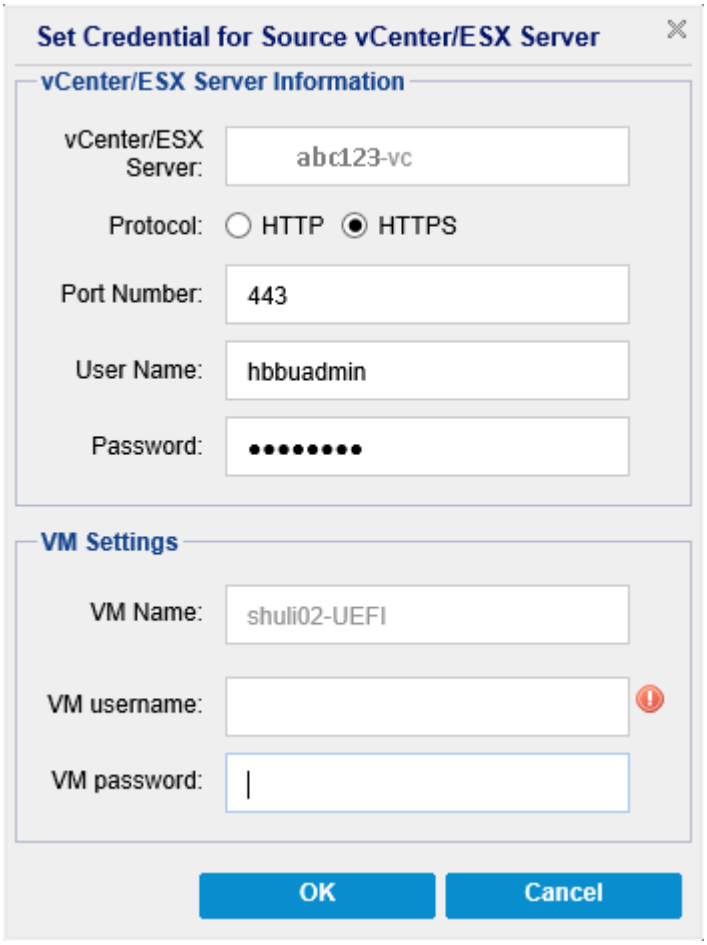
The available destination options are:

Restore to Original Location

Restores to the original location from where the backup image was captured.

Note: If you performed the recovery point backup using host-based agentless backup, restoring to original location is to restore the file back in to the virtual machine. In this case, a dialog box opens. You may enter the credentials of the hypervisor, and the operating system of the virtual machine.

For VMware VM:

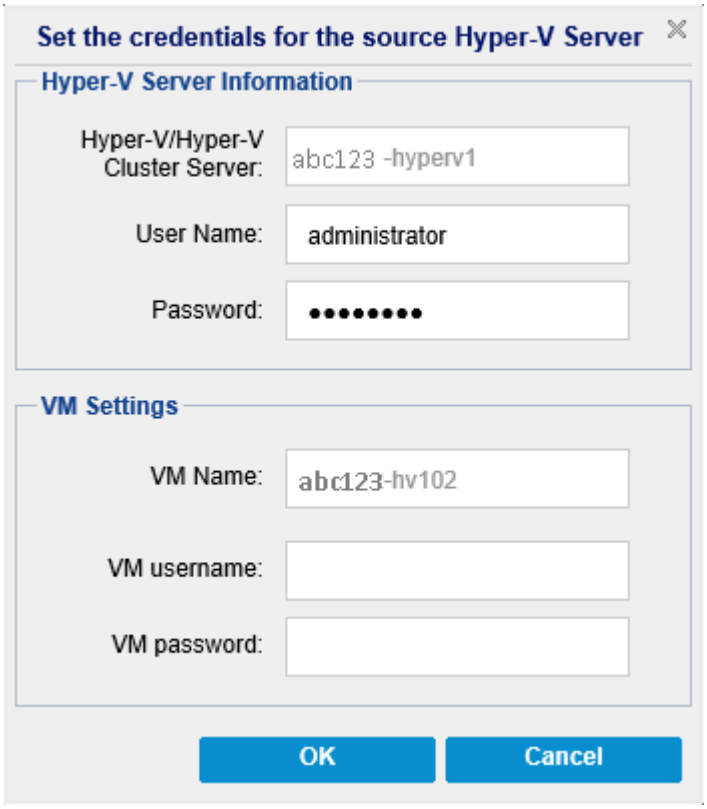


Note: To create or write files inside the VM, consider the following requirements for the settings and account permission of virtual machine:

- VMware Tools is installed and running.
- Firewall must allow File and Printer Sharing.
- The account is built-in local administrator, built-in domain administrator, or domain account that is member of the local Administrators group. If other accounts are used, do the following:
  - Disable the UAC remote access. To disable UAC remote access, see Import Virtual Machine Using Additional Administrative Account.
  - Disable UAC in the Local Security Policy by disabling the setting Run all administrator in Admin Approval Mode at secpol.msc -> Local Policies -> Security Options. (Secpol.msc is Microsoft's security policy editor).

Important! Do not attempt to disable the UAC in the User Account Control Settings dialog box that opens from the control panel.

For Hyper-V VM:



Note: To create or write files inside the VM, consider the following requirements for the settings and account permission of virtual machine:

- Hyper-V integration services are installed and running.
- Firewall must allow File and Printer Sharing.
- The account is built-in local administrator, built-in domain administrator, or domain account that is member of the local Administrators group. If other accounts are used, disable the UAC remote access. To disable UAC remote access, see Import Virtual Machine Using Additional Administrative Account.
- If virtual machine guest OS is Client version Windows (such as Windows 10), you need to manually configure firewall to allow Windows Management Instrumentation (WMI).

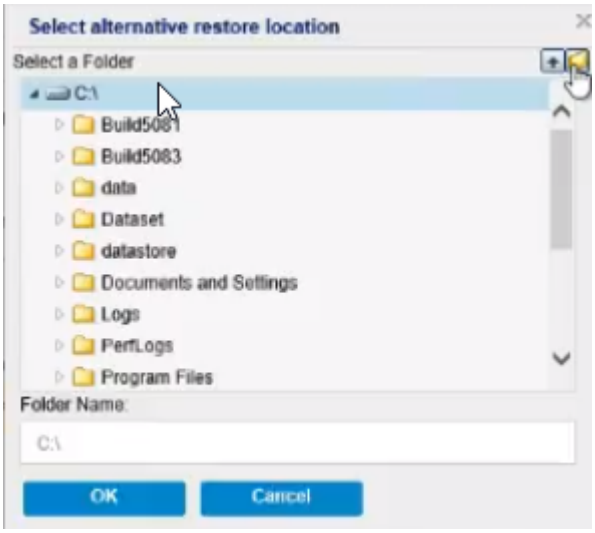
Restore to

Restores to the specified location. To restore data to the specified location, do the following:

- To provide the destination, click **Browse**.

  The Select alternative restore location dialog appears.

- Select the existing folder or create a new folder as needed, and then click **OK**.

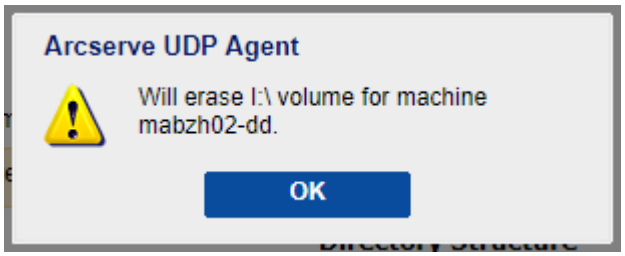


- To verify the connection to the specified location, click the green arrow button. If necessary, enter the Username and Password to gain access to that location.

2. To enable the communication between Windows Agent and Recovery Point Server, select the **Use selected network for restore traffic** check box, and then select the network from the drop-down list.

   Note: If the selected backup network is not accessible and to continue the backup job with the available network or with the default network, click the **Continue to run job even when unable to connect to the selected backup network** check box.

3. To improve the throughput when restoring a large volume with many small files, select a volume, and then under Block Level Restore, click the **Perform Block Restore** check box. Other options get disabled and a notification appears. To continue with the block level restore, click **OK**.



Note: The data gets restored to current UDP Agent machine, and the target volume is overwritten. The target volume is not accessible during the restore job.

4. Specify the Resolving Conflicts option that Arcserve UDP performs if conflicts are encountered during the restore process.

   The available options are:

   Overwrite existing files

   Overwrites (replaces) any existing files, which are at the restore destination. All the objects are restored from the backup files regardless of their current presence on your computer.

   Replace active files

   Replaces any active files upon reboot. If during the restore attempt Arcserve UDP Agent (Windows) discovers that the existing file is currently in use or being accessed, it will not immediately replace that file, but instead to avoid any problems will delay the replacement of the active files until the next time the machine is rebooted. (The restore occurs immediately, but the replacement of any active files is done during the next reboot).

   This option is only available if you select the **Overwrite existing files** option.

   Note: If you do not select this option, any active file is skipped from the restore.

   Rename files

   Creates a new file if the file name already exists. Selecting this option copies the source file to the destination with the same file name but a different extension. Data is then restored to the new file.

   Skip existing files

   Skips over and not overwrite (replace) any existing files that are located at the restore destination. Only objects that are not currently existing on your machine are restored from the backup files.

   Default: Skip existing files.

5. For Directory Structure, if you want to create a root directory during restore, select the **Create root directory** checkbox.

   The **Create root directory** checkbox specifies that if a root directory structure exists in the captured backup image, Arcserve UDP recreates that same root directory structure on the restore destination path.

   If this option not selected, the file or folder is restored directly to the destination folder.

For example, if during the backup you captured the files "C:\Folder-1\SubFolder2\A.txt" and "C:\Folder1\SubFolder2\B.txt", and during the restore you specified the restore destination as "D:\Restore".

- If you select to restore the "A.txt" and "B.txt" files individually, the destination for the restored files will be "D:\Restore\A.txt" and "D:\Restore\B.txt" (the root directory above the specified file level will not be recreated).
- If you select to restore from the "SubFolder2" level, the destination for the restored files will be "D:\Restore\SubFolder2\A.txt" and "D:\Restore\SubFolder2\B.txt" (the root directory above the specified folder level will not be recreated).

If this option selected, the entire root directory path for the files/folders (including the volume name) is recreated to the destination folder. If the files/folders to be restored are from the same volume name, then the destination root directory path does not include that volume name. However, if the files/folders to be restored are from different volume names, then the destination root directory path does include the volume name.

For example, if during the backup you captured the files "C:\Folder-1\SubFolder2\A.txt", "C:\Folder1\SubFolder2\B.txt", and also E:\Folder-3\SubFolder4\C.txt", and during the restore you specified the restore destination as "D:\Restore".

- If you select to restore just the "A.txt" file, the destination for the restored file will be "D:\Restore\ Folder1\SubFolder2\A.txt" (the entire root directory without the volume name will be recreated).
- If you select to restore both the "A.txt" and "C.txt" files, the destination for the restored files will be "D:\Restore\C\Folder1\SubFolder2\A.txt" and "D:\Restore\E\Folder3\SubFolder4\C.txt" (the entire root directory with the volume name will be recreated).

6. For Recovering ACL, to skip the original permission for the restored files/folders, select the **skip recovering ACL of files / folders** checkbox. Selecting this option lets you inherit the permissions of target folder instead. If you do not select this option, the original permissions are kept.
7. For **Backup Encryption or Password Protection**, if necessary, specify the backup encryption password when the data you are trying to restore is encrypted.

   Password is not required if you are attempting to restore from the same Arcserve UDP Agent (Windows) computer from where the encrypted backup was performed. However, if you are attempting to restore from a different Arcserve UDP Agent (Windows) computer, a password is required.

Note: If a clock icon appears with a lock symbol, it indicates that the recovery point contains encrypted information and may require a password for restore.

8. Click **Next**.

   The Restore Summary page opens.

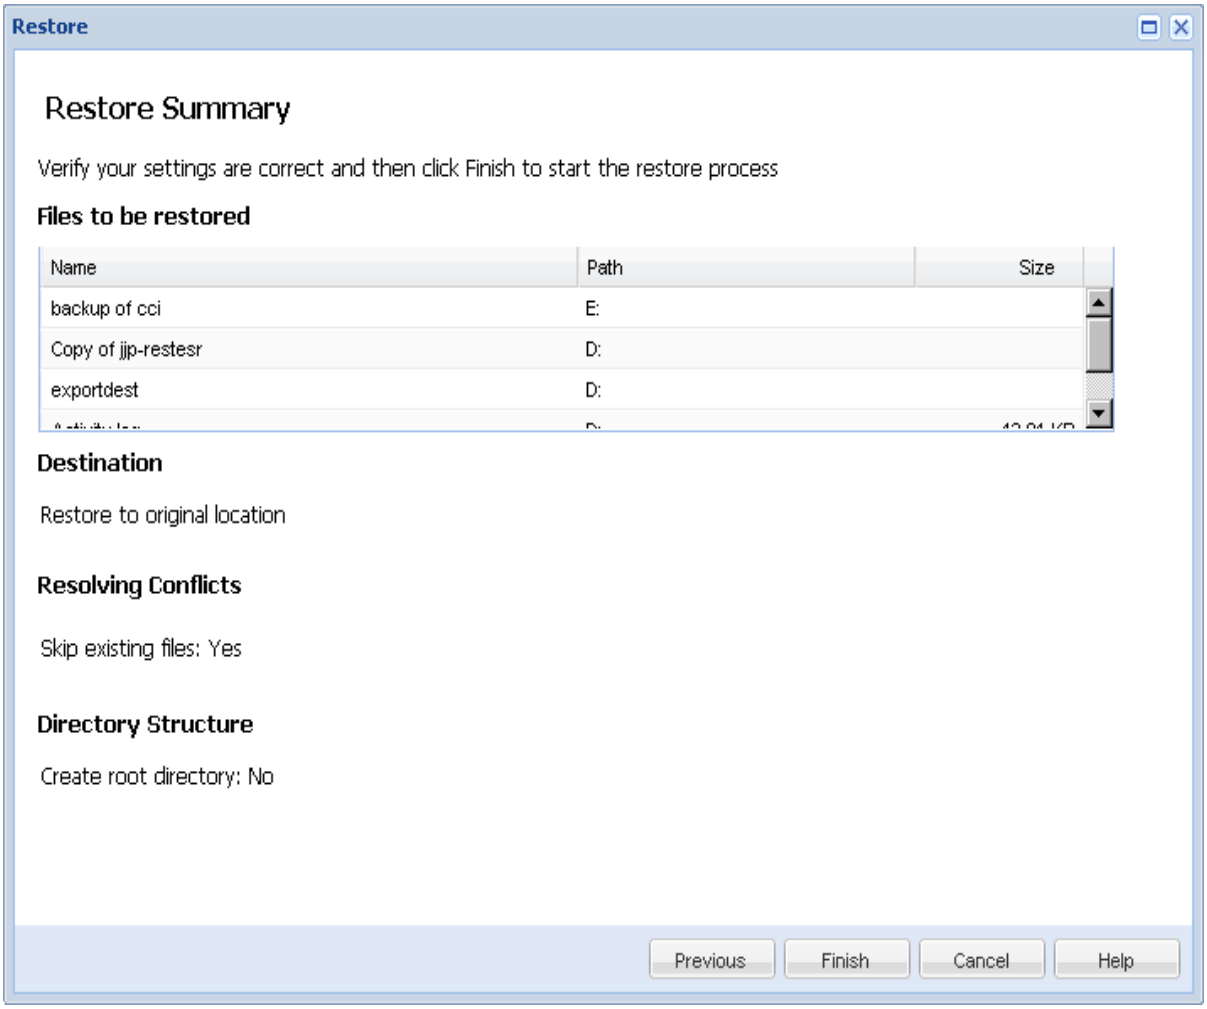
   The restore options are defined to restore from a recovery point.

Restore the Recovery Point Content

After you define the restore options, verify that your settings are correct and confirm the restore process. Restore Summary page helps you to review all the restore options that you defined and modify them if necessary.

Follow these steps:

1. On the Restore Summary dialog, review the displayed information to verify that all the restore options and settings are correct.



2. Do one of the following:
   - If the summary information is incorrect, click **Previous** and go back to the applicable dialog to change the incorrect setting.
   - If the summary information is correct, click **Finish** to launch the restore process.

The recovery point content is restored.

Verify that Content is Restored

After the completion of the restore process, verify that content is restored to the specified destination.

Follow these steps:

1. From the Arcserve UDP console, navigate to **jobs** > **Jobs Successfully Completed**.

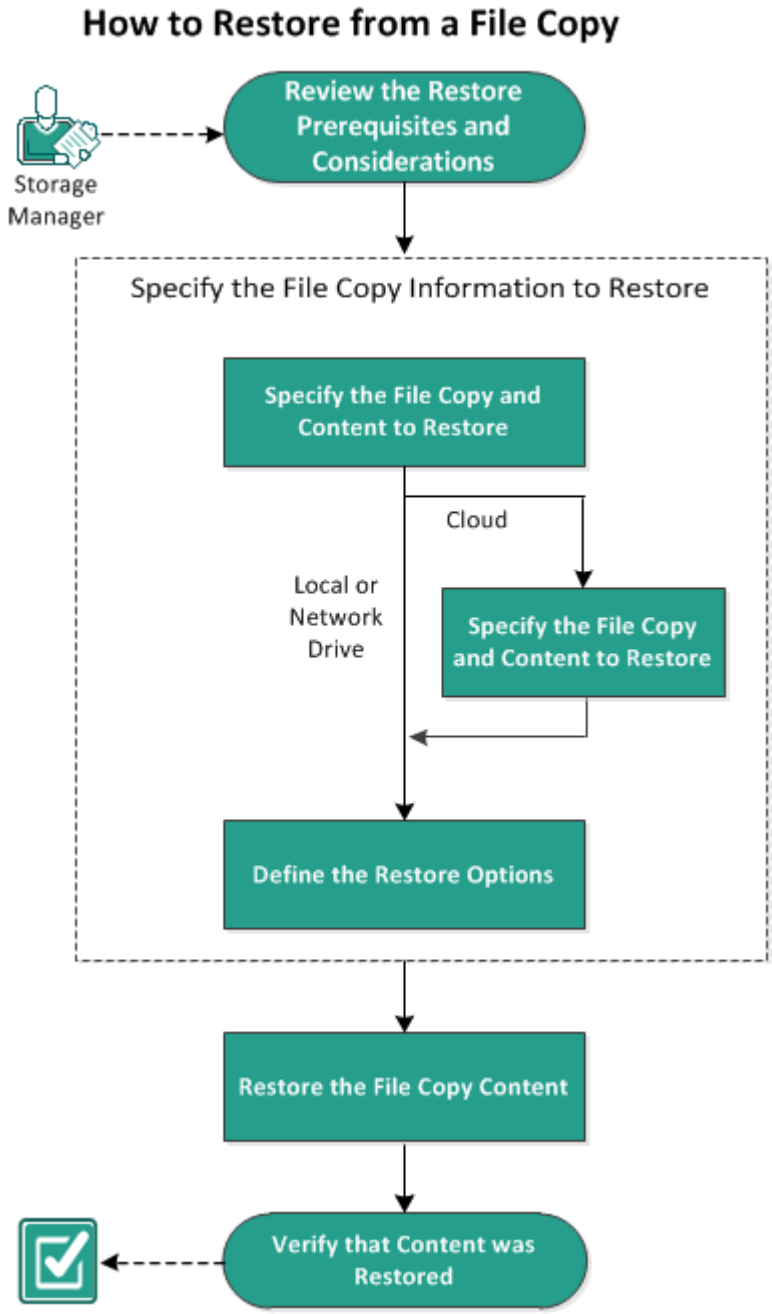
   The list of backup jobs completed appear on the center pane.
2. Click the finished reverse replication job. The right pane refreshes.
3. Under Job Details, click the **View Logs** hyperlink.
4. Verify the content to confirm the restore job.

   The restored content is successfully verified.

How to Restore From a File Copy

Each time Arcserve UDP performs a successful file copy job, it backs up all files that have changed since the last successful file copy job. This restore method allows you to browse the file copied data and specify exactly which file you want to restore.

The following diagram illustrates the process to restore from a file copy:



Perform the following tasks to restore from a File Copy:

1. Review the Restore Prerequisites and Considerations
2. Specify the File Copy Information to Restore
    a. Specify the File Copy and Content to Restore
        - Specify Cloud Configuration for Restore
    b. Define the Restore Options
3. Restore the Recovery Point Content
4. Verify that Content was Restored

Review the Restore Prerequisites and Considerations

Verify that the following prerequisites exist before performing a restore:

- You have at least one file copy available to restore.
- You have a valid and accessible file copy destination to restore the file copy content from.
- You have a valid and accessible target location to restore the file copy content to.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Review the following restore considerations:

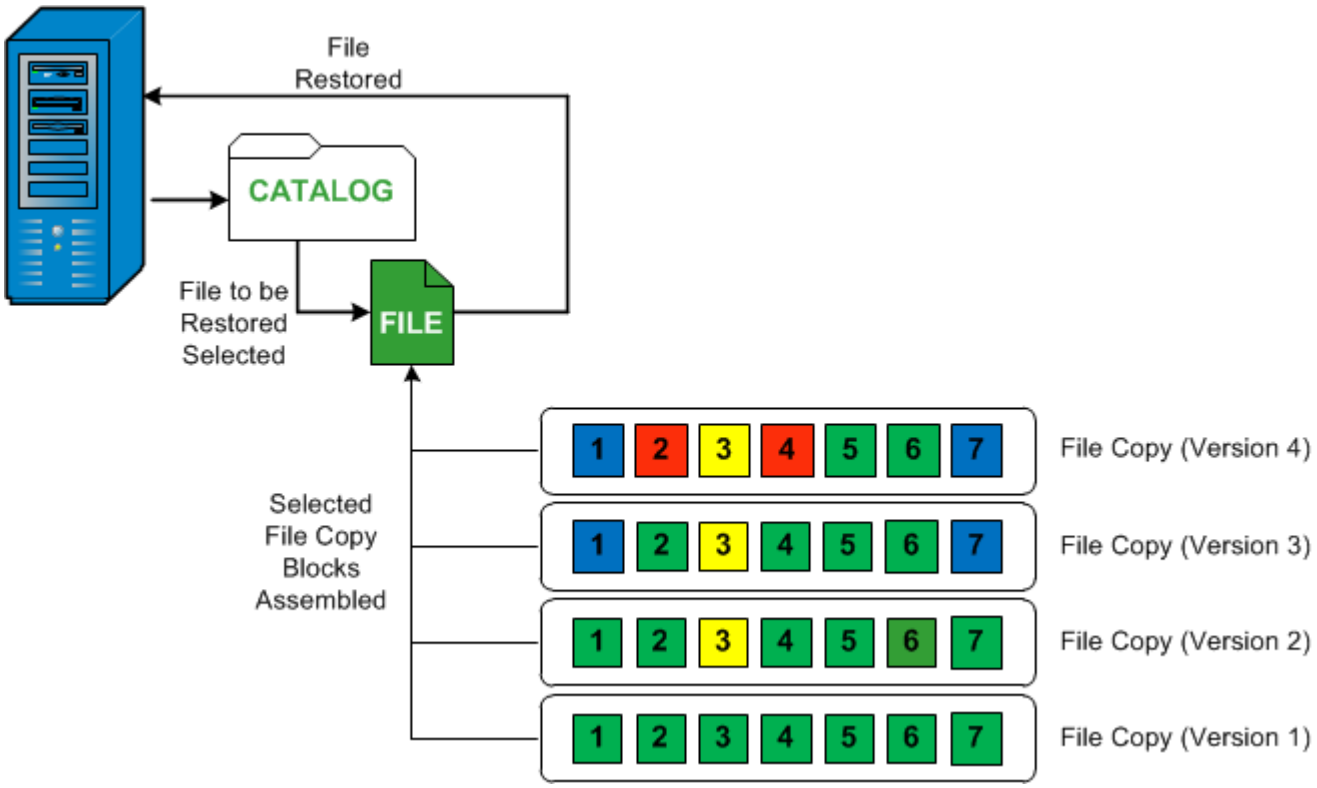
- Arcserve UDP only allows one restore job to run at the same time. If you attempt to launch a restore job manually, while another restore job is running, an alert message opens informing that another job is running and requests to try again later.
- If the restore is to a remote destination and if all the drive letters (A - Z) are occupied, the restore to a remote path will not succeed. Arcserve UDP Agent (Windows) needs to use one drive letter to mount the remote destination path.
- Enhance file copy to optimize performance:
  - File Copy can send multiple chunks simultaneously to the destination (ArchMultChunkIO)
  - File Copy can copy more than one file at a time from the destination (ThreadsForArchive).
  - Restore from a File Copy can download more than one file at a time (ThreadsForRestore).
  - Catalog Synchronization uses multiple threads (ThreadForCatalogSync).

  You can change the default File Copy Registry values by modifying the appropriate DWORD value. For more information, see Configure File Copy Settings to Optimize Performance in the *Agent for Windows online help*.
- (Optional) Understand how the restore process works. For more information, see How File Level Restores Work.

How File Level Restores Work

During a File Copy, each backed up file is made up of a collection of blocks that define the particular file. A catalog file is created for every version of the backed up file, along with the individual blocks that were used for these files. When you need to restore a particular file, you can browse and select the file you want to restore and the file copy versions you want to restore from. Then, Arcserve UDP collects the version of the blocks that were used for the file copy of the specified file, which reassembles and restores the file.

The following flow diagram shows the process of how Arcserve UDP restores a specific file:



Specify the File Copy Information to Restore

Arcserve UDP provides you with an option to restore data from a file copy. The goal of running a successful restore job is to identify quickly the data you need and to retrieve it from the appropriate backup location. Each restore job requires a source and destination.

The process involved in restoring from a file copy is as follows:

- Specify the File Copy and Content to Restore
- Define the Restore Options

Specify the File Copy and Content to Restore

Use the **Browse File Copies** option to restore from a file copy. This restore method allows you to browse the file copied data and specify exactly which file you want to restore.

Follow these steps:

1. Access the restore method selection dialog in one of the following ways:

   From Arcserve UDP:

   a. Log into Arcserve UDP.

   b. Click the **resources** tab.

   c. Select **All Nodes** in the left pane.

      All the added nodes are displayed in the center pane.

   d. In the center pane, select the node and click **Actions**.

   e. Click **Restore** from the **Actions** drop-down menu.

      The restore method selection dialog opens.

      Note: You are automatically logged in to the agent node and the restore method selection dialog is opened from the agent node.

   From Arcserve UDP Agent (Windows):

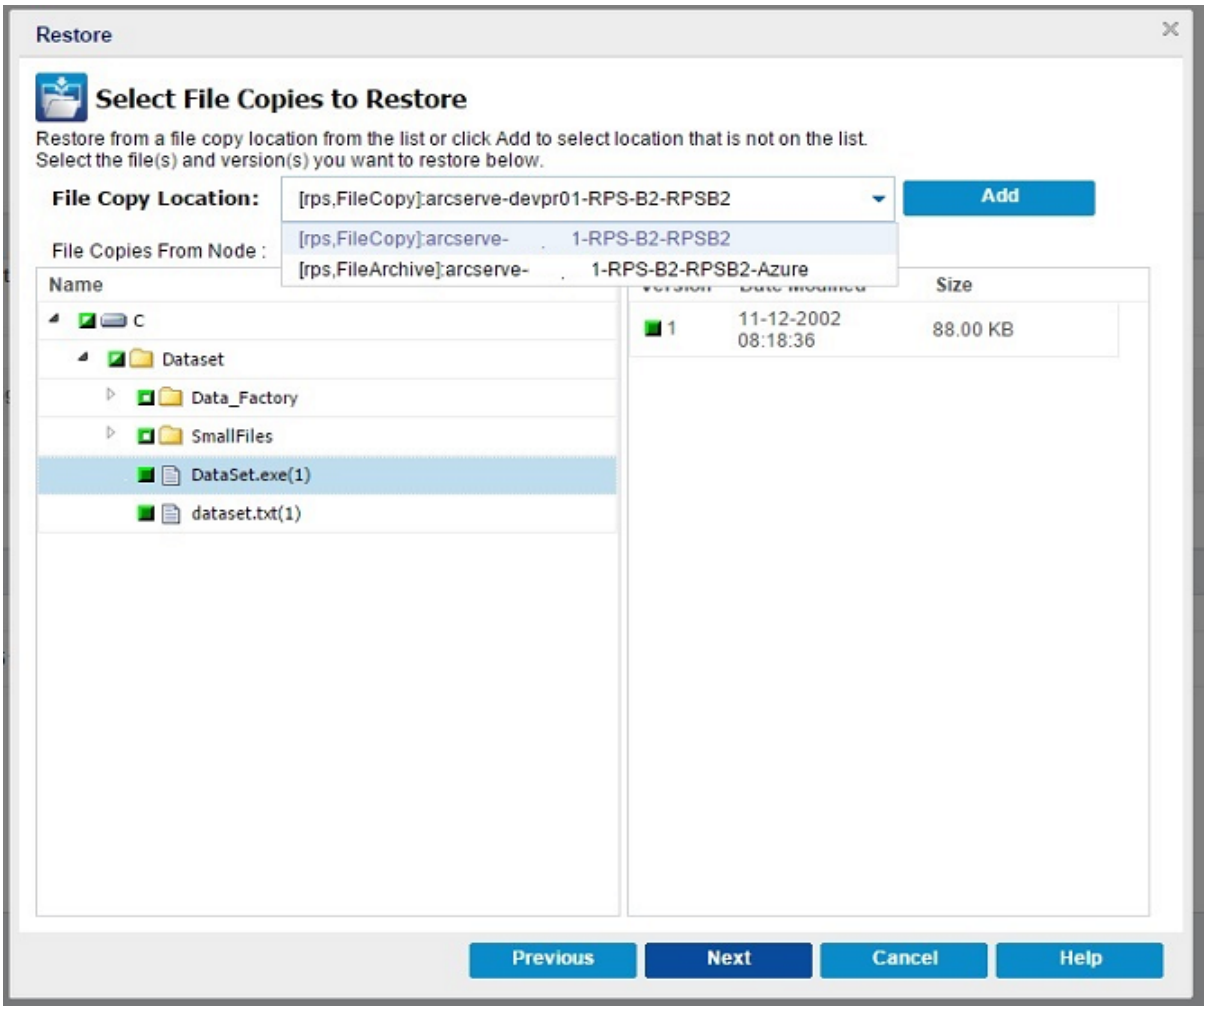
   a. Log into Arcserve UDP Agent (Windows).

   b. From the home page, select **Restore**.

   The restore method selection dialog opens.

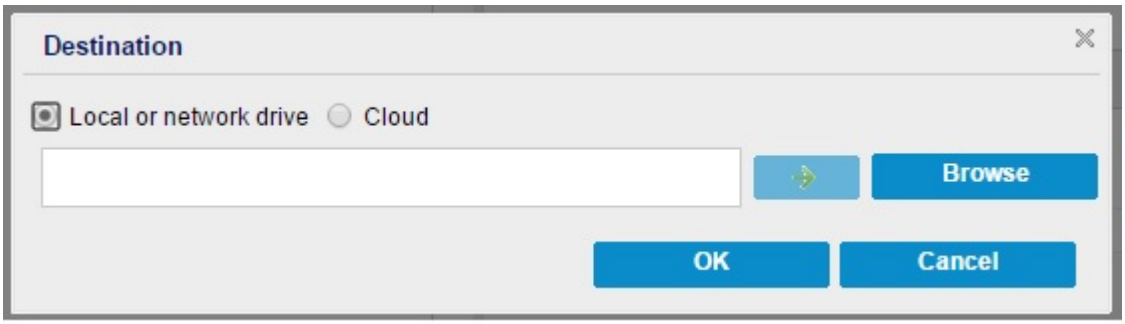
2. Click the **Browse File Copies** option.

   The **Restore** dialog opens. The destination that is currently showing in the **Restore From** field is the configured default **File Copy** destination.



3. If necessary, you can click **Add** to browse to an alternate location where your file copy images are stored.

   The **Destination** dialog opens displaying the available alternate destination options.



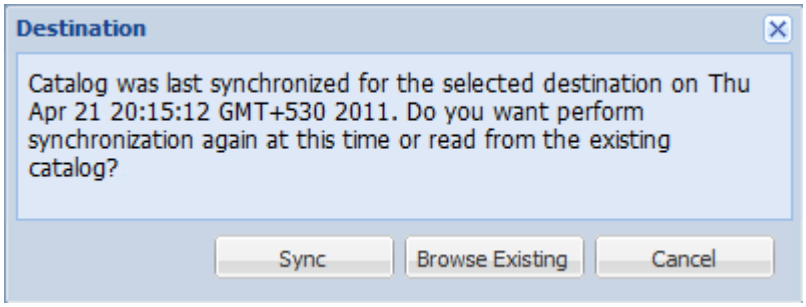
Local or network drive

The **Select a Backup Location** dialog opens, allowing you to browse to and select an alternate local or network drive location.

Cloud

The **Cloud Configuration** dialog opens, allowing you to access and select an alternate cloud location. For more information about this dialog, see Specify Cloud Configuration for Restore.

Regardless of whether you selected to restore from **Local or network drive** or from **Cloud**, when you change the destination to an alternate location a pop-up dialog will appear, asking if you want to perform a new catalog synchronization or read from the existing catalog.



- If you are performing a catalog synchronization for the first time, the **Browse Existing** button are disabled because is file copy catalog does not exist locally.
- If a catalog synchronization has been previously performed, this dialog will display details of the last time the catalog was synchronized from this destination. If more file copy jobs run since that displayed time, your catalog may not be currently synchronized and you can select the **Sync** option to ensure your file copy catalog is up-to-date.

  1. Click **Sync** to download the file copy catalog from the specified file copy destination to your local machine to provide faster browsing.

  2. Click **Browse Existing** to use the file copy catalog that is available locally and not download/sync it again.

4. On the left pane, specify the file copy data to be restored. You can select file copied folders or files to be restored.

   When you select an individual file to be restored, all file copied versions of that file are displayed in the right pane. If multiple versions are available, you must select which file copied version you want to restore.

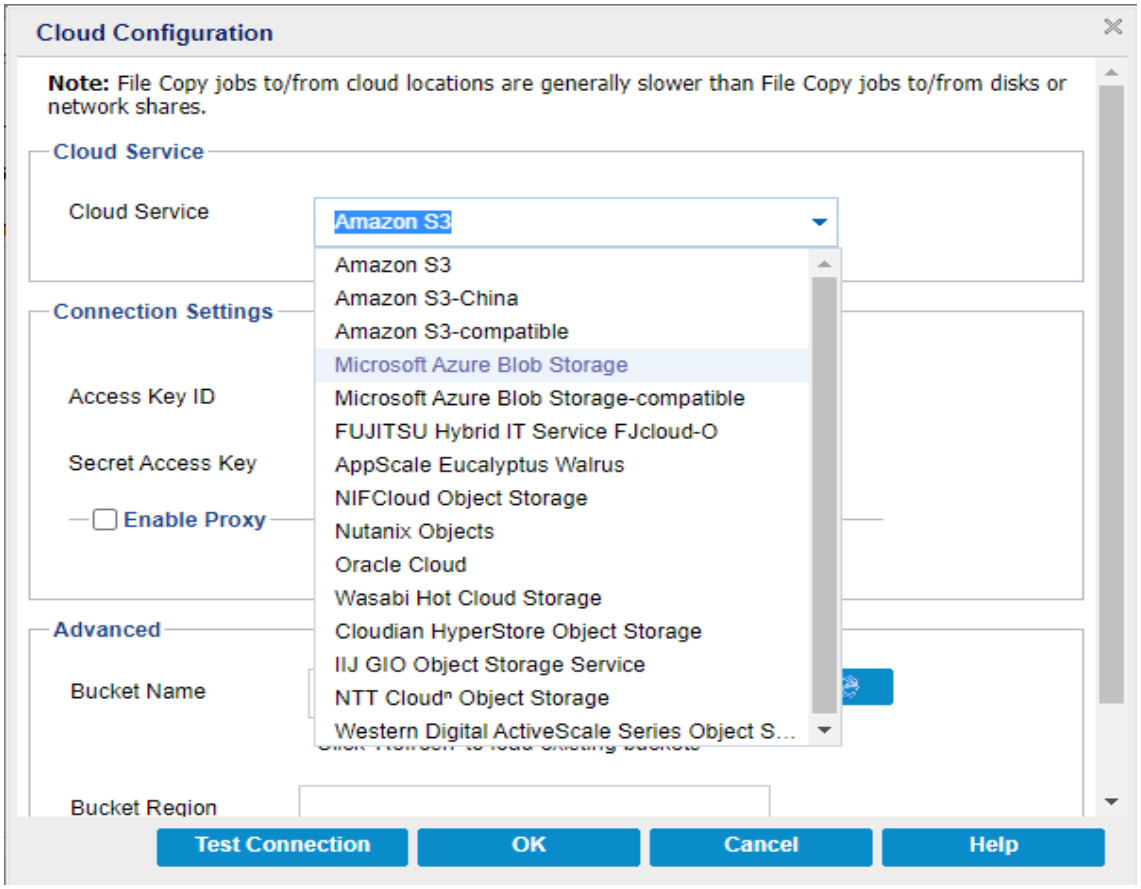
5. After selecting the file copied folder or file version to restore, click **Next**.

   The **Restore Options** dialog opens.

   The **File Copy and Content to restore** is specified.

Specify Cloud Configuration for Restore

Note: The following procedure only applies if you are restoring a file/folder from a file copy or file archive cloud location.



The available options are Amazon S3, Amazon S3-compatible, Windows Azure, Windows Azure-compatible, FUJITSU Hybrid IT Service FJcloud-O, and Eucalyptus-Walrus. (Amazon S3 is the default vendor).

Note: If you are using Eucalyptus-Walrus as your file copy cloud vendor, you will not be able to copy files whose entire path length is greater than 170 characters.

The configuration options for each cloud vendor are similar (with some different terminology), and any differences are described.

1. From the **Browse File Copies** option or the **Find Files/Folders to Restore** option, click **Add**.

   The **Destination** dialog opens.

2. Select **Cloud** and click **Browse**.

   The **Cloud Configuration** dialog opens.

3. Enter the following details:

Storage Name

Specifies the name of the cloud storage. This name will be added to Console for identifying the cloud account. Each cloud account must have a unique storage name.

Storage Service

Select the service from the drop-down list. The configuration option varies depending on the storage service that is selected.

Access Key ID/Account Name/Query ID

Identifies the user who is requesting access this location.

For this field, Amazon S3 uses Access Key ID, Windows Azure and FUJITSU Hybrid IT Service FJcloud-O use Account Name, and Eucalyptus-Walrus uses Query ID.

Secret Access Key/Secret Key

Because your Access Key is not encrypted, this Secret Access Key is a password that is used to verify the authenticity of the request to access this location.

Important! This Secret Access Key is crucial for maintaining the security of your accounts. You should keep your keys and your account credentials in a secure location. Do not embed your Secret Access Key in a web page or other publicly accessible source code and do not transmit it over insecure channels.

For this field, Amazon S3 uses Secret Access Key. Windows Azure, FUJITSU Hybrid IT Service FJcloud-O and Eucalyptus-Walrus use Secret Key.

Proxy Settings

Specifies the proxy server settings. Select **Connect using a proxy server** to enable this option. If you select this option, you must also include the IP address (or machine name) of the proxy server and the corresponding port number that is used by the proxy server for internet connections. You can also select this option if your proxy server requires authentication. You then must provide the corresponding authentication information (Domain Name\User-name and Password) that is required to use the proxy server.

Note: Proxy capability is not available for Eucalyptus-Walrus.

Bucket Name / Container

All files and folders moved or copied to the cloud vendor are stored and organized in your buckets (or containers). Buckets are like a container for your files

and are used to group and organize objects together. Every object stored at the cloud vendor is placed in a bucket.

For this field, Amazon S3 and Eucalyptus-Walrus use Bucket Name. Windows Azure and FUJITSU Hybrid IT Service FJcloud-O use Container.

Note: For the remainder of this step, all references to Buckets can also be applied to Containers unless specified.

Bucket Region

Refers to the region of bucket in Amazon and FUJITSU Hybrid IT Service FJcloud-O.

Contract Number

Refers to the number of contract that FUJITSU Hybrid IT Service FJcloud-O provides.

Project ID

Refers to the ID of project that FUJITSU Hybrid IT Service FJcloud-O generates.

Enable Reduced Redundancy Storage

For Amazon S3 only, this option lets you select to enable Reduced Redundancy Storage (RRS). RRS is a storage option within Amazon S3 that helps you reduce cost by storing non-critical, reproducible data at lower levels of redundancy than Amazon S3s standard storage. Both the standard and reduced redundancy storage options store data in multiple facilities and on multiple devices, but with RRS the data is replicated fewer times, so the cost is less. You should expect the same latency and throughput using either the Amazon S3 standard storage or RRS. By default this option is not selected (Amazon S3 uses the standard storage option).

4. Click **Test Connection** to verify the connection to the specified cloud location.
5. Click **OK**.

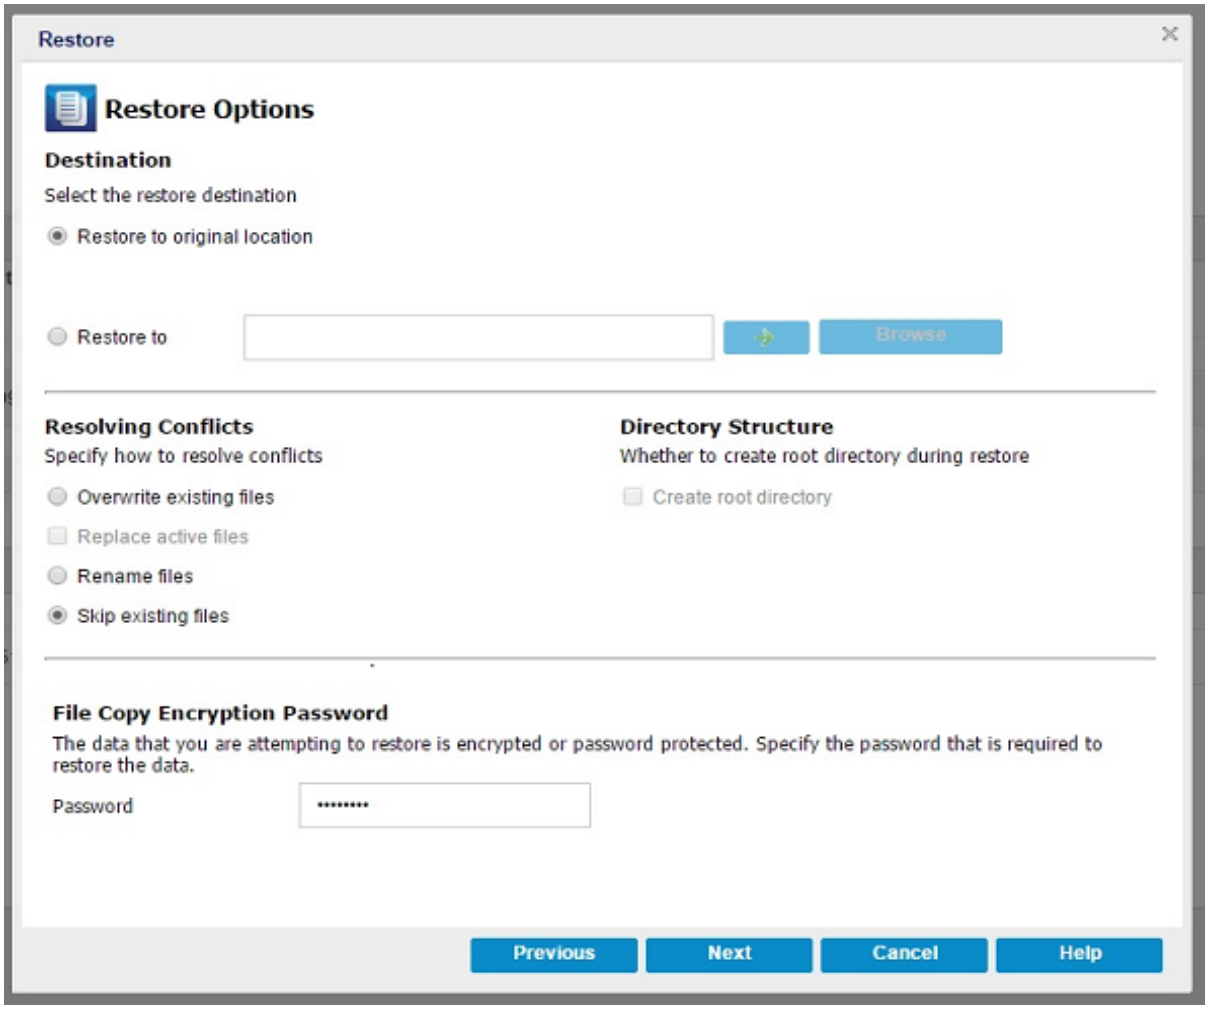
The cloud account is added to the Console.

Define the Restore Options

After you specify the file copy information to restore, define the copy options for the selected file copy and content.

Follow these steps:

1. On the **Restore Options** dialog, select the restore destination.



The available destination options are:

Restore to Original Location

Restores to the original location from where the backup image was captured.

Restore to

Restores to the specified location. You can click the green arrow button to verify the connection to the specified location. If necessary, enter the User-name and Password credentials to gain access to that location.

2. Specify the **Resolving Conflicts** option that Arcserve UDP performs if conflicts are encountered during the restore process.

   The available options are:

   Overwrite existing files

   Overwrites (replaces) any existing files, which are at the restore destination. All objects are restored from the backup files regardless of their current presence on your computer.

   Replace active files

   Replaces any active files upon reboot. If during the restore attempt Arcserve UDP Agent (Windows) discovers that the existing file is currently in use or being accessed, it will not immediately replace that file, but instead to avoid any problems will delay the replacement of the active files until the next time the machine is rebooted. (The restore occurs immediately, but the replacement of any active files is done during the next reboot).

   This option is only available if you select the **Overwrite existing files** option.

   Note: If you do not select this option, any active file is skipped from the restore.

   Rename files

   Creates a new file if the file name already exists. Selecting this option copies the source file to the destination with the same filename but a different extension. Data is then restored to the new file.

   Skip existing files

   Skips over and not overwrite (replace) any existing files that are located at the restore destination. Only objects that are not currently existing on your machine are restored from the backup files.

   Default: Skip existing files.

3. Specify the **Directory Structure** to create a root directory during restore.

   Create root directory

Specifies that if a root directory structure exists in the captured backup image, Arcserve UDP recreates that same root directory structure on the restore destination path.

With this option not selected, the file or folder is restored directly to the destination folder.

For example, if during the backup you captured the files "C:\Folder-1\SubFolder2\A.txt" and "C:\Folder1\SubFolder2\B.txt" and during the restore you specified to the restore destination as "D:\Restore".

- If you select to restore the "A.txt" and "B.txt" files individually, the destination for the restored files will be "D:\Restore\A.txt" and "D:\Restore\B.txt" (the root directory above the specified file level will not be recreated).
- If you select to restore from the "SubFolder2" level, the destination for the restored files will be "D:\Restore\SubFolder2\A.txt" and "D:\Restore\SubFolder2\B.txt" (the root directory above the specified folder level will not be recreated).

With this option selected, the entire root directory path for the files/folders (including the volume name) is recreated to the destination folder. If the files/-folders to be restored are from the same volume name, then the destination root directory path does not include that volume name. However, if the files/-folders to be restored are from different volume names, then the destination root directory path does include the volume name.

For example, if during the backup you captured the files "C:\Folder-1\SubFolder2\A.txt", "C:\Folder1\SubFolder2\B.txt", and also E:\Folder-3\SubFolder4\C.txt" and during the restore you specified to the restore destination as "D:\Restore".

- If you select to restore just the "A.txt" file, the destination for the restored file will be "D:\Restore\ Folder1\SubFolder2\A.txt" (the entire root directory without the volume name will be recreated).
- If you select to restore both the "A.txt" and "C.txt" files, the destination for the restored files will be "D:\Restore\C\Folder1\SubFolder2\A.txt" and "D:\Restore\E\Folder3\SubFolder4\C.txt" (the entire root directory with the volume name will be recreated).

4. Specify the encryption password in **File Copy Encryption Password.**
5. Click **Next**.

   The **Restore Summary** dialog opens.

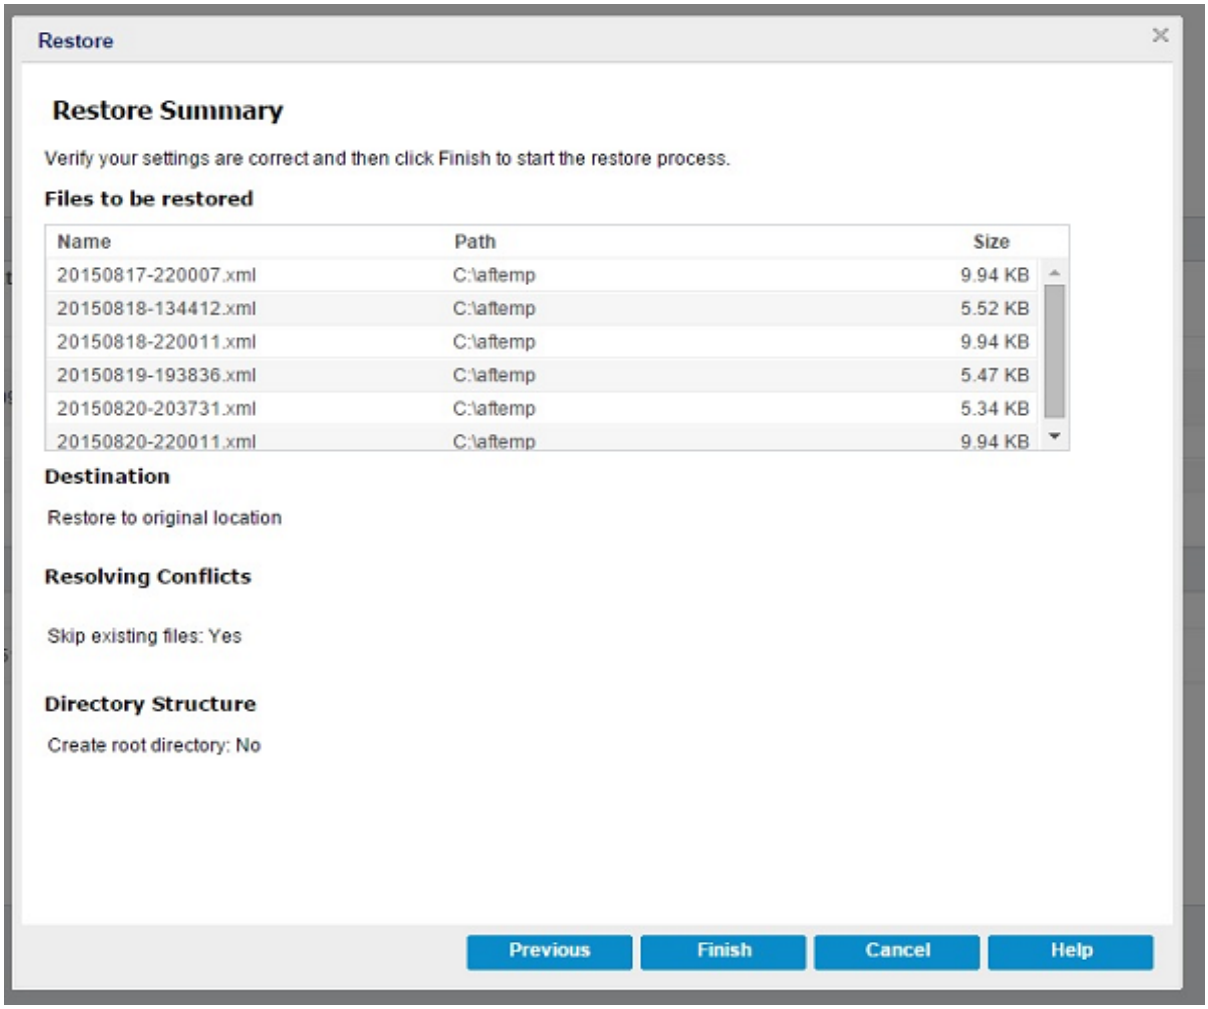
   The restore options are defined to restore from a file copy.

Restore the File Copy Content

After you define the restore options, verify that your settings are correct and confirm the restore process. **Restore Summary** helps you to review all the restore options that you defined and modify them if necessary.

Follow these steps:

1. On the **Restore Summary** dialog, review the displayed information to verify that all the restore options and settings are correct.



- If the summary information is incorrect, click **Previous** and go back to the applicable dialog to change the incorrect setting.
- If the summary information is correct, click **Finish** to launch the restore process.

The file copy content is restored.

Verify that Content was Restored

After the completion of the restore process, verify that content was restored to the specified destination.

Follow these steps:

1. Navigate to the restore destination you specified.

   A list of folders appears.

2. Locate the file to which you have restored the content.

   For example, if you select to restore the **A.txt** file to the restore destination as "D:\Restore, then navigate to the following location:

   D:\Restore\A.txt

3. Verify the content to confirm the restore job.

   The restored content is successfully verified.

How to Restore From a File Archive

Each time Arcserve UDP performs a successful file archive copy job, it archives all files that have changed since the last successful file archive job. This restore method allows you to browse the archived files and specify exactly which file you want to restore.

The file archive restore process is identical to file copy restore.

Perform the following tasks to restore from a File Archive:

1. Review the Restore Prerequisites and Considerations
2. Specify the File Copy Information to Restore
    a. Specify the File Copy and Content to Restore
        - Specify Cloud Configuration for Restore
    b. Define the Restore Options
3. Restore the Recovery Point Content
4. Verify that Content was Restored

Review the Restore Prerequisites and Considerations

Verify that the following prerequisites exist before performing a restore:

- You have at least one file copy available to restore.
- You have a valid and accessible file copy destination to restore the file copy content from.
- You have a valid and accessible target location to restore the file copy content to.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Review the following restore considerations:

- Arcserve UDP only allows one restore job to run at the same time. If you attempt to launch a restore job manually, while another restore job is running, an alert message opens informing that another job is running and requests to try again later.
- If the restore is to a remote destination and if all the drive letters (A - Z) are occupied, the restore to a remote path will not succeed. Arcserve UDP Agent (Windows) needs to use one drive letter to mount the remote destination path.
- Enhance file copy to optimize performance:
  - File Copy can send multiple chunks simultaneously to the destination (ArchMultChunkIO)
  - File Copy can copy more than one file at a time from the destination (ThreadsForArchive).
  - Restore from a File Copy can download more than one file at a time (ThreadsForRestore).
  - Catalog Synchronization uses multiple threads (ThreadForCatalogSync).

  You can change the default File Copy Registry values by modifying the appropriate DWORD value. For more information, see Configure File Copy Settings to Optimize Performance in the *Agent for Windows online help*.
- (Optional) Understand how the restore process works. For more information, see How File Level Restores Work.

Specify the File Copy Information to Restore

Arcserve UDP provides you with an option to restore data from a file copy. The goal of running a successful restore job is to identify quickly the data you need and to retrieve it from the appropriate backup location. Each restore job requires a source and destination.

The process involved in restoring from a file copy is as follows:

- Specify the File Copy and Content to Restore
- Define the Restore Options

Specify the File Copy and Content to Restore

Use the **Browse File Copies** option to restore from a file copy. This restore method allows you to browse the file copied data and specify exactly which file you want to restore.

Follow these steps:

1. Access the restore method selection dialog in one of the following ways:

   From Arcserve UDP:

   a. Log into Arcserve UDP.

   b. Click the **resources** tab.

   c. Select **All Nodes** in the left pane.

      All the added nodes are displayed in the center pane.

   d. In the center pane, select the node and click **Actions**.

   e. Click **Restore** from the **Actions** drop-down menu.

      The restore method selection dialog opens.

      Note: You are automatically logged in to the agent node and the restore method selection dialog is opened from the agent node.

   From Arcserve UDP Agent (Windows):

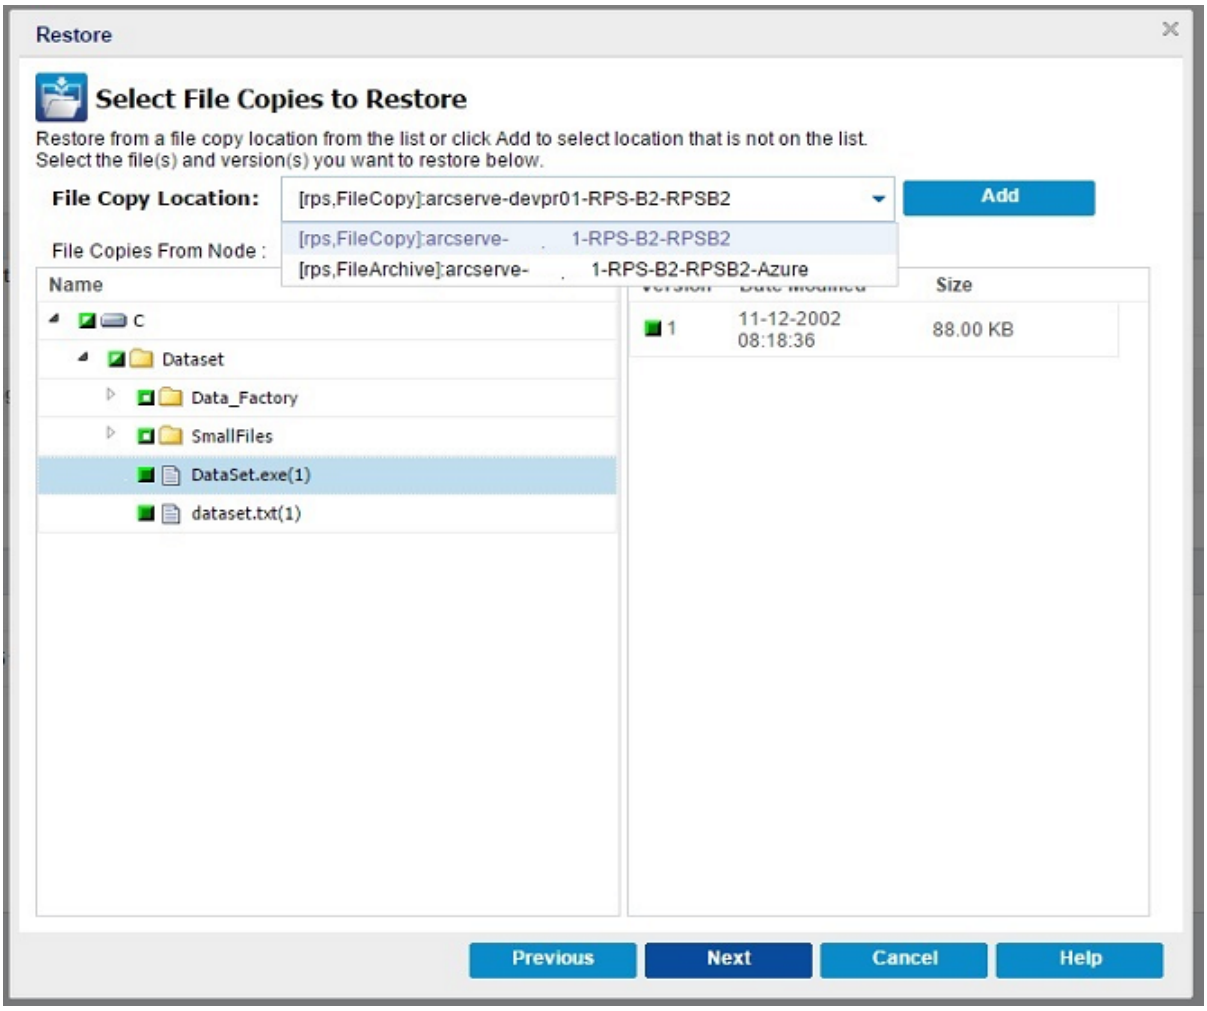
   a. Log into Arcserve UDP Agent (Windows).

   b. From the home page, select **Restore**.

   The restore method selection dialog opens.

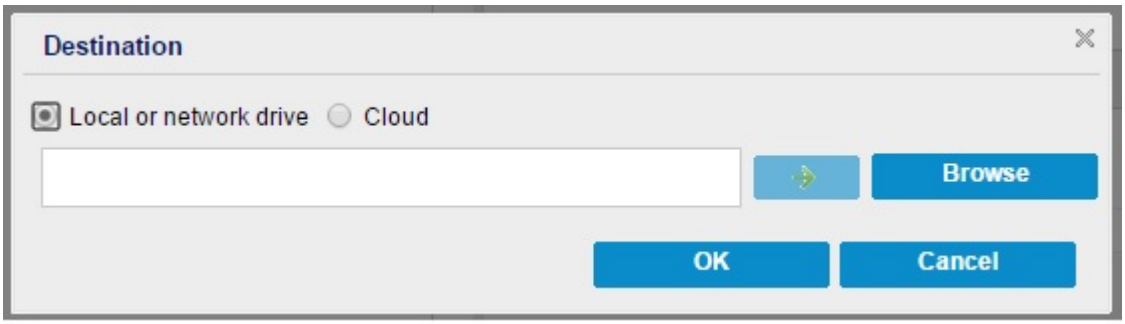
2. Click the **Browse File Copies** option.

   The **Restore** dialog opens. The destination that is currently showing in the **Restore From** field is the configured default **File Copy** destination.



3. If necessary, you can click **Add** to browse to an alternate location where your file copy images are stored.

   The **Destination** dialog opens displaying the available alternate destination options.



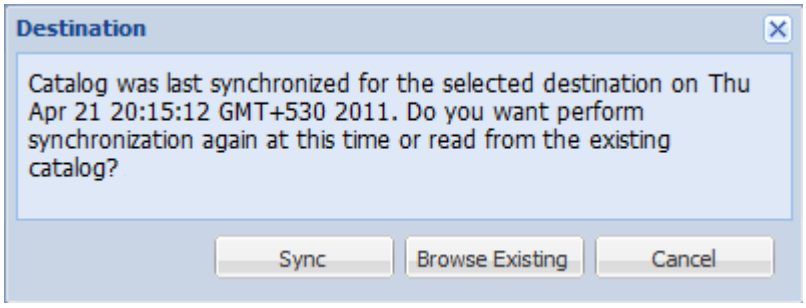
Local or network drive

The **Select a Backup Location** dialog opens, allowing you to browse to and select an alternate local or network drive location.

Cloud

The **Cloud Configuration** dialog opens, allowing you to access and select an alternate cloud location. For more information about this dialog, see Specify Cloud Configuration for Restore.

Regardless of whether you selected to restore from **Local or network drive** or from **Cloud**, when you change the destination to an alternate location a pop-up dialog will appear, asking if you want to perform a new catalog synchronization or read from the existing catalog.



- If you are performing a catalog synchronization for the first time, the **Browse Existing** button are disabled because is file copy catalog does not exist locally.
- If a catalog synchronization has been previously performed, this dialog will display details of the last time the catalog was synchronized from this destination. If more file copy jobs run since that displayed time, your catalog may not be currently synchronized and you can select the **Sync** option to ensure your file copy catalog is up-to-date.

  1. Click **Sync** to download the file copy catalog from the specified file copy destination to your local machine to provide faster browsing.

  2. Click **Browse Existing** to use the file copy catalog that is available locally and not download/sync it again.

4. On the left pane, specify the file copy data to be restored. You can select file copied folders or files to be restored.

   When you select an individual file to be restored, all file copied versions of that file are displayed in the right pane. If multiple versions are available, you must select which file copied version you want to restore.

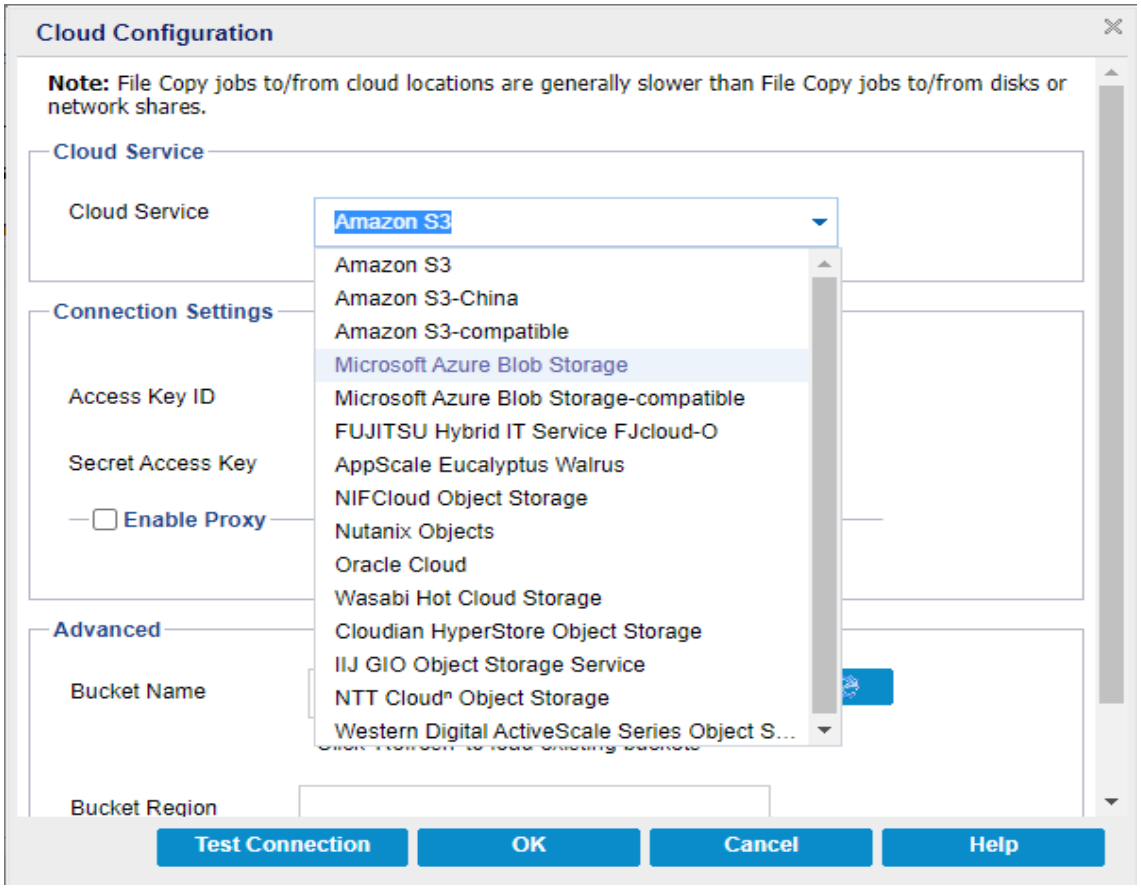
5. After selecting the file copied folder or file version to restore, click **Next**.

   The **Restore Options** dialog opens.

   The **File Copy and Content to restore** is specified.

Specify Cloud Configuration for Restore

Note: The following procedure only applies if you are restoring a file/folder from a file copy or file archive cloud location.



The available options are Amazon S3, Amazon S3-compatible, Windows Azure, Windows Azure-compatible, FUJITSU Hybrid IT Service FJcloud-O, and Eucalyptus-Walrus. (Amazon S3 is the default vendor).

Note: If you are using Eucalyptus-Walrus as your file copy cloud vendor, you will not be able to copy files whose entire path length is greater than 170 characters.

The configuration options for each cloud vendor are similar (with some different terminology), and any differences are described.

1. From the **Browse File Copies** option or the **Find Files/Folders to Restore** option, click **Add**.

   The **Destination** dialog opens.

2. Select **Cloud** and click **Browse**.

   The **Cloud Configuration** dialog opens.

3. Enter the following details:

Storage Name

Specifies the name of the cloud storage. This name will be added to Console for identifying the cloud account. Each cloud account must have a unique storage name.

Storage Service

Select the service from the drop-down list. The configuration option varies depending on the storage service that is selected.

Access Key ID/Account Name/Query ID

Identifies the user who is requesting access this location.

For this field, Amazon S3 uses Access Key ID, Windows Azure and FUJITSU Hybrid IT Service FJcloud-O use Account Name, and Eucalyptus-Walrus uses Query ID.

Secret Access Key/Secret Key

Because your Access Key is not encrypted, this Secret Access Key is a password that is used to verify the authenticity of the request to access this location.

Important! This Secret Access Key is crucial for maintaining the security of your accounts. You should keep your keys and your account credentials in a secure location. Do not embed your Secret Access Key in a web page or other publicly accessible source code and do not transmit it over insecure channels.

For this field, Amazon S3 uses Secret Access Key. Windows Azure, FUJITSU Hybrid IT Service FJcloud-O and Eucalyptus-Walrus use Secret Key.

Proxy Settings

Specifies the proxy server settings. Select **Connect using a proxy server** to enable this option. If you select this option, you must also include the IP address (or machine name) of the proxy server and the corresponding port number that is used by the proxy server for internet connections. You can also select this option if your proxy server requires authentication. You then must provide the corresponding authentication information (Domain Name\User-name and Password) that is required to use the proxy server.

Note: Proxy capability is not available for Eucalyptus-Walrus.

Bucket Name / Container

All files and folders moved or copied to the cloud vendor are stored and organized in your buckets (or containers). Buckets are like a container for your files

and are used to group and organize objects together. Every object stored at the cloud vendor is placed in a bucket.

For this field, Amazon S3 and Eucalyptus-Walrus use Bucket Name. Windows Azure and FUJITSU Hybrid IT Service FJcloud-O use Container.

Note: For the remainder of this step, all references to Buckets can also be applied to Containers unless specified.

Bucket Region

Refers to the region of bucket in Amazon and FUJITSU Hybrid IT Service FJcloud-O.

Contract Number

Refers to the number of contract that FUJITSU Hybrid IT Service FJcloud-O provides.

Project ID

Refers to the ID of project that FUJITSU Hybrid IT Service FJcloud-O generates.

Enable Reduced Redundancy Storage

For Amazon S3 only, this option lets you select to enable Reduced Redundancy Storage (RRS). RRS is a storage option within Amazon S3 that helps you reduce cost by storing non-critical, reproducible data at lower levels of redundancy than Amazon S3s standard storage. Both the standard and reduced redundancy storage options store data in multiple facilities and on multiple devices, but with RRS the data is replicated fewer times, so the cost is less. You should expect the same latency and throughput using either the Amazon S3 standard storage or RRS. By default this option is not selected (Amazon S3 uses the standard storage option).

4. Click **Test Connection** to verify the connection to the specified cloud location.
5. Click **OK**.

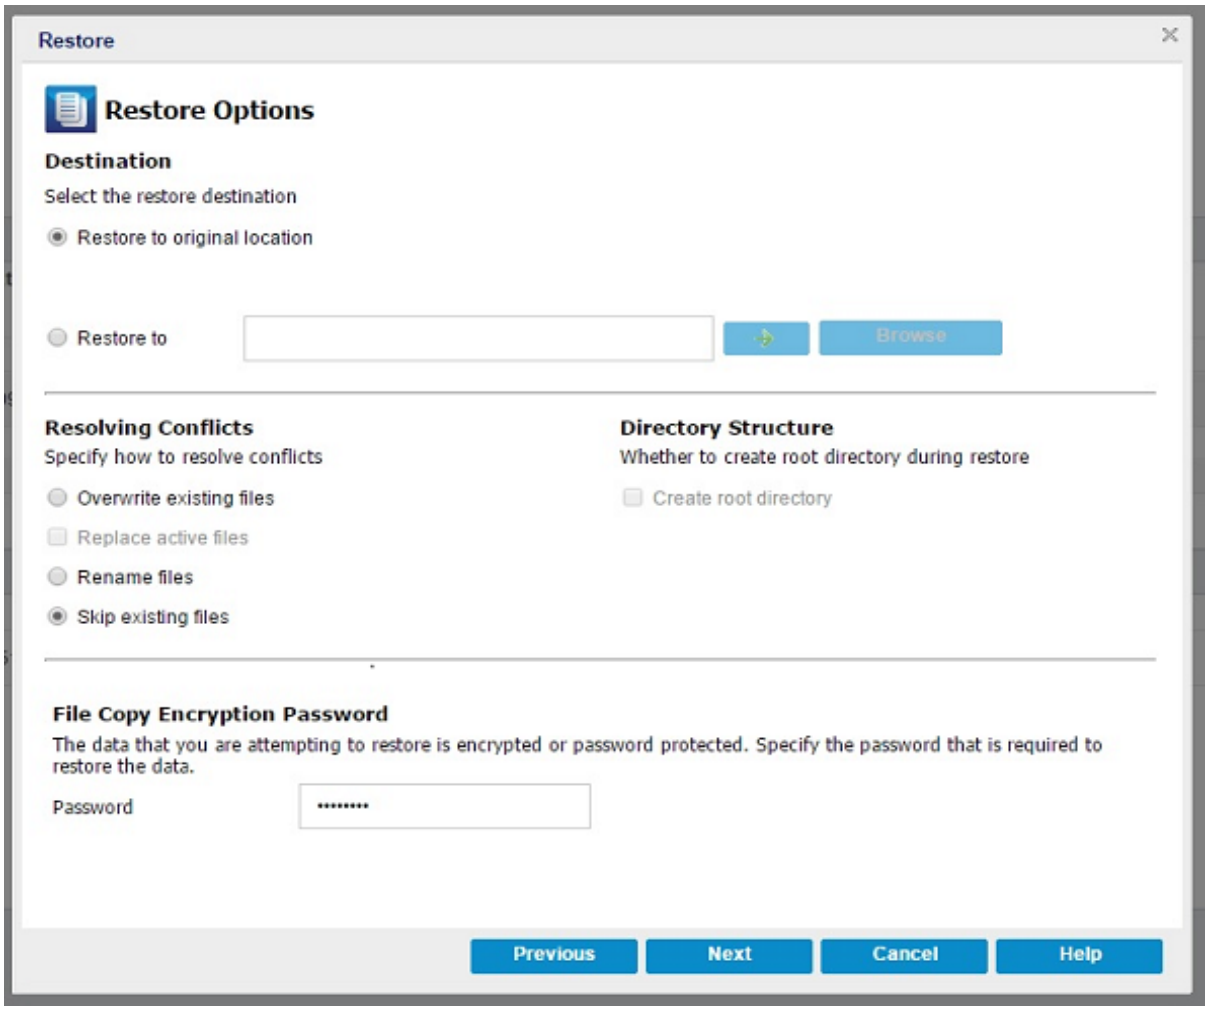
The cloud account is added to the Console.

Define the Restore Options

After you specify the file copy information to restore, define the copy options for the selected file copy and content.

Follow these steps:

1. On the **Restore Options** dialog, select the restore destination.



The available destination options are:

Restore to Original Location

Restores to the original location from where the backup image was captured.

Restore to

Restores to the specified location. You can click the green arrow button to verify the connection to the specified location. If necessary, enter the User-name and Password credentials to gain access to that location.

2. Specify the **Resolving Conflicts** option that Arcserve UDP performs if conflicts are encountered during the restore process.

   The available options are:

   Overwrite existing files

   Overwrites (replaces) any existing files, which are at the restore destination. All objects are restored from the backup files regardless of their current presence on your computer.

   Replace active files

   Replaces any active files upon reboot. If during the restore attempt Arcserve UDP Agent (Windows) discovers that the existing file is currently in use or being accessed, it will not immediately replace that file, but instead to avoid any problems will delay the replacement of the active files until the next time the machine is rebooted. (The restore occurs immediately, but the replacement of any active files is done during the next reboot).

   This option is only available if you select the **Overwrite existing files** option.

   Note: If you do not select this option, any active file is skipped from the restore.

   Rename files

   Creates a new file if the file name already exists. Selecting this option copies the source file to the destination with the same filename but a different extension. Data is then restored to the new file.

   Skip existing files

   Skips over and not overwrite (replace) any existing files that are located at the restore destination. Only objects that are not currently existing on your machine are restored from the backup files.

   Default: Skip existing files.

3. Specify the **Directory Structure** to create a root directory during restore.

   Create root directory

Specifies that if a root directory structure exists in the captured backup image, Arcserve UDP recreates that same root directory structure on the restore destination path.

With this option not selected, the file or folder is restored directly to the destination folder.

For example, if during the backup you captured the files "C:\Folder-1\SubFolder2\A.txt" and "C:\Folder1\SubFolder2\B.txt" and during the restore you specified to the restore destination as "D:\Restore".

- If you select to restore the "A.txt" and "B.txt" files individually, the destination for the restored files will be "D:\Restore\A.txt" and "D:\Restore\B.txt" (the root directory above the specified file level will not be recreated).
- If you select to restore from the "SubFolder2" level, the destination for the restored files will be "D:\Restore\SubFolder2\A.txt" and "D:\Restore\SubFolder2\B.txt" (the root directory above the specified folder level will not be recreated).

With this option selected, the entire root directory path for the files/folders (including the volume name) is recreated to the destination folder. If the files/-folders to be restored are from the same volume name, then the destination root directory path does not include that volume name. However, if the files/-folders to be restored are from different volume names, then the destination root directory path does include the volume name.

For example, if during the backup you captured the files "C:\Folder-1\SubFolder2\A.txt", "C:\Folder1\SubFolder2\B.txt", and also E:\Folder-3\SubFolder4\C.txt" and during the restore you specified to the restore destination as "D:\Restore".

- If you select to restore just the "A.txt" file, the destination for the restored file will be "D:\Restore\ Folder1\SubFolder2\A.txt" (the entire root directory without the volume name will be recreated).
- If you select to restore both the "A.txt" and "C.txt" files, the destination for the restored files will be "D:\Restore\C\Folder1\SubFolder2\A.txt" and "D:\Restore\E\Folder3\SubFolder4\C.txt" (the entire root directory with the volume name will be recreated).

4. Specify the encryption password in **File Copy Encryption Password.**
5. Click **Next**.

   The **Restore Summary** dialog opens.

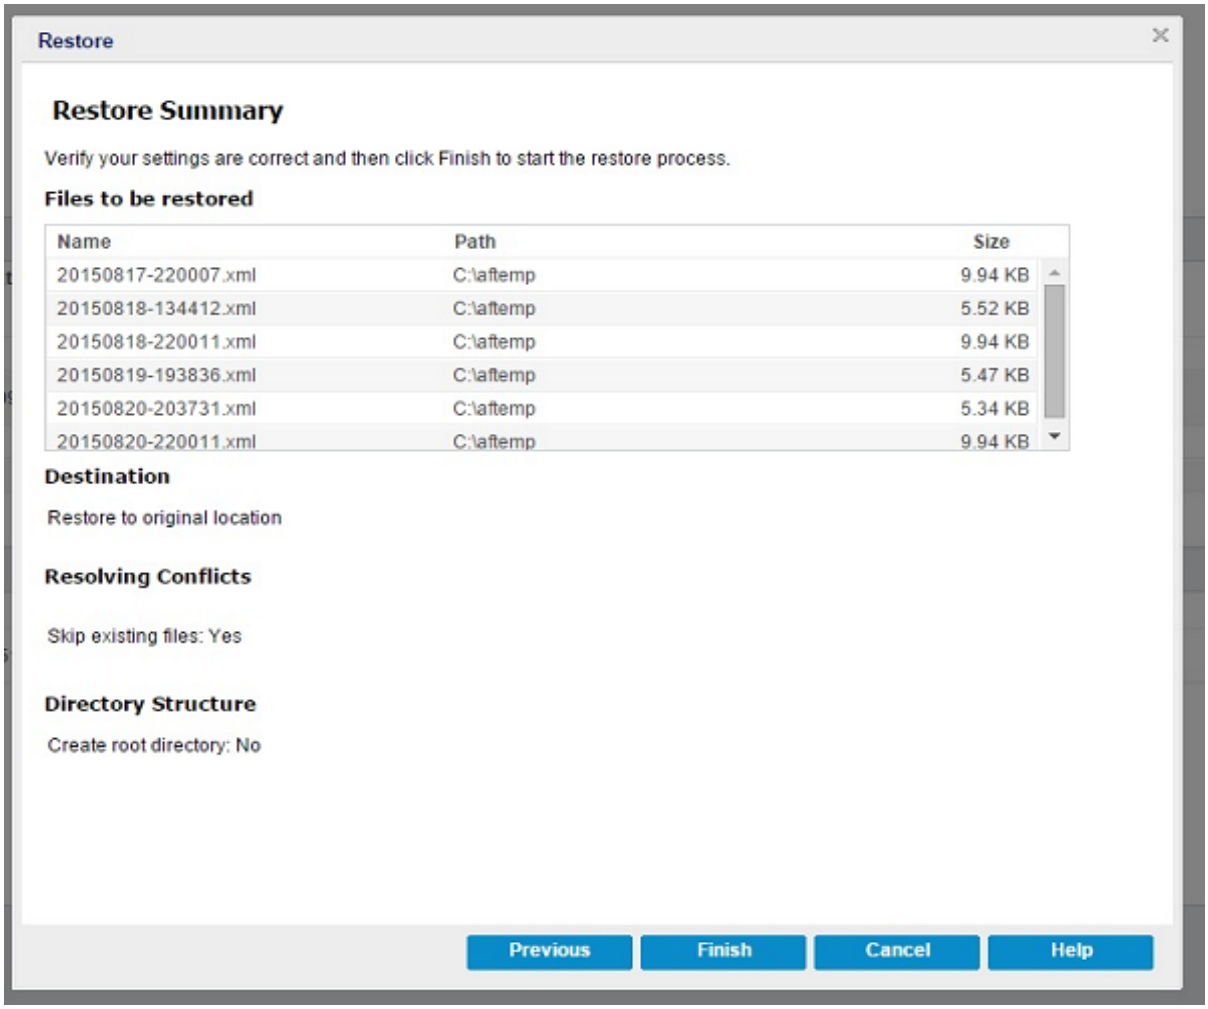
   The restore options are defined to restore from a file copy.

Restore the Recovery Point Content

After you define the restore options, verify that your settings are correct and confirm the restore process. **Restore Summary** helps you to review all the restore options that you defined and modify them if necessary.

Follow these steps:

1. On the **Restore Summary** dialog, review the displayed information to verify that all the restore options and settings are correct.



- If the summary information is incorrect, click **Previous** and go back to the applicable dialog to change the incorrect setting.
- If the summary information is correct, click **Finish** to launch the restore process.

The file copy content is restored.

Verify that Content is Restored

After the completion of the restore process, verify that content was restored to the specified destination.

Follow these steps:

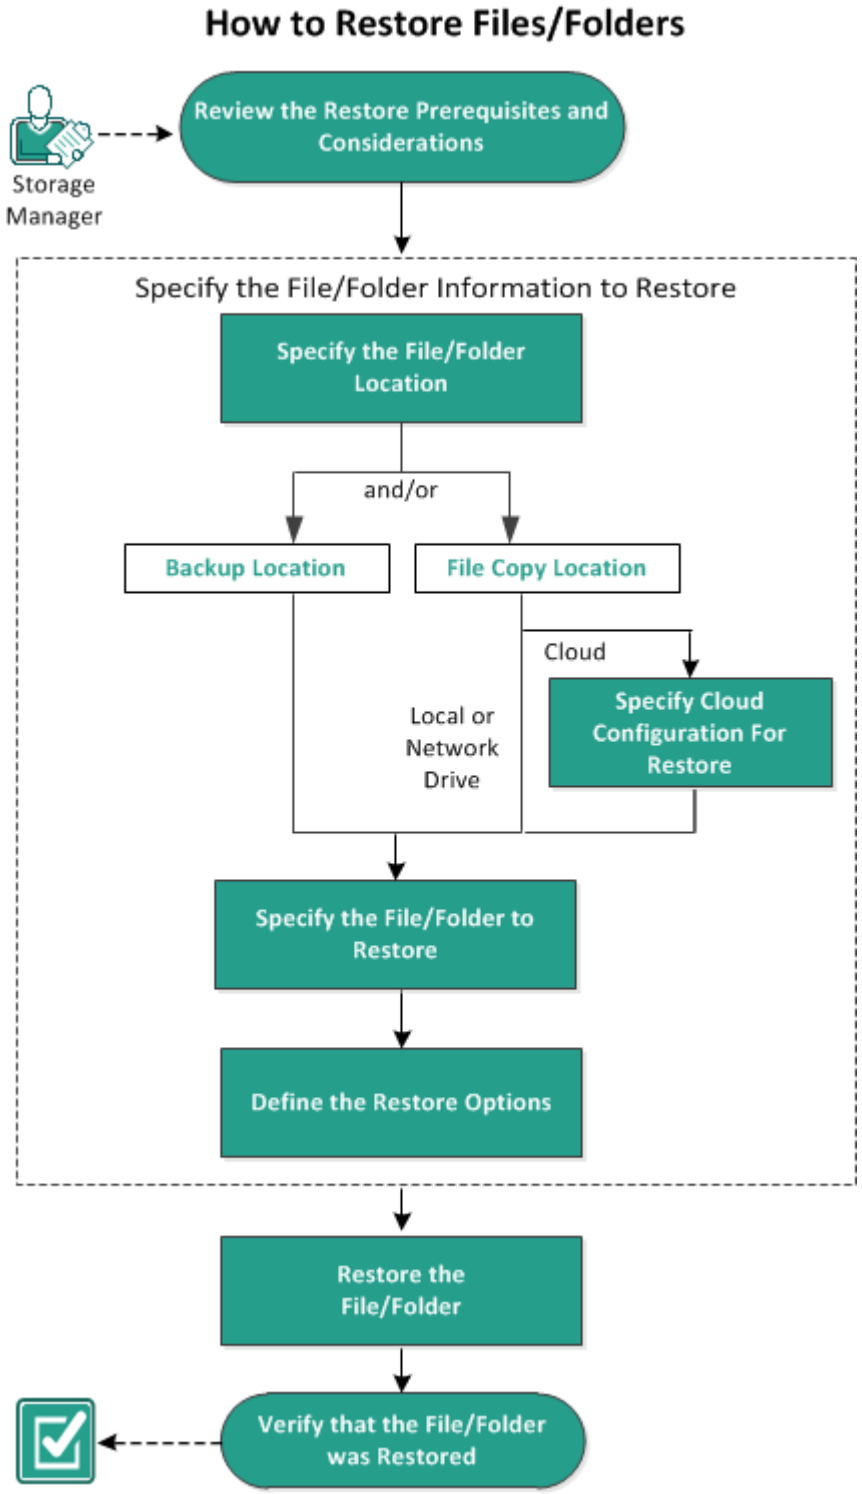
1. Log into the destination mailbox.
2. Check the mailbox item that you restored.
3. Verify the restored content.

   The restored content is successfully verified.

How to Restore Files/Folders

Each time Arcserve UDP performs a successful backup, all backed up files/folders are included in the snapshot image of your backup. This restore method allows you to specify exactly which file/folder you want to restore.

The following diagram illustrates the process to restore specific files/folders:



Perform the following tasks to restore files/folders:

1. Review the Restore Prerequisites and Considerations
2. Specify the File/Folder Information to Restore
    a. Specify the File/Folder Location
        - Specify Cloud Configuration for Restore
    b. Specify the File/Folder to Restore
    c. Define the Restore Options
3. Restore the File/Folder
4. Verify that the File/Folder was Restored

Review the Restore Prerequisites and Considerations

Verify that the following prerequisites exist before performing a restore:

- You have at least one backup or file copy version available to restore.
- You have a valid and accessible backup or file copy destination to restore the backup or file copy content from.
- You have a valid and accessible target location to restore the backup or file copy content to.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Review the following restore considerations:

- For a recovery point without a file system catalog created, to ensure you can browse and select files/folders to restore from the UI, the account/group should be granted access to all the folders/files on all volumes with read/list access before the backup is taken.

  The local system (SYSTEM) or built-in administrators group (BUILTIN\Administrators) needs to be added to the ACL of the folders for Arcserve UDP Agent (Windows) to be able to browse a backup without a file system catalog created. Otherwise, Arcserve UDP Agent (Windows) will not be able to browse the folders from the restore UI.

- (Optional) Understand how the restore process works. For more information, see How File Level Restores Work.

  Note: The process for restoring from a file copy location is similar to restoring from a backup location.

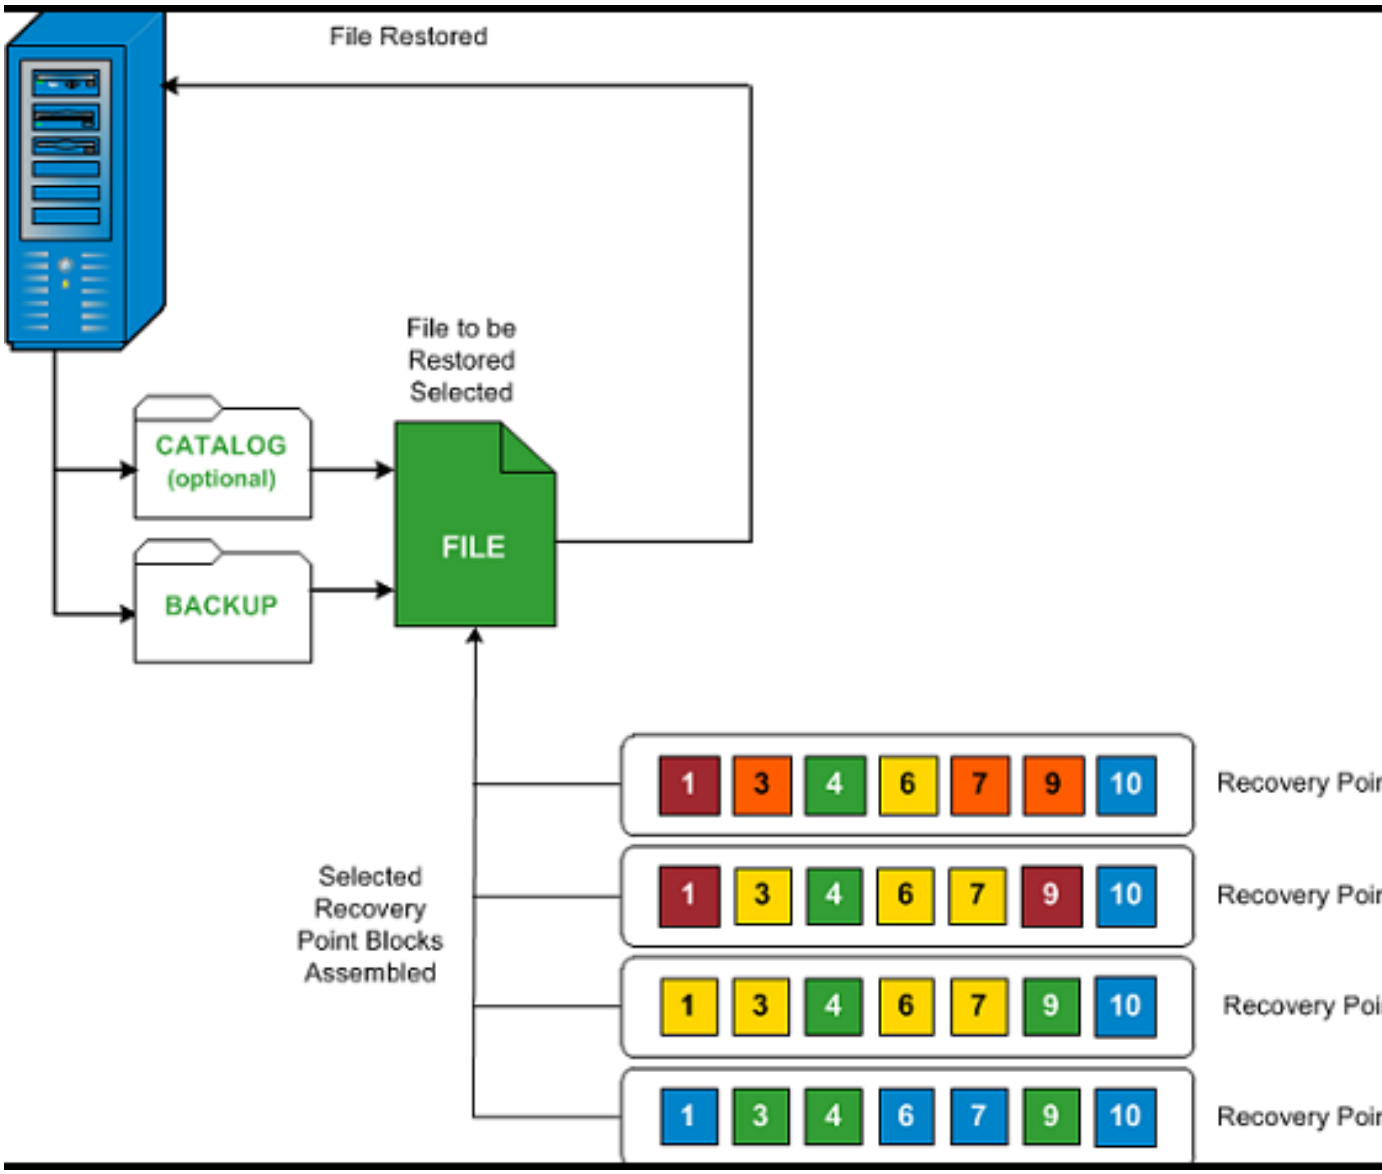
- (Optional) Review the files skipped during restore. For more information, see Files Skipped During Restore.

How File Level Restores Work

During a block-level backup, each backed up file is made up of a collection of blocks that define that particular file. When you need to restore a particular file, you can search your backup and select the file you want to restore and the recovery point you want to restore from. The Arcserve UDP Agent (Windows) then collects the version of the blocks that were used for the recovery point of the specified file, and reassembles and restores the file.

Note: When you specify your backup settings, you have an option to create a file catalog during backup. This file catalog lets you browse the backup sessions faster during restore. If you choose not to create the catalog during backup, it can still be created at a later time.

The following flow diagram shows the process of how Arcserve UDP restores a specific file.


File Restored
File to be Restored Selected
CATALOG (optional)
BACKUP
FILE
Selected Recovery Point Blocks Assembled
1 3 4 6 7 9 10
Recovery Poir
1 3 4 6 7 9 10
Recovery Poir
1 3 4 6 7 9 10
Recovery Poi
1 3 4 6 7 9 10
Recovery Poir

Files Skipped During Restore

While performing a restore by Arcserve UDP Agent (Windows) some files may be skipped intentionally.

The files and folders in the following table are skipped during a restore if the following two conditions exist:

- Files are skipped when such files exist before the restore and the conflict option is "skip existing files".
- Files and folders listed in the following table are skipped because they are an important component for Windows or Arcserve UDP Agent (Windows).

OS	Folder or Location	File or Folder Name	Remarks
All	Root folder of each volume	CAVolTrc.dat	Used by the Arcserve UDP tracking Driver.
		cavoltrcsnapshot.dat	
		System Volume Information*	Used to save files/folders by a Windows system, for example, volume shadow copy files.
		RECYCLER*	Used only on NTFS partitions. It contains a Recycle Bin for each user that logs on to the computer, sorted by their security identifier (SID).
		$Recycle.Bin*	When you delete a file in Windows NT Explorer or My Computer, the file is stored in the Recycle Bin until you empty the Recycle Bin or restore the file.
	Any folder contain picture files	Thumbs.db	Stores thumbnail images for Windows Explorer thumbnail view.
	Root folder of volume	PageFile.Sys	Windows virtual memory swap file.
		Hiberfil.sys	Hibernate file, used to save the system data when a computer goes into hibernate mode.

The following files and folders are skipped only when you restore to the original or alternate location:

OS	Folder or Location	File or Folder Name	Remark
All	Folder specified in value record under: HKLM\Software\Microsoft\Windows NT\CurrentVersion\WinLogon\SfcDllCache	All files/folders (recursively)	Folder contains a cached dll file which is used for System File Checker

			(SFC) and contents of the system dll cache directory are rebuilt by using SFC.
	%SystemRoot%\SYSTEM32\dllCache		
	Root folder of quorum_device	MSCS*	Used for Microsoft Cluster Server.
	%SystemRoot%\SYSTEM32\	perf?00?.dat perf?00?.bak	Per-formance data used by the Win-dows per-formance counter.
		CATROOT*	Used for Windows File Pro-tection (WFP) records digital sig-natures of the oper-ating sys-tem installs (such as DLL, EXE, SYS, OCX, and so on) to protect them from deletion or from replace-ment by older ver-sions.
	%SystemRoot%\inetsrv\	metabase.bin	Metabase binary file of earlier IIS versions before 6.0.

	File or folder specified in value except "SIS Common Store" under HKLM\SYSTEM\CurrentControlSet\Control\BackupRestore\FilesNotToBackup	All files/folders (recursively)	Files and folders should not be backed up and restored. For more information, see link.
XP W2003	System volume	NTLDR	The main boot loader.
		BOOT.INI	Contains boot configuration (if missing, NTLDR will default to \Windows on the first partition of the first hard drive).
		NTDETECT.COM	Required for booting an NT-based OS. Detects basic hardware information needed for a successful boot.
Vista and later	Root folder of system volume	boot*	Boot folder for Windows.
		bootmgr	Windows boot manager file.

		EFI\Microsoft\Boot*	Used for EFI boot.
	%SystemRoot%\SYSTEM32\	LogFiles\WMI\RTBackup*	Stores ETW trace files (extension .etl) for real time event trace sessions.
		config\RegBack*	Backup of current registry table.
Win-8 and later	System volume	swapfile.sys	System controller file, normally around 256 MB. It is used by Metro style applications that do not fit the traditional paging characteristics (such as usage pattern, growth, space reservation) of pagefile.sys.
		BOOTNXT	Used to boot from OS, other than Windows 8. Created when enabling the startup

			options, and updated by Windows.

The Activity log provides the following information:

- Date Time Information: jobxxxx System Files skipped. You can use Bare-Metal Recovery Option (BMR) to restore them.
- Date Time Information: jobxxxx Files or Directories skipped. Skipped files or directories are available at: C:\Program Files\Arcserve\Unified Data Protection\Engine\Logs\Restore-<YYYYMMDD>-<hhmmss>-<Process ID>-<Job ID>.log.

Specify the File/Folder Information to Restore

Arcserve UDP provides you with an option to find and restore a specific file or folder. The goal of running a successful restore job is to identify quickly the data you need and to retrieve it from the appropriate backup location. Each restore job requires a source and destination.

The process involved in restoring by find files/folders is as follows:

1. Specify the File/Folder Location
   - Specify Cloud Configuration for Restore
2. Specify the File/Folder to Restore
3. Define the Restore Options

Specify the File/Folder Location

Use the **Find Files/Folders** option to restore files and folders. This restore method allows you to specify exactly which file or folder you want to restore.

Follow these steps:

1. Access the restore method selection dialog in one of the following ways:

   From Arcserve UDP:

   a. Log into Arcserve UDP.

   b. Click the **resources** tab.

   c. Select **All Nodes** in the left pane.

      All the added nodes are displayed in the center pane.

   d. In the center pane, select the node and click **Actions**.

   e. Click **Restore** from the **Actions** drop-down options.

      The restore method selection dialog opens.

      Note: You are automatically logged into the agent node and the restore method selection dialog is opened from the agent node.

   From Arcserve UDP Agent (Windows):

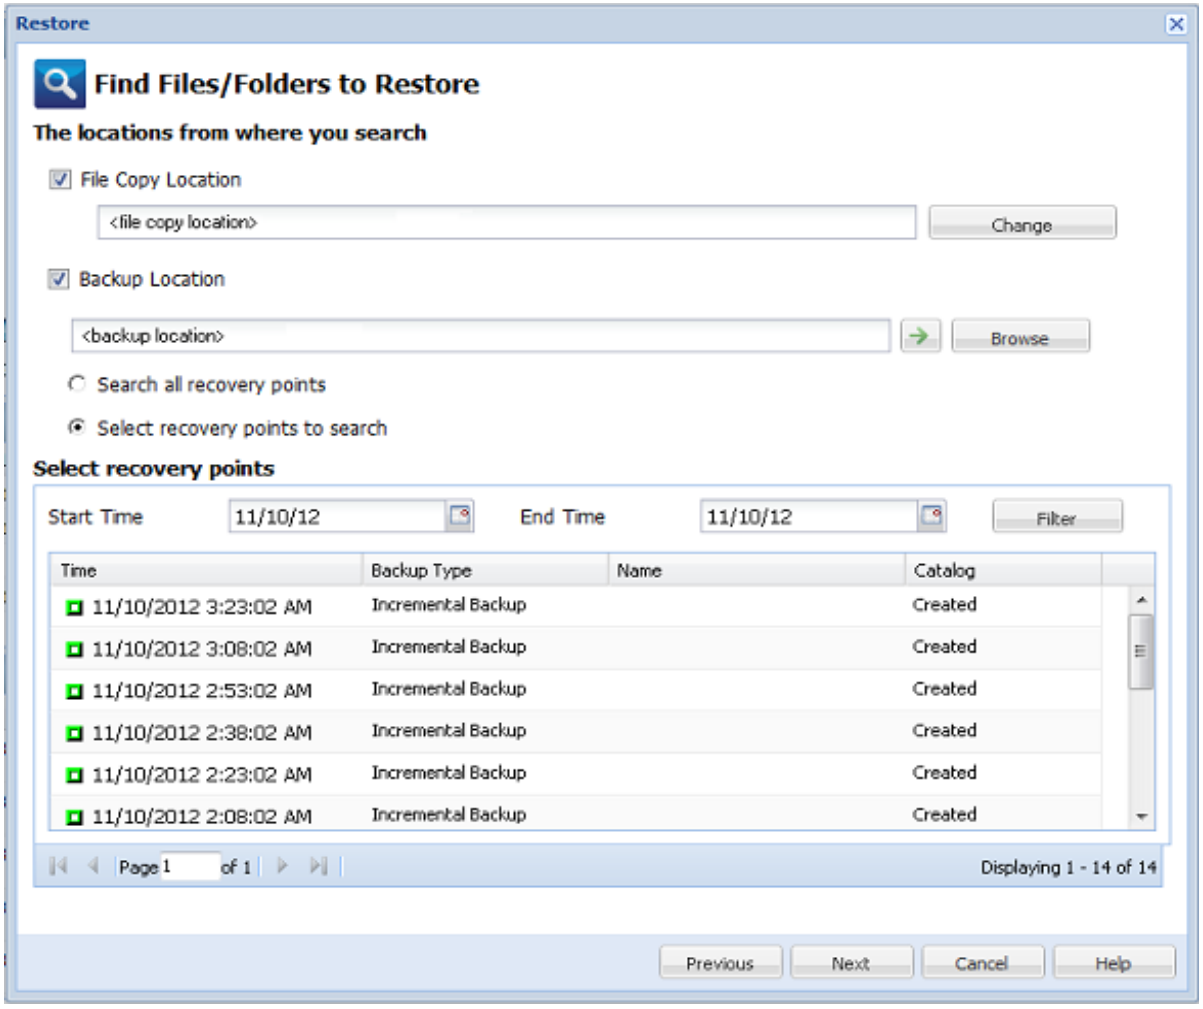
   a. Log into Arcserve UDP Agent (Windows).

   b. From the home page, select **Restore**.

   The restore method selection dialog opens.

2. Click the **Find Files/Folders to Restore** option.

   The **Find Files/Folders to Restore** dialog opens.



3. Select **File Copy Location** check box and click **Change** to modify the location to the destination where your file copy images are stored.

   The **Destination** dialog opens and you can select **Local or network drive** or **Cloud**.

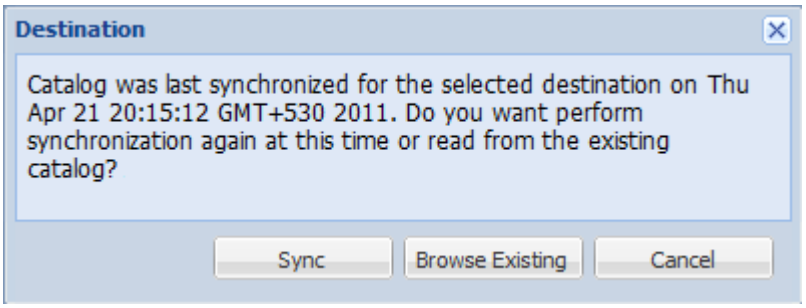
   Note: By default, the **Backup Location** and **File Copy Location** fields display the corresponding path used for the most recent backup/file copy destinations.



   - If you select **Local or network drive**, either specify a location or browse to the location where your file copy images are stored.

- You can click green arrow validation icon to verify proper access to the source location.
- If you select **Cloud**, either specify a cloud location or click the **Configure** button to display the **Cloud Configuration** dialog. For more information, see Specify Cloud Configuration for Restore.

  Regardless of whether you selected to restore from **Local or network drive** or from **Cloud**, when you change the destination to an alternate location a pop-up dialog will appear, asking if you want to perform a new catalog synchronization or read from the existing catalog.



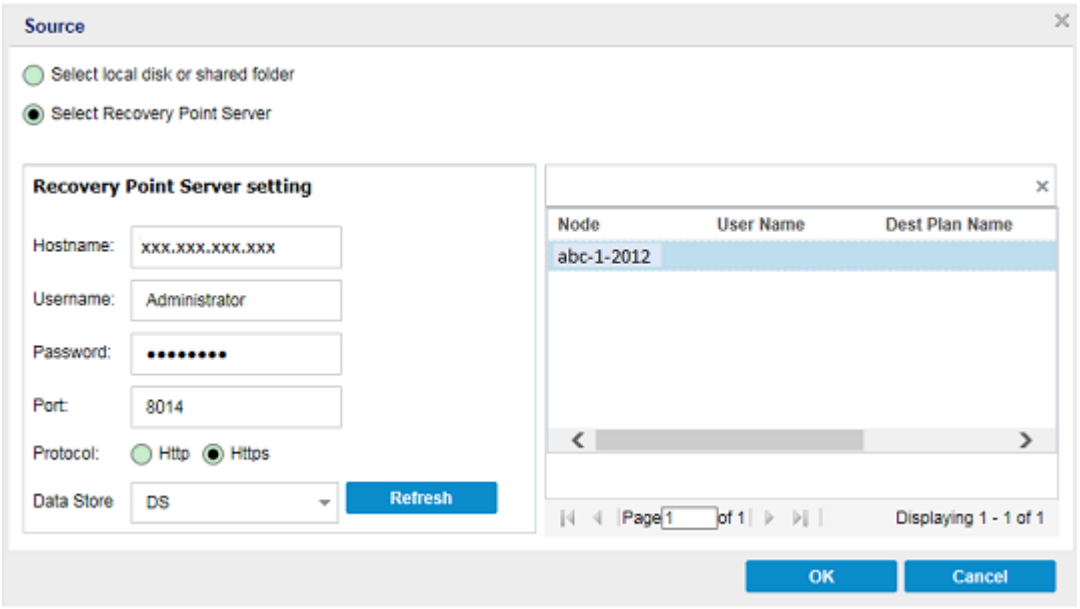
  - If you are performing a catalog synchronization the first time, the **Browse Existing** button is disabled because the file copy catalog does not exist locally.
  - If a catalog synchronization has been previously performed, this dialog will display details about the last time the catalog was synchronized from this destination. If more file copy jobs run since that displayed time, your catalog may not be currently synchronized and you can select the **Sync** option to ensure your file copy catalog is up-to-date.

    1. Click **Sync** to download the file copy catalog from the specified file copy destination to your local machine to provide faster browsing.

    2. Click **Browse Existing** to use the file copy catalog that is available locally and do not download/sync again.

4. Select the **Backup Location** check box and click **Change** to modify the Backup Location.

   The **Source** dialog opens where you can select the backup location.



5. Select one of the following options on the **Source** dialog:

   Select local disk or shared folder

   a. Specify or browse to the location where your backup images are stored and select the appropriate backup source.

      You can click the green arrow button to verify the connection to the specified location. If necessary, enter the Username and Password credentials to gain access to that source location.

      The **Select backup location** dialog opens.

   b. Select the folder where the recovery points are stored and click **OK**.

      The **Select backup location** dialog closes and you can see the backup location in the **Source** dialog.

   c. Click **OK**.

      The recovery points are listed in the **Find Files/Folders to Restore** dialog.

   Select Recovery Point Server

   a. Specify the **Recovery Point Server setting** details and click **Refresh**.

      All the agents are listed in the **Data Protection Agent** column in the **Source** dialog.

   b. Select the agent from the displayed list and click **OK**.

      The recovery points are listed in the **Find Files/Folders to Restore** dialog.

      Note: If you select a different agent and if the recovery points are encrypted, then you have to provide the encryption password when prompted.

6. Select one of the following options to search recovery points:

   Search all recovery points

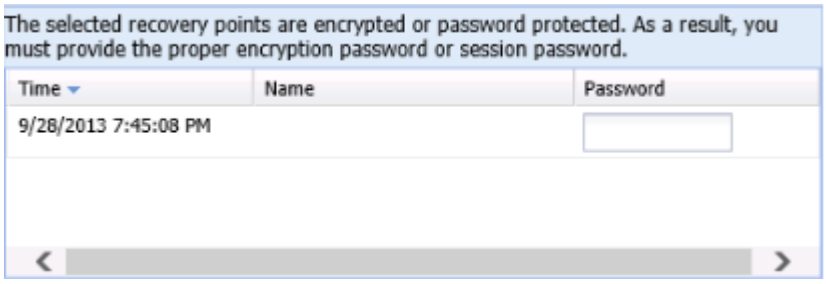
   Searches the file or folder in all the recovery points stored in the provided location. You have to specify the file or folder that you want to search on the **Find Files/Folders to Restore** dialog.

   Select recovery points to search

   Displays the recovery points between the specified time period. You can specify the start time and end time and then select the recovery point from the specified time period.

7. Select the recovery point and click **Next**.

   Note: If you have selected a different agent in the **Source** dialog and if the recovery points are encrypted, then the encryption dialog opens. Provide the password and click **OK**.



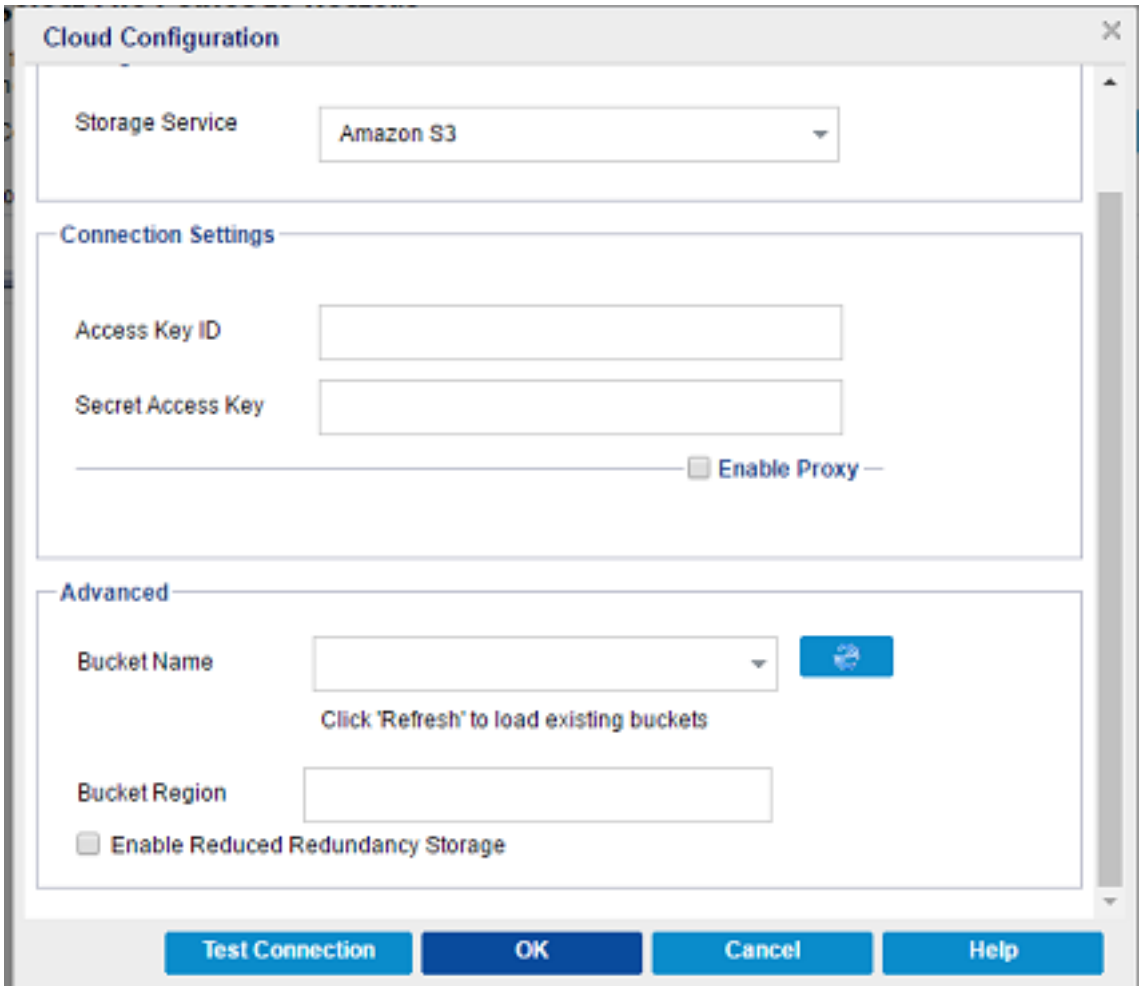
   The **Find Files/Folders to Restore** dialog opens.

   The **Backup or File Copy** location is specified.

Specify Cloud Configuration for Restore

Note: The following procedure only applies if you are restoring a file/folder from a file copy cloud location.

From the **Browse File Copies** option or the **Find Files/Folders to Restore** option, click the **Configure** button to display the **Cloud Configuration** dialog.



Follow these steps:

1. From the **Cloud Configuration** dialog, use the drop-down menu to select which cloud vendor type you want to restore from. The available options are **Amazon S3**, **Windows Azure**, **FUJITSU Hybrid IT Service FJcloud-O**, and **Eucalyptus-Walrus**. (**Amazon S3** is the default vendor). For more information about Fujitsu Cloud (Windows Azure) and FUJITSU Hybrid IT Service FJcloud-O, see the Overview and Registration.

   Note: After encoding the bucket name, if the path length is greater than 170 characters, Eucalyptus-Walrus will not be able to copy files.

2. Specify the **Configuration Options**.

The configuration options for each cloud vendor are similar (with some different terminology), and any differences are described.

a. Specify the **Connection Settings**:

Vendor URL

Identifies the URL address of the cloud provider.

For Amazon S3 and Windows Azure, the Vendor URL is automatically pre-populated. For Eucalyptus-Walrus, the Vendor URL must be manually entered using the specified format.

Access Key ID/Account Name/Query ID

Identifies the user who is requesting access this location.

For this field, Amazon S3 uses Access Key ID, Windows Azure and FUJITSU Hybrid IT Service FJcloud-O use Account Name, and Eucalyptus-Walrus uses Query ID.

Secret Access Key/Secret Key

Because your Access Key is not encrypted, this Secret Access Key is a password that is used to verify the authenticity of the request to access this location.

Important! This Secret Access Key is crucial for maintaining the security of your accounts. You should keep your keys and your account credentials in a secure location. Do not embed your Secret Access Key in a web page or other publicly accessible source code and do not transmit it over insecure channels.

For this field, Amazon S3 uses Secret Access Key. Windows Azure, FUJITSU Hybrid IT Service FJcloud-O, and Eucalyptus-Walrus use Secret Key.

Enable Proxy

If you select this option, you must also include the IP address (or machine name) of the proxy server and the corresponding port number that is used by the proxy server for internet connections. You can also select this option if your proxy server requires authentication. You then must provide the corresponding authentication information (Username and Password) that is required to use the proxy server.

(Proxy capability is not available for Eucalyptus-Walrus).

b. Specify the **Advanced Settings**:

Bucket Name / Container

All files and folders moved or copied to the cloud vendor are stored and organized in your buckets (or containers). Buckets are like a container for your files and are used to group and organize objects together. Every object stored at the cloud vendor is placed in a bucket.

Select a bucket name from the drop-down list. If necessary, you can click the **Refresh** button to update the list of available buckets.

For this field, Amazon S3 and Eucalyptus-Walrus use Bucket Name. Windows Azure and FUJITSU Hybrid IT Service FJcloud-O use Container.

Bucket Region

Refers to the region of bucket in Amazon and FUJITSU Hybrid IT Service FJcloud-O.

Contract Number

Refers to the number of contract that FUJITSU Hybrid IT Service FJcloud-O provides.

Project ID

Refers to the ID of project that FUJITSU Hybrid IT Service FJcloud-O generates.

Enable Reduced Redundancy Storage

For Amazon S3 only, this option lets you select to enable Reduced Redundancy Storage (RRS). RRS is a storage option within Amazon S3 that helps you reduce cost by storing non-critical, reproducible data at lower levels of redundancy than Amazon S3s standard storage. Both the standard and reduced redundancy storage options store data in multiple facilities and on multiple devices, but with RRS the data is replicated fewer times, so the cost is less. You should expect the same latency and throughput using either the Amazon S3 standard storage or RRS. By default this option is not selected (Amazon S3 uses the standard storage option).

3. Click **Test Connection** to verify the connection to the specified cloud location.
4. Click **OK** to exit the **Cloud Configuration** dialog.

Specify the File/Folder to Restore

After you specify the backup or file copy location, search for the file or folder name to restore. If a file has multiple file copy versions, all versions are listed and sorted by date (with the most recent listed first).

Follow these steps:

1. From the **Find Files/Folders to Restore** dialog, specify what to search for (file or folder name to restore).

   Note: The **File Name** field supports full name searching and wildcard searching. If you do not know the complete file name, you can simplify the results of the search by specifying the wildcard characters "*" and "?" in the File Name field.

   The wildcard characters supported for the file or folder name are as follows:

   - Use the asterisk to substitute zero or more characters in a file or folder name.
   - Use the question mark to substitute a single character in a file or folder name.

   For example, if you specify *.txt, all files with a .txt file extension appear in the search results.

2. (Optional) Specify a path to further filter your search and select whether to include or not include any subdirectories.
3. Click **Find** to launch search results.

   The search results are displayed. If the searched file has multiple file copy versions, all versions will be listed, sorted by date (with the most recent listed first). It also indicates if the searched file was backed up or file copied.

4. Select the version (occurrence) of the file/folder that you want to restore and click **Next**.

   The **Restore Options** dialog opens.

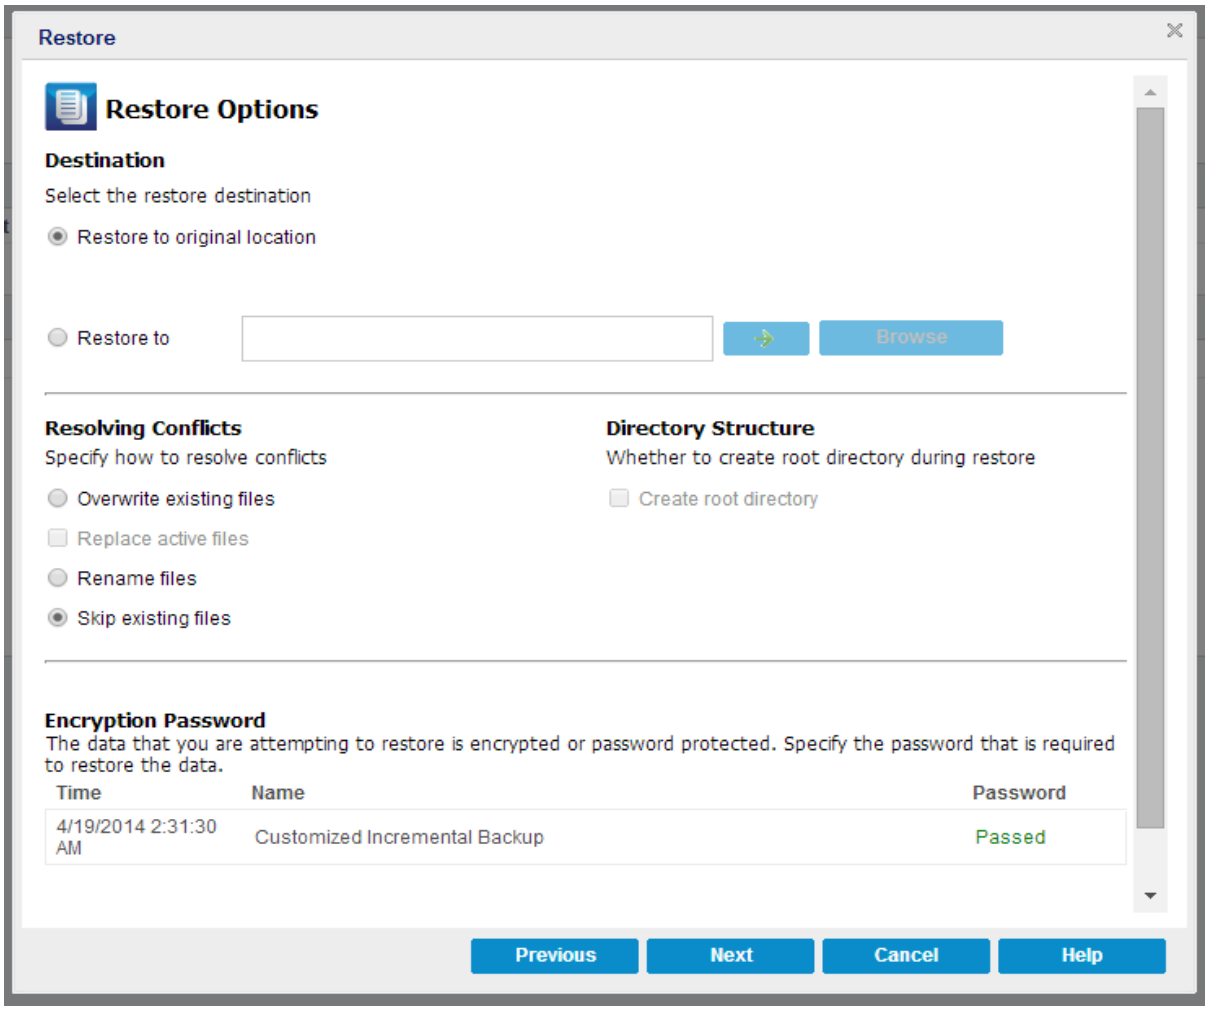
   The file/folder name to be restored is specified.

Define the Restore Options

After you specify the file or folder to restore, define the restore options for the selected file or folder.

Follow these steps:

1. From the **Restore Options** dialog, select the restore destination.



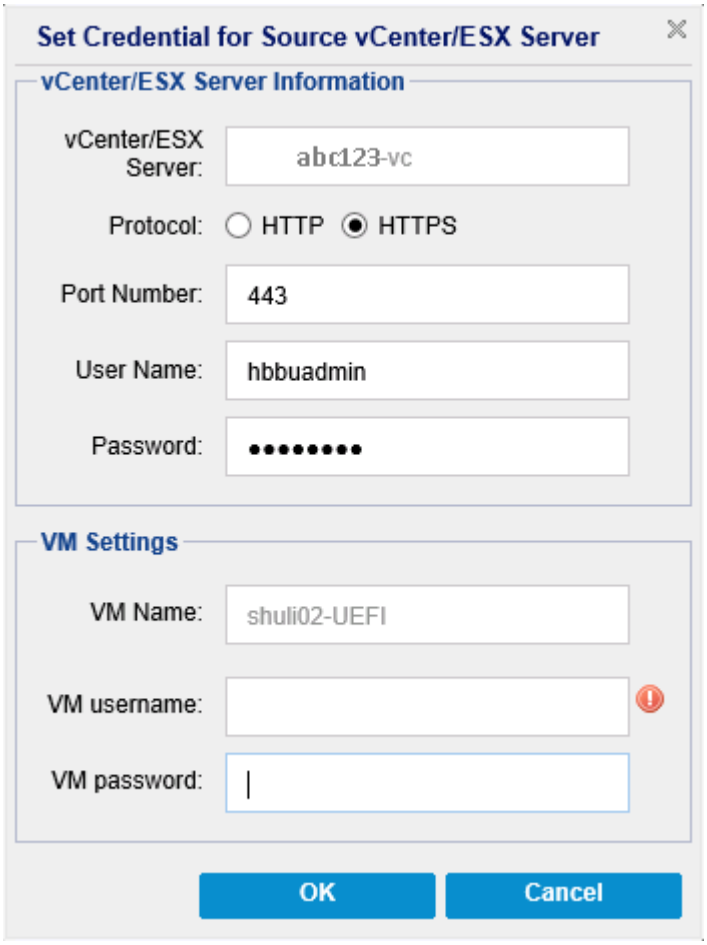
The available destination options are:

Restore to Original Location

Restores to the original location from where the backup image was captured.

Note: If you performed the recovery point backup using host-based agentless backup, restoring to original location is to restore the file back in to the virtual machine. In this case, a dialog box opens. You may enter the credentials of the hypervisor, and the operating system of the virtual machine.

For VMware VM:



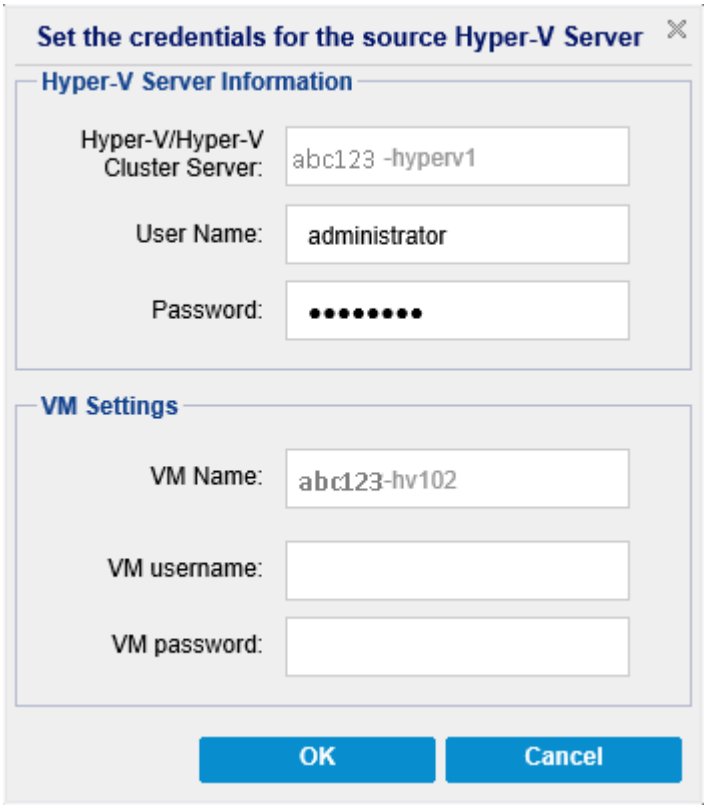
Note: To be able to create or write files inside the VM, consider the following requirements for the settings and account permission of virtual machine:

- VMware Tools is installed and running.
- Firewall must allow File and Printer Sharing.
- The account is built-in local administrator, built-in domain administrator, or domain account that is member of the local Administrators group. If other accounts are used:
  - Disable the UAC remote access. To disable UAC remote access, see Import Virtual Machine Using Additional Administrative Account.

- Disable UAC in the Local Security Policy by disabling the setting Run all administrator in Admin Approval Mode at secpol.msc -> Local Policies -> Security Options. (Secpol.msc is Microsoft's security policy editor).

  Important: Do not attempt to disable the UAC in the User Account Control Settings dialog box that opens from the control panel.

For Hyper-V VM:



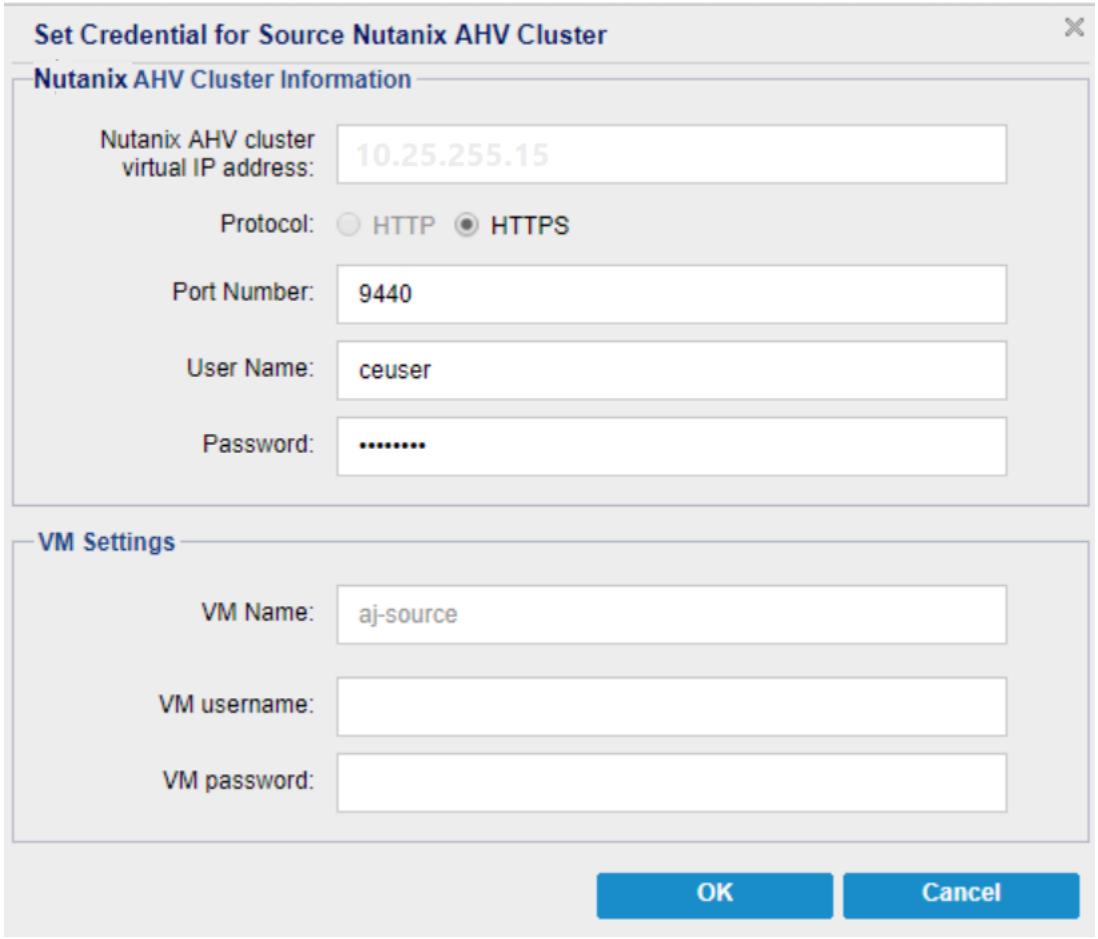
Note: To be able to create or write files inside the VM, consider the following requirements for the settings and account permission of virtual machine:

- Hyper-V integration services are installed and running.
- Firewall must allow File and Printer Sharing.
- The account is built-in local administrator, built-in domain administrator, or domain account that is member of the local Administrators group. If other accounts are used:

  Disable the UAC remote access. To disable UAC remote access, see Import Virtual Machine Using Additional Administrative Account.

- If virtual machine guest OS is Client version Windows (such as Windows 10), you need to manually configure firewall to allow Windows Management Instrumentation (WMI).

For Nutanix VM



Notes: To be able to create or write files inside the VM, consider the following requirements for the settings and account permission of virtual machine:

- Firewall must allow File and Printer Sharing.
- The account is built-in local administrator, built-in domain administrator, or domain account that is member of the local Administrators group. If other accounts are used:

  Disable the UAC remote access. To disable UAC remote access, see Import Virtual Machine Using Additional Administrative Account.
- Firewall must allow Windows Management Instrumentation(WMI).

Restore to

Restores to the specified location. You can click the green arrow button to verify the connection to the specified location. If necessary, enter the User-name and Password credentials to gain access to that location.

2. Specify the **Resolving Conflicts** option that Arcserve UDP performs if conflicts are encountered during the restore process.

   The available options are:

   Overwrite existing files

   Overwrites (replaces) any existing files, which are at the restore destination. All objects are restored from the backup files regardless of their current presence on your computer.

   Replace active files

   Replaces any active files upon reboot. If during the restore attempt Arcserve UDP Agent (Windows) discovers that the existing file is currently in use or being accessed, it will not immediately replace that file, but instead to avoid any problems will delay the replacement of the active files until the next time the machine is rebooted. (The restore occurs immediately, but the replacement of any active files is done during the next reboot).

   This option is only available if you select the **Overwrite existing files** option.

   Note: If you do not select this option, any active file is skipped from the restore.

   Rename files

   Creates a new file if the file name already exists. Selecting this option copies the source file to the destination with the same filename but a different extension. Data is then restored to the new file.

   Skip existing files

   Skips over and not overwrite (replace) any existing files that are located at the restore destination. Only objects that are not currently existing on your machine are restored from the backup files.

   Default: Skip existing files.

3. Specify the **Directory Structure** to create a root directory during restore.

   Create root directory

   Specifies that if a root directory structure exists in the captured backup image, Arcserve UDP recreates that same root directory structure on the restore destination path.

With this option not selected, the file or folder is restored directly to the destination folder.

For example, if during the backup you captured the files "C:\Folder-1\SubFolder2\A.txt" and "C:\Folder1\SubFolder2\B.txt" and during the restore you specified to the restore destination as "D:\Restore".

- If you select to restore the "A.txt" and "B.txt" files individually, the destination for the restored files will be "D:\Restore\A.txt" and "D:\Restore\B.txt" (the root directory above the specified file level will not be recreated).
- If you select to restore from the "SubFolder2" level, the destination for the restored files will be "D:\Restore\SubFolder2\A.txt" and "D:\Restore\SubFolder2\B.txt" (the root directory above the specified folder level will not be recreated).

With this option selected, the entire root directory path for the files/folders (including the volume name) is recreated to the destination folder. If the files/-folders to be restored are from the same volume name, then the destination root directory path does not include that volume name. However, if the files/-folders to be restored are from different volume names, then the destination root directory path does include the volume name.

For example, if during the backup you captured the files "C:\Folder-1\SubFolder2\A.txt", "C:\Folder1\SubFolder2\B.txt", and also E:\Folder-3\SubFolder4\C.txt" and during the restore you specified to the restore destination as "D:\Restore".

- If you select to restore just the "A.txt" file, the destination for the restored file will be "D:\Restore\ Folder1\SubFolder2\A.txt" (the entire root directory without the volume name will be recreated).
- If you select to restore both the "A.txt" and "C.txt" files, the destination for the restored files will be "D:\Restore\C\Folder1\SubFolder2\A.txt" and "D:\Restore\E\Folder3\SubFolder4\C.txt" (the entire root directory with the volume name will be recreated).

4. The **Encryption Password** for file copy destination is loaded automatically. If you select an alternate destination for the restore, you will need to enter the password manually.
5. Click **Next**.

   The **Restore Summary** dialog opens.

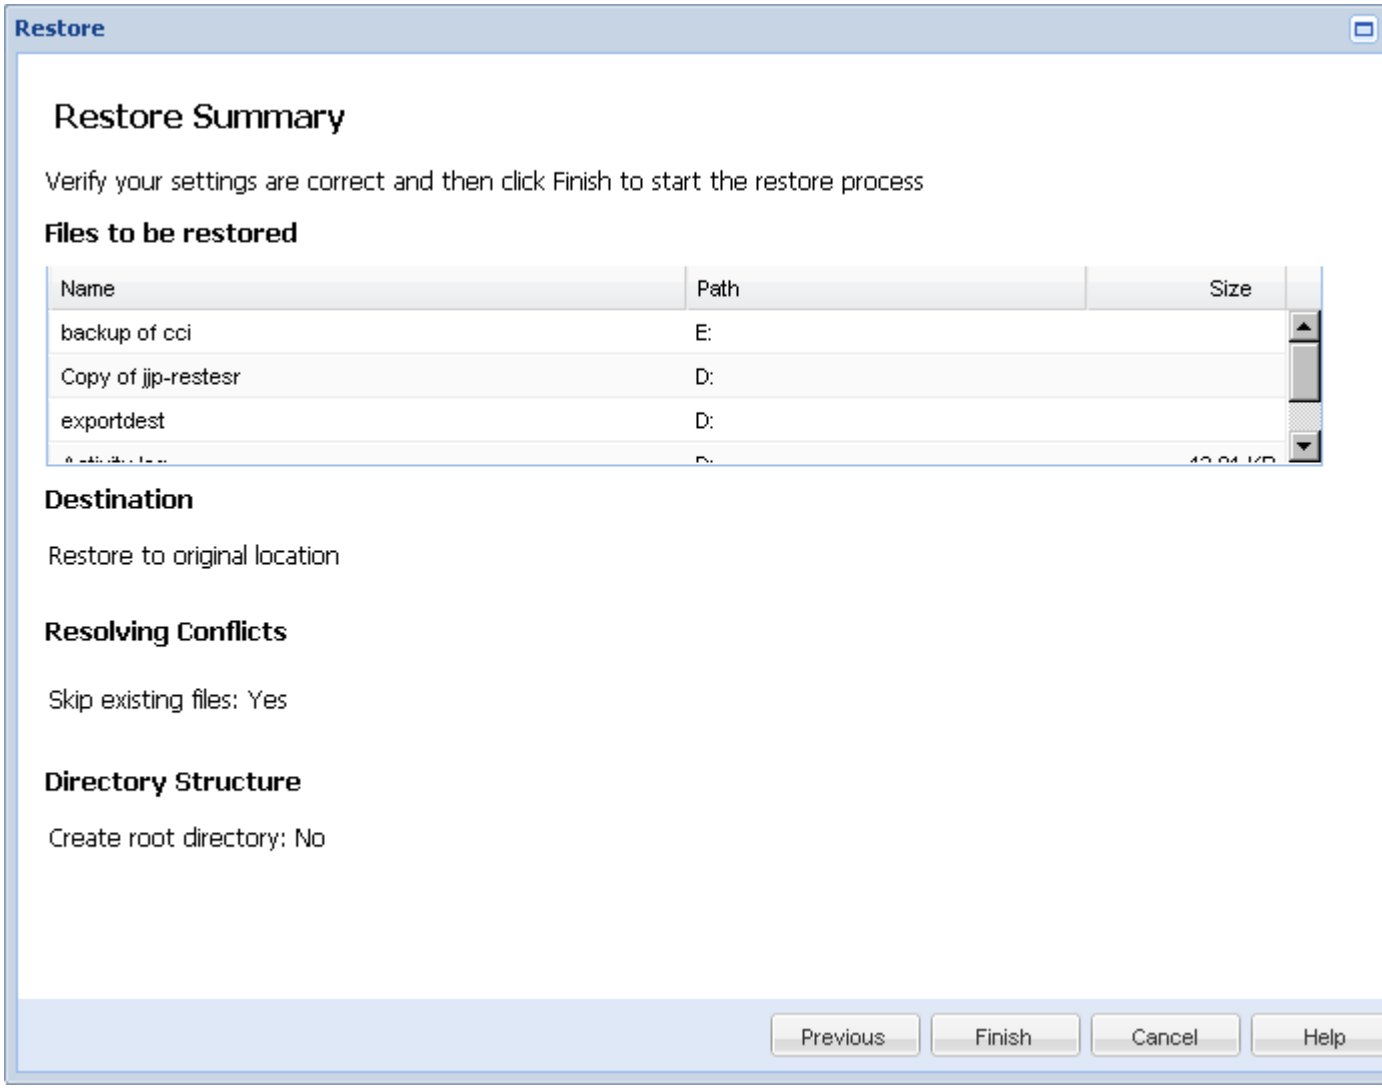
   The restore options are defined to restore the specified file/folder.

Restore the File/Folder

The **Restore Summary** dialog helps you to review all the restore options that you previously defined and lets you modify them if necessary.

Follow these steps:

1. From the **Restore Summary** dialog, review the displayed information to verify that all the restore options and settings are correct.



- If the summary information is incorrect, click **Previous** and go back to the applicable dialog to change the incorrect setting.
- If the summary information is correct, click **Finish** to launch the restore process.

The specified file/folder is restored.

Verify that the File/Folder was Restored

After the completion of the restore process, verify that the file/folder was restored to the specified destination.

Follow these steps:

1. Navigate to the restore destination you specified.

   A list of folders appears.

2. Locate the file to which you have restored the content.

   For example, if you select to restore the "A.txt'' file to the restore destination as "D:\Restore, then navigate to the following location:

   D:\Restore\A.txt.

3. Verify the content of the restored file/folder.

   The restored content is successfully verified.

How to Perform a Bare Metal Recovery Using a Virtual Standby VM or Instant VM

Bare Metal Recovery (BMR) is the process of restoring a computer system from "bare metal" including reinstalling the operating system and software applications, and then restoring the data and settings. The BMR process lets you restore a full computer with minimal effort, even to different hardware. BMR is possible because during the block-level backup process, Arcserve UDP Agent (Windows) not only captures the data, but also all information that is related to the following applications:

- Operating system
- Installed applications
- Configuration settings
- Necessary drivers

  All relevant information that is necessary to perform a complete rebuild of the computer system from "bare metal" is backed up into a series of blocks and stored on the backup location.

To perform a BMR from a virtual machine, use one of the following ways:

- Connect to the ESX server directly using the IP address
- Add the correct DNS setting in your BMR machine and resolve the hostname to the IP address

Complete the following tasks to perform a BMR using a Virtual Standby VM or Instant VM:

1. Review the BMR Prerequisites and Considerations
2. Define BMR Options
   - Recover Using a Hyper-V Virtual Standby VM or Instant VM
   - Recover Using a VMware Virtual Standby VM or Instant VM
   - Perform BMR in Express Mode
   - Perform BMR in Advanced Mode
3. Verify that the BMR was Successful
4. BMR Reference Information
5. Troubleshooting BMR Issues

Review the BMR Prerequisites and Considerations

Verify that the following prerequisites exist before performing a BMR:

- You must have one of the following images:
  - A created BMR ISO image burned onto a CD/DVD
  - A created BMR ISO image burned onto a portable USB stick

  Note: Arcserve UDP Agent (Windows) utilizes a Boot Kit Utility to combine a WinPE image and Arcserve UDP Agent (Windows) image to create a BMR ISO image. This ISO image is then burned onto a bootable media. You can then use either of these bootable media (CD/DVD or USB stick) to initialize the new computer system and allow the bare metal recovery process to begin. To ensure your saved image is always the most up-to-date version, create a new ISO image every time you update Arcserve UDP Agent (Windows).
- At least one full backup available.
- At least 1-GB RAM installed on the virtual machine and the source server that you are recovering.
- To recover VMware virtual machines to VMware virtual machines that are configured to behave as physical servers, verify the VMware Tools application is installed on the destination virtual machine.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Review the following restore considerations:

- Regardless of which method you used to create the Boot Kit image, the BMR process is basically the same.

  Note: The BMR process cannot create storage spaces. If the source machine had storage spaces, during BMR you cannot create storage spaces at the destination machine. You can either restore those volumes to regular disks/volumes or manually create storage spaces before performing the BMR, and then restore the data into those created storage spaces.
- Dynamic disks are restored at the disk level only. If your data is backed up to a local volume on a dynamic disk, you cannot to restore this dynamic disk during BMR. In this scenario, to restore during BMR you must perform one of the following tasks and then perform BMR from the copied Recovery Point:
  - Back up to a volume on another drive.
  - Back up to a remote share.
  - Copy a recovery point to another location.

Note: If you perform BMR with multiple dynamic disks, the BMR may fail because of some unexpected errors (such as fail to boot, unrecognized dynamic volumes, and so on). If this occurs, you should restore only the system disk using BMR, and then after the machine reboot you can restore the other dynamic volumes on a normal environment.

- If you attempt to perform a BMR on a Hyper-V VM with a 4 KB disk, add this 4 KB disk to the SCSI controller. If you add it to the IDE controller, the disk will not be detected in the Windows PE system.
- (Optional) Review the BMR Reference Information. For more information, see the following topics:
  - How Bare Metal Recovery Works
  - Operating Systems that Support UEFI/BIOS Conversion
  - Managing the BMR Operations Menu

Review the following considerations:

- If you upgrade to a newer version or update of Arcserve UDP, you must re-create the BMR ISO using the proper Windows AIK or ADK level to include support for latest features and bug fixes. However, once a BMR ISO is created, the ISO file can be used for the same OS level. The following OS levels can use the same ISO:
  - ISO created using Windows 7 WAIK – works for Windows 2003, Vista, 2008, 2008 R2
  - ISO create using Windows 8/8.1 ADK – works for Windows 8, 8.1, Server 2012, Server 2012 R2
  - ISO created using Windows 10 ADK – works for Windows 10

Define BMR Options

Prior to initiating the BMR process, you must specify some preliminary BMR options.

Follow these steps:

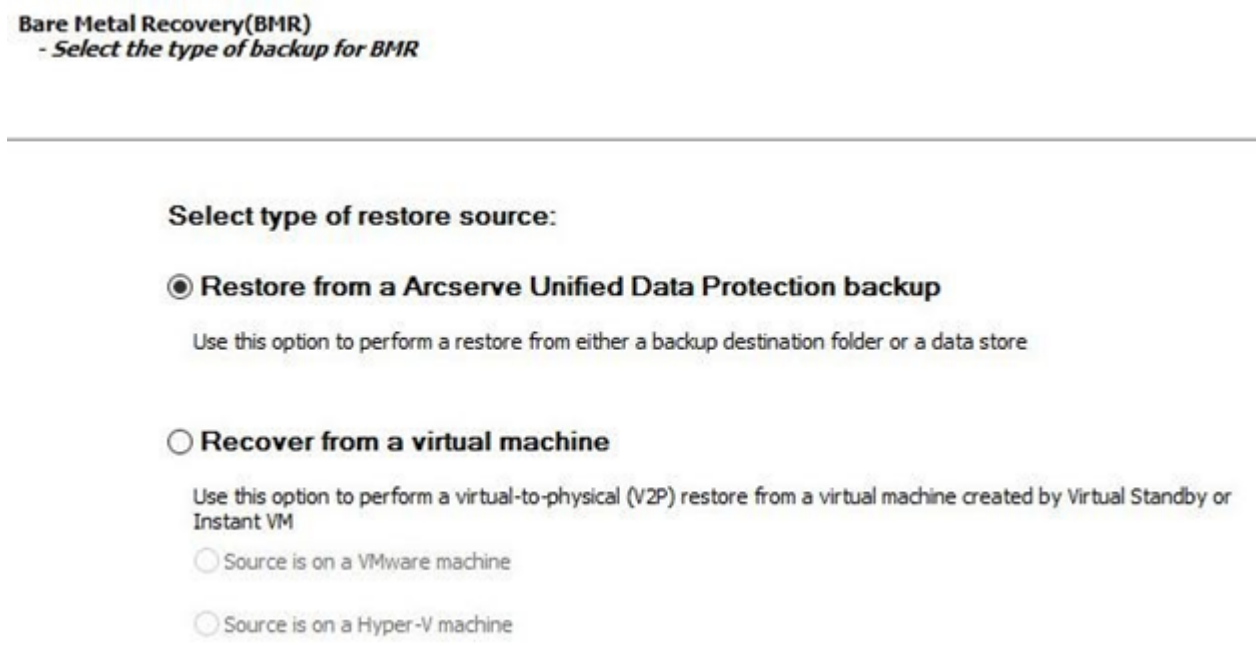
1. Insert the saved Boot Kit image media and boot the computer.
   - If you are using a BMR ISO image burned onto a CD/DVD, insert the saved CD/DVD.
   - If you are using a BMR ISO image burned onto a USB stick, insert the saved USB stick.

   The **BIOS Setup Utility** screen is displayed.
2. From the **BIOS Setup Utility** screen, select the CD-ROM Drive option or the USB option to launch the boot process. Select an architecture (x86/x64) and press **Enter** to continue.

   The Arcserve UDP Agent (Windows) language select screen is displayed.
3. Select a language and click **Next** to continue.



The Bare Metal Recovery process is initiated and the initial BMR wizard screen is displayed.



The BMR wizard screen allows you to select the type of BMR you want to perform:

- **Restore from an Arcserve Unified Data Protection backup**

  Use this option to perform a restore from either a backup destination folder or a data store.

  This option lets you recover data that was backed up using Arcserve UDP Agent (Windows). This option is used in connection with backup sessions performed with Arcserve UDP Agent (Windows) or with the Arcserve UDP host-based VM backup application.

  For more information, see How to Perform a Bare Metal Recovery Using a Backup in the online help.

- **Recover from a Virtual Standby VM**

  Use this option to perform a virtual-to-physical (V2P) restore from a virtual standby VM or Instant VM. Virtual-to-physical (V2P) is a term that refers to the migration of an operating system (OS), application programs and data from a virtual machine or disk partition to a computer's main hard disk. The target can be a single computer or multiple computers.

  - **Source is on a VMware machine**

Lets you recover data for a machine for which virtual conversion is done to a VMware virtual machine. This option is used in connection with the Arcserve Central Virtual Standby or Instant VM application.

Note: For this option, you can only recover data if the virtual conversion to a VMDK file (for VMware) was performed using Arcserve Central Virtual Standby or Instant VM.

If you select this option, see Recover using a VMware Virtual Standby VM or Instant VM to continue this procedure.

- **Source is on a Hyper-V machine**

  Lets you recover data for a machine for which virtual conversion is performed to a Hyper-V virtual machine. This option is used in connection with the Arcserve Central Virtual Standby or Instant VM application.

  Note: For this option, you can only recover data if the virtual conversion to a VHD file (for Hyper-V) was performed using Arcserve Central Virtual Standby or Instant VM.

  If you select this option, see Recover using a Hyper-V Virtual Standby VM or Instant VM to continue this procedure.

4. Select **Recover from a Virtual Standby VM**. Then select one of the sources.

   If you select the **Source is on a VMware machine** option, see Recover using a VMware Virtual Standby VM or Instant VM to continue this procedure.

   If you select the **Source is on a Hyper-V machine** option, see Recover using a Hyper-V Virtual Standby VM or Instant VM to continue this procedure.

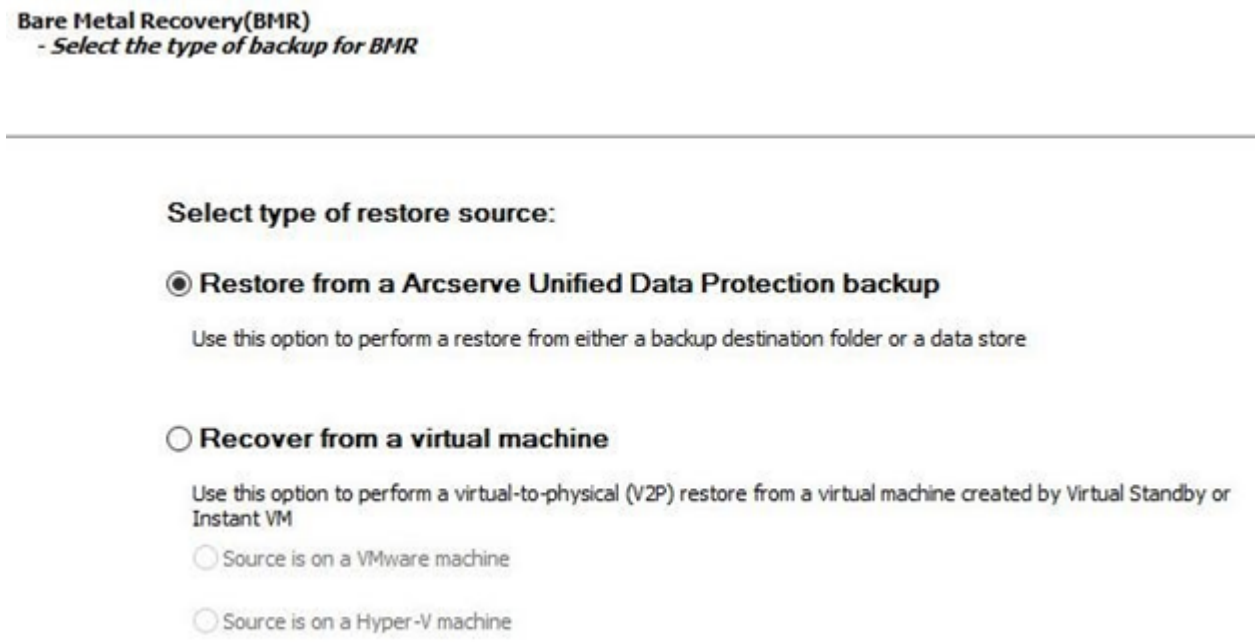
Recover using a Hyper-V Virtual Standby VM or Instant VM

Arcserve UDP Agent (Windows) provides the capability to perform Bare Metal Recovery for virtual-to-physical (V2P) machines. This feature lets you perform virtual-to-physical recovery from the latest state of a standby or instant virtual machine and helps you reduce the loss of your production machine.

Follow these steps:

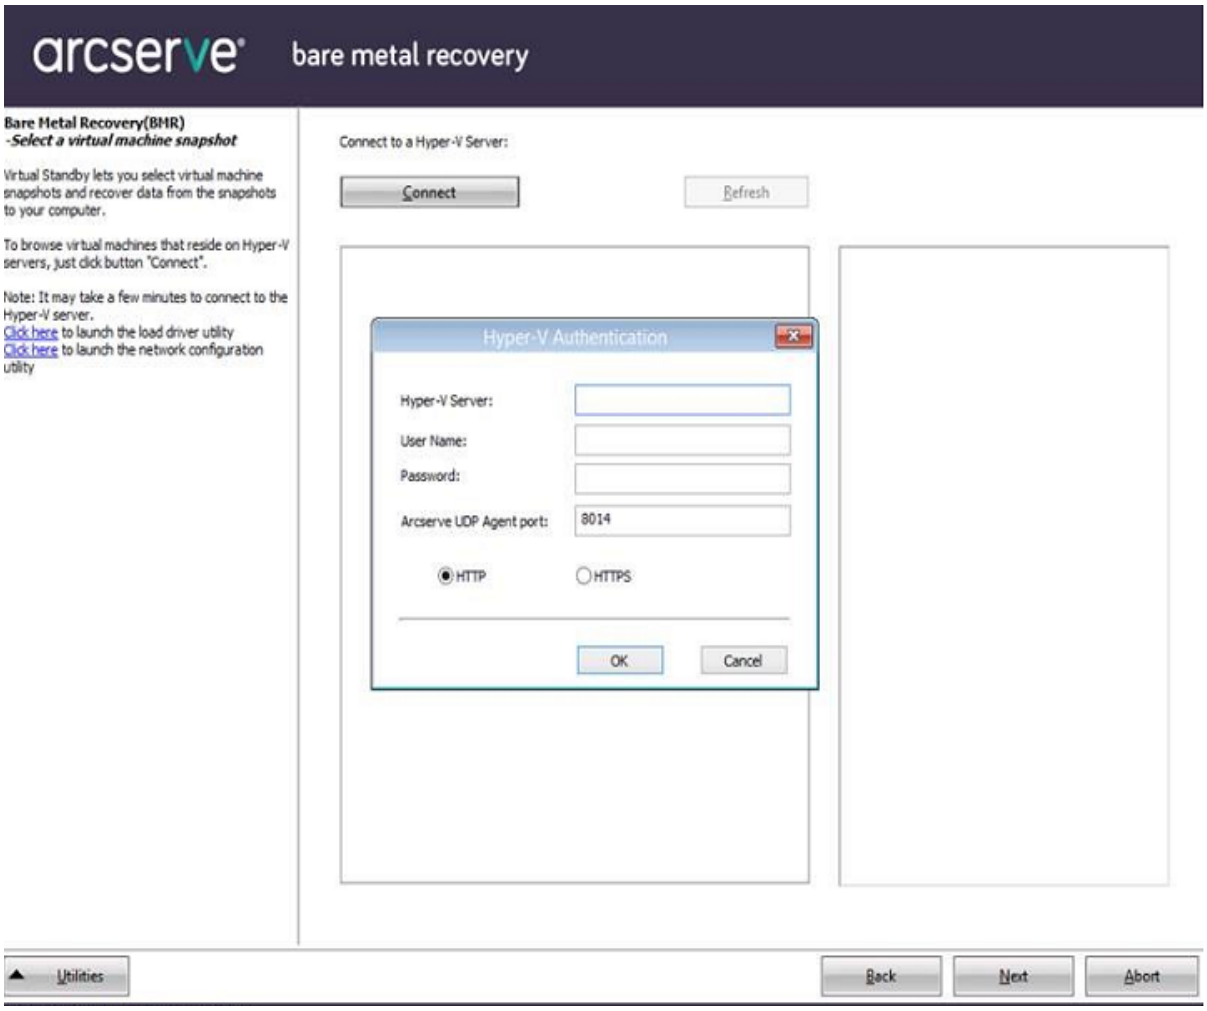
1. From the select the Type of Bare Metal Recovery (BMR) wizard screen, select the **Recover from a Virtual Standby VM** and select **Source is on a Hyper-V machine** option.

   Use this option to perform a virtual-to-physical restore from a virtual standby VM or Instant VM. The term virtual-to-physical refers to the migration of an operating system (OS), application programs and data from a virtual machine or disk partition to a computer's main hard disk. The target can be a single computer or multiple computers.



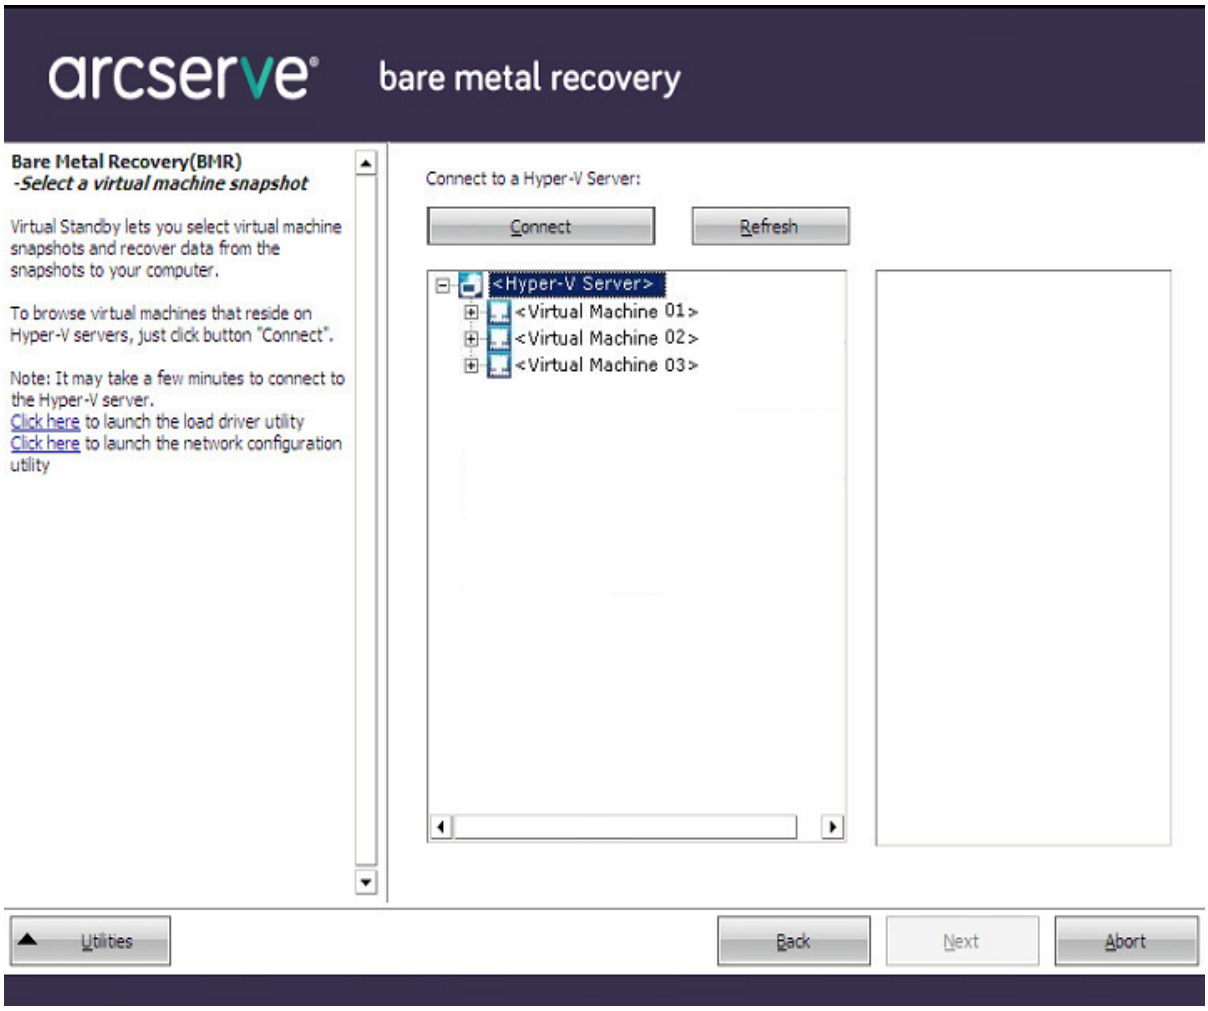
2. Click **Next**.

   The Select a virtual machine snapshot screen is displayed, with the Hyper-V Authentication dialog, prompting you for Hyper-V server details.



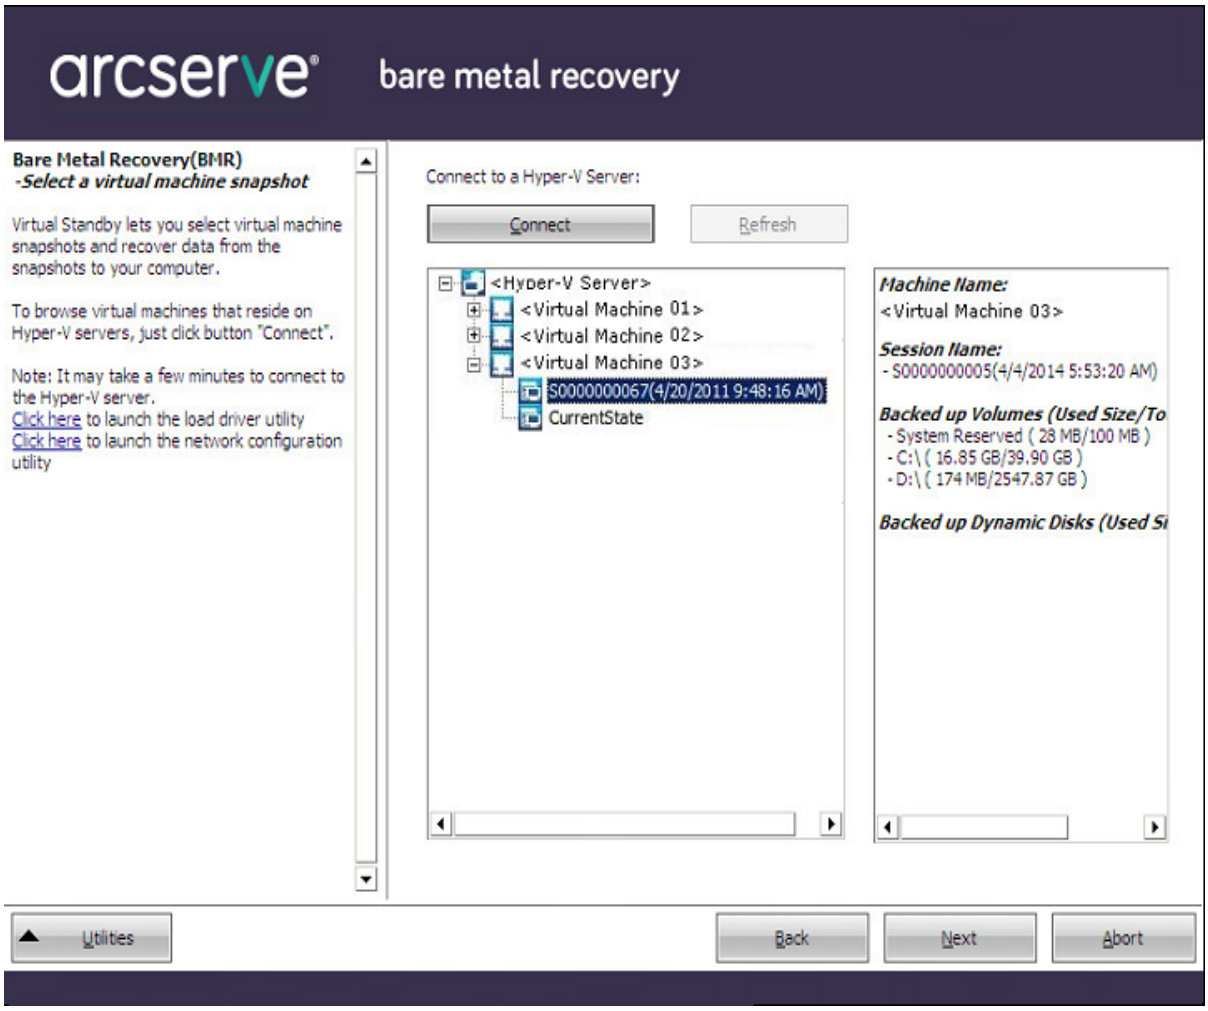
3. Enter the authentication information and click **OK**.

   Arcserve Unified Data Protection Agent for Windows detects and displays the Hyper-V Server with a listing of all the virtual machines that are converted to the specified Hyper-V server using Arcserve Central Virtual Standby or Instant VM.



4. Select the virtual machine that contains the recovery point snapshots for your backup image.

   The backup sessions (recovery point snapshots) for the selected virtual machine are displayed.



5. Select the virtual machine backup session (recovery point snapshot) that you want to recover.

   The corresponding details for the selected recovery point snapshot (virtual machine name, backup session name, backed up volumes) are displayed in the right pane.

   In addition to selecting one of the listed recovery points, you also have the option to select the **Current State** or the **Latest State** recovery point.

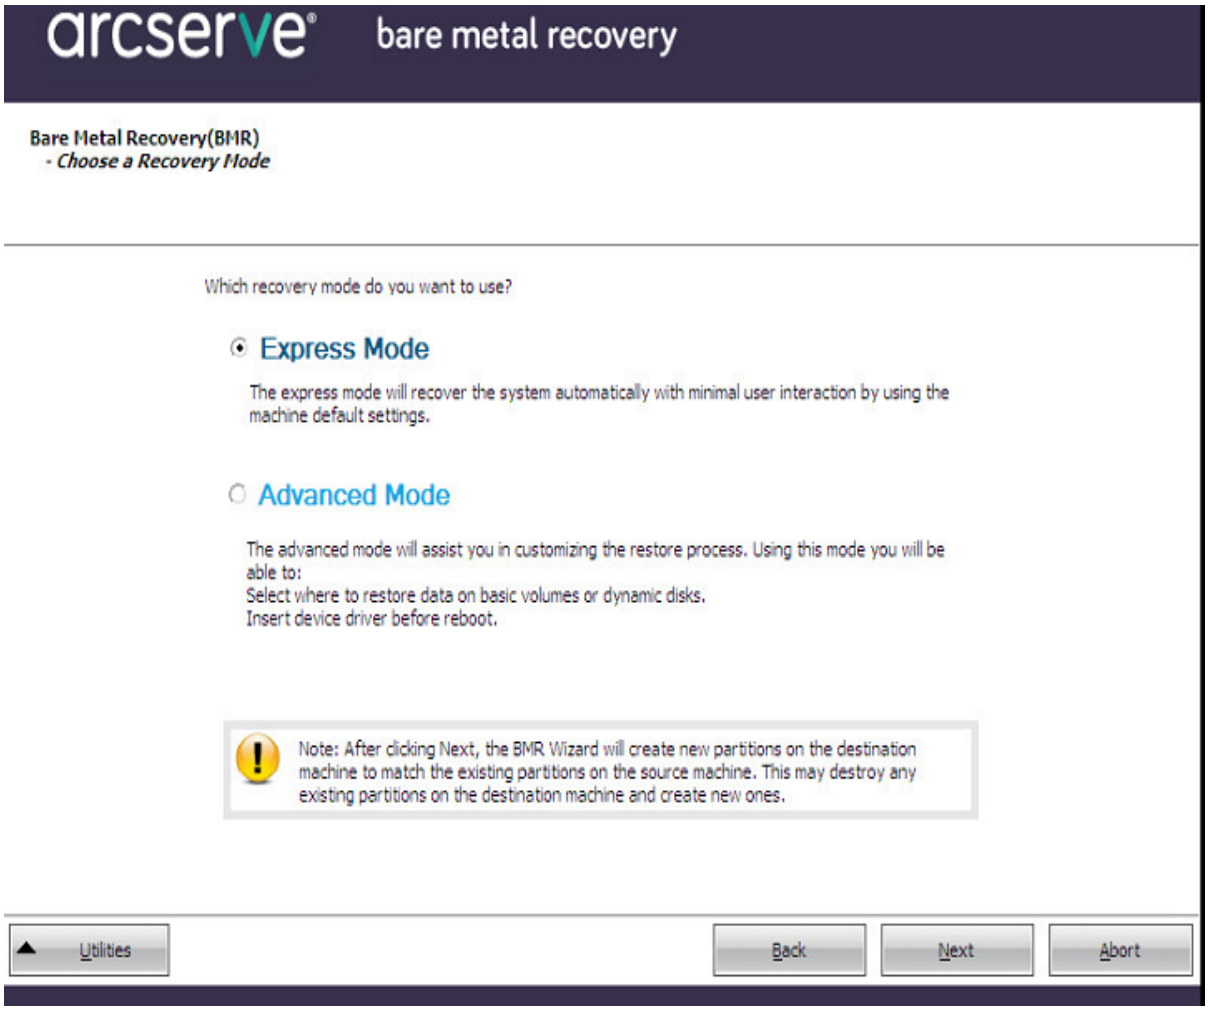
   - If the virtual machine that you are recovering from is powered on, the **Current State** recovery point is displayed.

     Note: If the virtual machine is powered on, then any data changes in the virtual machine after the BMR process started will not be recovered.

   - If the virtual machine that you are recovering from is powered off, the **Latest State** recovery point is displayed.

6. Verify this is the recovery point that you want to restore and click **Next**.

   A BMR wizard screen is displayed with the available recovery mode options.



The available options are **Advanced Mode** and **Express Mode**.

- Select **Express Mode** if you want minimal interaction during the recovery process. For more information see, Perform BMR in Express Mode.
- Select **Advanced Mode** if you want to customize the recovery process. For more information, see Perform BMR in Advanced Mode.

Default: Express Mode.

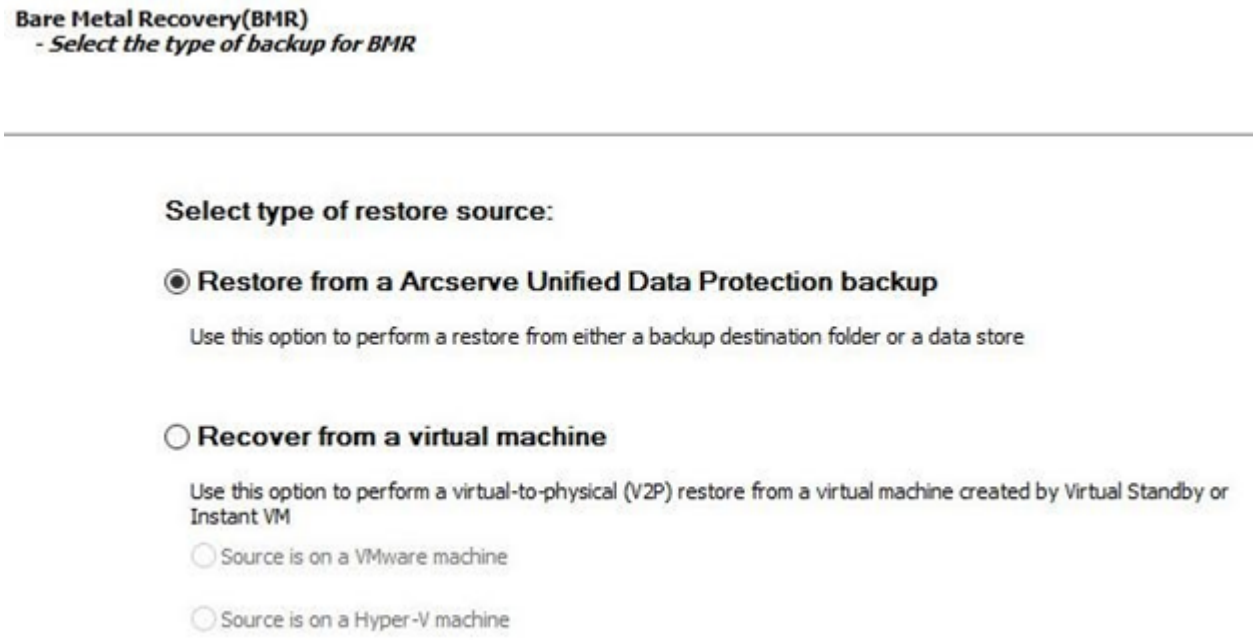
Recover using a VMware Virtual Standby VM or Instant VM

The Arcserve UDP Agent (Windows) provides the capability to perform Bare Metal Recovery for virtual-to-physical (V2P) machines. This feature lets you perform virtual-to-physical recovery from the latest state of a standby virtual machine and helps you reduce the loss of your production machine.

Follow these steps:

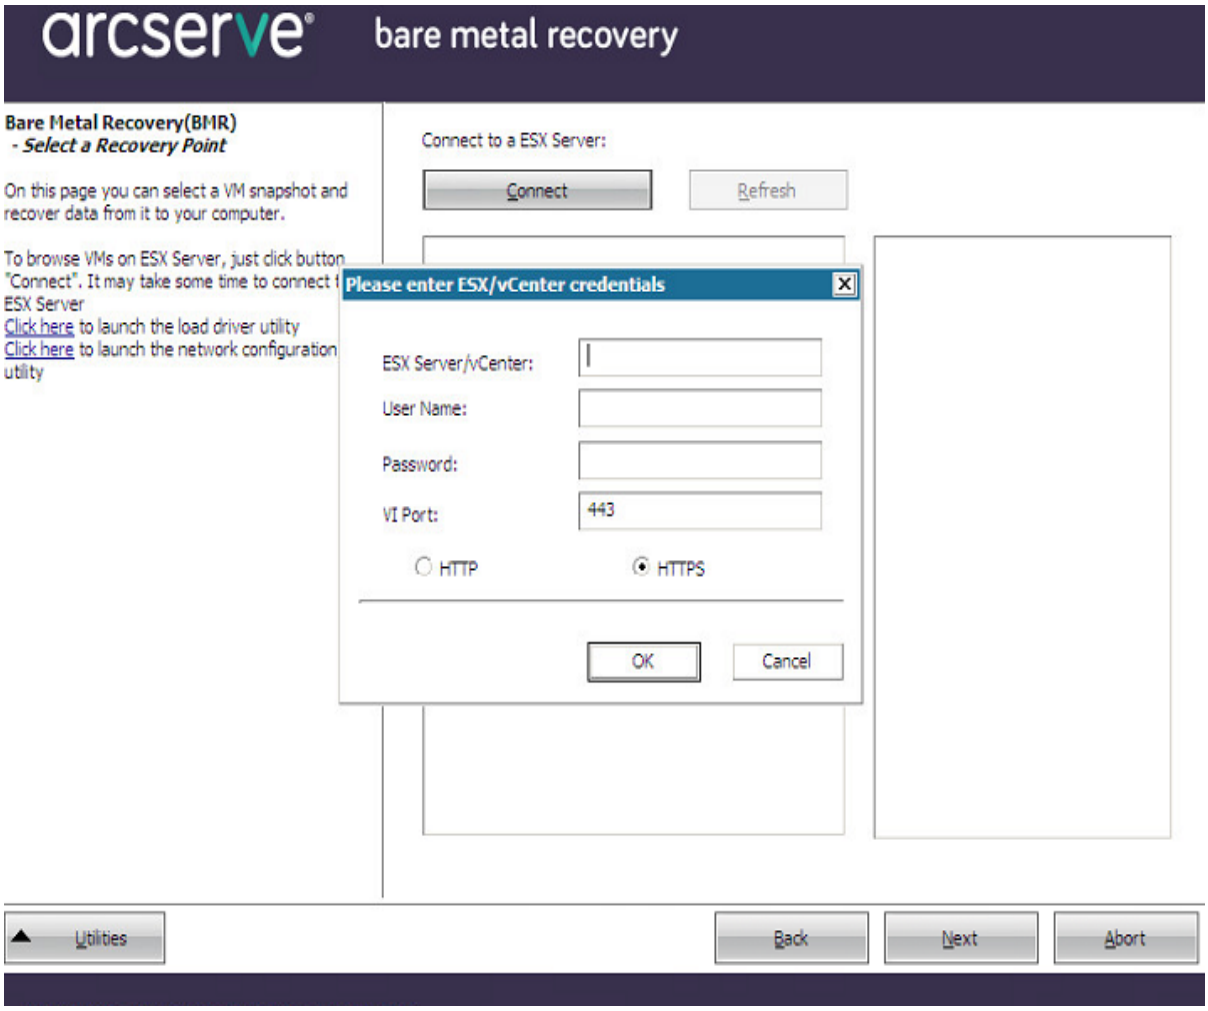
1. From the select the Type of Bare Metal Recovery (BMR) wizard screen, select the **Recover from a virtual machine** and select the **Source is on a VMware machine** option.

   Use this option to perform a virtual-to-physical restore from a virtual standby VM or Instant VM. The term virtual-to-physical refers to the migration of an operating system (OS), application programs and data from a virtual machine or disk partition to a computer's main hard disk. The target can be a single computer or multiple computers.



2. Click **Next**.

   The **Select a Recovery Point** screen is displayed with the **ESX/VC Credentials** dialog.



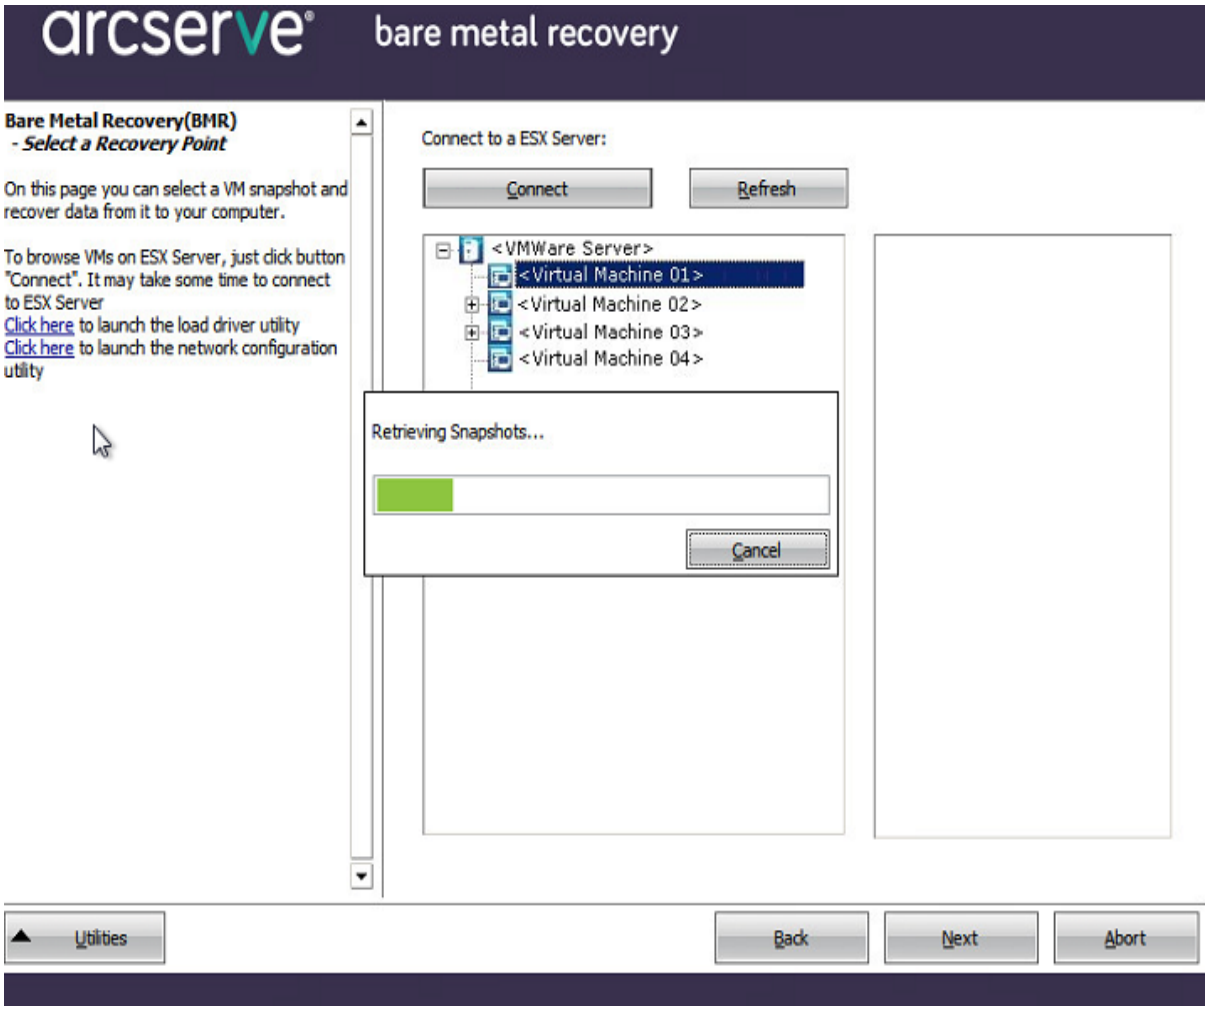
3. Enter the credential information and click **OK**.

   Note: If you are connecting to a vCenter, you do not need an Administrator permission at the vCenter Server level but you must have an Administrator permission at the Datacenter level. In addition, you must have the following permissions at the vCenter Server level:

   - Global, DisableMethods and EnableMethods
   - Global, License

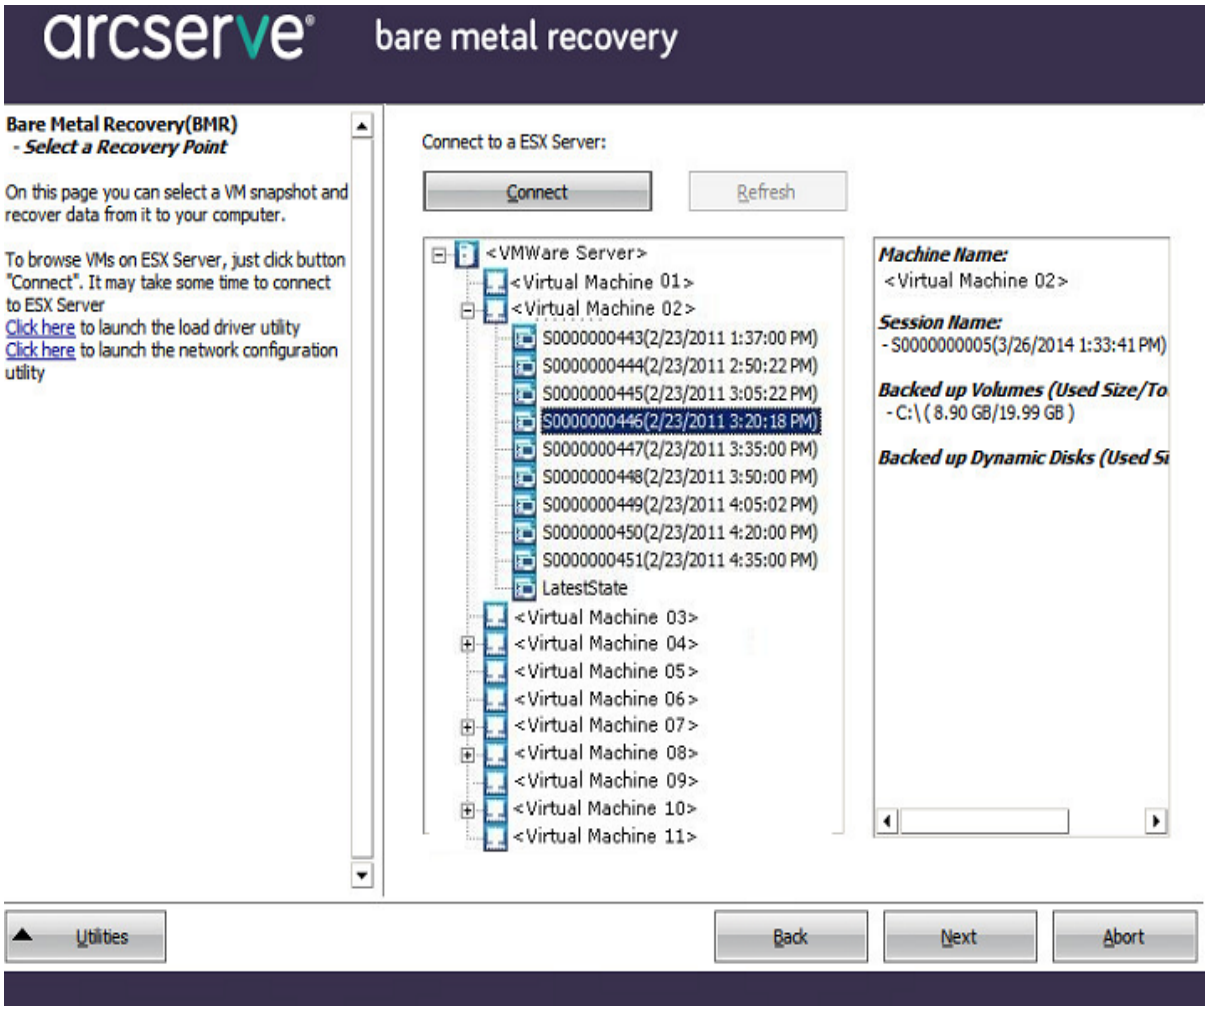
   The **Select a Recovery Point** screen is displayed.

   The Arcserve Unified Data Protection Agent for Windows then retrieves all the recovery point snapshots for the selected VMware server and displays the VMware Server in the left pane, with a listing of all the virtual machines that are hosted on the selected VMware server.



4. Select the virtual machine which contains recovery points for your backup image.

   The backup sessions (recovery point snapshots) for the selected virtual machine are displayed.



5. Select the virtual machine backup session (recovery point snapshots) that you want to recover.

   The corresponding details for the selected recovery point snapshot (virtual machine name, backup session name, backed up volumes, backed up dynamic disks) are displayed in the right pane.

   In addition to selecting one of the listed recovery points, you also have the option to select the **Current State** or the **Latest State** recovery point.

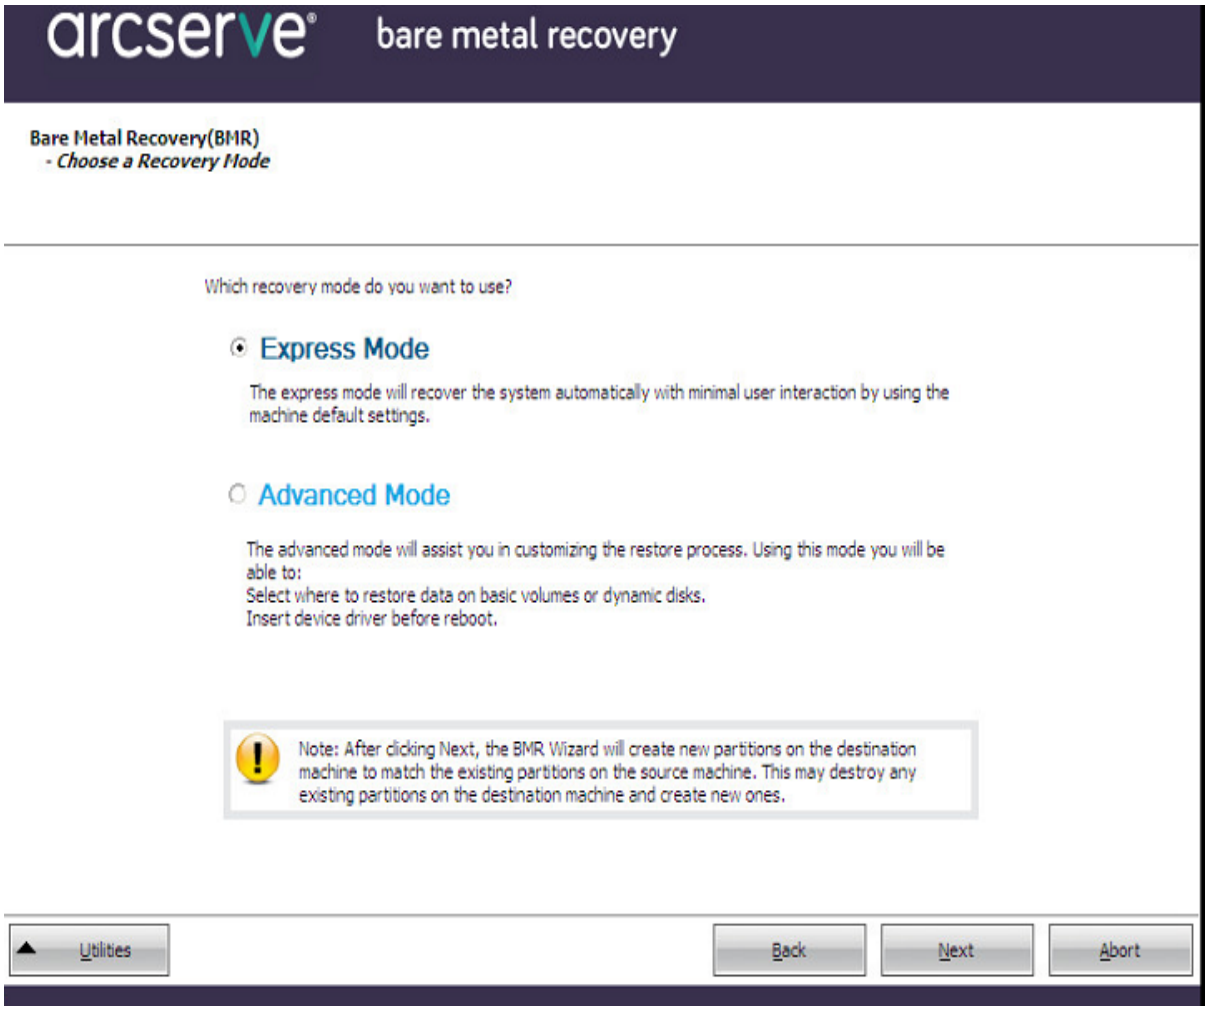
   - If the virtual machine that you are recovering from is powered on, the **Current State** recovery point is displayed.

     Note: If the virtual machine is powered on, then any data changes in the virtual machine after the BMR process started will not be recovered.

   - If the virtual machine that you are recovering from is powered off, the **Latest State** recovery point is displayed.

6. Verify this is the recovery point that you want to restore and click **Next**.

   A BMR wizard screen is displayed with the available recovery mode options.



The available options are **Advanced Mode** and **Express Mode**.

- Select **Express Mode** if you want minimal interaction during the recovery process. For more information see, Perform BMR in Express Mode.
- Select **Advanced Mode** if you want to customize the recovery process. For more information, see Perform BMR in Advanced Mode.

Default: Express Mode.

Perform BMR in Express Mode

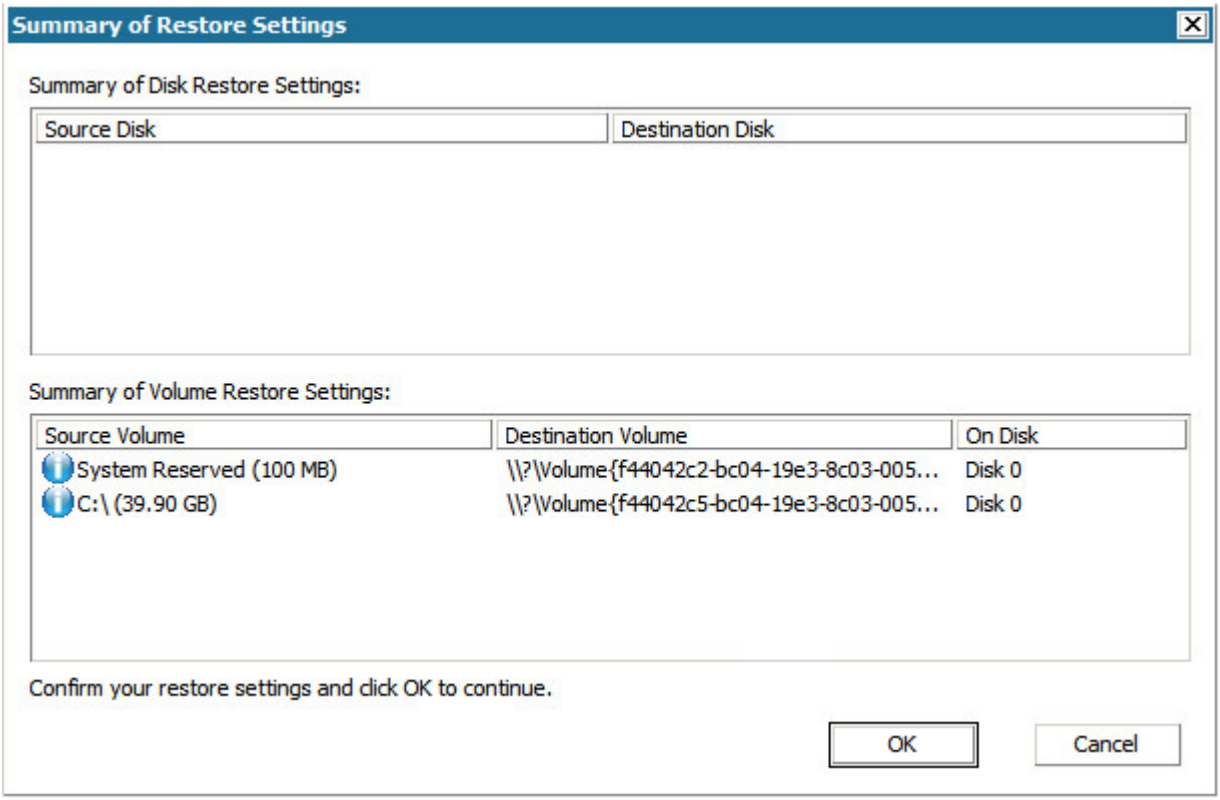
The **Express Mode** requires minimal interaction during the recovery process.

Follow these steps:

1. From the **Choose a Recovery Mode** dialog, select **Express Mode** and click **Next**.

   The **Summary of Disk Restore Settings** screen opens, displaying a summary of the volumes that are going to be restored.

   Note: On the bottom of restore summary window, the drive letters listed in **Destination Volume** column are automatically generated from the Windows Pre-installation Environment (WinPE). They can be different from the drive letters listed in **Source Volume** column. However, the data is still restored to proper volume even if drive letters are different.



2. After you have verified that the summary information is correct, click **OK**.

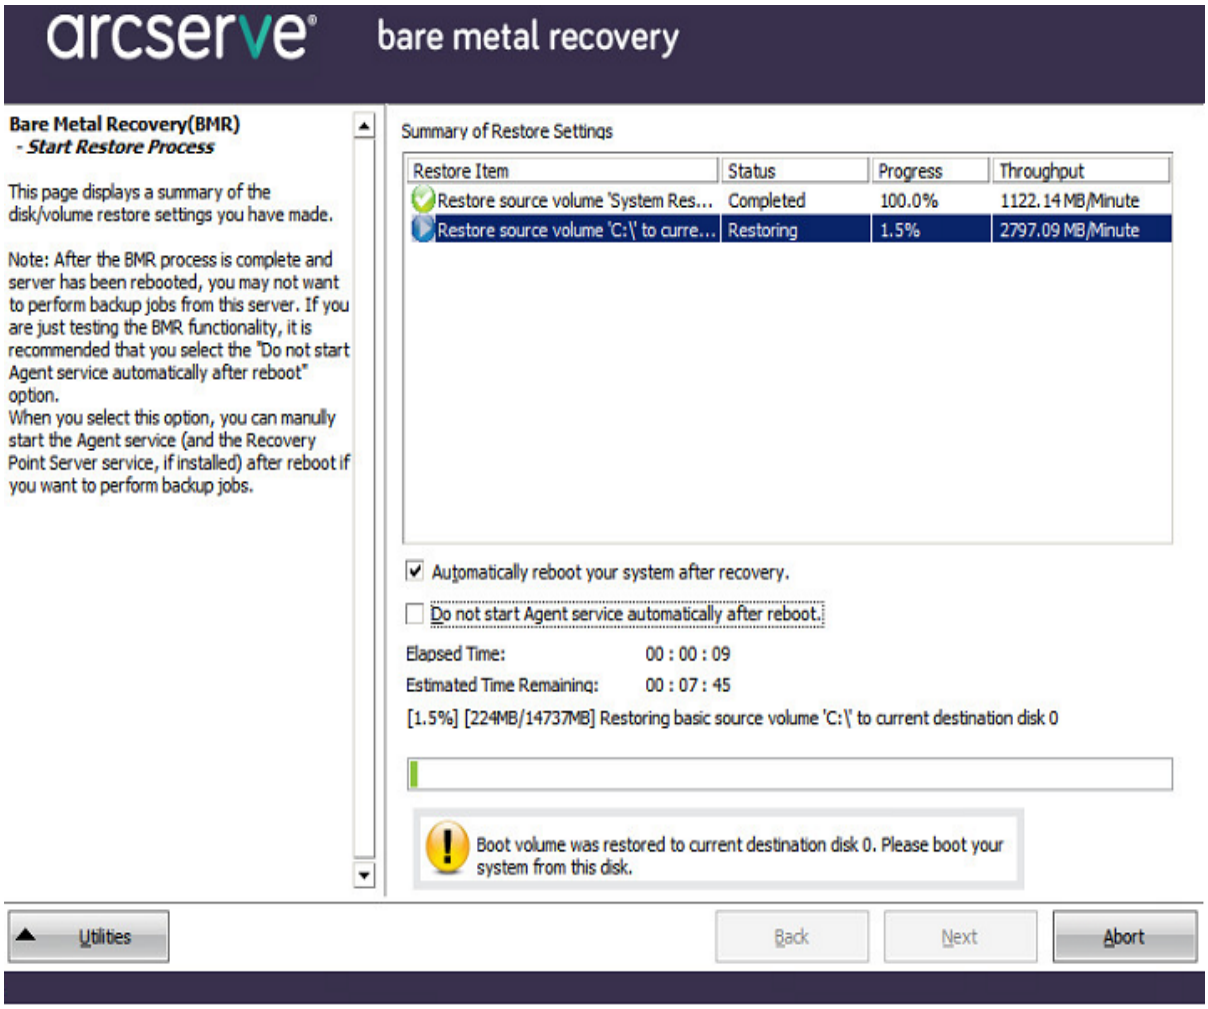
   The restore process starts. The BMR wizard screen displays the restore status for each volume.

   - Depending upon the size of the volume being restored, this operation can take some time

- During this process you are restoring, block-by-block whatever you had backed up for that recovery point and creating a replica of the source machine on the target machine.
- By default, the option to reboot your system automatically after recovery is selected. If necessary, you can clear this option and you can reboot manually at a later time.

  Important: If you are performing an authoritative restore of an active directory after a BMR, you must uncheck the option **Automatically reboot your system after recovery** and for more information, see How to Perform an Authoritative Restore of an Active Directory after a BMR.

  - If necessary, you can select Do not start Agent service automatically after reboot.
  - If necessary, you can cancel or abort the operation at any time.



3. From the **Utilities** menu, you can access the BMR **Activity Log** and you can use the **Save** option to save the Activity Log.

By default, the Activity Log is saved to the following location:

X:\windows\system32\dr\log.

Note: To avoid getting a Windows-generated error, do not save the Activity Log on your desktop or create a folder on your desktop using the **Save As** option from the BMR Activity Log window.

4. If you are restoring to dissimilar hardware (the SCSI/FC adapter which used to connect hard drives could have been changed) and no compatible driver is detected in your original system, a "driver injection" page is displayed to allow you to provide drivers for these devices.

   You can browse and select drivers to inject to the recovered system so that even if you are recovering to a machine with dissimilar hardware, you can still bring back the machine after BMR.

5. When the BMR process is completed, a confirmation notification is displayed.

Perform BMR in Advanced Mode

The **Advanced Mode** lets you customize the recovery process.

Follow these steps:

1. From the **Choose a Recovery Mode** dialog, select **Advanced Mode** and click **Next**.

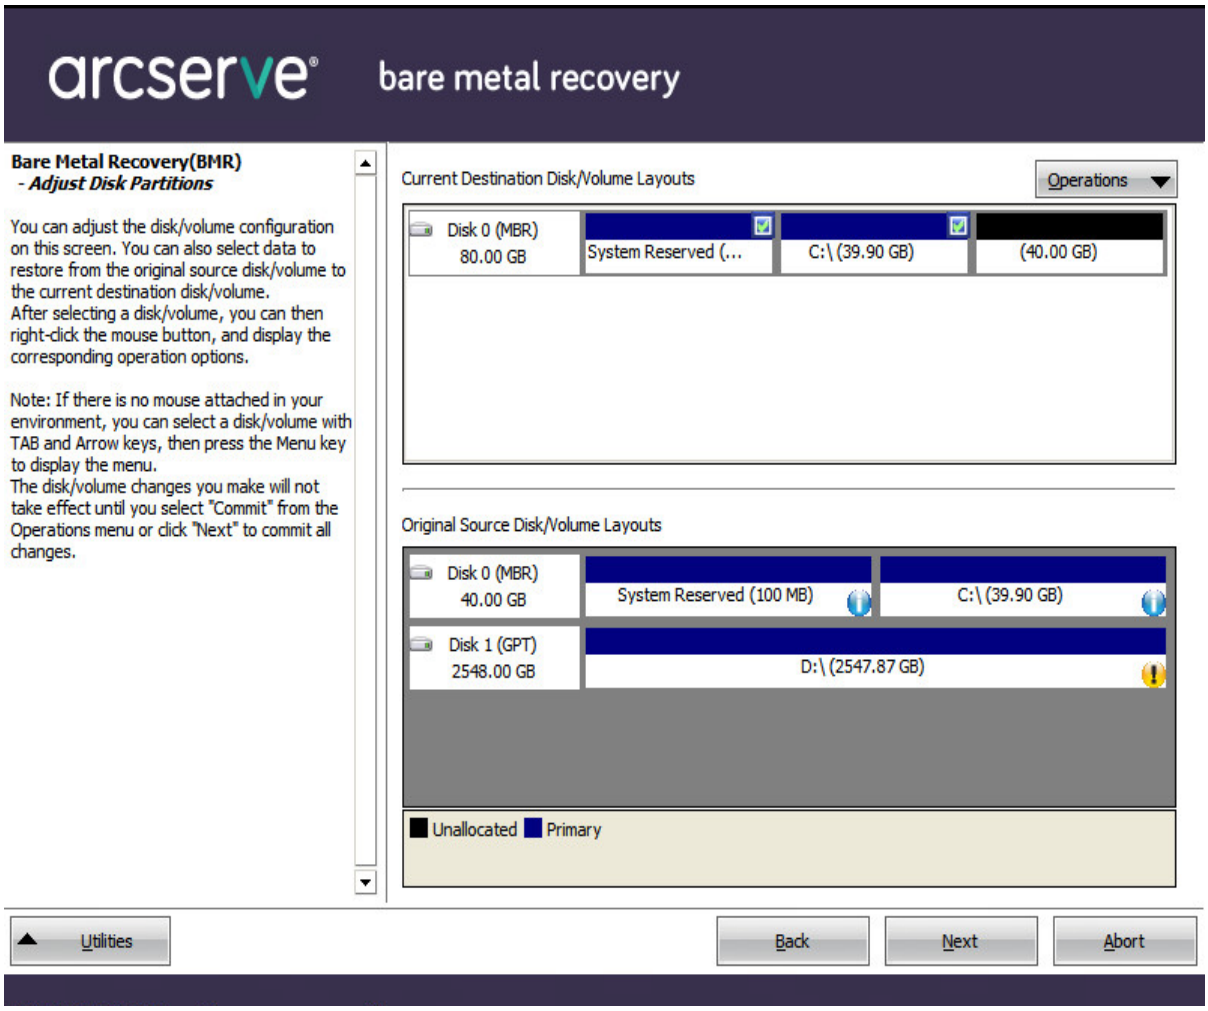
   The BMR utility starts locating the machine that is going to be recovered and displays the corresponding disk partition information.

   The upper pane shows the disk configuration that you have on the current (target) machine and the lower pane shows the disk partition information that you had on the original (source) machine.

   Important! A red X icon displaying for a source volume in the lower pane indicates that this volume contains system information and has not been assigned (mapped) to the target volume. This system information volume from the source disk must be assigned to the target disk and restored during BMR or the reboot fails.

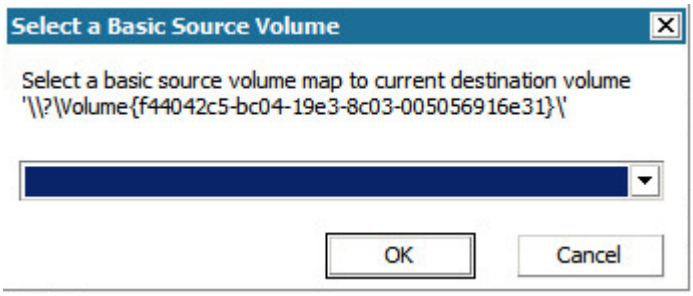
   Note: If you perform BMR and you restore the system volume to a disk which is not configured as the boot disk, it will fail to boot the machine after BMR is completed. Ensure that you are restoring the system volume to a properly configured boot disk.

   Note: When restoring to another disk/volume, the capacity of new disk/volume must be the same size or larger than original disk/volume. In addition, disk resizing is for basic disks only, and not for dynamic disks.



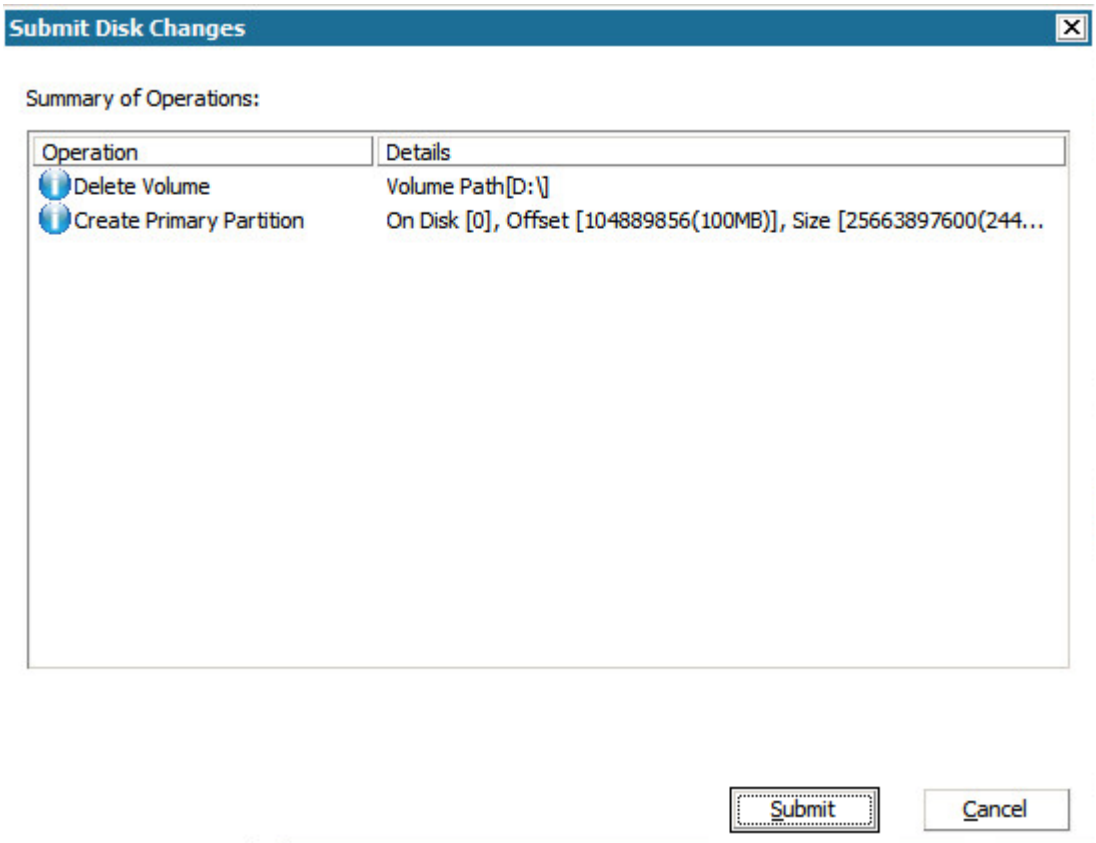
2. If the current disk information you are seeing does not appear correct, you can access the **Utilities** menu and check for missing drivers.
3. If necessary, on the target disk/volume pane you can click the **Operations** drop-down menu to display the available options. For more information about these options, see Managing the BMR Operations Menu.
4. Click on each target volume and from the pop-up menu, select the **Map Volume From** option to assign a source volume to this target volume.

   The **Select a Basic Source Volume** dialog opens.



5. From **Select a Basic Source Volume** dialog, click the drop-down menu and select the available source volume to assign to the selected target volume. Click **OK**.
   - On the target volume, a checkmark icon is displayed, indicating that this target volume has been mapped to.
   - On the source volume, the red X icon changes to a green icon, indicating that this source volume has been assigned to a target volume.
6. When you are sure all volumes that you want to restore and all volumes containing system information are assigned to a target volume, click **Next**.

   The Submit Disk Changes screen opens, displaying a summary of the selected operations. For each new volume being created, the corresponding information is displayed.



7. When you have verified the summary information is correct, click **Submit**. (If the information is not correct, click **Cancel**).

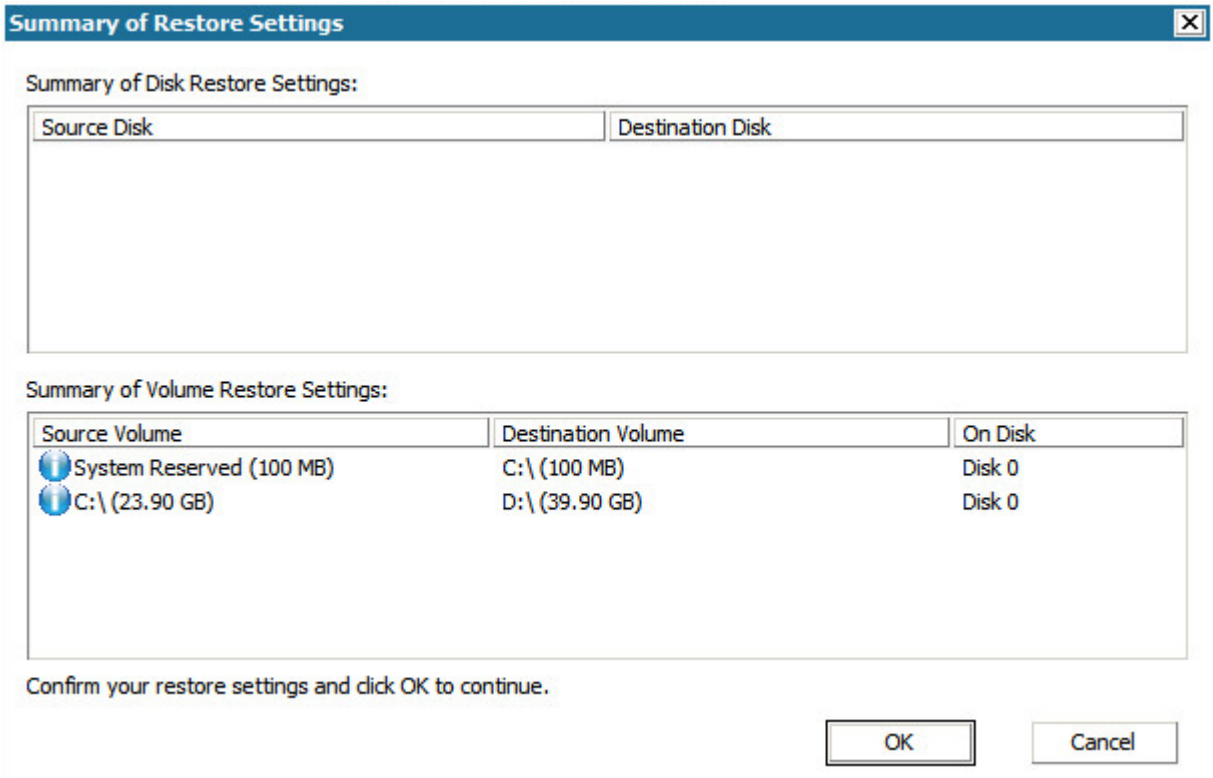
   Note: All operations to the hard drive do not take effect until you submit it.

   On the target machine, the new volumes are created and mapped to the corresponding source machine.

8. When the changes are completed, click **OK**.

   The Summary of Disk Restore Settings screen opens, displaying a summary of the volumes that are going to be restored.

   Note: On the bottom of restore summary window, the drive letters listed in "Destination Volume" column are automatically generated from the Windows Pre-installation Environment (WinPE). They can be different from the drive letters listed in "Source Volume" column. However, the data is still restored to proper volume even if drive letters are different.



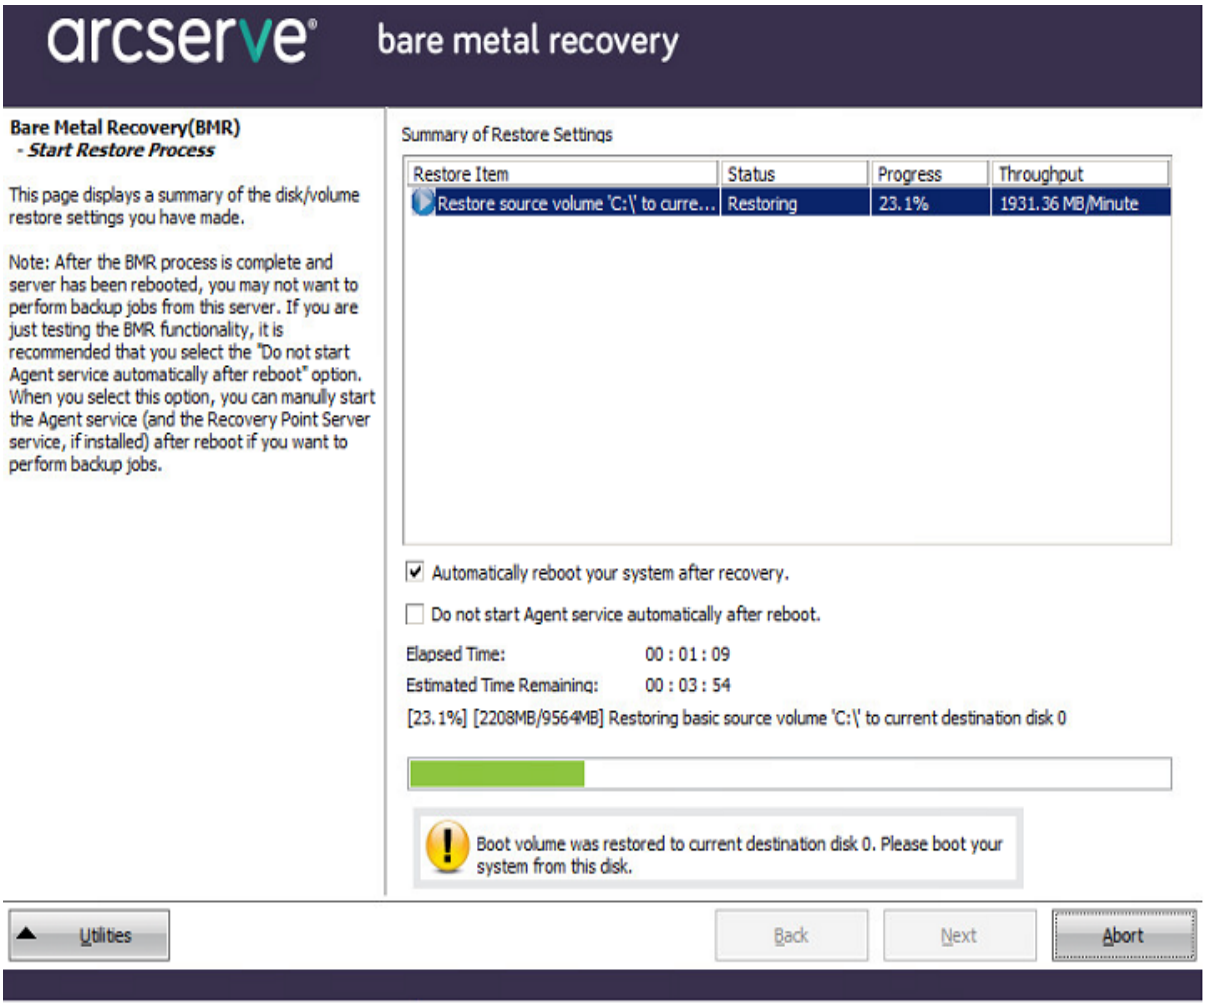
9. After you have verified that the summary information is correct, click **OK**.

   The restore process starts. The BMR wizard screen displays the restore status for each volume.

   - Depending upon the size of the volume being restored, this operation can take some time.

- During this process you are restoring, block-by-block whatever you had backed up for that recovery point and creating a replica of the source machine on the target machine.
- By default, the option to reboot your system automatically after recovery is selected. If necessary, you can clear this option and you can reboot manually at a later time.

  Important: If you are performing an authoritative restore of an active directory after a BMR, you must clear the selection for the option **Automatically reboot your system after recovery** and for more information, see How to Perform an Authoritative Restore of an Active Directory after a BMR.
- If necessary, you can select Do not start Agent service automatically after reboot.
- If necessary, you can cancel or abort the operation at any time.



10. From the **Utilities** menu, you can access the BMR **Activity Log** and you can use the **Save** option to save the Activity Log.

By default, the Activity Log is saved to the following location:

X:\windows\system32\dr\log

Note: To avoid getting a Windows-generated error, do not save the Activity Log on your desktop or create a folder on your desktop using the **Save As** option from the BMR Activity Log window.

11. If you are restoring to dissimilar hardware (the SCSI/FC adapter which used to connect hard drives could have been changed) and no compatible driver is detected in your original system, a "driver injection" page is displayed to allow you to provide drivers for these devices.

    You can browse and select drivers to inject to the recovered system so that even if you are recovering to a machine with dissimilar hardware, you can still bring back the machine after BMR.

12. When the BMR process is completed, a confirmation notification is displayed.

Verify that the BMR was Successful

To verify that the BMR was successful, perform the following tasks:

- Reboot the operating system.
- Verify all systems and applications function correctly.
- Verify all network settings are properly configured.
- Verify the BIOS is configured to boot from the disk on which the boot volume was restored to.
- When the BMR is completed, be aware of the following conditions:
  - The first backup that is performed after the BMR is a Verify Backup.
  - When the machine has been rebooted, you may need to configure the network adapters manually if you restored to dissimilar hardware.

    Note: When the machine is rebooting, a Windows Error Recovery screen may be displayed indicating that Windows did not shut down successfully. If this occurs, you can safely ignore this warning and continue to start Windows normally.
  - For dynamic disks, if the status of the disk is offline, you can manually change it to online from the disk management UI (accessed by running the Diskmgmt.msc control utility).
  - For dynamic disks, if the dynamic volumes are in a failed redundancy status, you can manually resynchronize the volumes from the disk management UI (accessed by running the Diskmgmt.msc control utility).

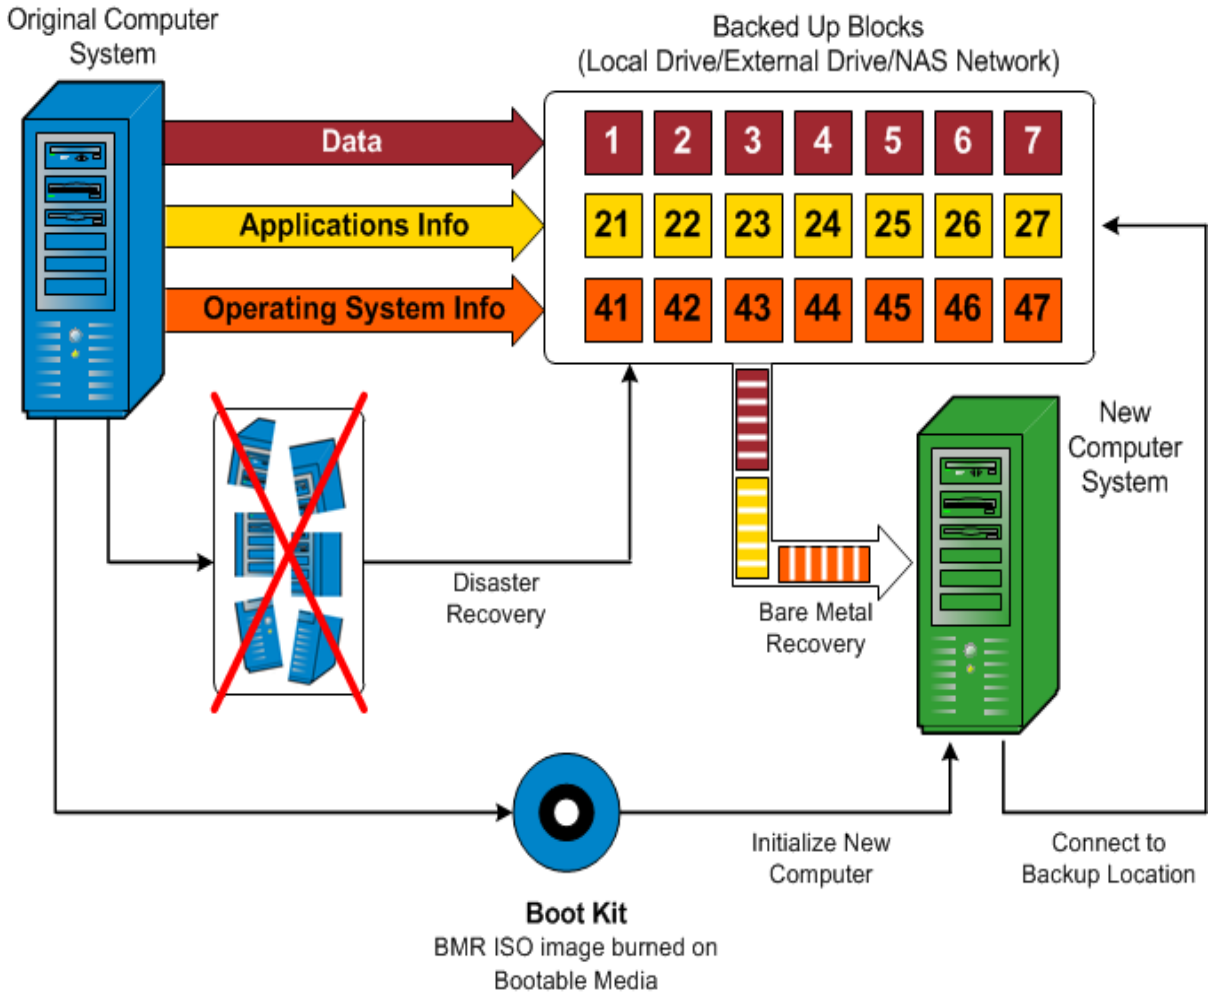
BMR Reference Information

- How Bare Metal Recovery Works
- Operating Systems that Support UEFI/BIOS Conversion
- Managing the BMR Operations Menu

How Bare Metal Recovery Works

Bare Metal Recovery is the process of restoring a computer system from "bare metal" by reinstalling the operating system and software applications, and then restoring the data and settings. The most common reasons for performing a bare metal recovery are because your hard drive either fails or becomes full and you want to upgrade (migrate) to a larger drive or migrate to newer hardware. Bare metal recovery is possible because during the block-level backup process, Arcserve UDP Agent (Windows) captures not only the data, but also all information related to the operating system, installed applications, configuration settings, necessary drivers, and so on. All relevant information that is necessary to perform a complete rebuild of the computer system from "bare metal" is backed up into a series of blocks and stored on the backup location.

Note: Dynamic disks are restored at disk level only. If your data is backed up to a volume on a dynamic disk, you will not be able to restore this dynamic disk (including all its volumes) during BMR.



When you perform a bare metal recovery, the Arcserve UDP Agent (Windows) boot disk is used to initialize the new computer system and allow the bare metal recovery process to begin. When the bare metal recovery is started, Arcserve UDP Agent (Windows) will prompt you to select or provide a valid location to retrieve these backed up blocks from, as well as the recovery point to be restored. You may also be prompted to provide valid drivers for the new computer system if needed. When this connection and configuration information is provided, Arcserve UDP Agent (Windows) begins to pull the specified backup image from the backup location and restore all backed up blocks to the new computer system (empty blocks will not be restored). After the bare metal recovery image is fully restored to the new computer system, the machine will be back to the state that it was in when the last backup was performed, and Arcserve UDP Agent (Windows) backups will be able to continue as scheduled. (After completion of the BMR, the first backup will be a Verify Backup).

Operating Systems that Support UEFI/BIOS Conversion

If it is detected that the operating system of your source machine is not the same firmware as your system, you will be asked if you want to convert UEFI to a BIOS-compatible system or BIOS to UEFI-compatible system. The following table lists each operating system and the type of conversion supported:

Operating System (OS)	CPU	uEFI to BIOS	BIOS to uEFI
Windows Vista (None SP)	x86	No	No
Windows Vista (None SP)	x64	No	No
Windows Vista SP1	x86	No	No
Windows Vista SP1	x64	Yes	Yes
Windows Server 2008	x86	No	No
Windows Server 2008	x64	Yes	Yes
Windows Server 2008 R2	x64	Yes	Yes
Windows 7	x86	No	No
Windows 7	x64	Yes	Yes
Windows 8	x86	No	No
Windows 8	x64	Yes	Yes
Windows Server 2012	x64	Yes	Yes
Windows 8.1	x86	No	No
Windows 8.1	x64	Yes	Yes
Windows 10	x86	No	No
Windows 10	x64	Yes	Yes
Windows Server 2012 R2	x64	Yes	Yes
Windows Server 2016	x64	Yes	Yes
Windows Server 2019	x64	Yes	Yes

Managing the BMR Operations Menu

The BMR Operations menu consists of the following three types of operations:

- Disk Specific Operations
- Volume/Partition Specific Operations
- BMR Specific Operation

Disk Specific Operations:

To perform disk specific operations, select the disk header and click **Operations**.

Clean Disk

This operation is used to clean all partitions of a disk and is:

- An alternate method to delete all volumes of a disk. With the **Clean Disk** operation, you do not have to delete each volume one by one.
- Used to delete the non-Windows partitions. Due to a VDS limitation, the non-Windows partition cannot be deleted from the UI, but you can use this operation to clean them all.

  Note: During BMR, when the destination disk has non-Windows partitions or OEM partitions, you cannot select this partition and delete it from the BMR UI. Usually this would occur if you ever installed Linux/Unix on the destination disk. To resolve this issue, perform one of the following tasks:
- Select the disk header on the BMR UI, click **Operations**, and use the **Clean Disk** operation to erase all partitions on the disk.
- Open a command prompt and type **Diskpart** to open the Diskpart command console. Then type "select disk x" , where 'x' is the disk number and "clean" to erase all partitions on the disk.

Convert to MBR

This operation is used to convert a disk to MBR (Master Boot Record). It is available only when the selected disk is a GPT (GUID Partition Table) disk and there are no volumes on this disk.

Convert to GPT

This operation is used to convert a disk to GPT. It is available only when the selected disk is an MBR disk and there are no volumes on this disk.

Convert to Basic

This operation is used to convert a disk to Basic. It is available only when the selected disk is a Dynamic disk and there are no volumes on this disk.

Convert to Dynamic

This operation is used to convert a disk to Dynamic Disk. It is available only when the selected disk is a Basic disk.

Online Disk

This operation is used to bring a disk online. It is available only when the selected disk is in the offline status.

Disk Properties

This operation is used to view detailed disk properties. It is always available and when you select this operation, a **Disk Properties** dialog appears.

Volume/Partition Specific Operations:

To perform volume/partition operations, select the disk body area and click **Operations**. From this menu, you can create new partitions to correspond to the disk partitions on the source volume.

Create Primary Partition

This operation is used to create a partition on a basic disk. It is available only when the selected area is an unallocated disk space.

Create Logical Partition

This operation is used to create a logical partition on a basic MBR disk. It is available only when the selected area is an extended partition.

Create Extended Partition

This operation is used to create an extended partition on a basic MBR disk. It is available only when the disk is an MBR disk and the selected area is an unallocated disk space.

Create System Reserved Partition

This operation is used to create the System Reserved Partition on a BIOS firmware system and builds a mapping relationship with the source EFI System Partition. It is only available when you restore a UEFI system to a BIOS system.

Note: If you previously converted from UEFI to a BIOS-compatible system, use the Create System Reserved Partition operation for destination disk resizing.

Create EFI System Partition

This operation is used to create the EFI System Partition on a basic GPT disk. It is available only when the target machine firmware is UEFI and the selected disk is a basic GPT disk.

Note: If you previously converted from BIOS to a UEFI-compatible system, use the Create EFI System Partition operation for destination disk resizing.

Note: Systems that support UEFI also require that the boot partition reside on a GPT (GUID Partition Table) disk. If you are using a MBR (Master Boot Record) disk, you must convert this disk to a GPT disk, and then use the Create EFI System Partition operation for disk resizing.

Resize Volume

This operation is used to resize a volume. It is an alternate method of Windows "Extend Volume/Shrink Volume". It is available only when the selected area is a valid disk partition.

Delete Volume

This operation is used to delete a volume. It is available only when the selected area is a valid volume.

Delete Extended Partition

This operation is used to delete the extended partition. It is available only when the selected area is the extended partition.

Volume Properties

This operation is used to view detailed volume properties. When you select this operation, a **Volume Properties** dialog appears.

BMR Specific Operations:

These operations are specific to BMR. To perform BMR operations, select the disk header or the disk body area and click **Operations**.

Map Disk From

This operation is used to build a mapping relationship between the source and target dynamic disks. It is available only when the selected disk is a Dynamic disk.

Note: When mapping to another disk, the capacity of each mapped target volume must be the same size or larger than the corresponding source volume.

Map Volume From

This operation is used to build a mapping relationship between the source and target basic volume. It is available only when the selected volume is a Basic volume.

Note: When mapping to another disk, the capacity of each mapped target volume must be the same size or larger than the corresponding source volume.

Commit

This operation is always available. All of the operations are cached in memory and they do not modify the target disks until you select the **Commit** operation.

Reset

This operation is always available. The **Reset** operation is used to relinquish your operations and restore the disk layout to the default status. This operation cleans all the cached operations. Reset means to reload the source and target disk layout information from the configure file and current OS, and discard any user changed disk layout information.

Troubleshooting BMR Issues

When a problem is detected, Arcserve UDP Agent (Windows) generates a message to help you identify and resolve the problem. These messages are contained in the Arcserve UDP Agent (Windows) **Activity Log**, which is accessed from the **View Logs** option on the home page UI. In addition, when an incorrect action is attempted, Arcserve UDP Agent (Windows) generally displays a pop-up message to help you identify and quickly resolve the problem.

- Slow throughput performance during BMR
- After BMR, dynamic volumes are not recognized by the operating system
- Unable to Reboot Hyper-V VM After BMR
- Unable to Reboot VMware VM After BMR
- Unable to boot the server after performing a BMR
- Failed to submit BMR job to Recovery Point Server

Slow throughput performance during BMR

This problem can be caused by SATA controllers with "AHCI" enabled.

During BMR, Arcserve UDP Agent (Windows) will install drivers for critical unknown devices. If the device already has a driver installed, Arcserve UDP Agent (Windows) will not update that driver again. For some devices, Windows 7PE may have the drivers for them, but these drivers may not be the best ones and this can cause the BMR to run too slow.

To resolve this problem, perform one of the following tasks:

- Check if the driver pool folder contains the newest disk drivers. If it does, and you are restoring to the original machine, please install the new driver from the driver pool folder. If you are restoring to alternate machine, download the latest disk drivers from the Internet, and load it before you start data recovery. To load the driver, you can use the "drvload.exe" utility, which is included in Windows PE.
- Change the device operating mode from "AHCI" (Advanced Host Controller Interface) to Compatibility mode. (Compatibility mode provides a better throughput).

If the problem persists, use Live Chat to contact Arcserve support. Live Chat lets you optimize intelligent conversation between you and the Technical Support team, allowing you to address your concerns and questions immediately, while still maintaining access to the product.

After BMR, Dynamic Volumes are Not Recognized by the Operating System

To keep dynamic disks in a consistent state, the Windows operating system automatically synchronizes the Logical Disk Manager (LDM) metadata on each dynamic disk. So when BMR restores one dynamic disk and brings it online, the LDM metadata on this disk is automatically updated by the operating system. This may result in a dynamic volume not being recognized by the operating system and missing after the reboot.

To resolve this problem, when you perform BMR with multiple dynamic disks, do not perform any pre-BMR disk operations such as cleaning, deleting volume, and so on.

If the problem persists, use Live Chat to contact Arcserve support. Live Chat lets you optimize intelligent conversation between you and the Technical Support team, allowing you to address your concerns and questions immediately, while still maintaining access to the product.

Unable to Reboot Hyper-V VM After BMR

If you performed BMR to a Hyper-V machine consisting of more than one disk connected to an Integrated Drive Electronics (IDE) controller and if the server does not reboot, perform the following troubleshooting procedure:

1. Verify that the disk that contains the system volume is the master disk.

   The Hyper-V BIOS searches for the system volume on the master disk (disk 1) which is connected to the master channel. If the system volume is not located on the master disk, the VM will not reboot.

   Note: Verify that the disk that contains the system volume is connected to an IDE controller. Hyper-V cannot boot from a SCSI disk.

2. If necessary, modify the Hyper-V settings, to connect the disk that contains the system volume to the IDE master channel and reboot the VM again.

   If the problem persists, use Live Chat to contact Arcserve support. Live Chat lets you optimize intelligent conversation between you and the Technical Support team, allowing you to address your concerns and questions immediately, while still maintaining access to the product.

Unable to Reboot VMware VM After BMR

If you performed BMR to a VMware machine consisting of more than one disk connected to an Integrated Drive Electronics (IDE) controller or a SCSI adapter and the server does not reboot, perform the following troubleshooting procedure:

1. Verify that the disk that contains the system volume is the master disk.

   The VMware BIOS searches for the system volume on the Master disk (disk 0) which is connected the master channel. If the system volume is not on the Master disk, the VM does not reboot.

2. If necessary, modify the VMware settings, to connect the disk that contains the system volume to the IDE master channel and reboot the VM again.

3. If the disk is a SCSI disk, verify the disk which contains boot volume is the first disk which connects to the SCSI adapter. If not, assign the boot disk from the VMware BIOS.

4. Verify the disk which contains boot volume is in the previous eight disks, because the VMware BIOS only detect eight disks during the boot. If there are more than seven disks ahead the disk which contains system volumes connected to the SCSI adapter, the VM cannot boot.

   If the problem persists, use Live Chat to contact Arcserve support. Live Chat lets you optimize intelligent conversation between you and the Technical Support team, allowing you to address your concerns and questions immediately, while still maintaining access to the product.

Unable to boot the server after performing a BMR

Symptom

When the source machine is an Active Directory server performing a BMR to a physical machine with different hardware or to a virtual machine on a Hyper-V server, the server does not boot and a blue screen displays with the following message:

STOP: c00002e2 Directory Services could not start because of the following error: a device attached to the system is not functioning. Error status: 0xc0000001.

Solution

Reboot the system to the BMR PE environment, rename all *.log files in the C:\Windows\NTDS folder, and restart the system. For example, rename the file edb.log to edb.log.old and restart the system.

If the problem persists, use Live Chat to contact Arcserve support. Live Chat lets you optimize intelligent conversation between you and the Technical Support team, allowing you to address your concerns and questions immediately, while still maintaining access to the product.

Failed to submit BMR job to Recovery Point Server

Only one BMR job is supported when restoring from same RPS server for the same node (Agent backup or Host-Based Backup). This is controlled by the job monitor on the RPS server.

If the machine where the BMR job is running is shut down or rebooted unexpectedly, the job monitor at the RPS server side will wait 10 minutes and then time out. During this time you cannot start another BMR for the same node from the same RPS server.

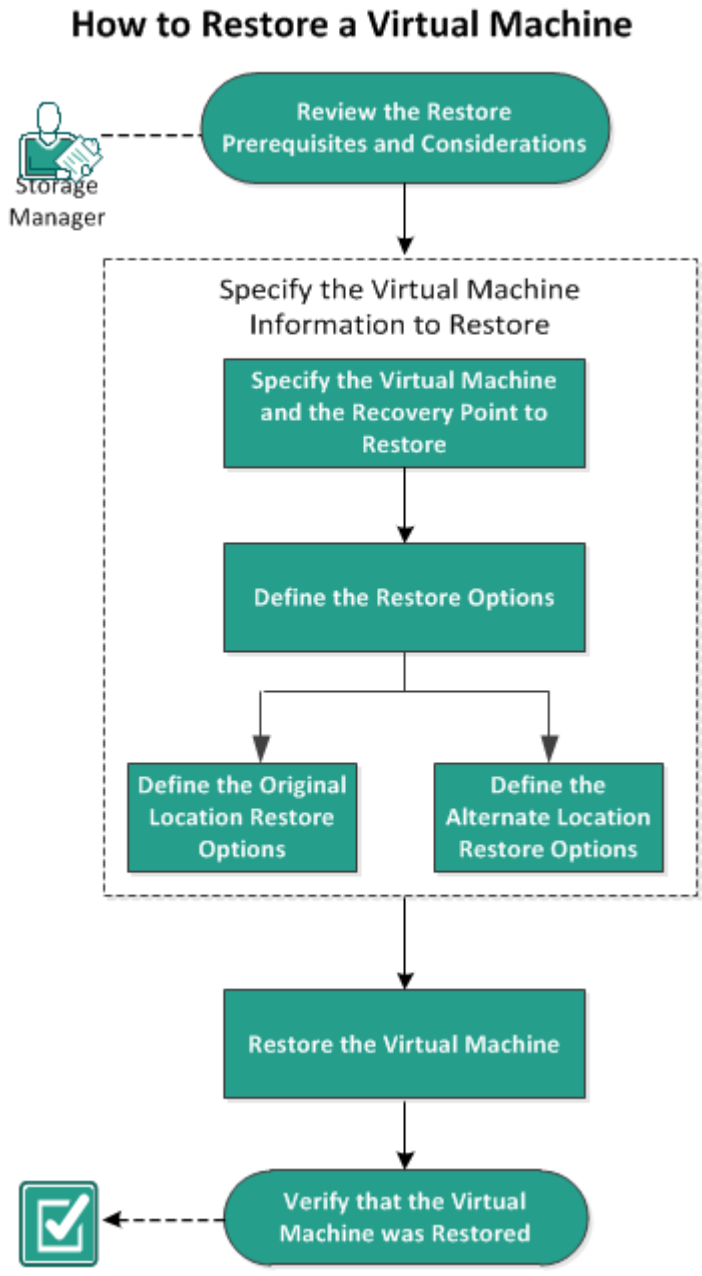
If you abort the BMR from the BMR UI, this problem does not exist.

If the problem persists, use Live Chat to contact Arcserve support. Live Chat lets you optimize intelligent conversation between you and the Technical Support team, allowing you to address your concerns and questions immediately, while still maintaining access to the product.

How to Restore a Virtual Machine

Arcserve UDP lets you use the **Recover VM** option to restore a virtual machine (VM) that you previously backed up using Host-Based Agentless backup. This method helps you restore the entire virtual machine to the original or to an alternate ESX or Hyper-V location. You can browse the available virtual machine recovery points from a calendar view and select which recovery point you want to restore.

The following diagram illustrates the process to restore from a virtual machine:



Perform the following tasks to restore a virtual machine:

1. Review the Restore Prerequisites and Considerations
2. Specify the Virtual Machine Information to Restore
    a. Specify the Virtual Machine and the Recovery Point to Restore
    b. Define the Restore Options
        - Define the Original Location Restore Options
        - Define the Alternate Location Restore Options
3. Restore the Virtual Machine
4. Verify that the Virtual Machine was Restored

Review the Restore Prerequisites and Considerations

Verify that the following prerequisites exist before performing a restore:

- You have a valid recovery point available to restore from.
- You have a valid and accessible target Virtual Center/ESX or Hyper-V server to recover the virtual machine.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Review the following restore considerations:

- If the Recover VM destination is Windows Server 2008 R2 then the source backup VM should not contain VHDx disks, which are not supported on the Hyper-V server (Windows Server 2008 R2).
- If the Recover VM destination is Windows Server 2008 R2 or Win2012 then the source backup VM's sub-system type should not be generation 2 (which was introduced in Windows Server 2012 R2), and is not supported on the Hyper-V server (Windows Server 2012/2008 R2).

Specify the Virtual Machine Information to Restore

You can recover an entire virtual machine from a recovery point.

The process involved in restoring virtual machine is as follows:

1. Specify the Virtual Machine and the Recovery Point to Restore
2. Define the Restore Options
   - Define the Original Location Restore Options
   - Define the Alternate Location Restore Options

Specify the Virtual Machine and the Recovery Point to Restore

Use the **Recover VM** option to restore a virtual machine that you previously backed up. This method quickly and consistently creates a virtual machine from an Arcserve UDP recovery point on an ESX or Hyper-V server. The recovered virtual machine can then simply be started to complete the recovery process.

Follow these steps:

1. Access the restore method selection dialog in one of the following ways:

   From Arcserve UDP:

   a. Log into Arcserve UDP.

   b. Click the **resources** tab.

   c. Select **All Nodes** in the left pane.

      All the added nodes are displayed in the center pane.

   d. In the center pane, select the node and click **Actions**.

   e. Click **Restore** from the **Actions** drop-down option.

      The restore method selection dialog opens.

      Note: You are automatically logged in to the agent node and the restore method selection dialog is opened from the agent node.

   From Arcserve UDP Agent (Windows):

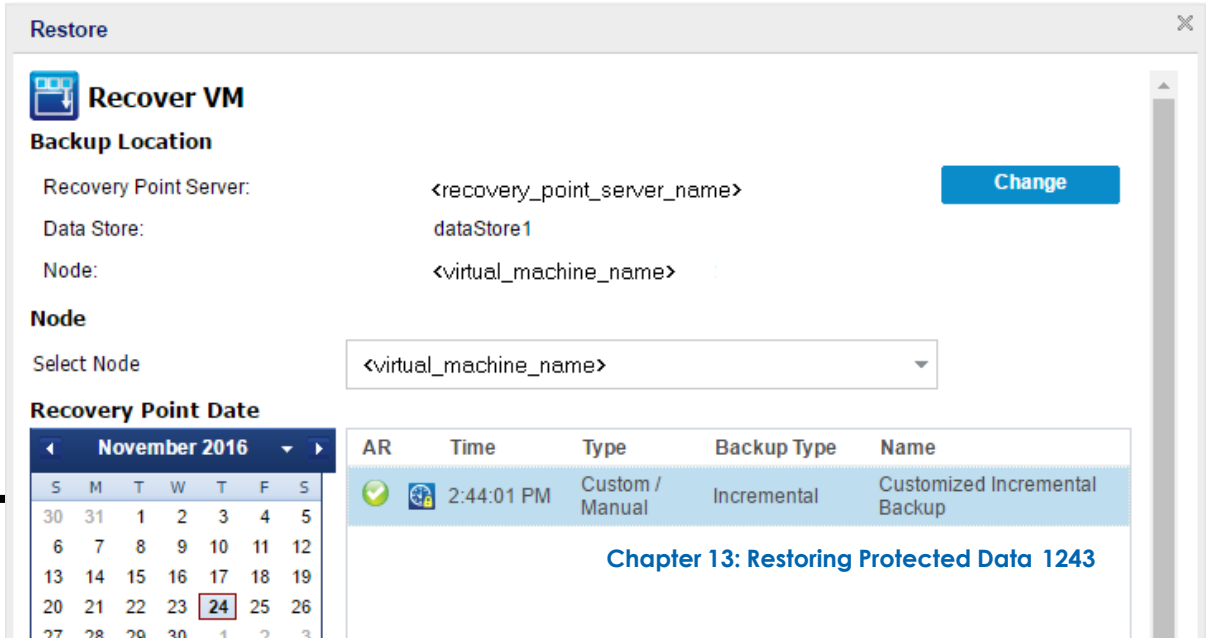
   a. Log into Arcserve UDP Agent (Windows).

   b. From the home page, select **Restore**.

   The restore method selection dialog opens.

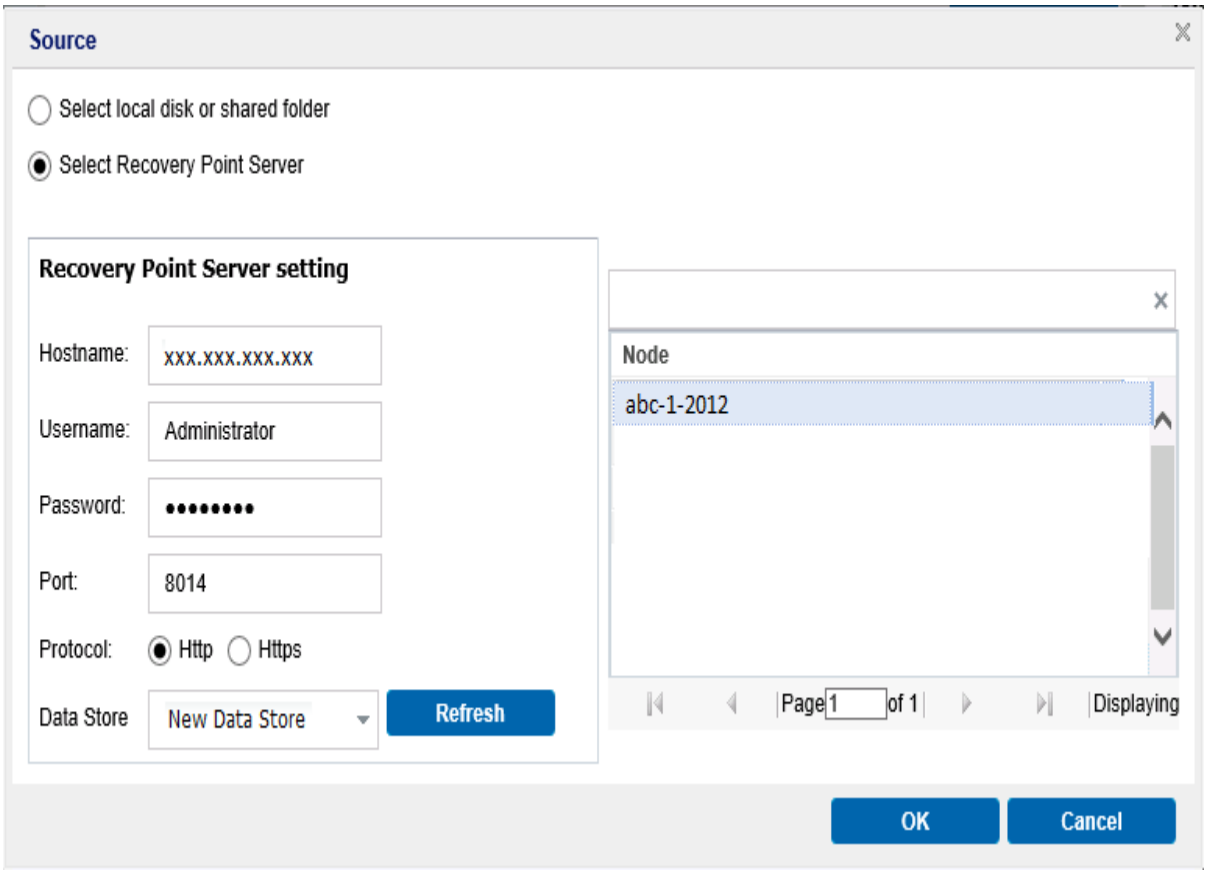
2. Click the **Recover VM** option.

   The **Recover VM** dialog opens.



3. Click **Change** to change the Backup Location.

   The **Source** dialog opens. You can select the backup location in this dialog.



4. Select one of the following options:

   Select local disk or shared folder

   a. Specify or browse to the location where your backup images are stored and select the appropriate backup source.

      You can click the green arrow button to verify the connection to the specified location. If necessary, enter the Username and Password credentials to gain access to that source location.

      The **Select backup location** dialog opens.

   b. Select the folder where the recovery points are stored and click **OK**.

      The **Select backup location** dialog closes and you can see the backup location in the **Source** dialog.

   c. Click **OK**.

   The recovery points are listed in the **Recover VM** dialog.

   Select Recovery Point Server

a. Specify the **Recovery Point Server setting** details and click **Refresh**.

   All the nodes (agents/virtual machines) are listed in the Node column in the **Source** dialog.

b. Select the node (agent/virtual machine) from the displayed list and click **OK**.

   The recovery points are listed in the **Recover VM** dialog.

5. From the **Virtual Machine** drop-down list, select the virtual machine to recover.

   The calendar view appears and all the dates containing recovery points for the specified backup source are highlighted in green.

6. Select the calendar date for the virtual machine image to restore.

   The corresponding recovery points for that date are displayed, with the time of the backup, the type of backup that was performed, and the name of the backup.

7. Select a recovery point to restore.

   The backup content (including any applications) for the selected recovery point is displayed. When restoring a virtual machine, the entire system is restored. As a result, you can view, but not select individual volumes, folders, or files from within the selected virtual machine.

   Note: A clock icon with a lock symbol indicates that the recovery point contains encrypted information and can require a password for restore.

8. Click **Next**.

   The **Restore Options** dialog opens.

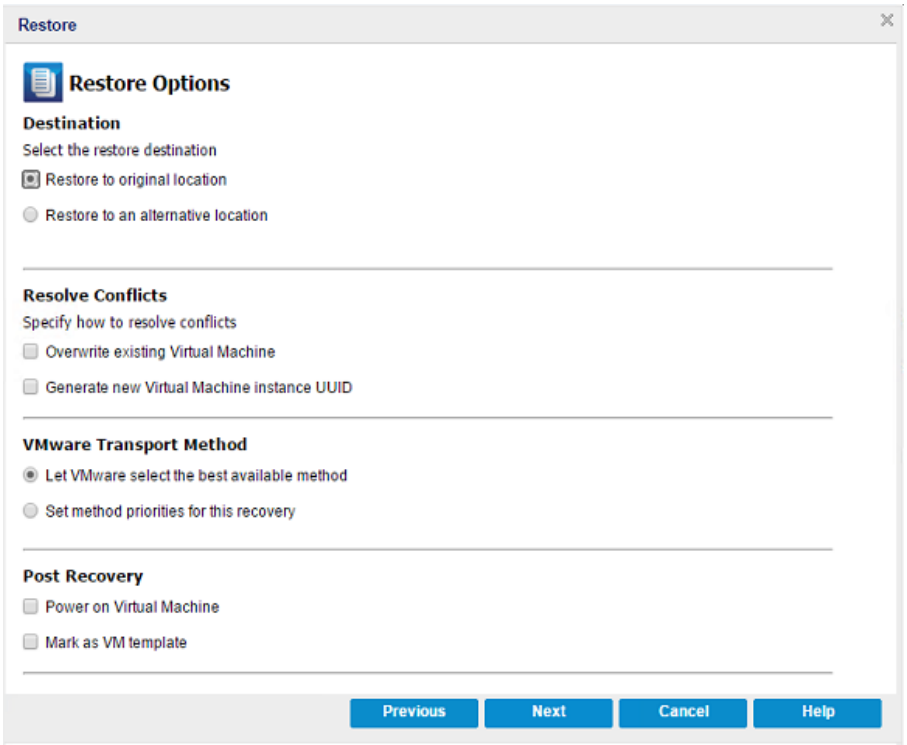
   The virtual machine and the recovery point to restore are specified.

Define the Restore Options

After you specify the virtual machine and the recovery point to restore, define the restore options for the selected virtual machine image.

Follow these steps:

1. From the **Restore Options** dialog, select the restore destination.



The available destination options are:

Restore to Original Location

Restores the virtual machine to the original location from where the backup image was captured. By default, this option is selected.

For more information, see Define the Original Location Restore Options.

Restore to an Alternative Location

Restores to the virtual machine to a different location from where the backup image was captured.

For more information, see Define the Alternate Location Restore Options.

2. Specify the **Resolving Conflicts** options that Arcserve UDP will perform if conflicts are encountered during the restore process

Overwrite existing Virtual Machine

This option is to specify whether to overwrite the existing virtual machine. By default, this overwrite option is not selected.

Note: For the **Overwrite existing Virtual Machine** option, an "existing virtual machine" is defined as a VM which has the same VM name and resides in the same ESXi host (for VMware VM), or a VM which has the same VM name and instance UUID and resides in the same Hyper-V host (for Hyper-V VM). For VMware VM, if there is a VM which has the same VM name but resides in a different ESXi host (which is under the same vCenter), the overwrite option does not work. In this case, VM recovery GUI detects that VM and displays an error message and blocks you from proceeding so that a VM is not overwritten by mistake. As a workaround, you need to either rename the existing VM or use the "Restore to alternative location" option and specify a different VM name.

- If you select this option, the restore process overwrites (replaces) any existing images of this virtual machine that are at the specified restore destination. The virtual machine image is restored from the backup files regardless of its current presence on your restore destination.
- If you do not select this option, VM recovery GUI displays an error message and blocks you from proceeding if the original VM still exists on the original location. You need to either rename the existing VM or use the "Restore to alternative location" option and specify a different VM name.

Generate new Virtual Machine instance UUID

This option is to specify whether to generate a new instance UUID for the restored VM or keep the original instance UUID.

Note: If you do not select this option, then the original instance UUID will be set to the restored VM. However, in case the destination vCenter/ESX or Hyper-V host already has a VM with the same instance UUID, new UUID will be used instead and a warning message is displayed in the activity log of VM recovery job.

3. Specify the **Post Recovery** option.

Power on Virtual Machine

Select whether power is applied to the virtual machine at the end of the restore process. By default, this option is not selected.

Mark as VM Template (available only for VMware VM)

Select whether to convert restored VM to template. If source node is VM when backed up, this option is not selected by default. If source node is template when backed up, this option is selected by default.

The restore options are defined to restore a virtual machine.

Define the Original Location Restore Options

During the Recover VM configuration process, you are required to select the option of where you want to restore the virtual machine to. The available selections are **Restore to the Original Location** and **Restore to an Alternative Location**.

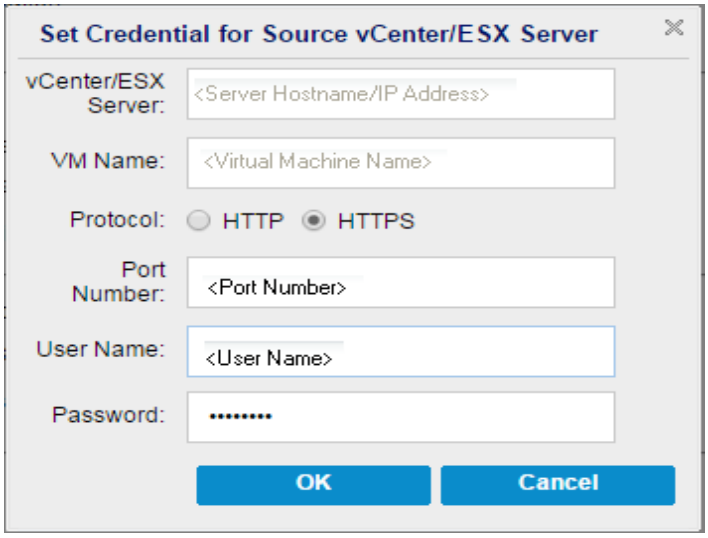
This procedure explains how to restore a virtual machine to the original location.

Follow these steps:

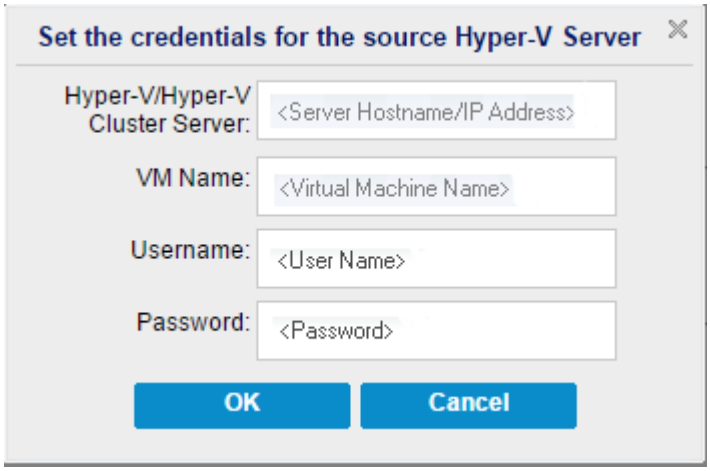
1. From the **Restore Options** dialog, after specifying the **Resolve Conflicts** and **Post Recovery** options, select **Restore to Original Location** and click **Next**.

   The appropriate dialog for VMware or Hyper-V is displayed.

   - For VMware the **Set Credential for Source vCenter/ESX Server** dialog is displayed.



   - For Hyper-V the **Set the credentials for the source Hyper-V Server** dialog is displayed.



2. Specify the credentials for accessing the virtual machine.

   - For VMware, complete the following fields.

     vCenter/ESX Server

     Displays the host name or IP address for the destination vCenter Server or ESX Server system.

     Note: You cannot edit this field. You can only view the details.

     VM Name

     Displays the virtual machine name that you are restoring.

     Note: You cannot edit this field. You can only view the details.

     Protocol

     Specifies the protocol that you want to use for communication with the destination server. The available selections are HTTP and HTTPS.

     Port Number

     Specifies the port that you want to use for data transfer between the source server and the destination.

     Default: 443.

     Username

     Specifies the user name that has access rights to log in to the vCenter/ESX server where you plan to restore the virtual machine.

     Password

     Specifies the corresponding password for the User Name.

   - For Hyper-V, complete the following fields.

     Hyper-V/Hyper-V Cluster Server

     Displays the host name or IP address for the destination Hyper-V Server or Hyper-V cluster server system.

     Note: You cannot edit this field. You can only view the details.

     VM Name

     Displays the virtual machine name that you are restoring.

     Note: You cannot edit this field. You can only view the details.

     Username

     Specifies the user name that has access rights to log in to the Hyper-V server where you plan to restore the virtual machine. For Hyper-V cluster

VM, specify the domain account which has administrative privilege of the cluster.

Password

Specifies the corresponding password for the User Name.

3. Click **OK**.

   The **Restore Summary** dialog opens.

   The restore options for original location are defined.

Define the Alternate Location Restore Options

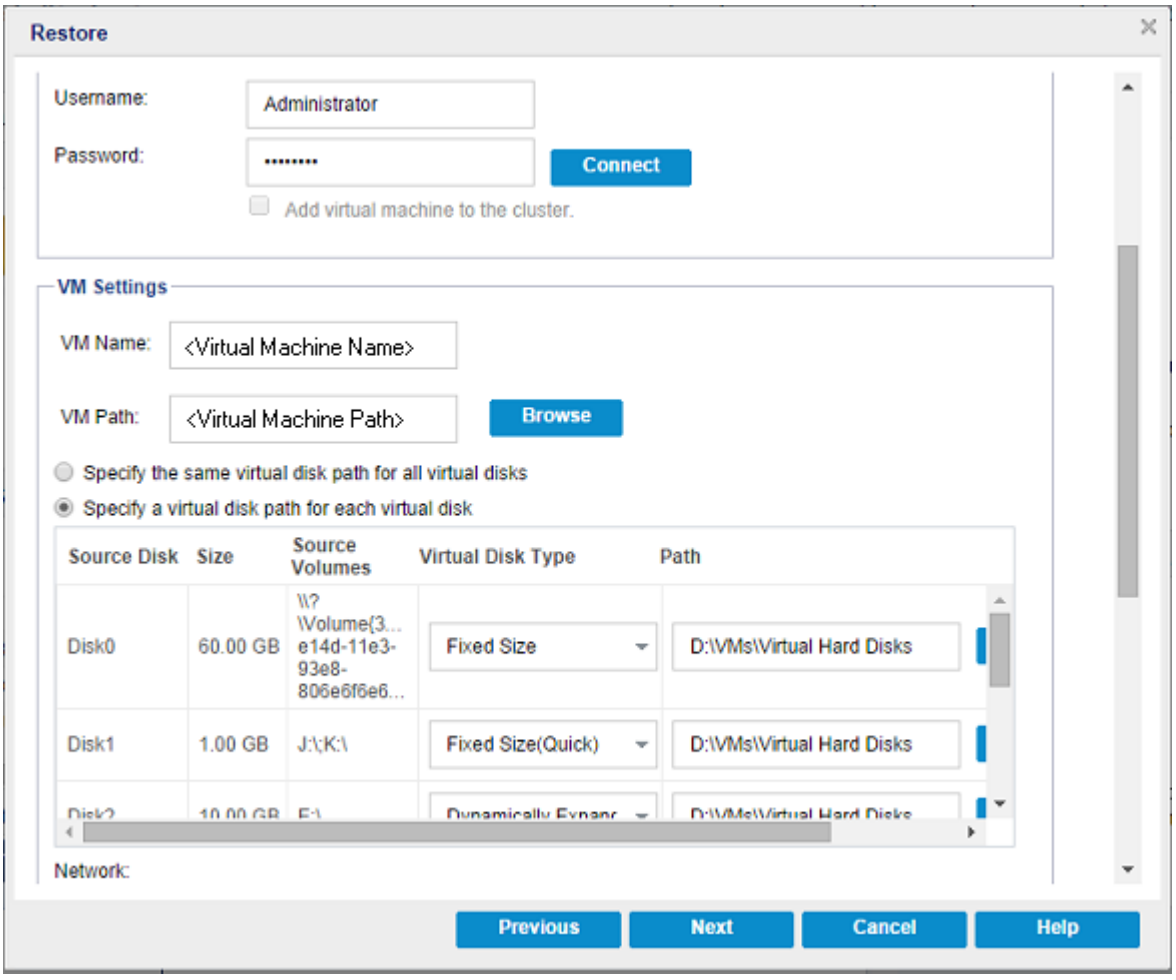
During the Restore VM configuration process, specify where the recovered virtual machine is stored. The available selections are **Restore to the Original Location** and **Restore to an Alternative Location**.

This procedure explains how to restore a virtual machine to alternate location or different data store.

Follow these steps:

1. From the **Restore Options** dialog, after specifying the **Resolve Conflicts** and **Post Recovery** options, select **Restore to an Alternative Location**.
   - For VMware, the **Restore Options** dialog expands to display additional restore to alternative options.
   - For Hyper-V, the **Restore Options** dialog expands to display additional restore to alternative options.

   If you select the **Specify a virtual disk path for each virtual disk** option, the following dialog appears:



2. Specify the appropriate server Information.

   - For VMware, enter the following fields:

     vCenter/ESX Server

     Specifies the host name or IP address for the destination vCenter or ESX server system.

     Username

     Specifies the user name that has access rights to log in to the vCenter/ESX server where you plan to restore the virtual machine.

     Password

     Specifies the corresponding password for the User Name.

     Protocol

     Specifies the protocol that you want to use for communication with the destination server. The available selections are HTTP and HTTPS.

     Default: HTTPS.

     Note: VMware Virtual Disk Development Kit (VDDK) 6.x.x is in-built with Arcserve UDP 6.5 but VDDK 6.x.x does not support HTTP. Make sure to select HTTPS, unless you manually replace the built-in VDDK 6.x.x with another version of VDDK.

     Port Number

     Specifies the port that you want to use for data transfer between the source server and the destination.

     Default: 443.

   - For Hyper-V, enter the following fields:

     Hyper-V Server

     Displays the host name or IP address for the destination Hyper-V Server system.

     Username

     Specifies the user name that has access rights to log in to the Hyper-V server where you plan to restore the virtual machine. For Hyper-V cluster VM, specify the domain account that has administrative privilege of the cluster.

     Password

     Specifies the corresponding password for the User Name.

Add virtual machine to the cluster

Select the option if you want to add the virtual machine that Arcserve UDP restores, into the cluster. Consider the following options:

- If you provide the cluster node name as the Hyper-V server name, the check box is disabled and checked by default. As a result, the virtual machine is automatically added into the cluster.
- If you provide the host name of a Hyper-V server that is part of the cluster the check box is enabled and you can select to add the virtual machine into the cluster.
- If you provide the host name of a standalone Hyper-V server that is not part of the cluster the check box is disabled and unchecked.

3. When the vCenter/ESX Server Information or Hyper-V Server Information is specified, click the **Connect to this vCenter/ESX Server** button or click the **Connect to this Hyper-V Server** button.

   If the alternative server access credential information is correct, the **VM Settings** fields become enabled.

4. Specify the **VM Settings**.

   - For VMware, enter the following fields.

     VM Name

     Specifies the virtual machine name that you are restoring.

     ESX Server

     Specifies the destination ESX server. The drop-down menu contains a listing of all ESX servers that are associated with a vCenter server.

     Resource Pool

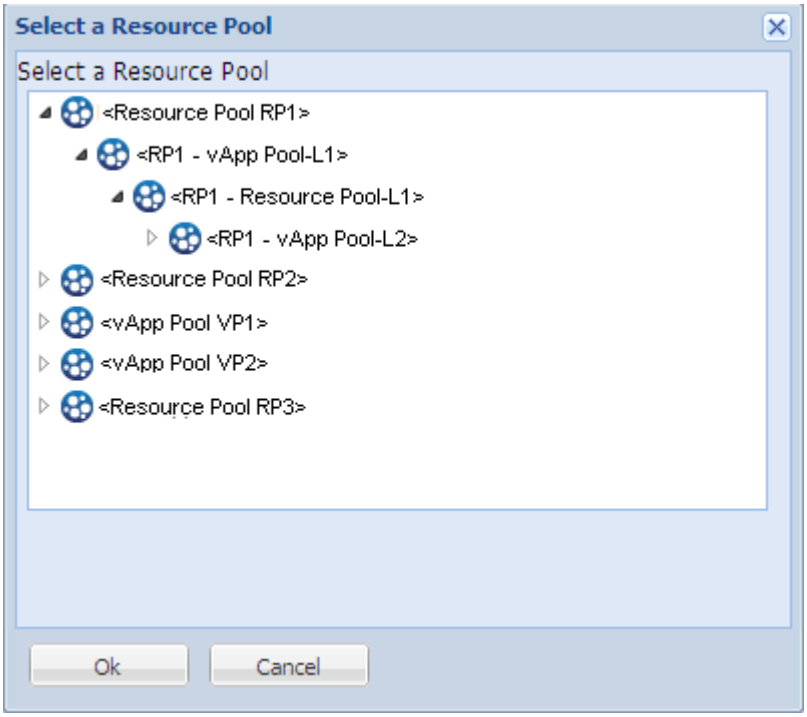
     Selects the **Resource Pool** or **vApp Pool** you want to use for the virtual machine recovery.

     Note: A Resource Pool is a configured collection of CPU and memory resources. A vApp Pool is a collection of one or more virtual machines that can be managed as a single object.

     Default: empty.

     Click the **Browse Resource Pool** button to display the **Select a Resource Pool** dialog. This dialog contains a listing of all Resource Pools and vApp Pools available for the destination ESX server. Select the pool to use for the virtual machine recovery. You can leave this field blank when you do

not want to assign a Resource Pool or vApp Pool to this virtual machine recovery.



Storage Policy

Specify the VM storage policy that is applied to VM home of restored VM. Select Datastore Default if you do not want to apply the VM storage policy.

Note: if you can see only Datastore Default but actually there are other storage policies defined in vCenter, then the account used to connect vCenter does not have enough permission to get the storage policy from vCenter. Verify if the account has the privilege Profile-driven storage view at the vCenter level.

VM DataStore

Specify the destination datastore for VM home of restored VM.

Note: By default, only those datastores that are compatible with selected storage policy are listed. If you want to see all datastores, clear selection of the checkbox **Show only compatible datastores for selected storage policy** that in under the Disk Datastore table.

Disk Datastore

Specify Virtual Disk Type, Storage Polic,y and Target Datastore for each of the virtual disk of the VM, respectively.

- Virtual Disk Type: Select one of the following options: Thin, Thick Lazy Zeroed, or Thick Eager Zeroed.
- Storage Policy: Select the VM storage policy that is applied to this virtual disk. Select Datastore Default if you do not want to apply the VM storage policy.
- Target Datastore: Select the datastore where the virtual disk is restored.

Note: By default only the datastores that are compatible with selected storage policy are listed. If you want to see all datastores, clear selection of the checkbox **Show only compatible datastores for selected storage policy** that in under the Disk Datastore table.

Network

Specifies the vSphere Standard Switch/vSphere Distributed Switch configuration details.

- For Hyper-V, enter the following fields.

VM Name

Specifies the virtual machine name that you are restoring.

VM Path

Specifies the destination path (on Hyper-V server) where to save the Hyper-V VM configuration file. The default folder of the VM configuration file for the Hyper-V server is shown by default. You can modify the path directly in the field or click **Browse** to select one.

Note: If you are restoring the virtual machine into Hyper-V cluster and you want the virtual machine to migrate among the cluster nodes, specify the cluster shared volume (CSV) for both- the VM path and the virtual disk path.

Specify the same virtual disk path for all virtual disks

Specify one path (on Hyper-V server) where to save all virtual disks of the VM together. The default folder of the VM disk file for the Hyper-V server is shown by default. You can modify the path directly in the field or click **Browse** to select one.

Note: If you are restoring the virtual machine into Hyper-V cluster and you want the virtual machine to migrate among the cluster nodes, specify the cluster shared volume (CSV) for both- the VM path and the virtual disk path.

Specify a virtual disk path for each virtual disks

Specify the path (on Hyper-V server) for each of the virtual disks of the VM respectively. The default folder of the VM disk file for the Hyper-V server is shown by default. You can modify the path directly in the field or click **Browse** to select one. To assign the virtual disk type, select one of the following options: Fixed Size, Fixed Size (Quick), Dynamically Expanding, and Keep same as Source disk.

Notes:

- If you are restoring the virtual machine into Hyper-V cluster and you want the virtual machine to migrate among the cluster nodes, specify the cluster shared volume (CSV) for both- the VM path and the virtual disk path.
- Do not use Fixed Size (Quick) option unless you are sure that earlier you have not saved sensitive information on the storage device where the virtual disk file resides.

Fixed Size (Quick)

Using this option, you can restore Fixed Size disk in a quicker way. You do not need to clear unused disk blocks to zero while restoring the disk. However, because of this, some fragments of original data remained on underlying storage. That situation creates risks of information leaks. After the disk is mounted into the virtual machine, the user of the virtual machine may use some disk tools to analyze the raw data in the disk and get the original data on Hyper-V server storage device where the file of virtual disk resides.

Network

Specifies the network configuration details for the VM.

5. Click **OK**.

   The **Restore Summary** dialog opens.

   The restore options for alternate location are defined.

Restore the Virtual Machine

The **Restore Summary** helps you to review all the restore options that you defined and modify them if necessary.

On the **Restore Summary** dialog, review the displayed information to verify that all the restore options and settings are correct.

- If the summary information is incorrect, click **Previous** and go back to the applicable dialog to change the incorrect setting.
- If the summary information is correct, click **Finish** to launch the restore process.

The virtual machine is restored.

Verify that the Virtual Machine was Restored

After the completion of the restore process, verify that the virtual machine was restored to the specified destination.

Follow these steps:

1. Navigate to the restore destination you specified.

   For example, if you select to restore the virtual machine to the restore destination as original location, then log in to the original vCenter/ESX or Hyper-V Server and check if the virtual machine exists.

   If you select to restore the virtual machine to the Alternate location, then log in to the alternate vCenter/ESX or Hyper-V Server provided in the restore options and check if the virtual machine exists.

2. Verify the virtual machine was restored.

   The virtual machine is restored successfully.

How to Use Exchange Granular Restore (GRT) Utility

This section contains the following information about the Exchange Granular Restore (GRT) Utility:

Introduction

Review the Prerequisites and Considerations

How to Restore Microsoft Exchange Data Using Exchange Granular Restore (GRT) Utility

Introduction

The Exchange Granular Restore utility is used to restore Microsoft Exchange email and non-email objects. The utility includes the injection capability for items, such as emails, from offline databases (*.EDB) and log files to the original live Exchange databases, as well as granular data extraction to Personal Storage File (.pst) files.

This utility includes the following key benefits:

- Supports non-email items (for example, Calendar, Contacts, Tasks) and public folders.
- Can work with just a database file as well. Logs are not mandatory, but having them will ensure more recent data available for restore.
- It does not need to generate a catalog and directly restores the mail from the mounted recovery point.
- Takes a minimum amount of time to restore a mailbox level item from a database or user mailbox of any size.
- Supports the command line options to process several databases.

Note: For more details on the supported specifications, functionalities, and other features, see the Exchange Granular Restore user guide (esr.pdf).

Review the Prerequisites and Considerations

Verify that the following prerequisites exist before performing a restore:

- The Exchange Granular Restore utility is available at the following location:

  The tool is installed with the Arcserve UDP Agent under the following directory:

  ```
  X:\Program Files\Arcserve\Unified Data Protection\Engine\Exchange GRT
  ```

  Note: The tool is installed with the Arcserve UDP Agent.

- The Restore job is set to run from the Exchange machine or HBBU proxy machine.

  Note: If you want to run the restore job on any other machine, search the recovery point from the backup destination.

- The database name, Server name, path to database (.edb), and the log files of the user are identified to perform the restore job.

  To identify, use the Exchange Management Console (EMC), Exchange Control Panel (ECP), or Exchange Management Shell.

  For example:

  ```
  Get-Mailbox -identity "username" | fl Database
  ```

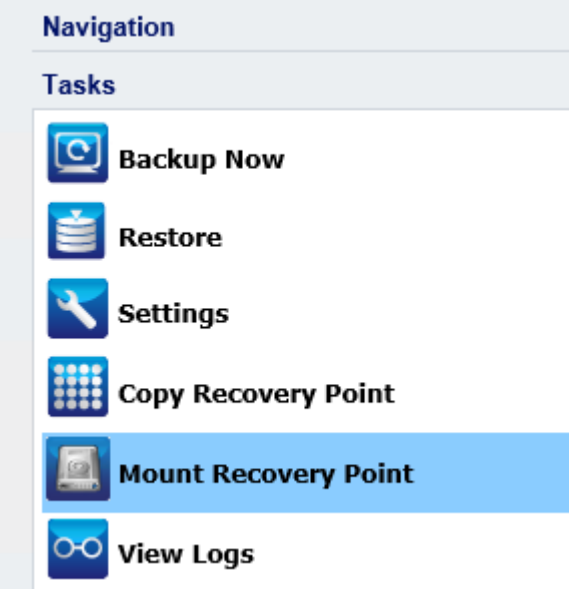
  ```
  Get-MailboxDatabase -identity "Databasename" | fl Name, Server, EdbFilePath,LogFolderPath
  ```

How to Restore Microsoft Exchange Data Using Exchange Granular Restore (GRT) Utility

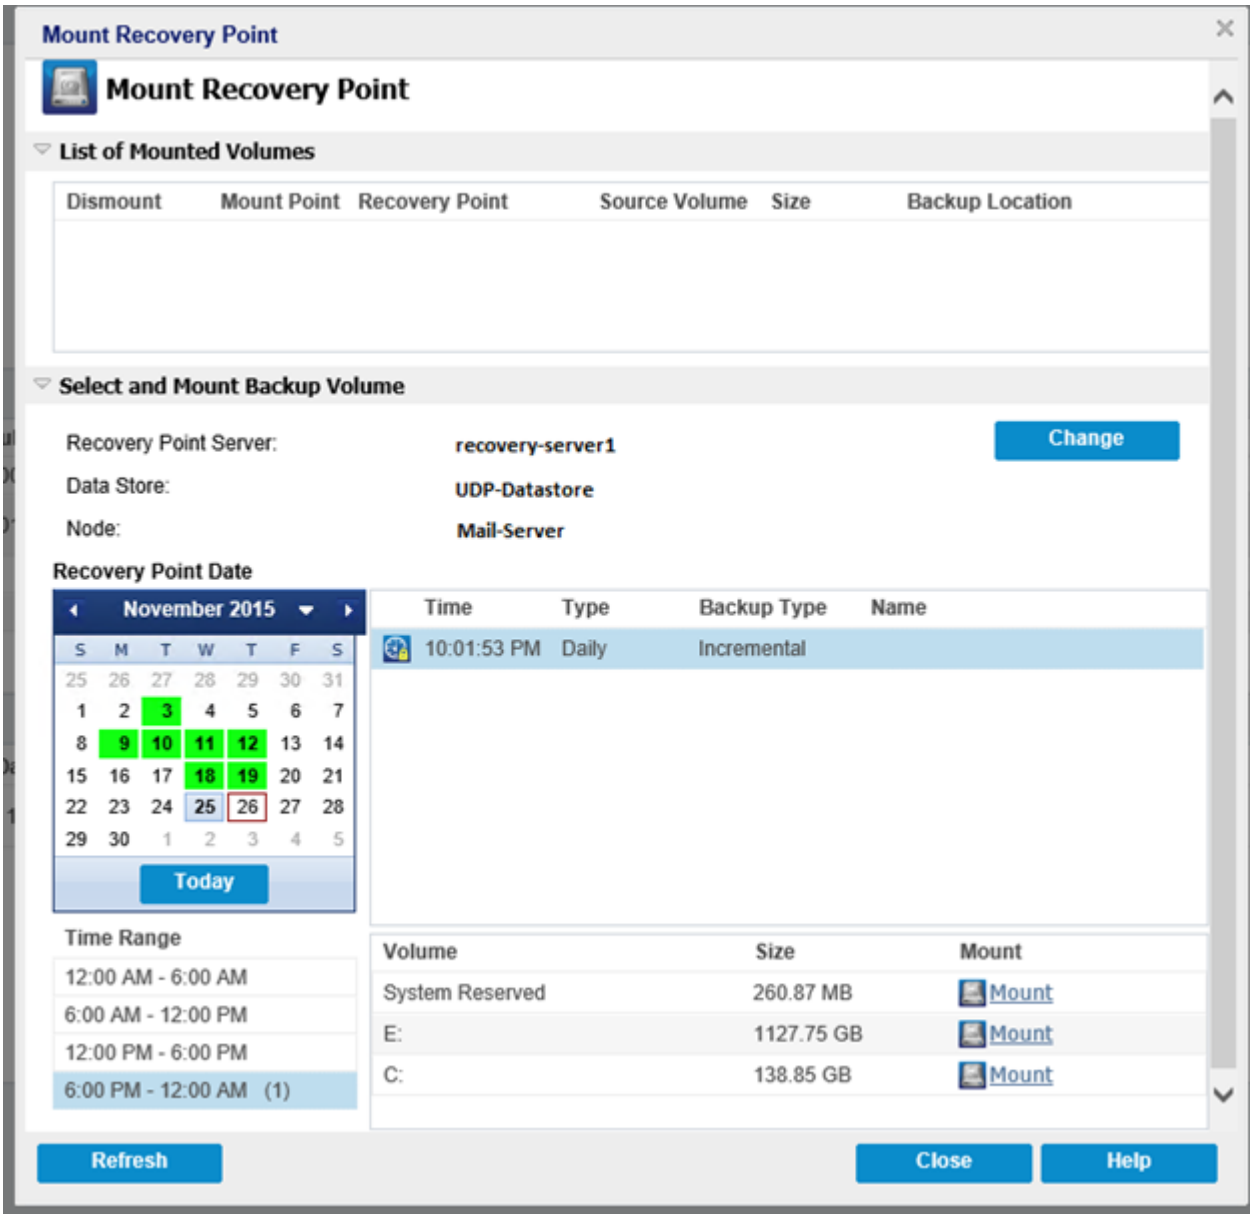
Before you begin, review the prerequisites and considerations.

Perform the following tasks to restore Microsoft Exchange mailbox items, using the Exchange Granular Restore utility:

1. From the Arcserve UDP Agent console, select the Mount Recovery Point task (recommended) or restore the Exchange database to the local drive. The Mount Recovery Point dialog opens.

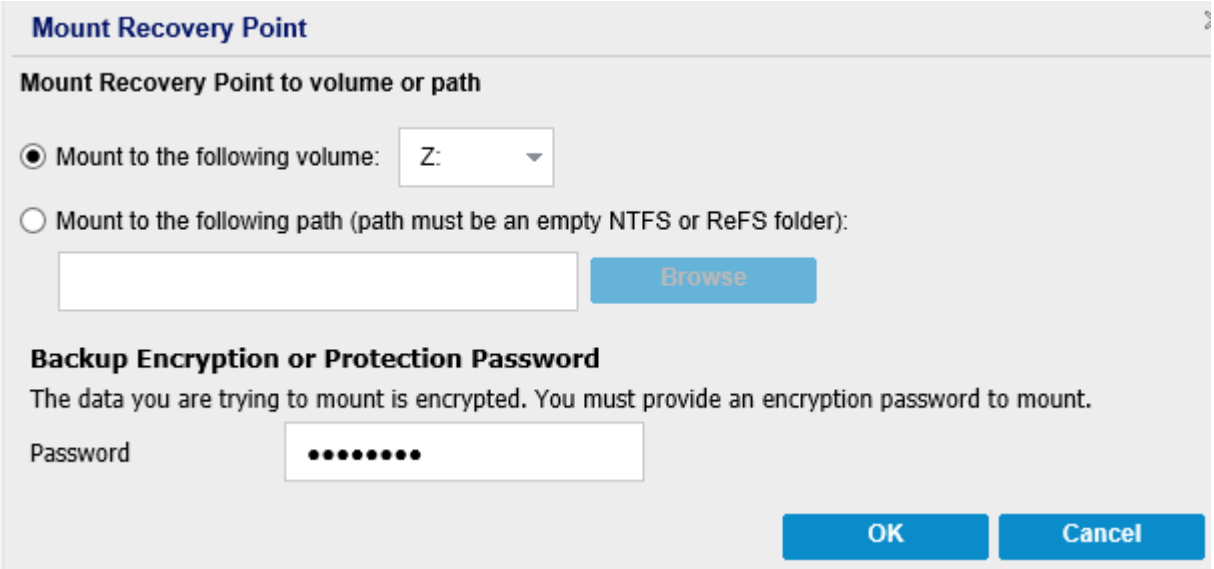


2. Select the recovery point date and click **Mount** for the volume(s) that contain Exchange Database and logs.



Note: If the server that is running the restore job is not the Exchange or HBBU proxy, click **Change** to select the appropriate Recovery Point Server, Data Store, and Exchange Server.

3. Select the drive letter to mount the volume and click **OK**.



4. Launch the Exchange Granular Restore utility from one of the following locations:

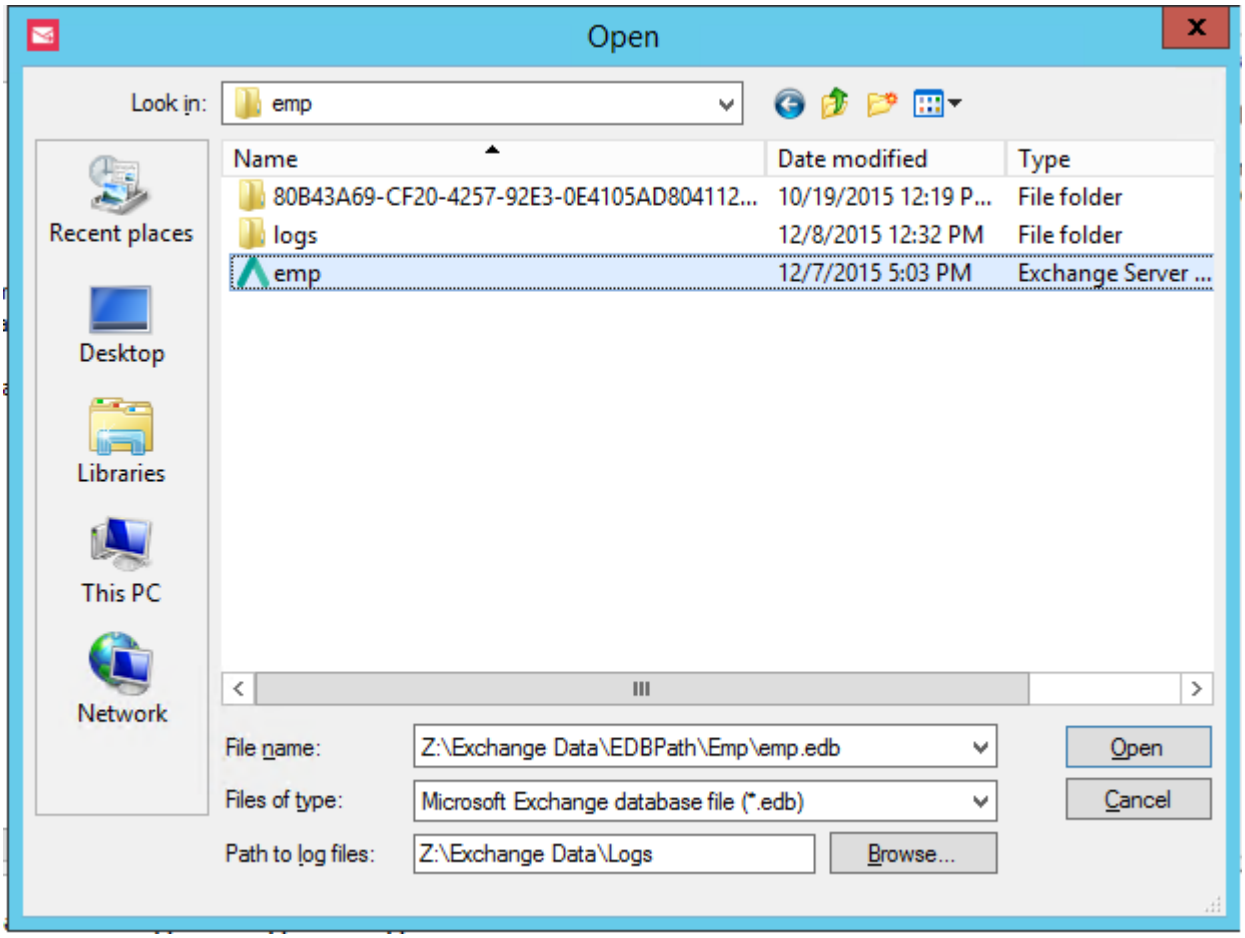
   Start > All Programs > Arcserve > Unified Data Protection > Arcserve UDP Exchange Granular Restore

   or

   X:\Program Files\Arcserve\Unified Data Protection\Engine\Exchange GRT\esr.exe

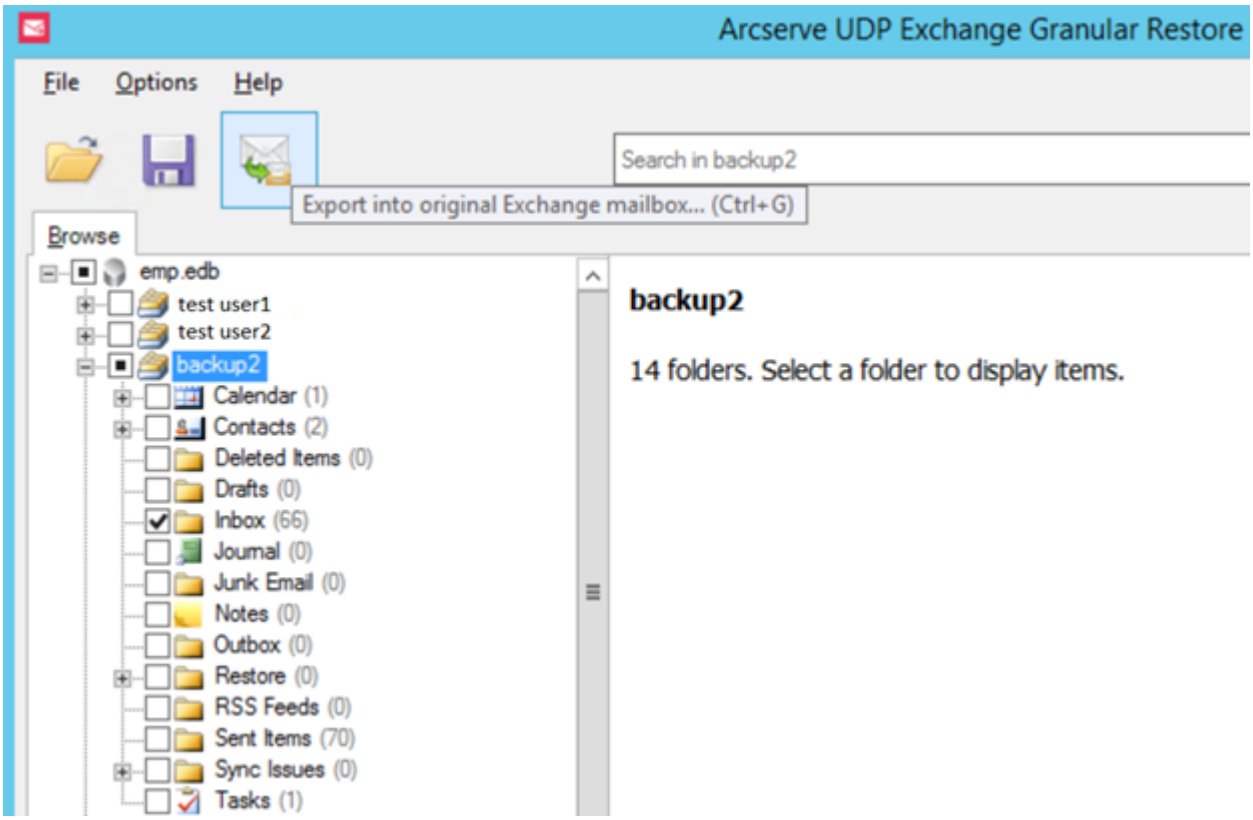
   A dialog appears to specify the path for the database and log files.

5. Specify the path to the mounted volume and click **Open**.



The Arcserve UDP Exchange Granular Restore utility opens.

6. Select the user data to restore and click **Export into original mailbox** or **Export into .PST**.



Notes:

- For more details on the supported specifications, features, user options and limitations, see the Exchange Granular Restore user guide (esr.pdf), located at:

  %ProgramFiles%\Arcserve\Unified Data Protection\Engine\Exchange GRT or Bookshelf.

- By default, the utility uses the current user who is logged in to Windows to establish the connection. If the current user does not have permissions to impersonate the selected user, an error message appears in the **Details** pane.

  If an error is reported, the recommended action to take is to log in to the machine with an account that has impersonation rights for the selected user or the account of the selected user.

7. When the restore job completes, dismount the volume that was used for the recovery.

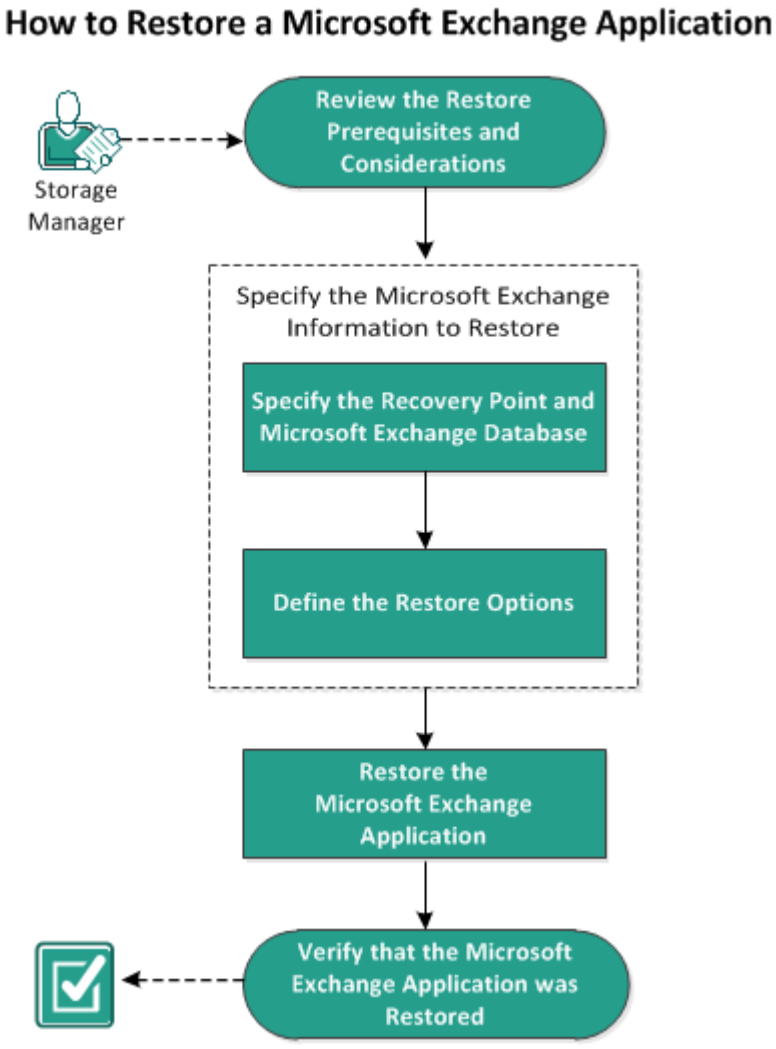
   To dismount the volume, from the Arcserve UDP Agent console, click **Mount Recovery Point** and then click **Dismount**.


Mount Recovery Point
Mount Recovery Point
List of Mounted Volumes
Dismount
Mount Point
Recovery Point
Source Volume
Size
Backup Location
Dismount
Z:
11/25/2015 10:01:53...
E:
1127.75 GB
\\BackupServer\CAU.....

How to Restore a Microsoft Exchange Application

Arcserve UDP Agent (Windows) allows you to not only protect and recover your data, but also helps you to get the applications that will use that data back up and running. All application recoveries can only be made using the Restore by Recovery Point method. During an application recovery, Arcserve UDP Agent (Windows) takes advantage of Windows Volume Shadow Copy Service (VSS) to help ensure data consistency for any VSS-aware application. With Arcserve UDP Agent (Windows), you can recover the Microsoft Exchange Server application without performing a full disaster recovery.

The following diagram illustrates the process to restore a Microsoft Exchange Application:



Perform the following tasks to restore a Microsoft Exchange Application:

1. Review the Restore Prerequisites and Considerations
2. Specify the Microsoft Exchange Information to Restore
    a. Specify the Recovery Point and Microsoft Exchange Database
    b. Define the Restore Options
3. Restore the Microsoft Exchange Application
4. Verify that the Microsoft Exchange Application was Restored

Review the Restore Prerequisites and Considerations

Arcserve UDP Agent (Windows) supports the following versions of Microsoft Exchange Server:

- Microsoft Exchange 2010 - Single Server Environment and Database Availability Group (DAG) environment.
- Microsoft Exchange 2013 and 2016 - Single Server Environment and Database Availability Group (DAG) environment.

  For Microsoft Exchange Server 2010, 2013, and 2016 DAG environment, Arcserve UDP Agent (Windows) must be installed on all member servers in the DAG group. A backup job can also be performed from any member server for both active and passive database copies, but restore can only be performed to an active database copy.
- Although all DAG members can be part of same or different backup plan, we recommend to use same deduplication data store to eliminate duplicate data.

You can restore Microsoft Exchange Server at the following levels:

Microsoft Exchange Writer Level

Defines if you want to restore all the Microsoft Exchange Server data, you can perform a restore at Microsoft Exchange Writer level.

Storage Group Level

Defines if you want to restore a specific Storage Group, you can perform a restore at this level.

Note: The Storage Group Level does not apply for Microsoft Exchange Server 2010, 2013, and 2016.

Mailbox Database Level (Microsoft Exchange 2010, 2013, and 2016)

Specifies if you want to restore a specific Mailbox Database, you can perform a restore at this level.

Mailbox Level (Microsoft Exchange 2010, 2013, and 2016)

Defines if you want to restore a specific Mailbox or mail object.

Verify that the following prerequisites exist before performing a Microsoft Exchange restore:

Database-level restore

- The target machine has the same name and the same version of Microsoft Exchange installed.

- The target database has the same database name and the same storage group name (Microsoft Exchange 200X) and be a part of the same Microsoft Exchange organization.

Granular-level restore

- To restore Microsoft Exchange data, use the Exchange Granular Restore utility.

Specify the Microsoft Exchange Information to Restore

Arcserve UDP Agent (Windows) allows you to not only protect and recover your data, but also helps you to get the Microsoft Exchange Server application that uses that data back up and running. The Microsoft Exchange Server recovery can only be made using the Restore by Recovery Point method.

The process involved in restoring a Microsoft Exchange Application is as follows:

1. Specify the Recovery Point and Microsoft Exchange Database
2. Define the Restore Options

Specify the Recovery Point and Microsoft Exchange Database

Use the **Browse Recovery Points** option to restore from a recovery point. When you select a recovery date, all the associated recovery points for that date are displayed. You can then browse and select the Microsoft Exchange database to be restored.

Follow these steps:

1. Access the restore method selection dialog in one of the following ways:

   From Arcserve UDP:

   a. Log into Arcserve UDP.

   b. Click the **resources** tab.

   c. Select **All Nodes** in the left pane.

      All the added nodes are displayed in the center pane.

   d. In the center pane, select the node and click **Actions**.

   e. Click **Restore** from the **Actions** drop-down option.

      The restore method selection dialog opens.

      Note: You are automatically logged in to the agent node and the restore method selection dialog is opened from the agent node.

   From Arcserve UDP Agent (Windows):

   a. Log into Arcserve UDP Agent (Windows).

   b. From the home page, select **Restore**.

   The restore method selection dialog opens.

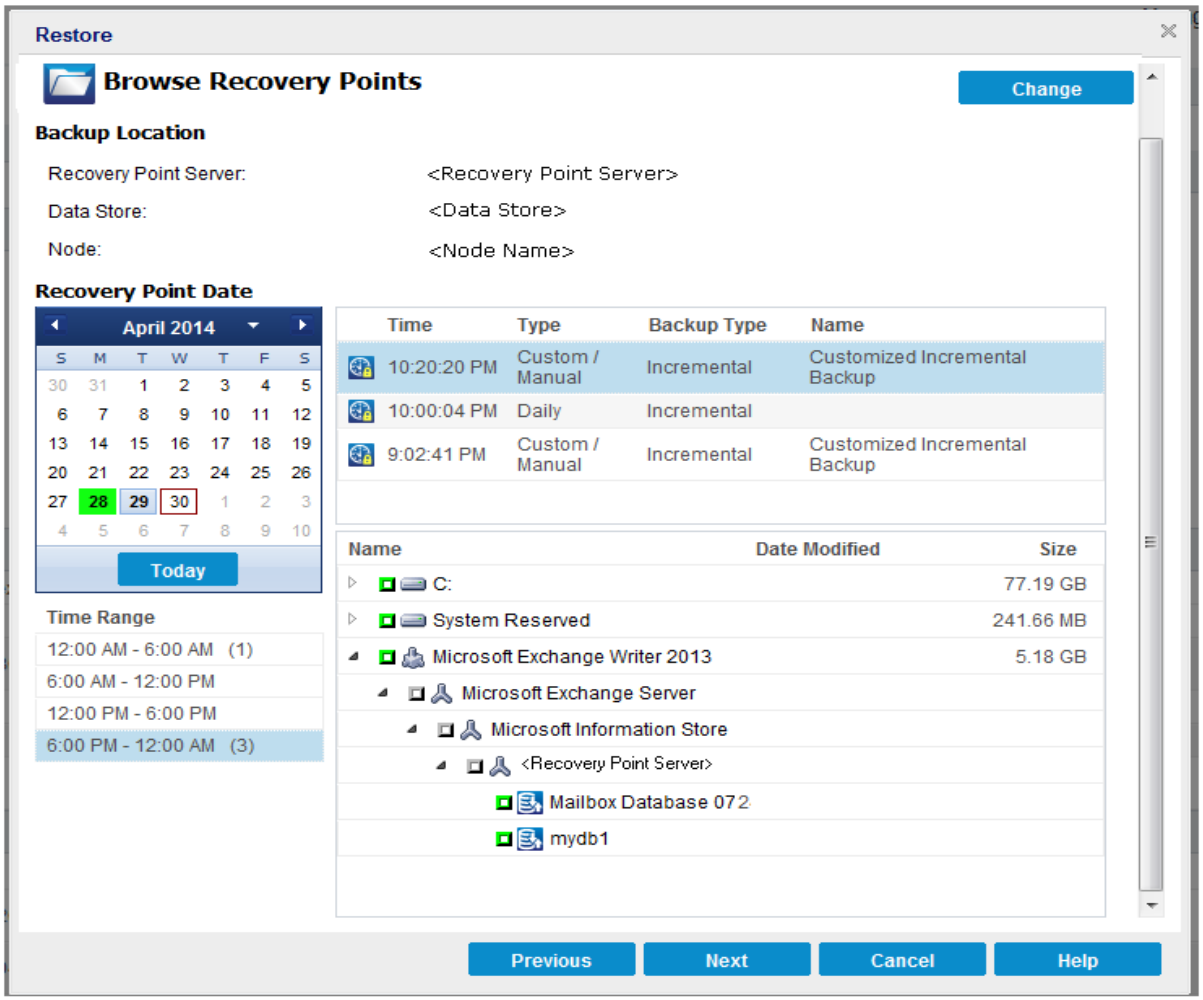
2. Click the **Browse Recovery Points** option.

   The **Browse Recovery Points** dialog opens.

3. Select the recovery point (date and time) and then select the Microsoft Exchange database to be restored.

   The corresponding marker box becomes filled (green) to indicate that the database has been selected for the restore.

   Note: If you do not want the transaction log files to be applied after the restore, you must manually delete it before the restore is performed. For more information about manually deleting transaction log files, refer to the Microsoft Exchange Server documentation.



4. Click **Next**.

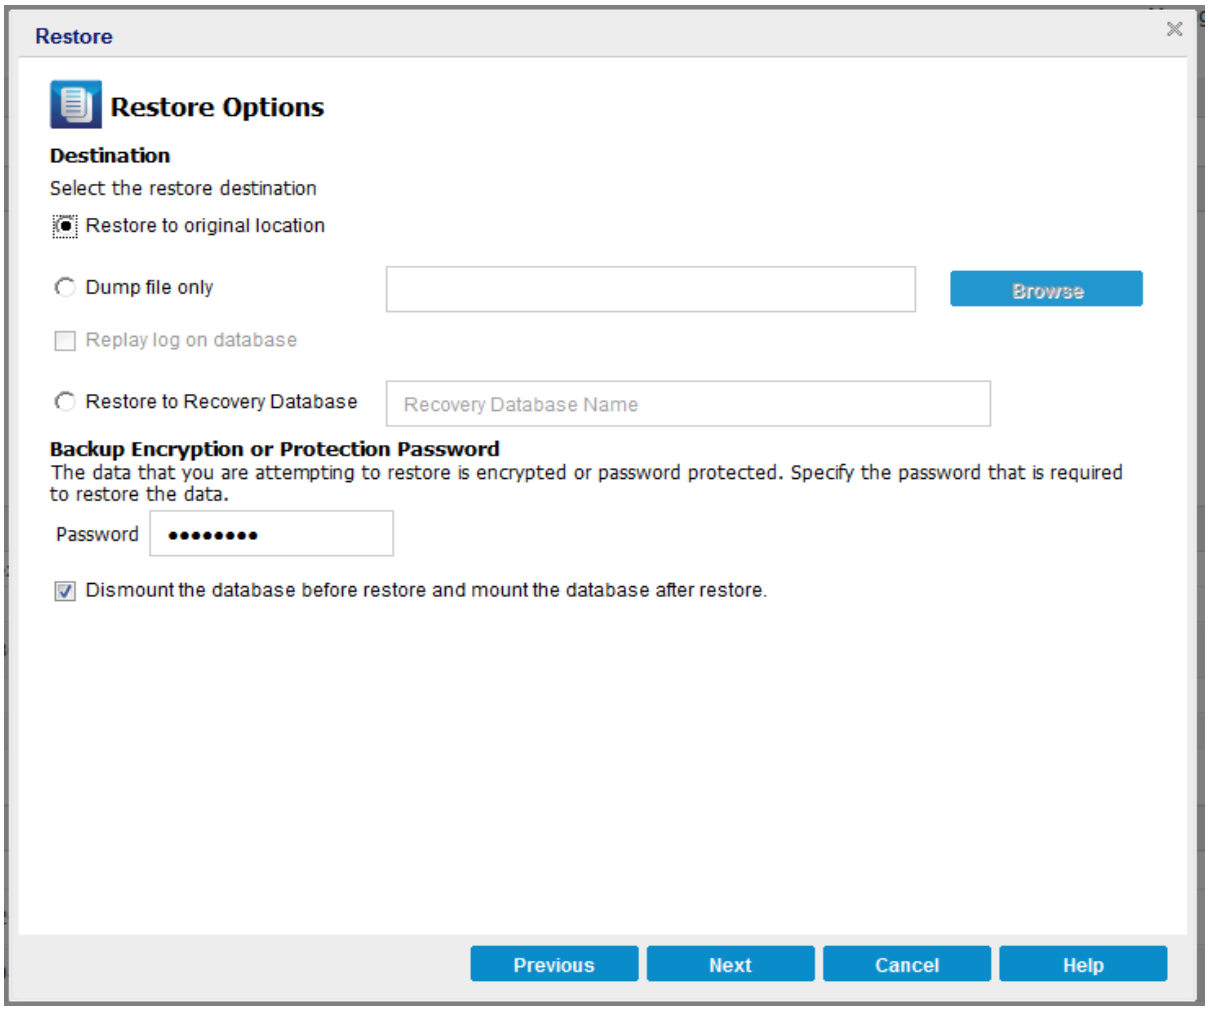
   The **Restore Options** dialog opens.

Define the Restore Options

After you specify a recovery point and content to restore, define the copy options for the selected recovery point.

Follow these steps:

1. From the **Restore Options** dialog, select the restore destination.



2. Select the destination for the restore.

The available options are to restore to the original location of the backup, restore the dump file only, or restore to a Recovery Storage Group/Recovery Mailbox Database.

Restore to original location

Restores to the original location from where the backup image was captured.

Dump file only

Restores the dump files only.

For this option, Arcserve UDP Agent (Windows) will restore the Microsoft Exchange database file to a specified folder, and will not bring it online after recovery. You can then use it to mount on Microsoft Exchange Server manually.

Note: When a Recovery Mailbox Database exists, restore with **Dump file only** option will fail.

Replay log on database

Specifies that when the database files are dumped to the destination folder, you can replay Microsoft Exchange transaction log files and commit them to the database.

Dismount the database before restore and mount the database after restore

Typically before a restore, Microsoft Exchange will perform some checks to help ensure the following:

- The database to be restored is in *Dismounted* status.
- The database is not restored unexpectedly.

To protect a Microsoft Exchange production database from being restored unexpectedly, a switch is added to allow the database to be overwritten during the restore process. Microsoft Exchange will refuse to restore a database if this switch is not set.

For Arcserve UDP Agent (Windows), these two options are controlled by this "Dismount the database before restore and mount the database after restore" option. With this option, Arcserve UDP Agent (Windows) lets you launch the restore process automatically without any manual operations. (You can also specify to dismount/mount database manually).

- If checked, specifies that the recovery process will automatically dismount the Microsoft Exchange database before the restore process and then mount the database after the restore process is completed. In addition, if checked, this option will also allow the Microsoft Exchange database to be overwritten during the restore.

- If unchecked, specifies that the recovery process will not automatically dismount the Microsoft Exchange database before recovery and mount the database after recovery.

  The Microsoft Exchange administrator would have to perform some manual operations such as dismount the Microsoft Exchange database, set the Allow Overwrite flag on the database, and mount the Microsoft Exchange database. (The recovery procedure is performed by Exchange during the mounting of the database).

  In addition, if unchecked, this option does not allow the Microsoft Exchange database to be overwritten during restore.

Restore to Recovery Database (Microsoft Exchange 2010 and 2013)

Restores the database to a Recovery Database. A Recovery Database is a database that can be used for recovery purposes. You can restore a Microsoft Exchange Mailbox Database from a backup to a Recovery Database and then recover and extract data from it, without affecting the production database that is being accessed by end users.

Before restoring a Microsoft Exchange 2010 or Exchange 2013 database to a Recovery Database, you must first create a Recovery Database.

3. Click **Next**.

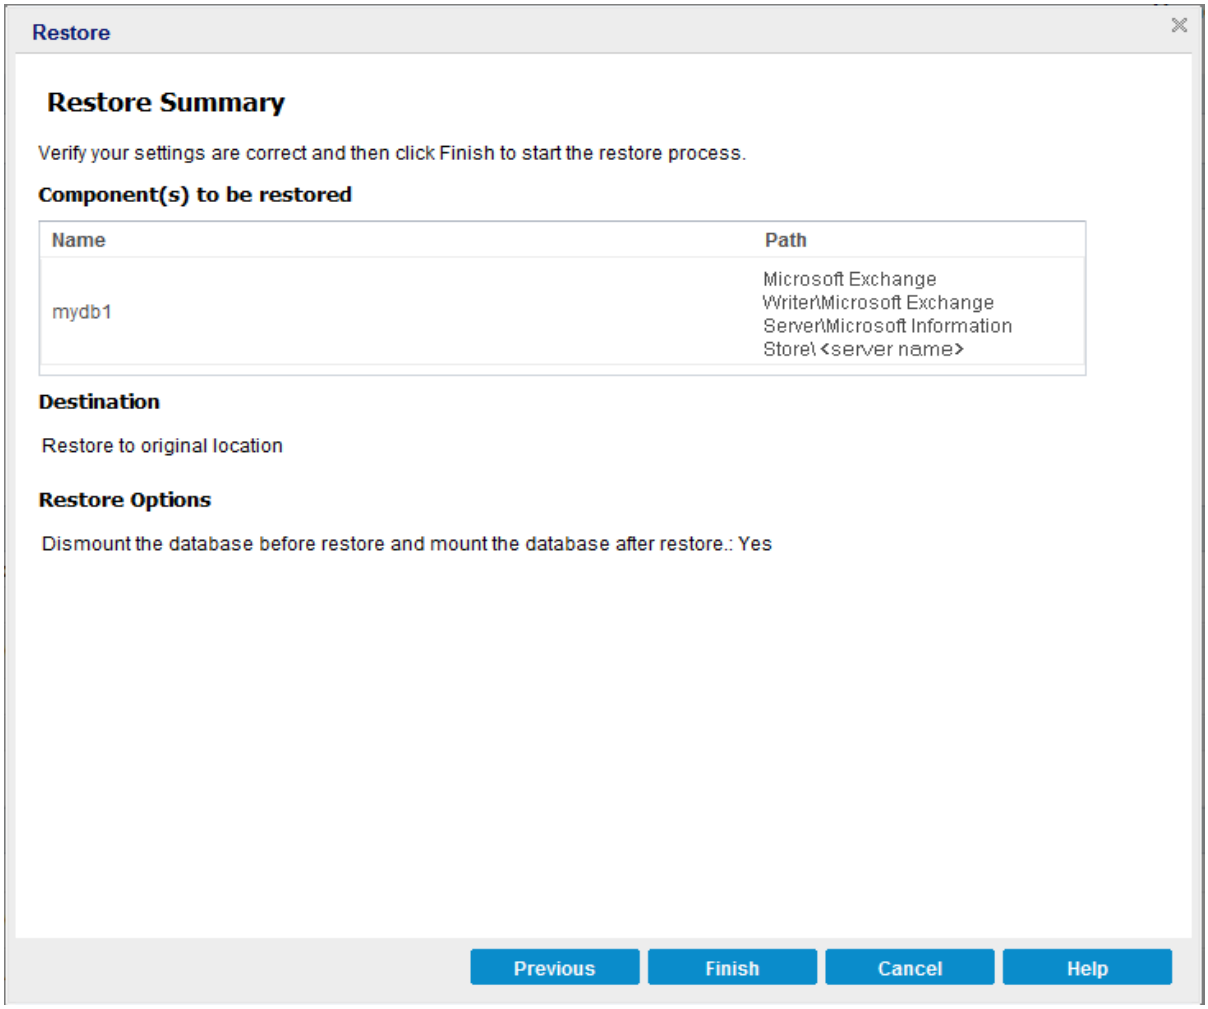
   The **Restore Summary** dialog opens.

Restore the Microsoft Exchange Application

After you define the restore options, verify that your settings are correct and confirm the restore process. **Restore Summary** helps you to review all the restore options that you defined and modify them if necessary.

Follow these steps:

1. From the **Restore Summary** dialog, review the displayed information to verify that all the restore options and settings are correct.



- If the summary information is not correct, click **Previous** and go back to the applicable dialog to change the incorrect setting.
- If the summary information is correct, click **Next** and then **Finish** to launch the restore process.

The Microsoft Exchange Application is restored.

Verify that the Microsoft Exchange Application was Restored

Follow these steps:

1. Navigate to the Arcserve UDP Agent (Windows) restore destination you specified.

   For example, if you select to restore the Microsoft Exchange database to the original location, after the restore is complete, then browse to the physical location to check if the Microsoft Exchange database and logs are restored.

   If you select to restore the Microsoft Exchange database to Dump File only location then Arcserve UDP Agent (Windows) will restore the Microsoft Exchange database and logs to a specified location.

2. Verify if the Microsoft Exchange Application was restored and check if the database is mounted and is accessible.

   The Microsoft Exchange Application is restored successfully.

How to Restore Exchange Data on a VMware Virtual Machine

Important! To restore Microsoft Exchange data on a VMware virtual machine, we recommend to use the Exchange Granular Restore utility.

How to Download File/Folders without Restore

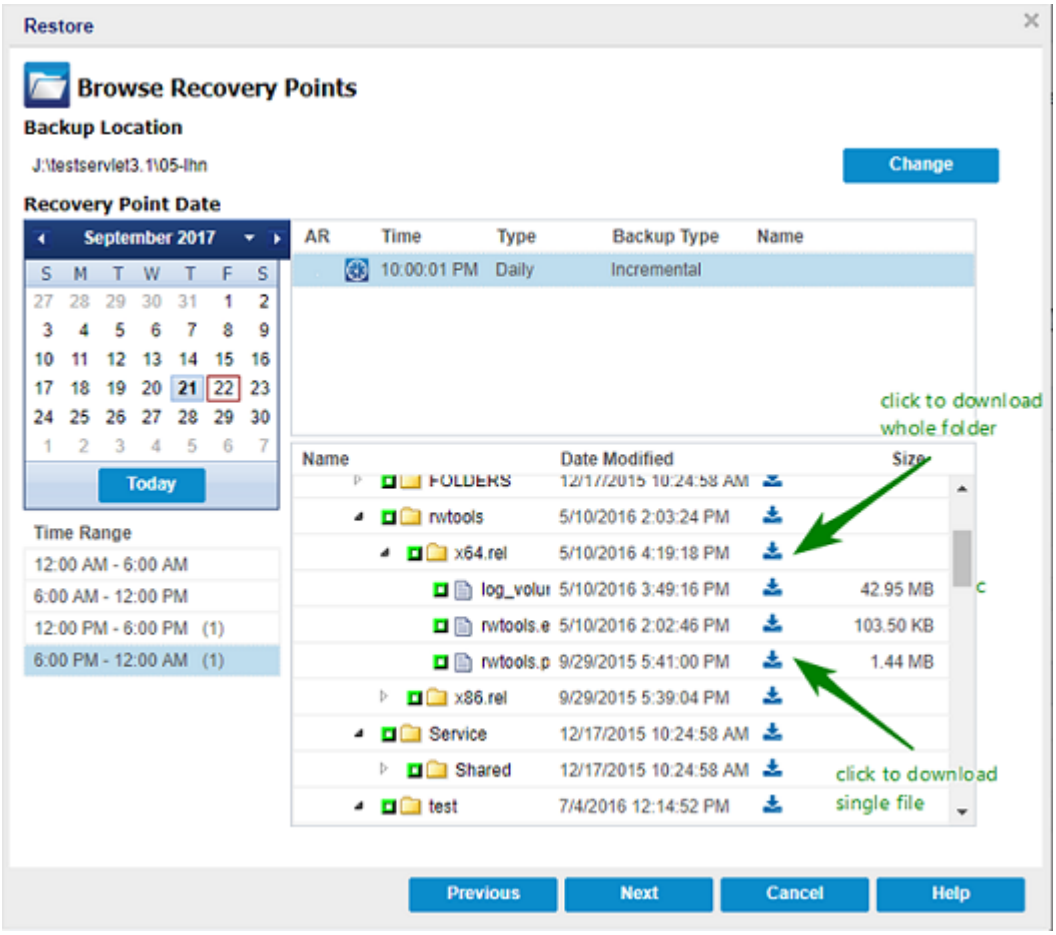
Arcserve UDP lets you download a file or complete folder without submitting for restore. From the Restore wizard, the Browse Recovery Points screen lets you directly download any file or a complete folder with all the files. Downloading before restore may help perform a quick check of files to avoid undesired files getting restored.

Note: Downloading the files does not maintain file or folder permissions.

A single file is downloaded directly in the same format, while a folder is downloaded as a zip file. The zip file has the following name format:

[nodename]_[sessionid]_[timestamp].zip

To download, you simply need to reach the Browse Recovery point screen in the Restore wizard. The below screenshot displays how to perform download of a file or folder:



Considerations for download:

- Downloading or packaging as zip file is not possible for some system file. The agent tomcat service does not have enough privileges to access system file or user files of other protected node.
- To avoid excess consumption of Tomcat memory and CPU usage, we recommend submitting a restore job to alternative path while downloading a huge file or folder.
- Using Windows Compressed Folder Tools to browse the downloaded zip files may fail as the tool finds some of the zip entry names too long to browse. We recommend using other zip tools to open the file. For example, WinZip, WinRAR,7-Zip.
- IE9 user using https in IE9 and agent web service to provide service may not be able to download the files. A known issue from IE9 in downloading resource from a dynamic page through https prevents such download. For more information and solution, click link for Microsoft article.

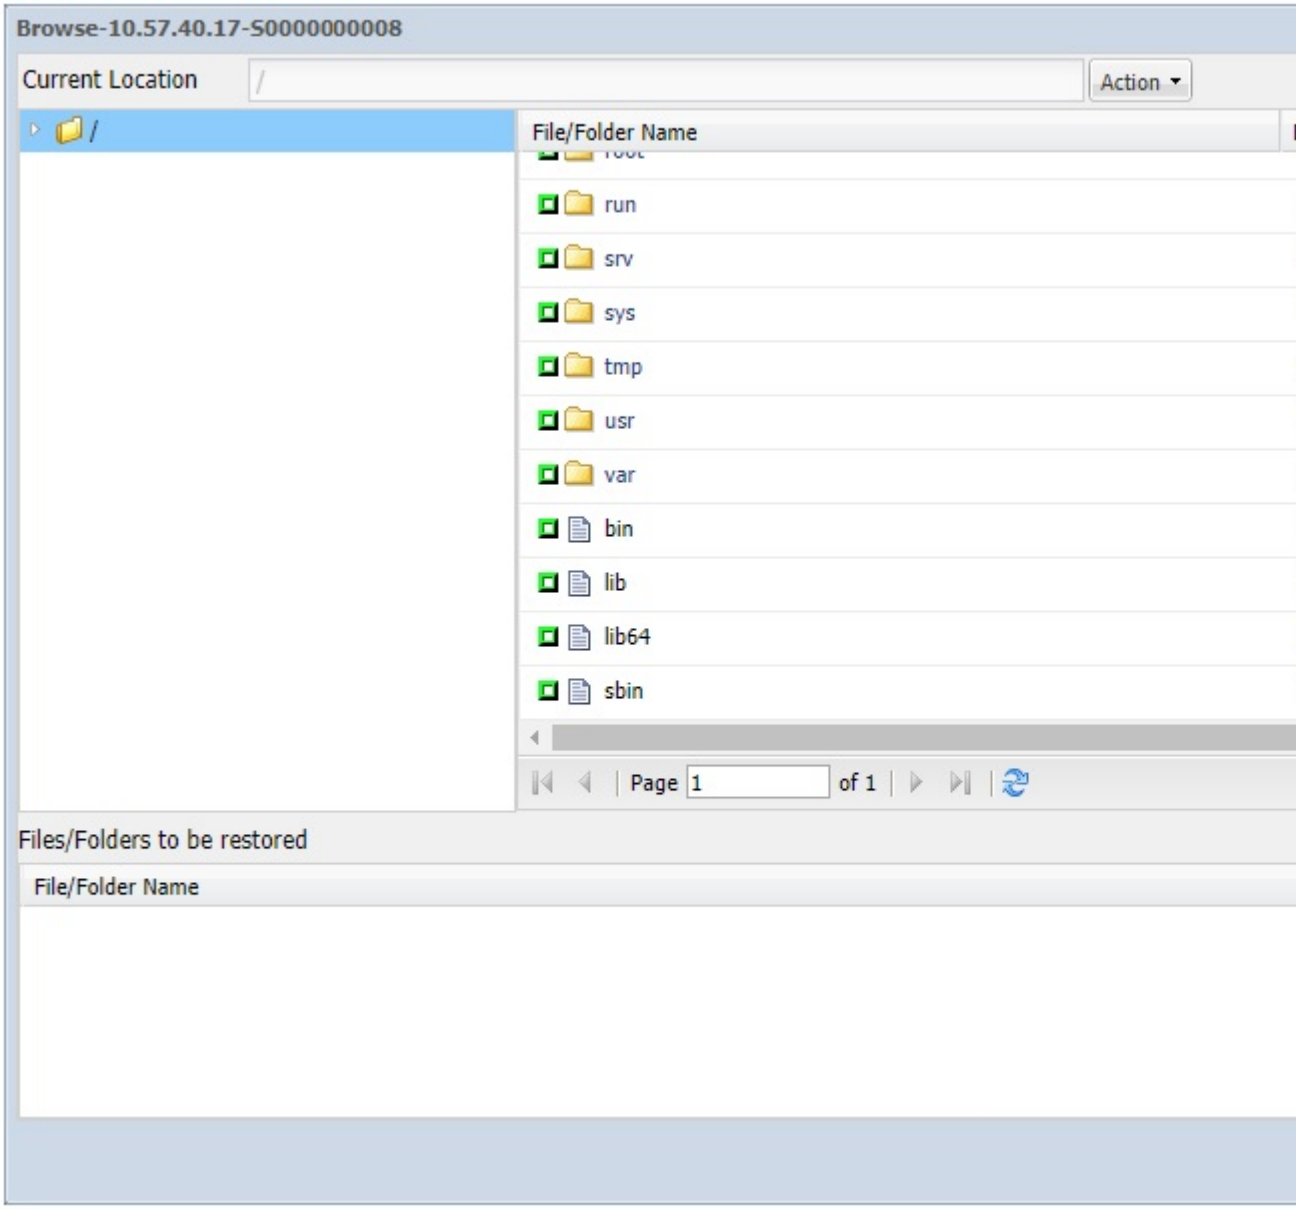
How to Download File/Folders without Restore for Linux Nodes

Arcserve UDP lets you download a file or complete folder without submitting for restore. From the Restore wizard, the Browse Recovery Points screen lets you directly download any file or a complete folder with all the files. Downloading before restore may help perform a quick check of files to avoid undesired files getting restored.

A single file is downloaded directly in the same format, while a folder is downloaded as a zip file. The zip file has the following name format:

[nodename]_[sessionid]_[timestamp].zip

To download, you simply need to reach the Browse Recovery point screen in the Restore wizard. The below screenshot displays how to perform download of a file or folder for linux nodes:

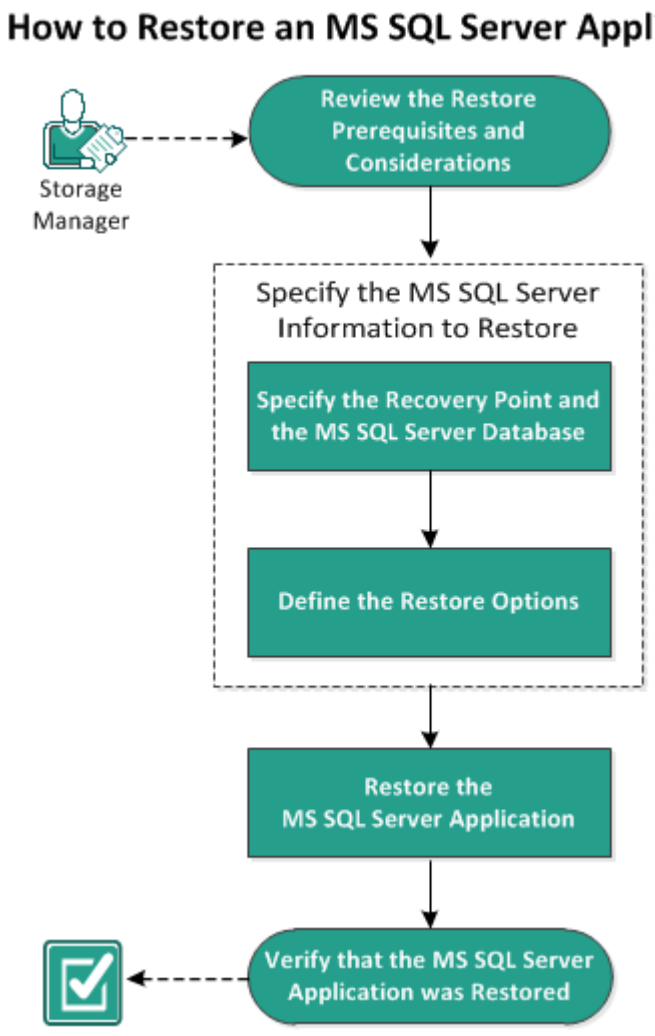


To open the downloaded files, use zip tools such as WinZip, WinRAR, 7-Zip, and so on.

How to Restore a Microsoft SQL Server Application

The Arcserve UDP Agent (Windows) allows you to not only protect and recover your data, but also helps you to get the applications that will use that data back up and running. All application recoveries can only be made using the Restore by Recovery Point method. During an application recovery, Arcserve UDP Agent (Windows) takes advantage of Windows Volume Shadow Copy Service (VSS) to help ensure data consistency for any VSS-aware application. With Arcserve UDP Agent (Windows), you can recover the Microsoft SQL Server application without performing a full disaster recovery.

The following diagram illustrates the process to restore a Microsoft SQL Server Application:



Perform the following tasks to restore a Microsoft SQL Server Application:

1. Review the Restore Prerequisites and Considerations
2. Specify the Microsoft SQL Server Information to Restore

a. Specify the Recovery Point and Microsoft SQL Server Database

b. Define the Restore Options

3. Restore the Microsoft SQL Server Application
4. Verify that the Microsoft SQL Server Application was Restored

Review the Restore Prerequisites and Considerations

Verify that the following prerequisites exist before performing a restore:

- You need Microsoft SQL Server instance before performing a SQL Application restore.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Review the following restore considerations:

- We cannot restore database across an instance. Restore to alternate location in Arcserve UDP Agent (Windows) means we can restore database and change its database name and file location. For more information, see Microsoft SQL Server Restore to Alternate Location Considerations.
- If the jobs are not for the same VM, Arcserve UDP allows multiple restore jobs to run at the same time. If you attempt to launch a restore job, while another restore job is running for the same VM, an alert message informs you that another job is running and requests you to try again later.
- Arcserve UDP_agt_windows> only allows one restore job to run at the same time. If you attempt to launch a restore job manually, while another restore job is running, an alert message opens informing you that another job is running and requests you to try again later.

Microsoft SQL Server Restore to Alternate Location Considerations

When you specify to restore a Microsoft SQL Server application to an alternate location, you can either restore it to an alternate location on the same machine or on a different machine.

Prior to performing an Arcserve UDP Agent (Windows) restore of a Microsoft SQL Server application to an alternate location, you should consider the following:

If alternate location is on the same machine

For this option, you can either restore a database to a new location (with the same name) or restore with a new name (to the same location):

- **Same Name - New Location**

  For example, if Database A is installed in the current SQL Server at "C:\DB_A" and has been backed up. You can use this option and specify "Alternate File Location" to restore Database A to an alternate location such as "D:\Alternate_A.

  After the database has been restored, the database file located at the new location "D:\Alternate_A" will then be used.

Important! During restore, if you change the database location but retain the database name, then the previous database will be deleted after the restore is complete. The restored database file will be pointed to the new location.

When you restore to an alternate location, the Instance Name section is unavailable because the Instance Name should always be the same and cannot be changed. As a result, you cannot restore a database to an alternate instance that is currently located on the same MS SQL Server.

- **Same Location - New Name**

  For example, if you have two databases (Database A and Database B) installed in the current SQL Server and both have been backed up. You can use this option and specify "New database Name" to restore Database A to same location as Database A_New.

  After the databases have been restored, this location will now have three databases (Database A, Database B, and Database A_New).

If alternate location is on the different machine

- The SQL Server installation path must be the same as the path that existed when the backup was performed.

  For example, if the backup of the SQL Server is installed at "C:\SQLServer", then the SQL Server on the new Arcserve UDP Agent (Windows) server must also be installed at C:\SQLServer.

- The same instance name for the database that existed when the backup was performed must be installed on Arcserve UDP Agent (Windows) server, otherwise the database associated with that instance will be skipped from the restore.

  For example, if the backup of the SQL Server contained "Instance_1" with Database A and Database B and "Instance_2" with Database C, but the Arcserve UDP Agent (Windows) server only has "Instance_1". After the restore is complete, Database A and Database B will be restored, but Database C will not be restored.

- The SQL Server version on the Arcserve UDP Agent (Windows) server must be backwards compatible to the version of the SQL Server used during the backup session.

  For example, you can restore a SQL Server 2008 machine to a SQL Server 2010 machine; however, you cannot restore a SQL Server 2010 machine to a SQL Server 2008 machine.

- Restoring a database of 64-bit instance to 32-bit instance is not supported.

Microsoft SQL Server 2012/2014 AAG Restore Considerations

When restoring a Microsoft SQL Server 2012/2014 database that is part of an AlwaysOn Availability Group (AAG), there are some considerations that you should be aware of.

If the MS SQL database is part of the MS SQL 2012/2014 AlwaysOn Availability Group (AAG), and restoring to the original location fails, perform the following tasks:

1. Remove the database to be restored away from the Availability Group. For more information, see the link.
2. Share the backup session to Arcserve UDP Agent (Windows)on every Availability Group node and then restore the session by Arcserve UDP Agent (Windows)on every Availability Group node.
3. Add the database back to an Availability Group. For more information, see the link.

Specify the Microsoft SQL Server Information to Restore

Arcserve UDP Agent (Windows) allows you to not only protect and recover your data, but also helps you to get the Microsoft SQL Server application that uses that data back up and running. The Microsoft SQL Server recovery can only be made using the Restore by Recovery Point method.

The process involved in restoring a Microsoft SQL Server Application is as follows:

1. Specify the Recovery Point and Microsoft SQL Server Database
2. Define the Restore Options

Specify the Recovery Point and Microsoft SQL Server Database

Use the **Browse Recovery Points** option to restore from a recovery point. When you select a recovery date, all the associated recovery points for that date are displayed. You can then browse and select the Microsoft SQL Server database to be restored.

Follow these steps:

1. Access the restore method selection dialog in one of the following ways:

   - From Arcserve UDP:

     a. Log into Arcserve UDP.

     b. Click the **resources** tab.

     c. Select **All Nodes** in the left pane.

     All the added nodes are displayed in the center pane.

     d. In the center pane, select the node and click **Actions**.

     e. Click **Restore** from the server name drop-down option.

     The restore method selection dialog opens.

     Note: You are automatically logged in to the agent node and the restore method selection dialog is opened from the agent node.

   - From Arcserve UDP Agent (Windows):

     a. Log into Arcserve UDP Agent (Windows).

     b. From the home page, select **Restore**.

     The restore method selection dialog opens.

   6. Click the **Browse Recovery Points** option.

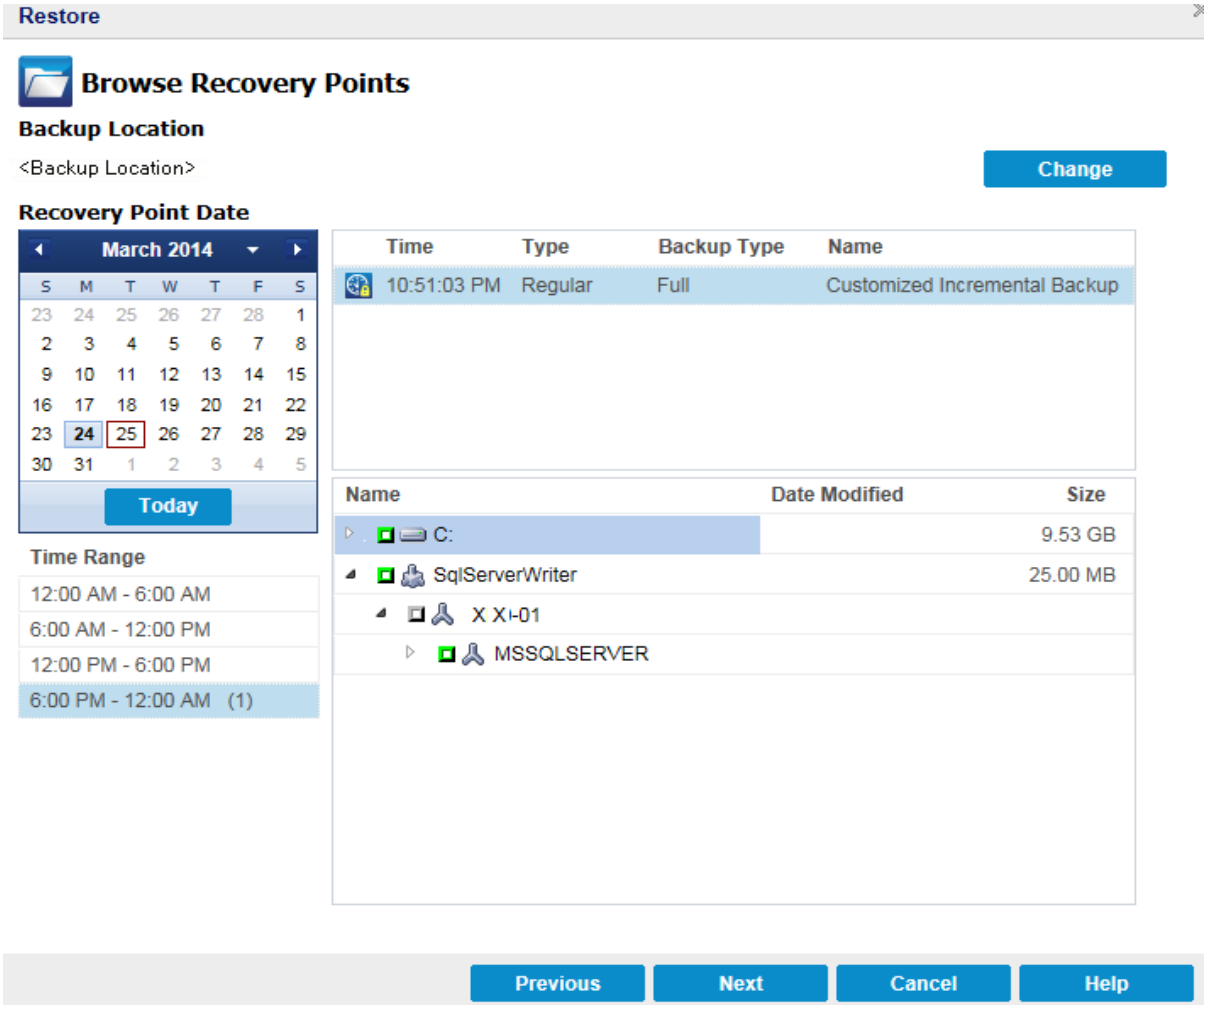
      The **Browse Recovery Points** dialog opens.

   7. Select the recovery point (date and time) and then select the Microsoft SQL Server database to be restored.

   8. The corresponding marker box becomes filled (green) to indicate that the database has been selected for the restore.

      Note: If you do not want the transaction log files to be applied after the restore, you must manually delete it before the restore is performed. For

more information about manually deleting transaction log files, refer to the Microsoft SQL Server documentation.



9. Click **Next**.

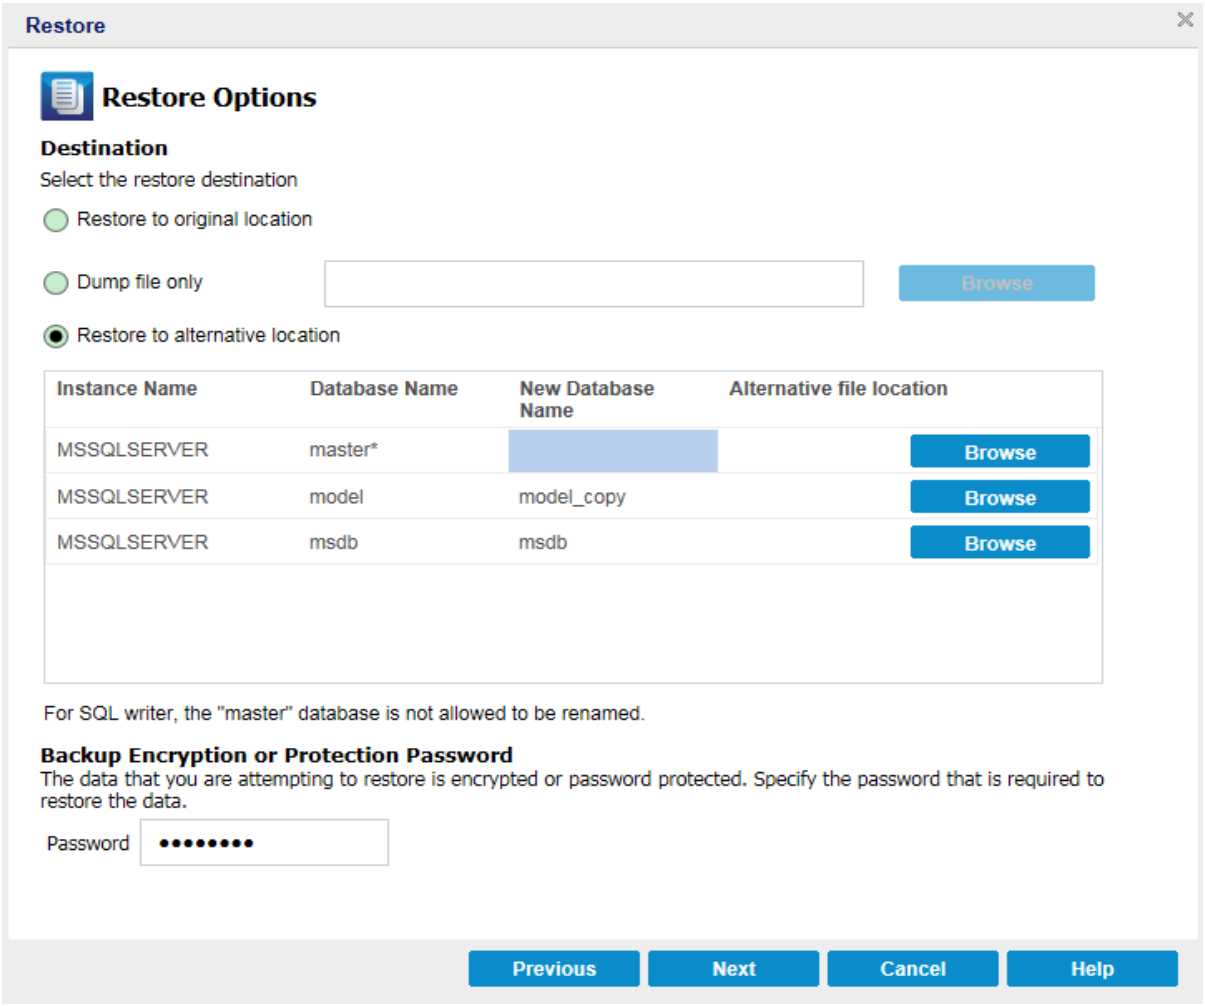
   The **Restore Options** dialog opens.

Define the Restore Options

After you specify a recovery point and content to restore, define the copy options for the selected recovery point.

Follow these steps:

1. From the **Restore Options** dialog, select the restore destination.



2. Select the destination for the restore.

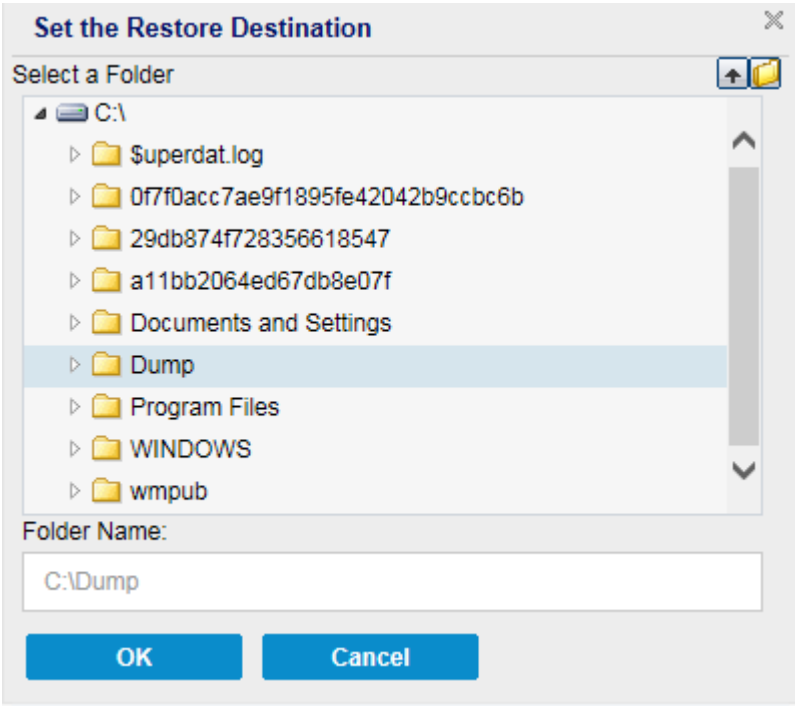
   The available options are to restore to the original location of the backup, restore the dump file only, or restore to alternative location.

Restore to original location

Restores to the original location from where the backup image was captured.

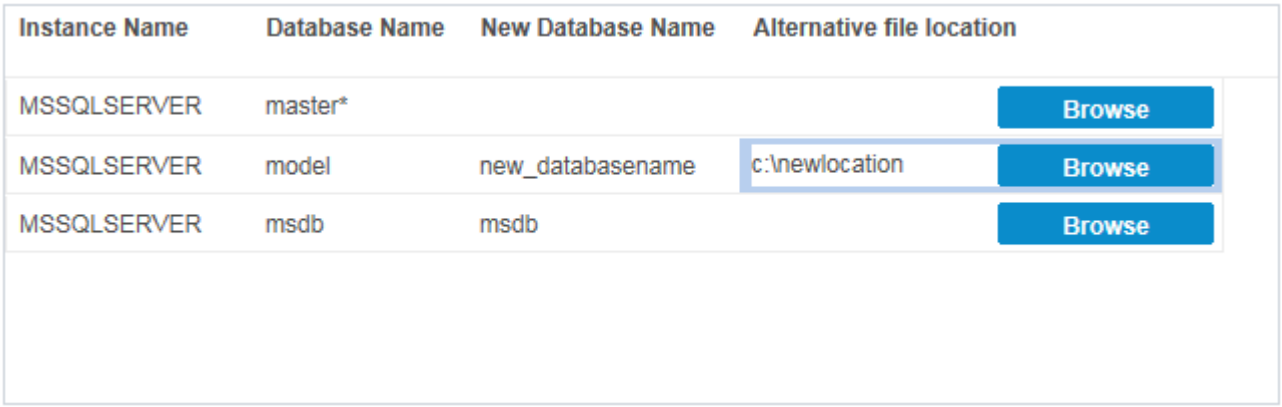
Dump file only

For this option, Arcserve UDP Agent (Windows) dumps the selected Microsoft SQL database files to the specified folder. When you select this option, you can then specify or browse to the folder location where the dump file will be restored to.



Restore to alternate location

Restores to an alternate location (not the original location).



Backups can be copied to network locations and they can be used by multiple SQL Server instances. You can perform a multiple database restore (simultaneously) from the instance level. From this listing, you can select the

database instance and specify a new database name and alternate location to restore the database to. In addition, you can also browse to the alternate location where the database will be restored to.

When restoring a Microsoft SQL Server application to an alternate location, there are some considerations that you should be aware of. For more information, see the **Microsoft SQL Server Restore to Alternate Location Considerations** section in the topic Review the Restore Prerequisites and Considerations.

3. Click **Next**.

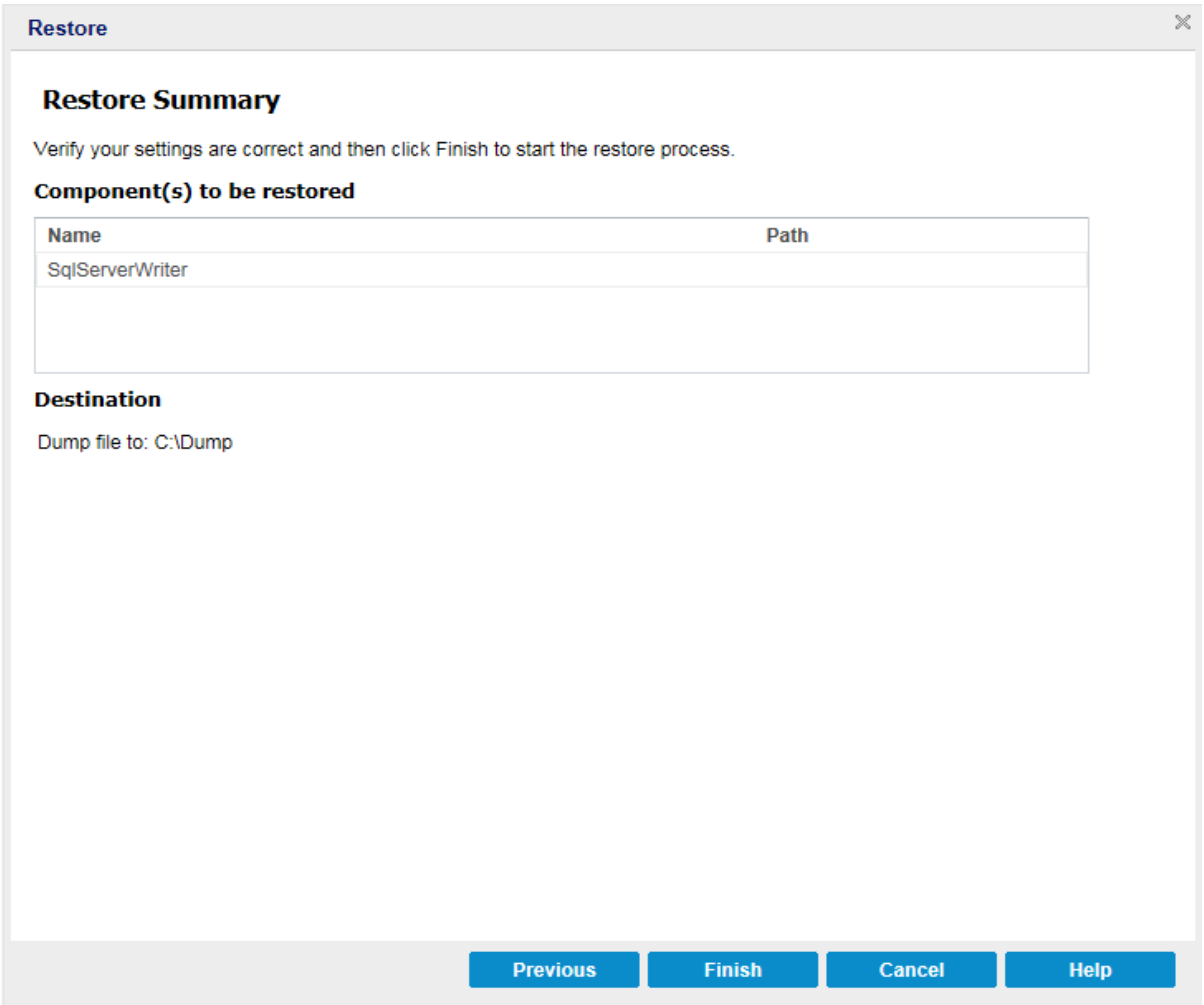
   The **Restore Summary** dialog opens.

Restore the Microsoft SQL Server Application

After you define the restore options, verify that your settings are correct and confirm the restore process. **Restore Summary** helps you to review all the restore options that you defined and modify them if necessary.

Follow these steps:

1. From the **Restore Summary** dialog, review the displayed information to verify that all the restore options and settings are correct.



- If the summary information is not correct, click **Previous** and go back to the applicable dialog to change the incorrect setting.
- If the summary information is correct, click **Finish** to launch the restore process.

The Microsoft SQL Server Application is restored.

Verify that the Microsoft SQL Server Application was Restored

Follow these steps:

1. Navigate to the Arcserve UDP Agent (Windows) restore destination you specified.

   For example, if you select to restore the Microsoft SQL Server database to the original location, after the restore is complete, then browse to the physical location to check if the Microsoft SQL Server database and logs are restored.

   If you select to restore the Microsoft SQL Server database to Dump File only location then Arcserve UDP Agent (Windows) will restore the Microsoft SQL Server database and logs to a specified location.

2. Verify if the Microsoft SQL Server Application was restored and check if the database is mounted and is accessible.

   The Microsoft SQL Server Application is restored successfully.

How to Restore From a UNC/NFS Path

Each time Arcserve UDP performs a successful UNC/NFS path backup, it backs up all files/folders that have changed after the last successful job. This restore method allows you to browse the archived files/folders and specify exactly which file you want to restore.

Perform the following tasks to restore from a recovery point:

1. Review the Restore Prerequisites and Considerations
2. Specify the Files/Directories on a UNC/NFS Path to Restore
    a. Specify the Files and Content to Restore
    b. Define the Restore Options
3. Restore the Files and Content
4. Verify that Content is Restored

Review the Restore Prerequisites and Considerations

Verify that the following prerequisites exist before performing a restore:

- You have at least one recovery point available to restore.
- You have a valid and accessible recovery point destination to restore the recovery point content from.
- You have a valid and accessible target location to restore the recovery point content to.

Specify the UNC/NFS Path Information to Restore

Arcserve UDP provides you with an option to restore data from a UNC/NFS path. The goal of running a successful restore job is to identify quickly the data you need and to retrieve it from the appropriate backup location. Each restore job requires a source and destination.

The process involved in restoring from a UNC/NFS path is as follows:

1. Specify the Files/Folders and Content to Restore
2. Define the Restore Options

Specify the UNC/NFS Files/Folders and Content to Restore

This section provides information about how to specify the UNC/NFS Files/Folders and content to restore.

Follow these steps:

1. Access the restore method selection dialog in one of the following ways:

   From Arcserve UDP Console:

   a. Log into Arcserve UDP Console.

   b. Navigate to **resources** > **Node** > **All Nodes**.

      All the added nodes are displayed in the center pane.

   c. Right-click the node, and then click **Restore**.

      You are automatically logged into the agent node console, and the Restore dialog is opened.

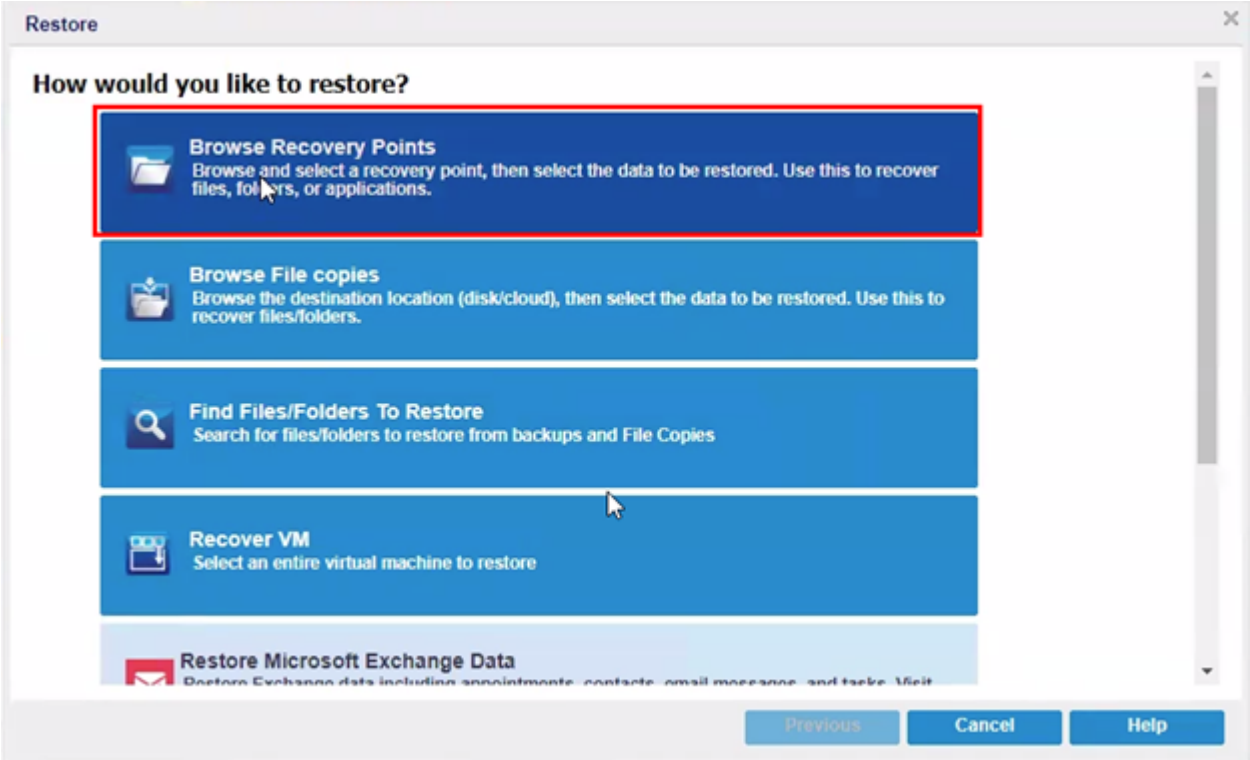
   From Arcserve UDP Agent (Windows):

   a. Log into Arcserve UDP Agent (Windows).

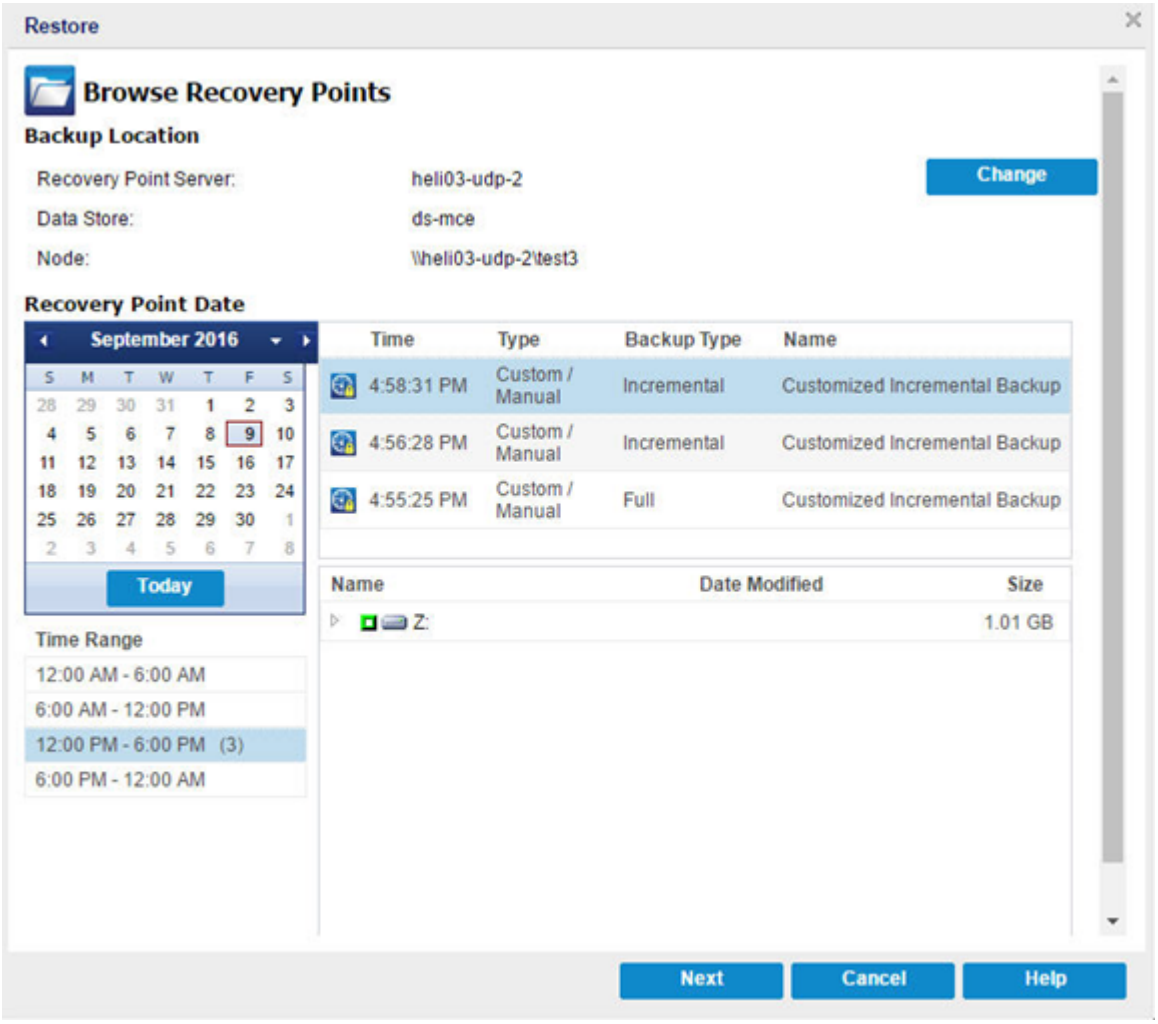
   b. From the home page, select **Restore**.

      The Restore dialog opens.

2. On the Restore dialog, click the **Browse Recovery Points** option.



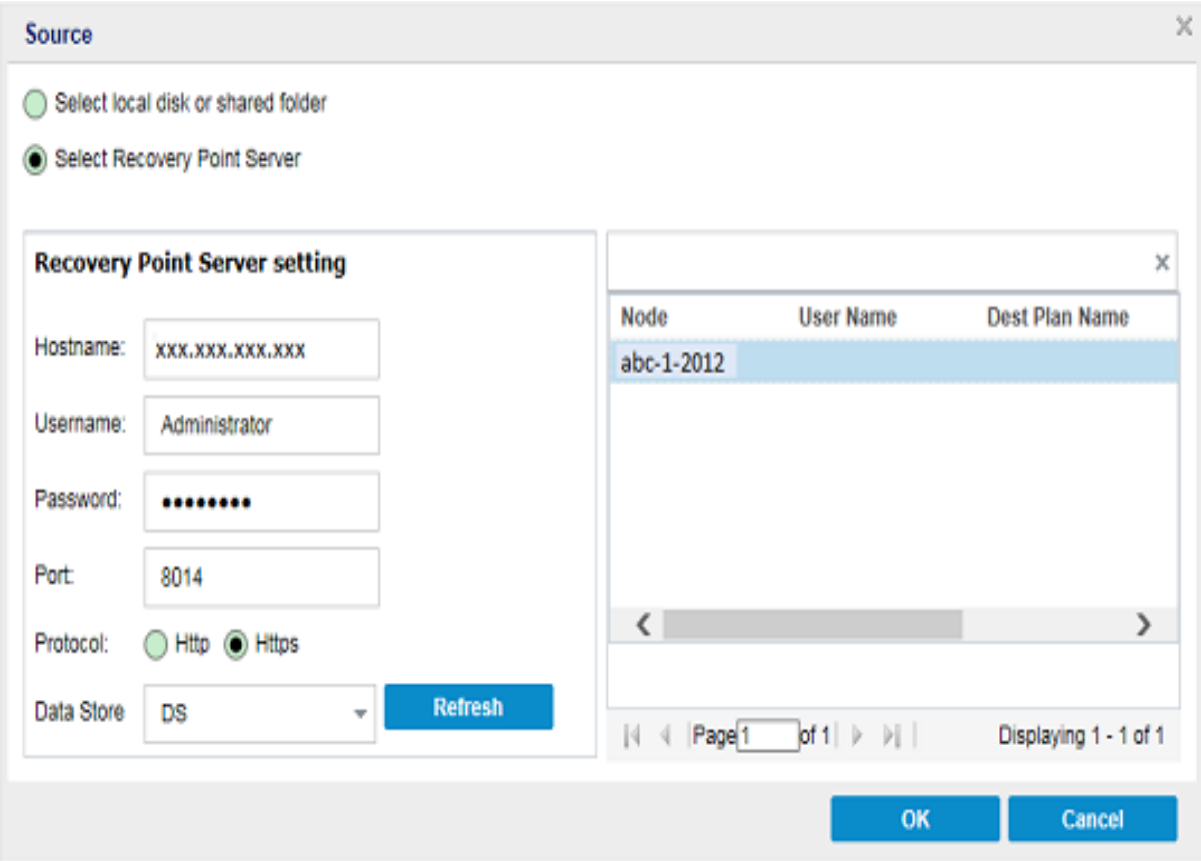
3. On the Browse Recovery Points page, do the following:



a. To update the backup location, click **Change**.

   The Source dialog opens.

b. On the Source dialog, select one of the following backup locations, and then click **OK**:



Select local disk or shared folder

1. Specify or browse to the location where your backup images are stored and select the appropriate backup source.
2. To verify the connection to the specified location, click the green arrow button. If required, enter the Username and Password to access to that source location.

   The Select backup location dialog opens.
3. Select the folder where the recovery points are stored, and then click **OK**.

   The Select backup location dialog closes and you can see the backup location in the Source dialog.

   The recovery points are listed in the Browse Recovery Points dialog.

Select Recovery Point Server

1. Specify the Recovery Point Server setting details, and then click **Refresh**.

   All the agents are listed in the Data Protection Agent column in the Source dialog.

2. Select the agent from the displayed list.

   The recovery points are listed in the Browse Recovery Points dialog.

c. Select the calendar date for the backup image to restore.

   All the dates containing recovery points for the specified backup source are highlighted in green.

   The corresponding recovery points for that date are displayed, with the time of the backup, the type of backup that was performed (Full, Incremental, or Verify), and the name of the backup.

d. Select a recovery point to restore.

   The backup content (including any applications) for the selected recovery point displays.

   Note: If a clock icon appears with a lock symbol, it indicates that the recovery point contains encrypted information and may require a password for restore.

e. Select the content to restore.

   - For a volume-level restore, you can specify to restore the entire volume or selected files/folders within the volume.
   - For an application-level restore, you can specify to restore the entire application or selected components, databases, instances, and so on, within the application.

4. Click **Next**.

   The Restore Options dialog opens.

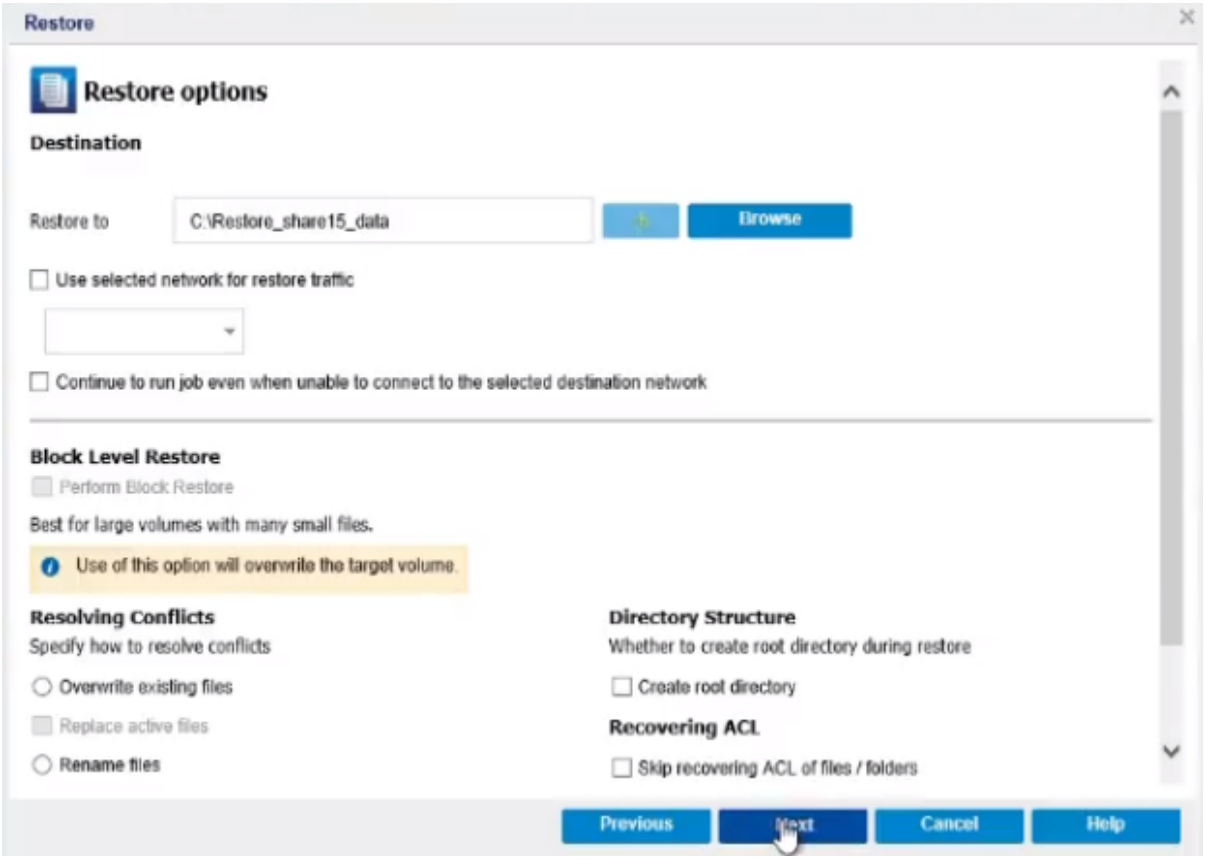
   The recovery point and content to restore is specified.

Define the Restore Options

After you specify a recovery point and content to restore, define the copy options for the selected recovery point.

Follow these steps:

1. On the Restore Options dialog, select the restore destination.



Available destination option is:

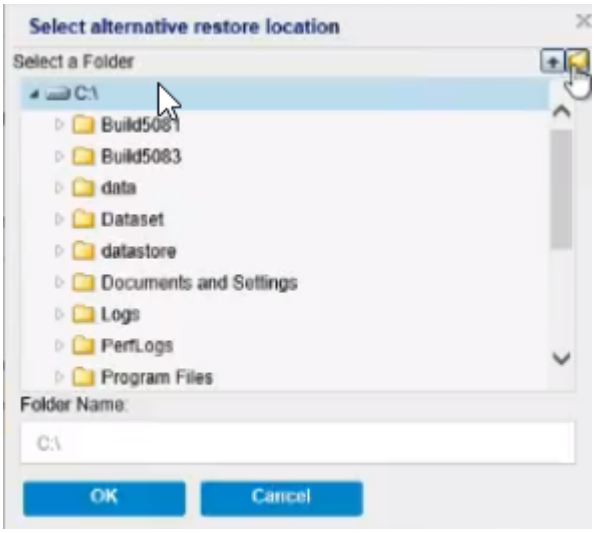
Restore to

Restores to the specified location. To restore data to the specified location, do the following:

a. To provide the destination, click **Browse**.

   The Select alternative restore location dialog appears.

b. Select the existing folder or create a new folder as needed, and then click **OK**.



c. To verify the connection to the specified location, click the green arrow button. If necessary, enter the Username and Password to gain access to that location.

Note: Restore to NFS share from an NFS Protection plan is not allowed.

2. To enable the communication between the selected network and the Recovery Point Server, select the **Use selected network for restore traffic** check box, and then select the network from the drop-down list.

   Note: If the selected backup network is not accessible and to continue the backup job with the available network or with the default network, click the **Continue to run job even when unable to connect to the selected backup network** check box.

3. Specify the **Resolving Conflicts** option that Arcserve UDP performs if conflicts are encountered during the restore process.

   The available options are:

   Overwrite existing files

   Overwrites (replaces) any existing files, which are at the restore destination. All the objects are restored from the backup files regardless of their current presence on your computer.

   Replace active files

   Replaces any active files upon reboot. If during the restore attempt Arcserve UDP Agent (Windows) discovers that the existing file is currently in use or being accessed, it will not immediately replace that file, but instead to avoid any problems will delay the replacement of the active files until the next time the machine is rebooted. (The restore occurs immediately, but the replacement of any active files is done during the next reboot).

   This option is only available if you select the **Overwrite existing files** option.

Note: If you do not select this option, any active file is skipped from the restore.

Rename files

Creates a new file if the file name already exists. Selecting this option copies the source file to the destination with the same file name but a different extension. Data is then restored to the new file.

Skip existing files

Skips over and not overwrite (replace) any existing files that are located at the restore destination. Only objects that are not currently existing on your machine are restored from the backup files.

Default: Skip existing files.

4. To create a root directory during restore, specify the **Directory Structure**.

Create root directory

Specifies that if a root directory structure exists in the captured backup image, Arcserve UDP recreates that same root directory structure on the restore destination path.

With this option not selected, the file or folder is restored directly to the destination folder.

For example, if during the backup you captured the files "C:\Folder-1\SubFolder2\A.txt" and "C:\Folder1\SubFolder2\B.txt" and during the restore you specified to the restore destination as "D:\Restore".

- If you select to restore the "A.txt" and "B.txt" files individually, the destination for the restored files will be "D:\Restore\A.txt" and "D:\Restore\B.txt" (the root directory above the specified file level will not be recreated).
- If you select to restore from the "SubFolder2" level, the destination for the restored files will be "D:\Restore\SubFolder2\A.txt" and "D:\Restore\SubFolder2\B.txt" (the root directory above the specified folder level will not be recreated).

With this option selected, the entire root directory path for the files/folders (including the volume name) is recreated to the destination folder. If the files/-folders to be restored are from the same volume name, then the destination root directory path does not include that volume name. However, if the files/-folders to be restored are from different volume names, then the destination root directory path does include the volume name.

For example, if during the backup you captured the files "C:\Folder-1\SubFolder2\A.txt", "C:\Folder1\SubFolder2\B.txt", and also E:\Folder-3\SubFolder4\C.txt" and during the restore you specified to the restore destination as "D:\Restore".

- If you select to restore just the "A.txt" file, the destination for the restored file will be "D:\Restore\ Folder1\SubFolder2\A.txt" (the entire root directory without the volume name will be recreated).
- If you select to restore both the "A.txt" and "C.txt" files, the destination for the restored files will be "D:\Restore\C\Folder1\SubFolder2\A.txt" and "D:\Restore\E\Folder3\SubFolder4\C.txt" (the entire root directory with the volume name will be recreated).

5. Specify if you want the Recovering ACL to skip recovering ACL of files/folders.

   If the option **Skips recovering ACL of files/foders** is selected, only source files/-folders are restored. The attribute of files/folders are not restored so that everyone can access.

   Default: Restores source files/folders with the attribute - Access Conrol List.

6. If necessary, specify the **Backup Encryption Password**, when the data you are trying to restore is encrypted.

   A password is not required if you are attempting to restore from the same Arcserve UDP Agent (Windows) computer from where the encrypted backup was performed. However, if you are attempting to restore from a different Arcserve UDP Agent (Windows) computer, a password is required.

   Note: If a clock icon appears with a lock symbol, it indicates that the recovery point contains encrypted information and may require a password for restore.

7. Click **Next**.

   The Restore Summary dialog opens.

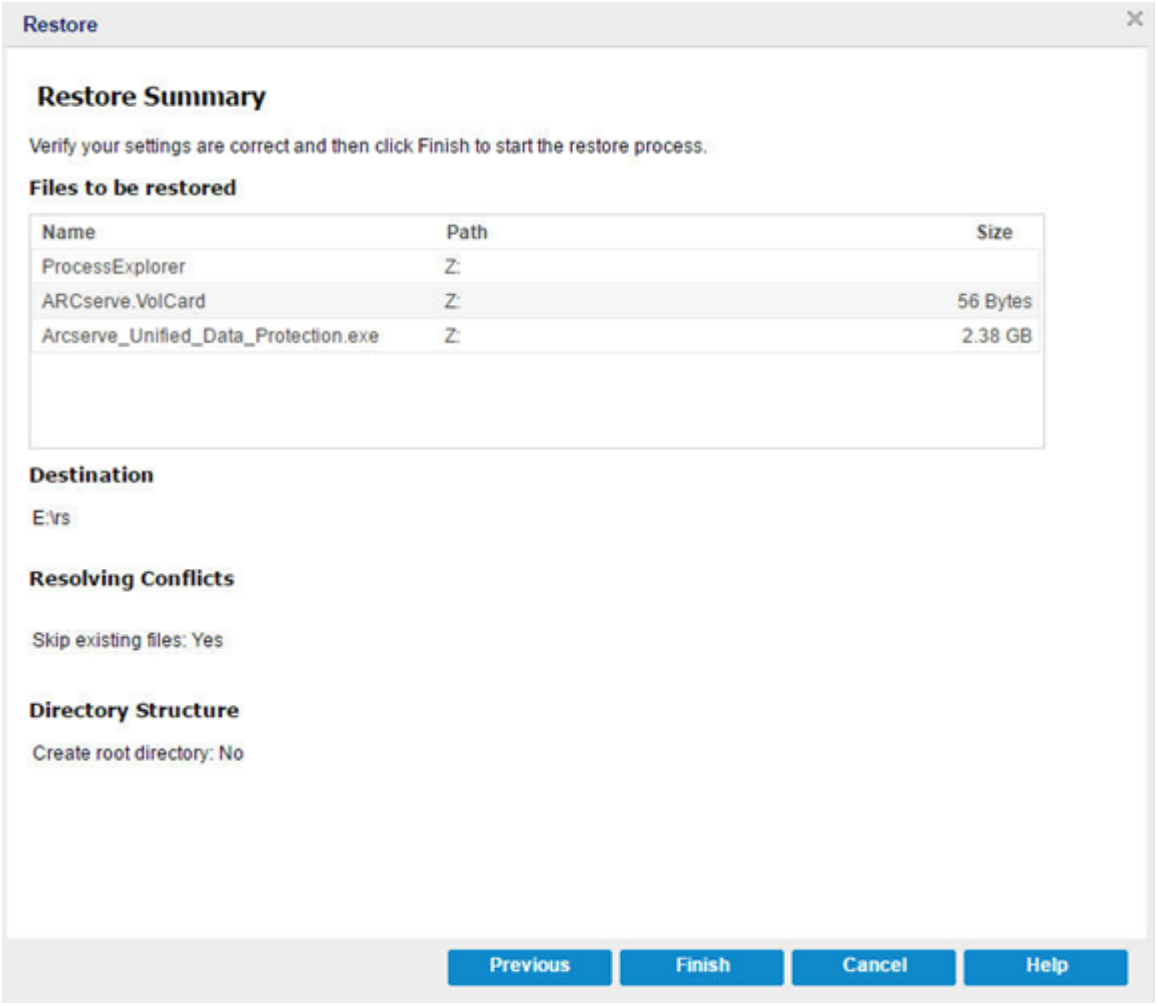
   The restore options are defined to restore from a recovery point.

Restore the Files/Folders and Content Located on UNC/NFS Path

After you define the restore options, verify that your settings are correct and confirm the restore process. Restore Summary page helps you to review all the restore options that you defined and modify them if necessary.

Follow these steps:

1. On the Restore Summary dialog, review the displayed information to verify that all the restore options and settings are correct.



2. Do one of the following:
   - If the summary information is incorrect, click **Previous** and go back to the applicable dialog to change the incorrect setting.
   - If the summary information is correct, click **Finish** to launch the restore process.

The recovery point content is restored.

Verify that Content is Restored

After the completion of restore job, verify that all the files are restored in the target node. Check the **Job History** and **Activity Log** tabs in the Status pane to monitor the progress of the restore process.

Follow these steps:

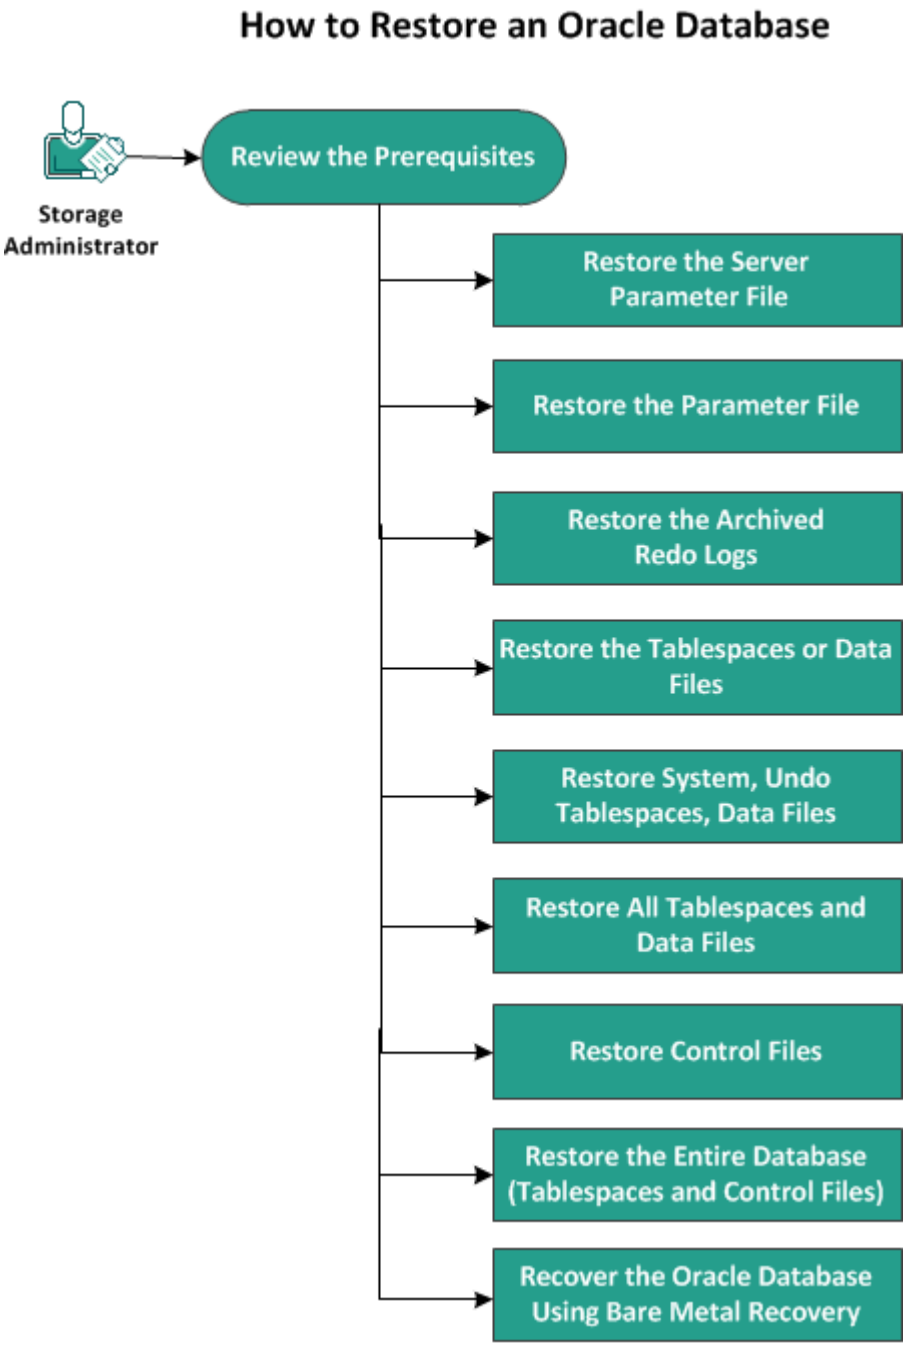
1. Navigate to the target machine where you restored data.
2. Verify that the required data from the recovery point is restored.

   The restored content is successfully verified.

How to Restore an Oracle Database

You can restore either certain files and tablespaces or the entire Oracle database using the restore wizard. To restore an Oracle database, locate the files or tablespace on the destination node. Then, you restore the files or tablespace using the restore wizard.

The following diagram illustrates the process to restore an Oracle database:



Perform the following tasks to restore an Oracle database:

- Review the Prerequisites
- Restore the Server Parameter File
- Restore the Parameter File
- Restore the Archived Redo Logs
- Restore the Tablespaces or Data Files
- Restore System, Undo Tablespaces, Data Files
- Restore All Tablespaces and Data Files
- Restore Control Files
- Restore the Entire Database (Tablespaces and Control Files)
- Recover the Oracle Database Using Bare Metal Recovery

Review the Prerequisites and Considerations

Review the following prerequisites before you restore the Oracle database:

- The Oracle VSS writer on the backup node is functioning properly. If the Oracle VSS writer does not function properly, you get a warning message in the Activity Log associated with the backup job.
- You have a valid recovery point.
- To avoid any restore failure problem, you have saved a duplicate copy of your system files before you overwrite the original files.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Restore the Server Parameter File

The server parameter file is a repository for initialization parameters. Before you restore, you must locate the file. When you locate the files, ensure that the database is in the Open state.

Follow these steps:

1. Log in to the computer where you want to restore the files.
2. Locate the server parameter file using the following command:

   SQL> SHOW PARAMETER SPFILE;
3. Shut down the database or the Oracle instance before you begin the restore process:

   SQL> SHUTDOWN IMMEDIATE;
4. Log in to the Arcserve UDP Console.
5. Restore the server parameter file using the Restore Wizard. For more information on the restore process, see How to Restore From a Recovery Point.
6. Log in to the destination computer.
7. Navigate to the specific folders and verify that the files are restored.
8. Connect to SQL*Plus to restart the Oracle instance with the restored server parameter file.

   The server parameter file is restored.

Restore the Parameter File

The parameter file includes a list of initialization parameters and values for each parameters. Before you restore, you must locate the file. When you locate the files, ensure that the database is in the Open state.

Follow these steps:

1. Log in to the computer where you want to restore the files.
2. Locate the parameter file (pfile).

   Typically, the pfile (INIT<SID>.ORA) is located in the %ORACLE_HOME/database directory. You can type "INIT<SID>.ORA" to locate the pfile.
3. Shut down the database or the Oracle instance before you begin the restore process:

   SQL> SHUTDOWN IMMEDIATE;
4. Log in to the Arcserve UDP Console.
5. Restore the parameter file using the Restore Wizard. For more information on the restore process, see How to Restore From a Recovery Point.
6. Log in to the destination computer.
7. Navigate to the specific folders and verify that the files are restored.
8. Connect to SQL*Plus to restart the Oracle instance with the restored parameter file.

   The parameter file is restored.

Restore the Archived Redo Logs

Archived redo logs are used to recover a database or update a standby database. Before you restore, you must locate the file. When you locate the files, ensure that the database is in the Open state.

Follow these steps:

1. Log in to the computer where you want to restore the files.
2. Locate the archived redo logs using the following command.

   SQL> ARCHIVE LOG LIST;

   SQL> SHOW PARAMETER DB_RECOVERY_FILE_DEST;
3. Log in to the Arcserve UDP Console.
4. Restore the archived redo logs using the Restore Wizard. For more information on the restore process, see How to Restore From a Recovery Point.
5. Log in to the destination computer.
6. Navigate to the specific folders and verify that the archived redo logs are restored.

   The archived redo logs are restored.

Restore the Tablespaces or Data Files

You can restore the tablespace or data files. Before you restore, you must locate the file. When you locate the files, ensure that the database is in the Open state. If the database is open, use the ALTER TABLESPACE. OFFLINE statement to take the tablespaces or datafiles offline before you begin the restore process.

Follow these steps:

1. Log in to the computer where you want to restore the tablespaces or datafiles.
2. Locate the user tablespaces or datafiles using the following command:

   SQL> SELECT FILE_NAME, TABLESPACE_NAME FROM DBA_DATA_FILES;
3. Change the state of the database to mount, or nomount, or shutdown before you restore the tablespaces or datafiles.

   SQL> STARTUP MOUNT;

   SQL> STARTUP NOMOUNT;

   SQL> SHUTDOWN IMMEDIATE;
4. Log in to the Arcserve UDP Console.
5. Restore the tablespaces or datafiles using the Restore Wizard. For more information on the restore process, see How to Restore From a Recovery Point.
6. Log in to the destination computer.
7. Navigate to the specific folders and verify that the tablespaces or datafiles are restored.
8. Recover the tablespace or data files.
   - To recover a tablespace, enter the following command at the SQL*Plus prompt screen:

     SQL> RECOVER TABLESPACE "tablespace_name";
   - To recover a data file, enter the following command at the SQL*Plus prompt screen:

     SQL> RECOVER DATAFILE 'path';

   Oracle checks for the archive redo log files that it needs to apply and displays the names of the files in a sequence.
9. Enter AUTO in the SQL*Plus prompt screen to apply the files.

   Oracle applies the log files to restore the data files. After Oracle finishes applying the redo log file, it displays the following messages:

   Applying suggested logfile

Log applied

After each log is applied, Oracle continues to apply the next redo log file until the recovery is complete.

10. Enter the following command to bring the tablespace online:

    SQL> ALTER TABLESPACE "tablespace_name" ONLINE;

    The tablespace is now recovered to the last available log file.

Restore System, or Undo Tablespaces or Data Files

You can restore system, or undo tablespaces or data files. Before you restore, you must locate the file. When you locate the files, ensure that the database is in the Open state.

Follow these steps:

1. Log into the computer where you want to restore system or undo tablespaces or datafiles.
2. Locate the user tablespaces or datafiles using the following command:

   SQL> SELECT TABLESPACE_NAME, FILE_NAME FROM DBA_DATA_FILES;
3. Change the state of the database to mount, or nomount, or shutdown before you restore the tablespaces or datafiles.

   SQL> STARTUP MOUNT;

   SQL> STARTUP NOMOUNT;

   SQL> SHUTDOWN IMMEDIATE;
4. Log into the Arcserve UDP Console.
5. Restore the tablespaces or datafiles using the Restore Wizard. For more information on the restore process, see How to Restore From a Recovery Point.
6. Log into the destination computer.
7. Navigate to the specific folders and verify that the system, or undo tablespaces or datafiles are restored.
8. Recover the tablespace or data files.
   - To recover a tablespace, enter the following command at the SQL*Plus prompt screen:

     SQL> RECOVER TABLESPACE "tablespace_name";
   - To recover a data file, enter the following command at the SQL*Plus prompt screen:

     SQL> RECOVER DATAFILE 'path';

   Oracle checks for the archive redo log files that it needs to apply and displays the names of the files in a sequence.
9. Enter AUTO in the SQL*Plus prompt screen to apply the files.

   Oracle applies the log files to restore the data files. After Oracle finishes applying the redo log file, it displays the following messages:

   Applying suggested logfile

Log applied

After each log is applied, Oracle continues to apply the next redo log file until the recovery is complete.

10. Enter the following command to bring the tablespace online:

    SQL> ALTER TABLESPACE "tablespace_name" ONLINE;

    The tablespace is now recovered to the last available log file.

Restore All Tablespaces and Data Files

You can restore all the tablespaces and data files. Before you restore, you must locate the file. When you locate the files, ensure that the database is in the Open state. If the database is open, use the ALTER TABLESPACE. OFFLINE statement to take the tablespaces or datafiles offline before you begin the restore process.

Follow these steps:

1. Log into the computer where you want to restore the tablespaces or datafiles.
2. Locate the user tablespaces or datafiles using the following command:

   SQL> SELECT FILE_NAME, TABLESPACE_NAME FROM DBA_DATA_FILES;
3. Change the state of the database to mount, or nomount, or shutdown before you restore the tablespaces or datafiles.

   SQL> STARTUP MOUNT;

   SQL> STARTUP NOMOUNT;

   SQL> SHUTDOWN IMMEDIATE;
4. Log into the Arcserve UDP Console.
5. Restore the tablespaces or datafiles using the Restore Wizard. For more information on the restore process, see How to Restore From a Recovery Point.
6. Log i to the destination computer.
7. Navigate to the specific folders and verify that the tablespaces or datafiles are restored.
8. Recover the database.

   SQL> RECOVER DATABASE;

   Oracle checks for the archive redo log files that it needs to apply and displays the names of the files in a sequence.
9. Enter AUTO in the SQL*Plus prompt screen to apply the files.

   Oracle applies the log files to restore the data files. After Oracle finishes applying the redo log file, it displays the following messages:

   Applying suggested logfile

   Log applied

   After each log is applied, Oracle continues to apply the next redo log file until the recovery is complete.

   Note: If Oracle displays an error indicating that the log file cannot be opened, the log file may not be available. In such cases, perform the incomplete media recovery

to recover the database again. After all the log files are applied, the database recovery is complete. For more information about incomplete media recovery, see the Oracle documentation.

10. Enter the following command to bring the database online:

    SQL> ALTER DATABASE OPEN;

    The database is now recovered to the last available log file.

    Note: If you perform an incomplete media recovery, enter the following command to change the database to the open state:

    SQL> ALTER DATABASE OPEN RESETLOGS;

Restore Control Files

You can restore the control files that stores the physical structure of the database. Before you restore, you must locate the file. When you locate the files, ensure that the database is in the Open state.

Follow these steps:

1. Log in to the computer where you want to restore the control files.
2. Locate the control files using the following command:

   SQL> SHOW PARAMETER CONTROL FILES;
3. Change the state of the database to nomount or shutdown before you restore the control files.

   SQL> STARTUP NOMOUNT;

   SQL> SHUTDOWN IMMEDIATE;
4. Log in to the Arcserve UDP Console.
5. Restore the Control file using the Restore Wizard. For more information on the restore process, see How to Restore From a Recovery Point.
6. Log in to the destination computer.
7. Navigate to the specific folders and verify that the control files are restored.
8. Mount the database to begin the database recovery:

   SQL> STARTUP MOUNT
9. Enter the RECOVER command with the USING BACKUP CONTROLFILE clause.

   SQL> RECOVER DATABASE USING BACKUP CONTROLFILE

   The database recovery process begins.
10. (Optional) Specify the UNTIL CANCEL clause to perform an incomplete recovery.

    SQL> RECOVER DATABASE USING BACKUP CONTROLFILE UNTIL CANCEL
11. Apply the prompted archived logs.

    Note: If the required archived log is missing, then it implies that a necessary redo record is located in the online redo logs. It occurs because unarchived changes are located in the online logs when the instance failed. You can specify the full path of an online redo log file and press Enter (you may have to try this a few times until you find the correct log).
12. Enter the following command to return the control file information about the redo log of a database:

    SQL>SELECT * FROM V$LOG;

13. (Optional) Enter the following command to see the names of all of the member of a group:

    SQL>SELECT * FROM V$LOGFILE;

    Example: After applying the prompted archived logs, you may see the following messages:

    ORA-00279: change 55636 generated at 24/06/2014 16:59:47 needed for thread 1

    ORA-00289: suggestion e:\app\Administrator\flash_recovery_ area\orcl\ARCHIVELOG\2014_06_24\ O1_MF_1_2_9TKXGGG2_.ARC

    ORA-00280: change 55636 for thread 1 is in sequence #24

    Specify log: {<RET>=suggested | filename | AUTO | CANCEL}

14. Specify the full path of the online redo log file and press Enter.

    Example: E:\app\Administrator\oradata\orcl\redo01.log

    Note: You have to specify the full path multiple times until you get the correct log.

    The following messages are displayed:

    Log applied

    Media recovery complete

15. Open the database with the RESETLOGS clause after completing the recovery process.

    SQL> ALTER DATABASE OPEN RESETLOGS;

    The lost control files are recovered.

Restore the Entire Database (Tablespaces and Control Files)

You can restore all the entire database (all tablespaces and control files). Before you restore, you must locate the file. When you locate the files, ensure that the database is in the Open state. If the database is open, use the ALTER TABLESPACE. OFFLINE statement to take the tablespaces or datafiles offline before you begin the restore process.

Follow these steps:

1. Log into the computer where you want to restore the tablespaces or datafiles.
2. Locate the user tablespaces or datafiles using the following command:

   SQL> SELECT TABLESPACE_NAME, FILE_NAME from DBA_DATA_FILES;

   SQL> SHOW PARAMETER CONTROL FILES;
3. Change the state of the database to nomount, or shutdown before you restore the tablespaces or datafiles.

   SQL> STARTUP NOMOUNT;

   SQL> SHUTDOWN IMMEDIATE;
4. Log in to the Arcserve UDP Console.
5. Restore the tablespaces or datafiles using the Restore Wizard. For more information on the restore process, see How to Restore From a Recovery Point.
6. Log in to the destination computer.
7. Navigate to the specific folders and verify that the tablespaces or datafiles are restored.
8. Recover the database.

   SQL> RECOVER DATABASE USING BACKUP CONTROLFILE UNTIL CANCEL;
9. Apply the prompted archived logs.

   Note: If the required archived log is missing, then it implies that a necessary redo record is located in the online redo logs. It occurs because unarchived changes are located in the online logs when the instance failed. You can specify the full path of an online redo log file and press Enter (you may have to try this a few times until you find the correct log).
10. Enter the following command to return the control file information about the redo log of a database:

    SQL>SELECT * FROM V$LOG;

11. (Optional) Enter the following command to see the names of all of the member of a group:

    SQL>SELECT * FROM V$LOGFILE;

    Example: After applying the prompted archived logs, you may see the following messages:

    ORA-00279: change 55636 generated at 24/06/2014 16:59:47 needed for thread 1

    ORA-00289: suggestion e:\app\Administrator\flash_recovery_ area\orcl\ARCHIVELOG\2014_06_24\ O1_MF_1_2_9TKXGGG2_.ARC

    ORA-00280: change 55636 for thread 1 is in sequence #24

    Specify log: {<RET>=suggested | filename | AUTO | CANCEL}

12. Specify the full path of the online redo log file and press Enter.

    Example: E:\app\Administrator\oradata\orcl\redo01.log

    Note: You have to specify the full path multiple times until you get the correct log.

    The following messages are displayed:

    Log applied

    Media recovery complete

13. Open the database with the RESETLOGS clause after completing the recovery process.

    SQL> ALTER DATABASE OPEN RESETLOGS;

    Note: For multitenant databases (CDB/PDB), need to open all pluggable database also.

    SQL> ALTER PLUGGABLE DATABASE <PDB_NAME> OPEN;

    The entire database is restored.

14. Reboot the Oracle server after performing steps 1 to 13 for multi-tenant databases.

    Note: This step is not required for server running only stand-alone databases.

Recover the Oracle Database Using Bare Metal Recovery

Bare metal recovery lets you recover and rebuild the entire computer system during a disaster. You can restore the original computer or you can restore another computer.

Follow these steps:

1. Restore the computer using one of the following methods:
   - If the recovery points are from an agent-based backup, perform a BMR to restore the computer.
   - If the recovery points are from a host-based agentless backup, then use Recover VM to restore the computer.
2. Log in to the restored computer.
3. Open the command prompt and connect to the Oracle instance (for example ORCL) as sysdba.
4. Verify the status of the Oracle instance.

   SQL> SELECT STATUS FROM V$INSTANCE;
5. Perform one of the following steps depending on the status of the Oracle instance:
   - If the status is Shutdown, then start and open the instance.

     SQL> STARTUP;

     SQL> ALTER DATABASE OPEN;
   - If the status is Nomount, then mount and open the instance.

     SQL> ALTER DATABASE MOUNT;

     SQL> ALTER DATABASE OPEN;
   - If the status is Mount, then open the Oracle instance.

     SQL> ALTER DATABASE OPEN;
6. Recovery by executing the RECOVER command if database need media recovery

   SQL> RECOVER DATABASE;
7. Open the Oracle instance after the media recovery is complete.

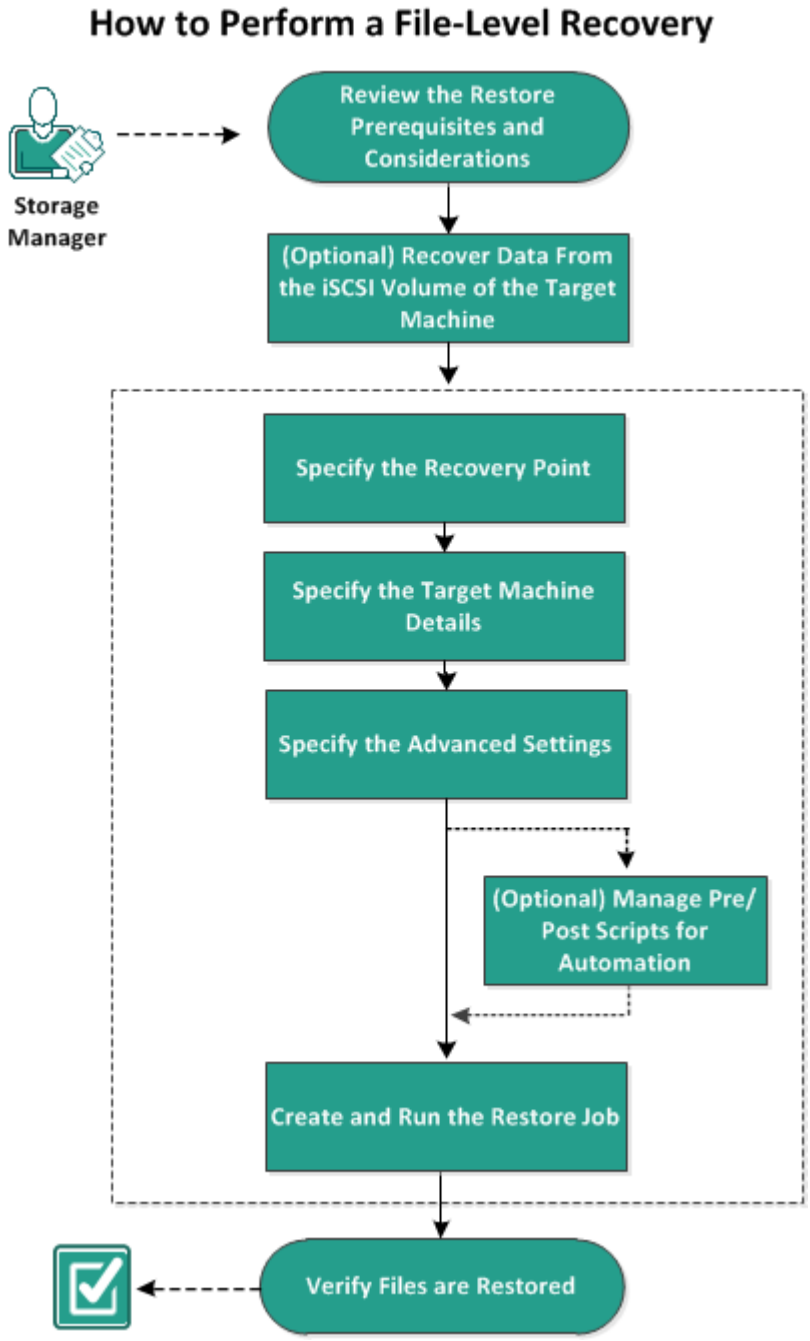
   SQL> ALTER DATABASE OPEN;

   The Oracle database is recovered using the bare metal recovery.

How to Perform a File-Level Recovery on Linux Nodes

A file-level recovery restores individual files and folders from a recovery point. You can restore as minimum as one file from the recovery point. This option is useful if you want to restore selected files and not the entire recovery point.

The following diagram displays the process to perform a file-level recovery:



Perform these tasks for a file-level recovery:

- Review the Restore Prerequisites
- (Optional) Recover Data from the iSCSI Volume to the Target Machine
- Specify the Recovery Point
- Specify the Target Machine Details
- Specify the Advanced Settings
- (Optional) Manage Pre/Post Scripts for Automation
- Create and Run the Restore Job
- Verify that Files are Restored

Review the Prerequisites

Consider the following options before you perform a file-level recovery:

- You have a valid recovery point and the encryption password, if any.
- You have a valid target node to recover data.
- You have verified that the Linux Backup Server supports the file system that you want to restore.

  For example, RedHat 7.x does not support the *reiserfs* file system. If the operating system of the Backup Server is RedHat 7.x and you want to restore the reiserfs file system, you must install the file system driver to support reiserfs. You can also use Arcserve Unified Data Protection Agent for Linux Live CD to perform the file-level restore because Live CD supports all types of file system.

- You have installed the following packages on the Linux Backup Server:
  - mdadm
  - kpartx
  - lvm2
  - Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

(Optional) Recover Data from the iSCSI Volume to the Target Machine

If you have stored your data in an iSCSI target volume, you can connect to the iSCSI volume and recover data. The iSCSI volume lets you manage data and transfer data over a network.

Verify that you have the latest release of the iSCSI-initiator software installed on your Backup Server. The initiator software on RHEL systems is packaged as iscsi-initiator-utils. The initiator software on SLES systems is packaged as open-iscsi.

Follow these steps:

1. Log into the shell environment of the Backup Server.
2. Run one of the following commands to start the iSCSI initiator daemon.
   - For RHEL systems:

     /etc/init.d/iscsid start

     The service on RHEL systems is named iscsid.
   - For SLES systems:

     /etc/init.d/open-iscsi start

     The service on SLES systems is named open-iscsi.
3. Run a discovery script to discover the iSCSI target host.

   iscsiadm -m discovery -t sendtargets -p <ISCSI-SERVER-IP-ADDRESS>:<Port_Number>

   The default port value of iSCSI target host is 3260.
4. Make a note of the iSCSI Qualified Name (IQN) of the iSCSI target host found by the discovery script before you manually log into the discovered target.
5. List the available block device of the Backup Server.

   #fdisk -l
6. Log in to the discovered target.

   iscsiadm -m node -T <iSCSI Target IQN name> -p <ISCSI-SERVER-IP-ADDRESS>:<Port_Number> -l

   You can see a block device in the /dev directory of the Backup Server.
7. Run the following command to obtain the new device name:

   #fdisk -l

   You can see an additional device named /dev/sd<x> on the Backup Server.

For example, consider the name of the device is /dev/sdc. This device name is used to create a partition and a file system in the following steps.

8. Mount the iSCSI volume using the following commands:

   # mkdir /iscsi

   # mount /dev/sdc1 /iscsi

   Note: When you specify the session location in the Restore Wizard, you need to select Local and enter the path /iscsi.

   Example: <path>/iscsi

9. (Optional) Add the following record to the /etc/fstab file so that the iSCSI volume automatically connects with the Backup Server after you restart the server.

   /dev/sdc1 /iscsi ext3 _netdev 0 0

   The Backup Server can now connect to the iSCSI volume and can recover data from the iSCSI volume.

Specify the Recovery Point for Agent-Based Backup

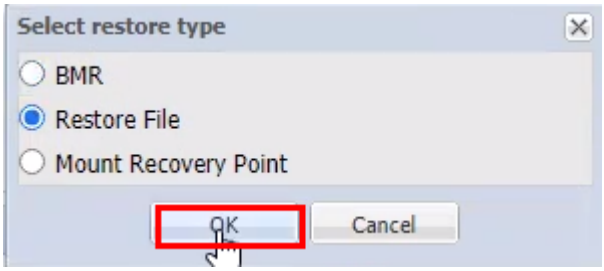
Each time you perform a backup, a recovery point is created. Specify the recovery point information in the Restore Wizard so that you can recover the exact data that you want. You can restore specific files or all files depending on your requirement.

Follow these steps:

1. Access the Restore Wizard in one of the following ways:
   - **From Arcserve UDP:**
     a. Log into Arcserve UDP.
     b. Navigate to **resources** > **Node** > **All Nodes**.

        All the added nodes are displayed on the center pane.
     c. Right-click the node, and then click **Restore**.

        The Arcserve UDP Agent (Linux) web interface opens and displays the Select restore type dialog.
     d. On the Select restore type dialog, click the **Restore File** option, and then click **OK**.

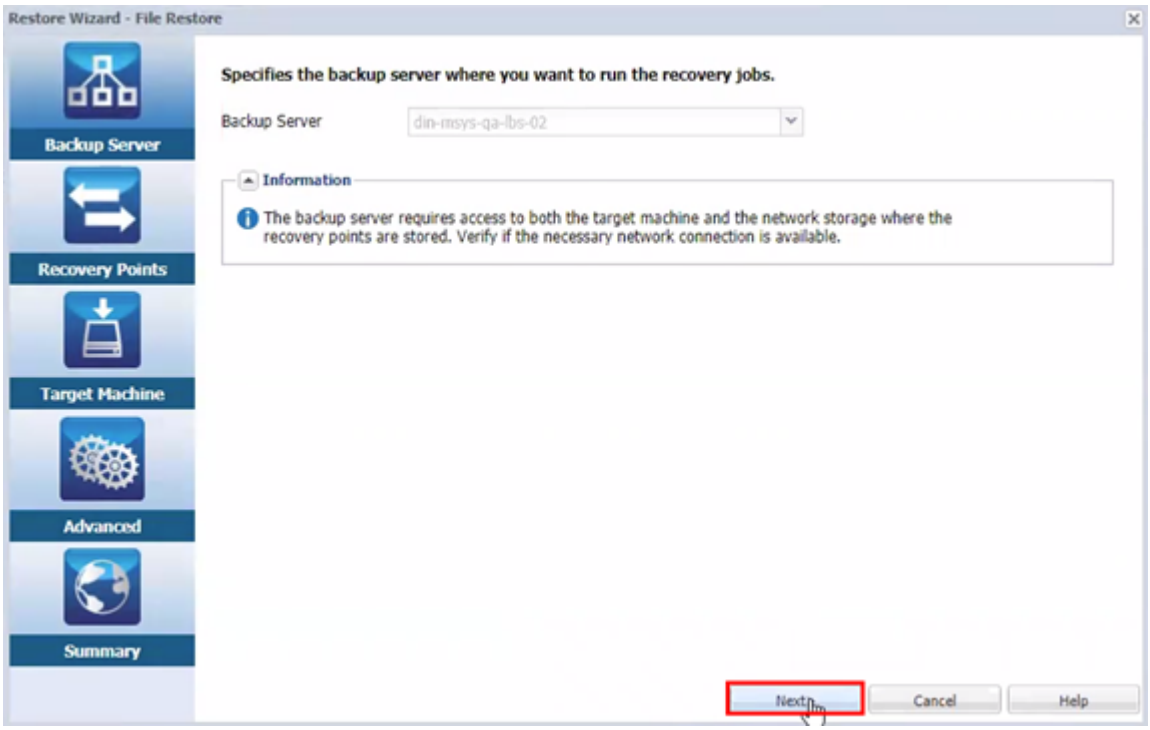
        

        Note: You are automatically logged into the agent node and the Restore Wizard opens from the agent node.
   - **From Arcserve Unified Data Protection Agent for Linux:**
     a. Open the Arcserve Unified Data Protection Agent for Linux web interface.

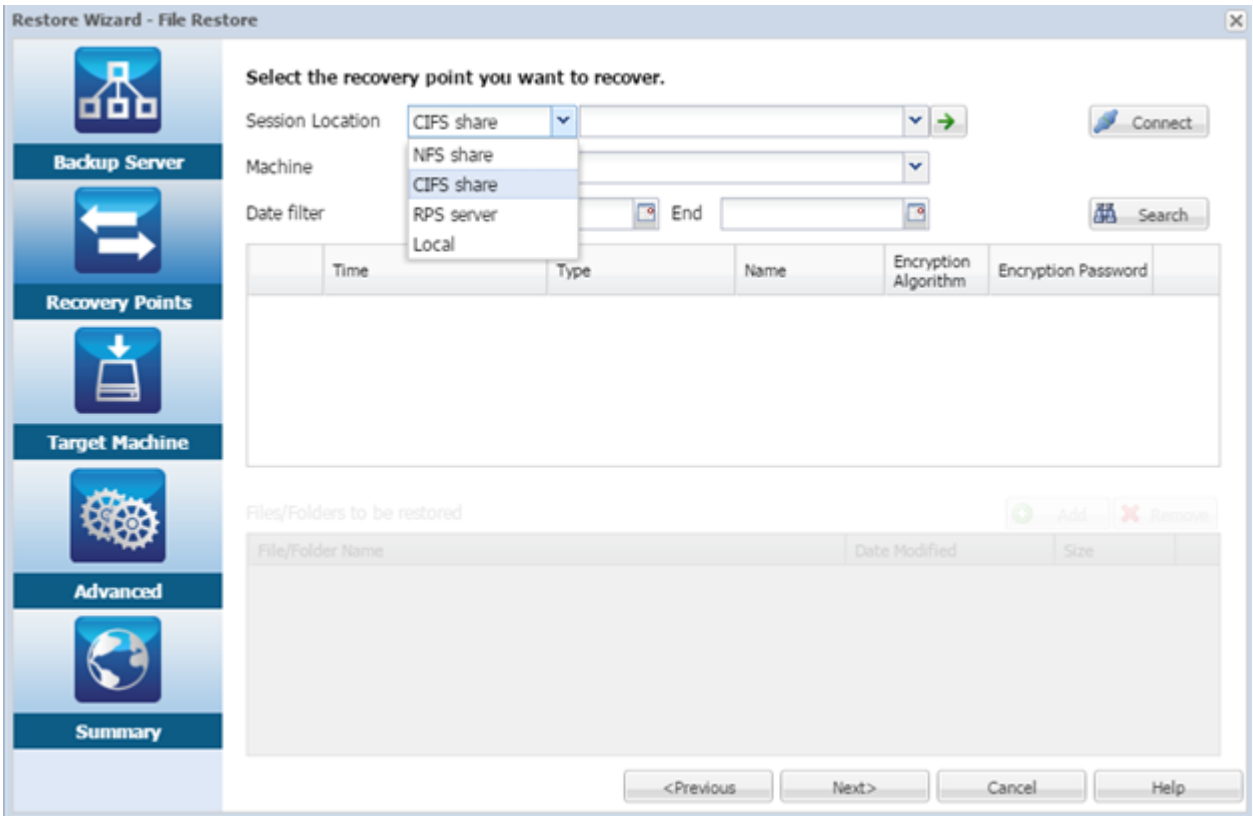
        Note: During the installation of Arcserve Unified Data Protection Agent for Linux, you received the URL to access and manage the server. Log into Arcserve Unified Data Protection Agent for Linux.
     b. From the Wizard menu, click **Restore**, and then select **Restore File**.

        The Restore Wizard - File Restore dialog opens.

2. On the Backup Server page of the Restore Wizard, you can see the backup server. You cannot select any option from the Backup Server drop-down list. Click **Next**.



3. On the Recovery Points page of the Restore Wizard, do the following:



Important! If you have opened the Wizard from the Console, the session location and machine details are automatically displayed. You can skip to Step 4.

a. Select the **CIFS share/NFS share/RPS server/Local** from the Session Location drop-down list.

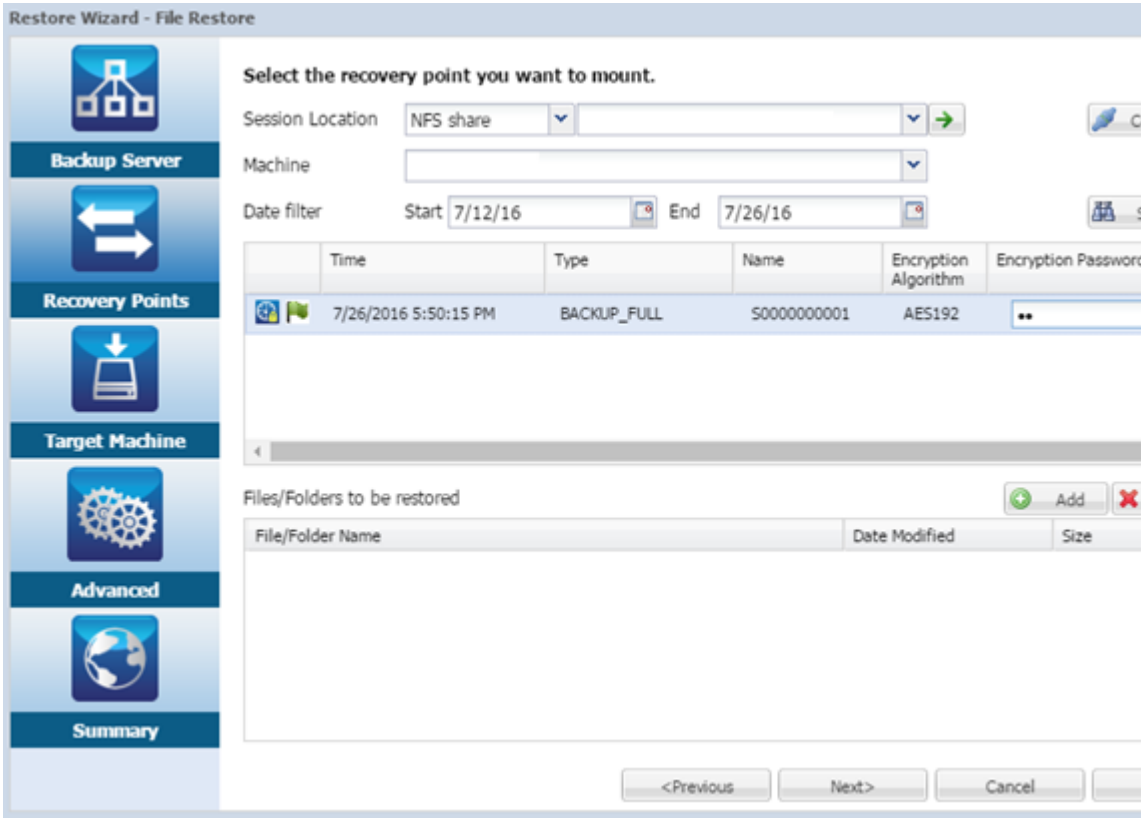
b. Perform one of the following steps depending on your session location:

For CIFS share/NFS share/Local

Specify the full path of the CIFS share/NFS share/Local, and then click **Connect**.

All the machines are listed in the Machine drop-down list.

Note: If you select the **CIFS share** option, specify the user name and password.



For RPS server

a. Select the RPS server, and then click **Add**.

The Recovery Point Server Information dialog opens.

b. Provide the RPS details, and then click **Yes**.

c. Select the data store from the drop-down list.

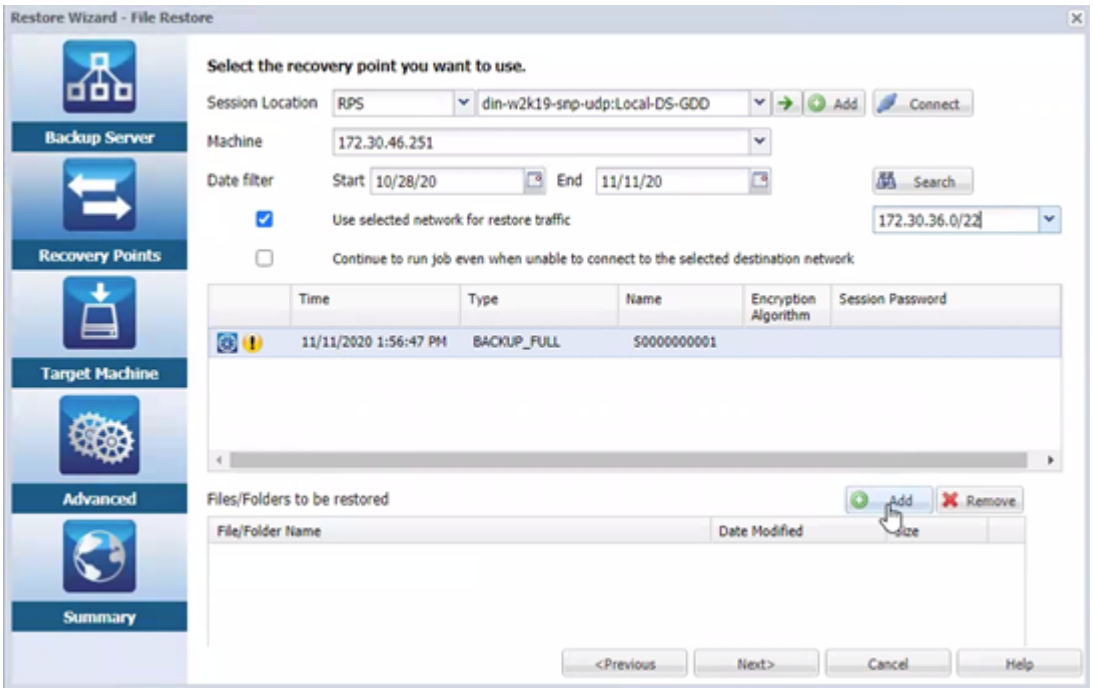
The Recovery Point Server Information dialog closes and you see the wizard.

d. Click **Connect**.

All the nodes that have been backed up to this location get listed in the Machine drop-down list.

e. From the Machine drop-down list, select the node that you want to restore.

All the recovery points of the selected node get listed.



4. Apply the date filter to display the recovery points that are generated between the specified date, and then click **Search**.

Default: Recent two weeks.

All the recovery points available between the specified dates are displayed.

5. To enable the communication between Linux Agent and Recovery Point Server, select the **Use selected network for restore traffic** check box, and then select the network from the drop-down list.

Note: If the selected backup network is not accessible and to continue the backup job with the available network or with the default network, click the **Continue to run job even when unable to connect to the selected backup network** check box.

6. Select the recovery point that you want to restore. If the recovery point is encrypted, enter the encryption password to restore data.

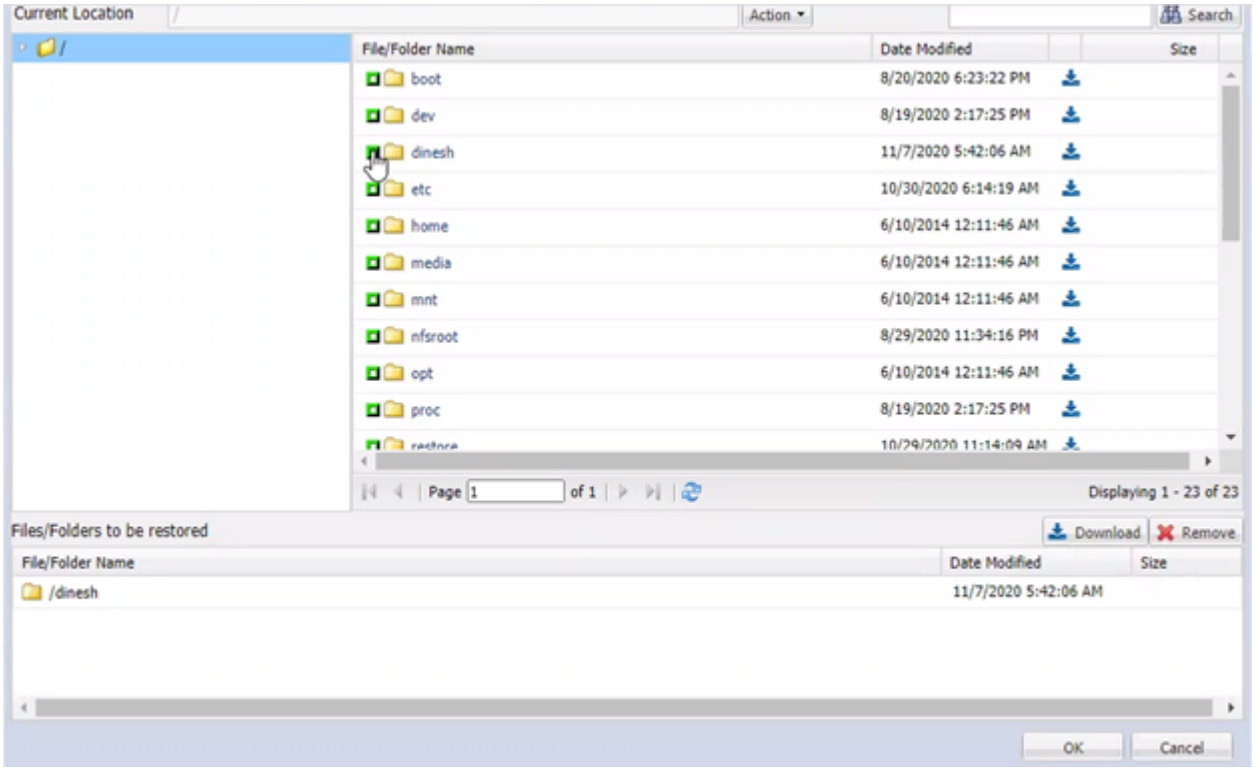
7. For Files/Folders to be restored, click **Add**.

   The Browse-<node name> dialog opens.

   Important! If you see the warning message, "The files/folders are displayed under device file. Click for more information." on the Console, refer to the following Note for resolution.

   Note: For some complex disk layout, the file system is shown by the device file. The change in the file system display behavior does not affect the function of host-based Linux VM file-level restore. You can browse the file system under the device file. Also, you can use the search function to search specific file or directory.

8. On the Browse-<node name> dialog, select the file or folder that you want to restore, and then click **OK**.



   Note: If you try to locate a file or folder using the **Search** field, ensure that you select the highest folder in the hierarchy. The search is conducted on all the child folders of the selected folder.

   The Browse-<node name> dialog closes and you return to the Recovery Points page. The selected files and folders are listed under Files/Folders to be restored.

9. Click **Next**.

   The Target Machine page opens.

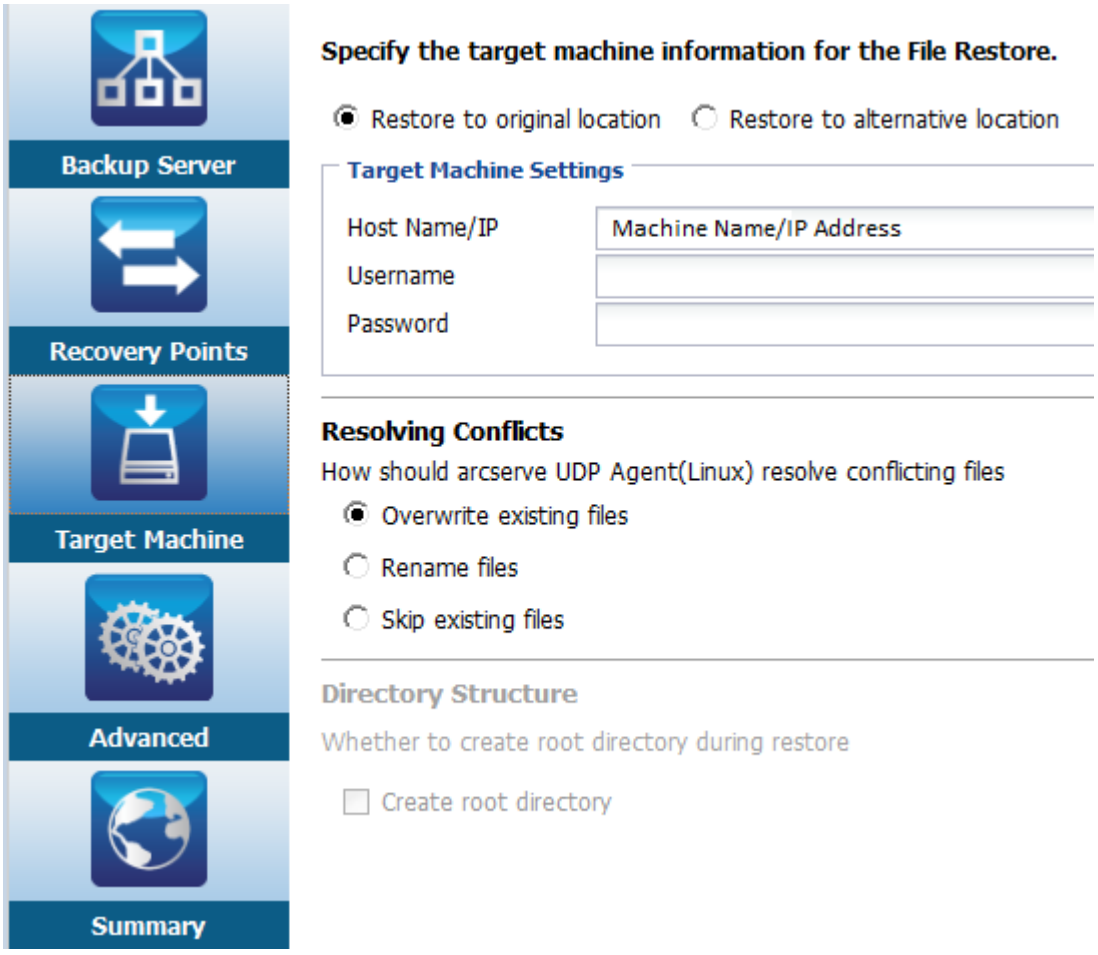
The recovery point is specified.

Specify the Target Machine Details

Specify the target node details so that data is restored to that node. You can restore the selected files or folders to the source node or to a new node.

To restore to the node from where the data was backed up, follow these steps:

1. On the Target Machine page, select **Restore to original location**.



2. Enter the user name and the password of the node.
3. Select one of the following options to resolve conflicting files:

   Overwrite existing files

   Specifies that if the file exists in the target machine then the backup file from the recovery point replaces the existing file.

   Rename files

   Specifies that if the file exists in the target machine, then a new file is created with the same file name and *.d2dduplicate<x>* file extension. <x>

specifies the number of times the file is restored. All the data is restored to the new file.

Skip existing files

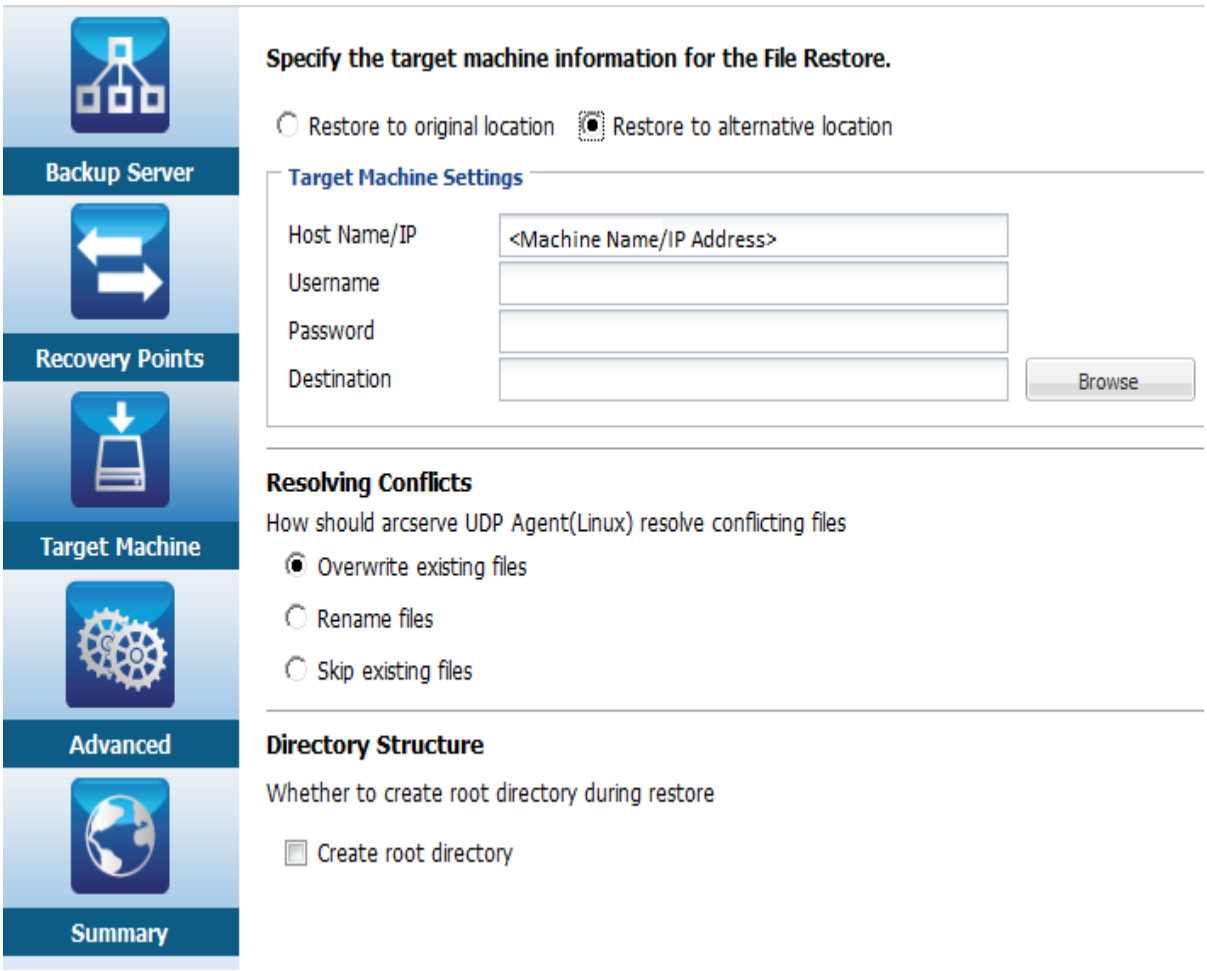
Specifies that if the same file exists in the target machine, then those files are not restored from the recovery point.

4. (Optional) Select **Create root directory**.
5. Click **Next**.

   The Advanced page opens.

To restore to a new node, follow these steps:

1. On the Target Machine page, select **Restore to alternative location**.



2. Enter the host name or IP address of the target node.
3. Enter the user name and the password of the node.

4. Enter the path where the data is restored, or click **Browse** to select the folder where the data is restored, and then click **OK**.
5. Select one of the following options to resolve conflicting files:

   Overwrite existing files

   Specifies that if the file exists in the target machine then the backup file from the recovery point replaces the existing file.

   Rename files

   Specifies that if the file exists in the target machine, then a new file is created with the same file name and *.d2dduplicate<x>* file extension. <x> specifies the number of times the file is restored. All the data is restored to the new file.

   Skip existing files

   Specifies that if the same file exists in the target machine then those files are not restored from the recovery point.
6. (Optional) Select **Create root directory**.
7. Click **Next**.

   The Advanced page opens.

The target machine details are specified.

Specify the Advanced Settings

Specify the advanced settings to perform a scheduled recovery of your data. Scheduled recovery ensures that your data is recovered at the specified time even in your absence.

Follow these steps:

1. Set the start date and time by selecting one of the following options:

   Run Now

   Starts the file-level restore job as soon as you submit the job.

   Set Starting Date and Time

   Starts the file-level restore job at the specified date and time after submitting the job.

2. (Optional) Select **Estimate file size**.
3. (Optional) Select a script from the **Pre/Post Scripts Settings** option.

   These scripts run script commands for actions to take before the start of the job and/or upon the completion of the job.

   Note: The **Pre/Post Script Settings** fields are populated only if you already created a script file and placed it at the following location:

   /opt/Arcserve/d2dserver/usr/prepost

   Note: For more information about creating the pre/post scripts, see *Manage Pre/Post Scripts for Automation*.

4. Click **Next**.

   The **Summary** page opens.

   The advanced settings are specified.

(Optional) Manage Pre/Post Scripts for Automation

Pre/Post scripts let you run your own business logic at specific stages of a running job. You can specify when to run your scripts in **Pre/Post Script Settings** of the **Backup Wizard** and the **Restore Wizard** in the UI. The scripts can be run on the Backup Server depending on your setting.

Managing the pre/post script is a two part process, consisting of creating the pre/-post script and placing the script in the prepost folder.

Create Pre/Post Scripts

Follow these steps:

1. Log into the Backup Server as a root user.
2. Create a script file using the environment variables in your preferred scripting language.

Pre/Post Script Environment Variables

To create your script, use the following environment variables:

D2D_JOBNAME

Identifies the name of the job.

D2D_JOBID

Identifies the job ID. Job ID is a number provided to the job when you run the job. If you run the same job again, you get a new job number.

D2D_TARGETNODE

Identifies the node that is being backed up or restored.

D2D_JOBTYPE

Identifies the type of the running job. The following values identify the D2D_ JOBTYPE variable:

backup.full

Identifies the job as a full backup.

backup.incremental

Identifies the job as an incremental backup.

backup.verify

Identifies the job as a verify backup.

restore.bmr

Identifies the job as a bare-metal recovery (BMR). This is a restore job.

restore.file

Identifies the job as a file-level restore. This is a restore job.

D2D_SESSIONLOCATION

Identifies the location where the recovery points are stored.

D2D_PREPOST_OUTPUT

Identifies a temp file. The content of the first line of the temp file is displayed in the activity log.

D2D_JOBSTAGE

Identifies the stage of the job. The following values identify the D2D_JOBSTAGE variable:

pre-job-server

Identifies the script that runs on the Backup Server before the job starts.

post-job-server

Identifies the script that runs on the Backup Server after the job completes.

pre-job-target

Identifies the script that runs on the target machine before the job starts.

post-job-target

Identifies the script that runs on the target machine after the job completes.

pre-snapshot

Identifies the script that runs on the target machine before capturing the snapshot.

post-snapshot

Identifies the script that runs on the target machine after capturing the snapshot.

D2D_TARGETVOLUME

Identifies the volume that is backed up during a backup job. This variable is applicable for pre/post snapshot scripts for a backup job.

D2D_JOBRESULT

Identifies the result for a post job script. The following values identify the D2D_JOBRESULT variable:

success

Identifies the result as successful.

fail

Identifies the result as unsuccessful.

D2DSVR_HOME

Identifies the folder where Backup Server is installed. This variable is applicable for the scripts that run on the Backup Server.

The script is created.

Note: For all scripts, a return value of zero indicates success and a nonzero return value indicates failure.

Place the Script in the Prepost Folder and Verify

All the pre/post scripts for a Backup Server are centrally managed from the prepost folder at the following location:

/opt/Arcserve/d2dserver/usr/prepost

Follow these steps:

1. Place the file in the following location of the Backup Server:

   /opt/Arcserve/d2dserver/usr/prepost
2. Provide the execution permission to the script file.
3. Log into the Arcserve Unified Data Protection Agent for Linux web interface.
4. Open the **Backup Wizard** or the **Restore Wizard** and navigate to the **Advanced** tab.
5. Select the script file in the **Pre/Post Script Settings** drop-down list and then submit the job.
6. Click **Activity Log** and verify that the script is executed to the specified backup job.

   The script is executed.

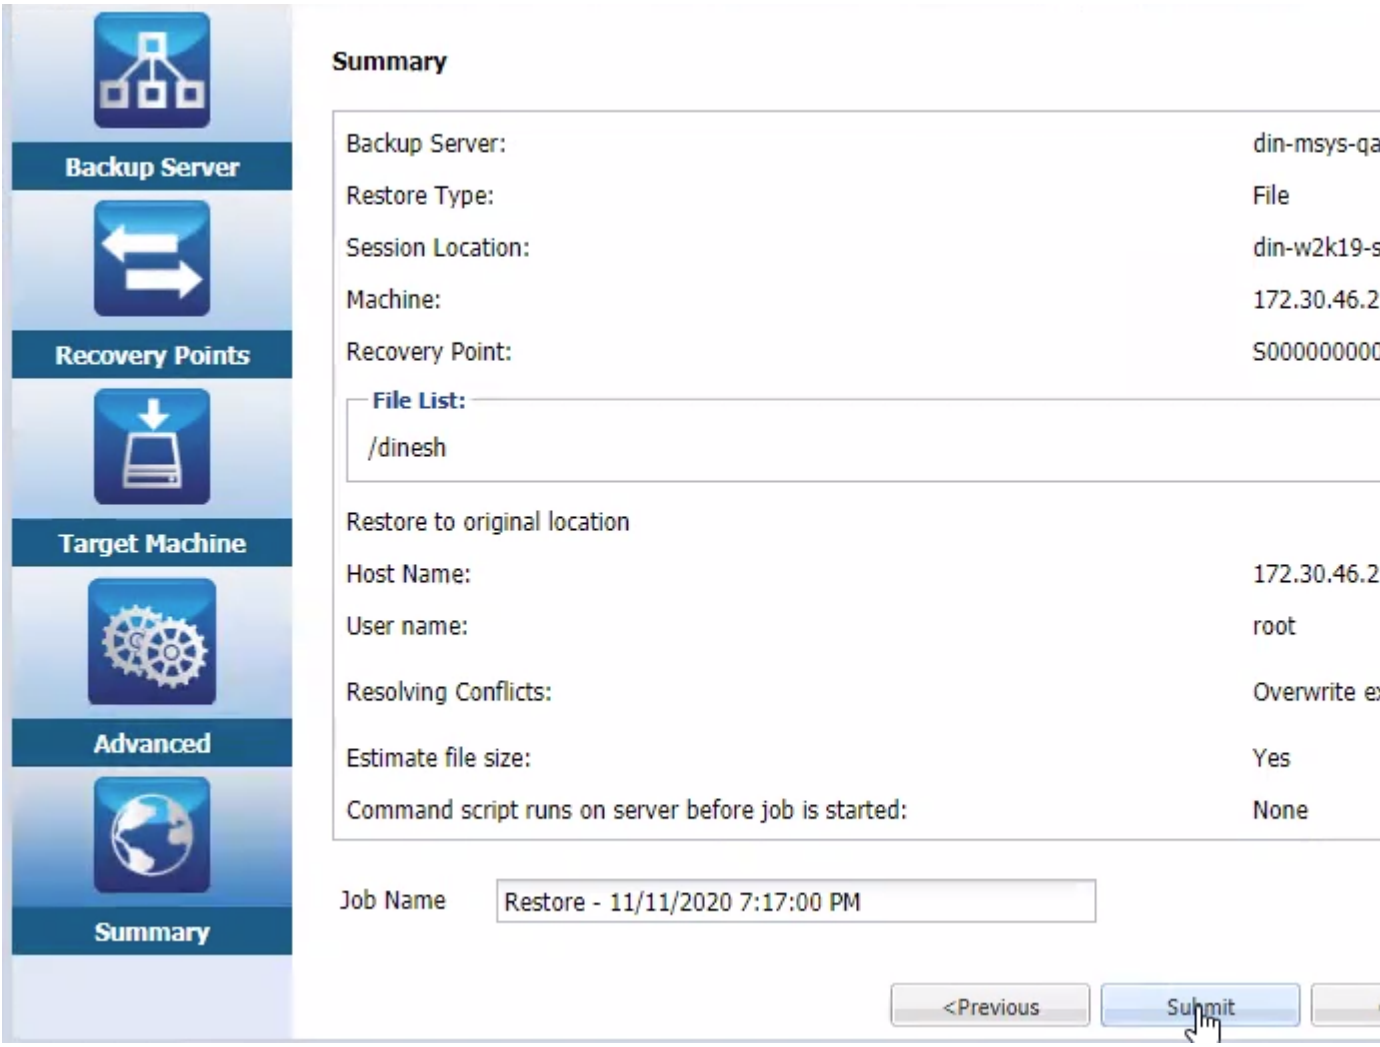
   The pre/post scripts are successfully created and placed in the prepost folder.

Create and Run the Restore Job

Create and run the restore job so that you can initiate the file-level recovery. Verify the recovery point information before you restore the files. If needed, go back and change the restore settings on the wizard.

Follow these steps:

1. On the Summary page of the Restore Wizard, verify the restore details.



2. Do one of the following:
   - If the summary information is incorrect, click **Previous** and go back to the applicable dialog to change the incorrect setting.
   - If the summary information is correct, enter a job name, and then click **Submit** to launch the restore process.

Note: The Job Name field has a default name initially. You can enter a new job

name of your choice, but you cannot leave this field empty.

The Restore Wizard closes. You can see the status of the job on the Job Status page.

The restore job is successfully created and run.

Verify that Files are Restored

After the completion of restore job, verify that all the files are restored in the target node. Check the **Job History** and **Activity Log** tabs in the **Status** pane to monitor the progress of the restore process.

Follow these steps:

1. Navigate to the target machine where you restored data.
2. Verify that the required data from the recovery point is restored.

   The files are successfully verified.

   The file-level recovery is successfully performed.

How to Perform a File-Level Recovery from Host-Based Agentless Backup for Linux Nodes

A file-level recovery restores individual files and folders from a recovery point. You can restore as minimum as one file from the recovery point. This option is useful if you want to restore selected files and not the entire recovery point.

Perform these tasks for a file-level recovery:

- Review the Restore Prerequisites
- Specify the Recovery Point
- Specify the Target Machine Details
- Specify the Advanced Settings
- (Optional) Manage Pre/Post Scripts for Automation
- Create and Run the Restore Job
- Verify that Files are Restored

Review the Prerequisites

Consider the following options before you perform a file-level recovery:

- You have a valid recovery point and the encryption password, if any.
- You have a valid target node to recover data.
- You have verified that the Linux Backup Server supports the file system that you want to restore.

  For example, RedHat 7.x does not support the *reiserfs* file system. If the operating system of the Backup Server is RedHat 7.x and you want to restore the reiserfs file system, you must install the file system driver to support reiserfs. You can also use Arcserve Unified Data Protection Agent for Linux Live CD to perform the file-level restore because Live CD supports all types of file system.

- You have installed the following packages on the Linux Backup Server:
  - mdadm
  - kpartx
  - lvm2
  - Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Specify the Recovery Point

Each time that you perform a backup, a recovery point is created. Specify the recovery point information in the **Restore Wizard** so that you can recover the exact data that you want. You can restore specific files or all files depending on your requirement.

Follow these steps:

1. Access the Restore Wizard in one of the following ways:

   From Arcserve UDP:

   a. Log into Arcserve UDP.

   b. Click the **resources** tab.

   c. Select **All Nodes** in the left pane.

      All the added nodes are displayed in the center pane.

   d. In the center pane, select the node and click **Actions**.

   e. Click **Restore** from the **Actions** drop-down menu.

      The Arcserve Unified Data Protection Agent for Linux web interface opens. The restore type selection dialog is displayed in the Agent UI.

   f. Select the restore type and click **OK**.

      Note: You are automatically logged in to the agent node and the **Restore Wizard** opens from the agent node.

   From Arcserve Unified Data Protection Agent for Linux:

   a. Open the Arcserve Unified Data Protection Agent for Linux web interface.

      Note: During the installation of Arcserve Unified Data Protection Agent for Linux, you received the URL to access and manage the server. Log in to Arcserve Unified Data Protection Agent for Linux.

   b. Click **Restore** from the **Wizard** menu and select **Restore File**.

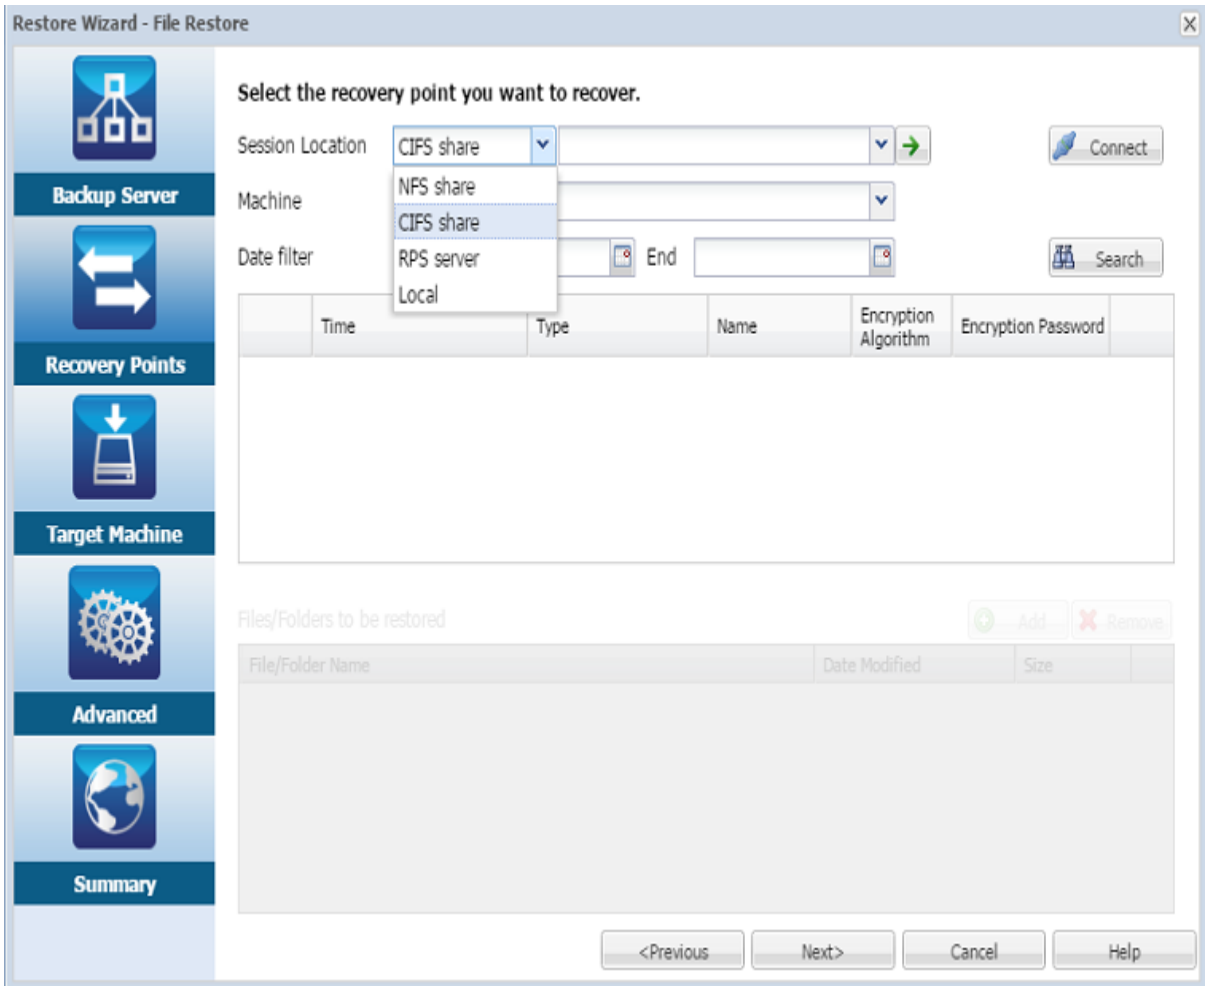
      Restore Wizard - File Restore opens.

      You can see the Backup Server in the **Backup Server** page of the **Restore Wizard**. You cannot select any option from the **Backup Server** drop-down list.

2. Click **Next**.

   The **Recovery Points** page of the **Restore Wizard** opens.

   Important! If you have opened the Wizard from the Console, the session location and machine details are automatically displayed. You can skip to Step 5.



3. Select either a **CIFS share** or **RPS server** from the **Session Location** drop-down list.

   Note: You cannot select NFS share or Local for restoring host-based agentless backup sessions.

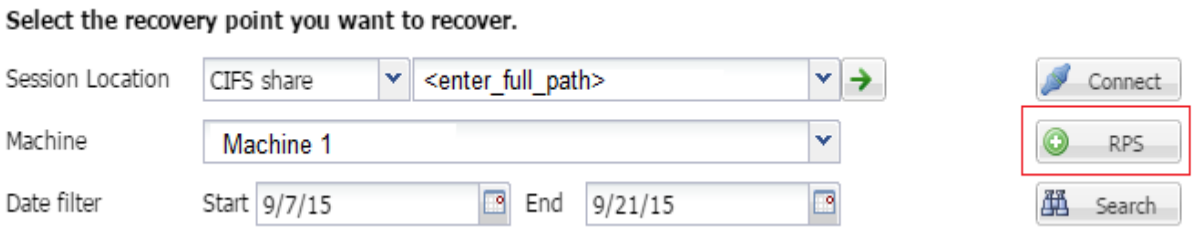
4. Follow one of the following steps depending on your session location:

   For CIFS share

   a. Specify the full path of the CIFS share and click **Connect**.

   b. Specify the username and password to connect to the CIFS share and click **OK**.

      All the machines are listed in the Machine drop-down list and an RPS button appears beside Machine.



c. Select the machine from the drop-down list and click RPS.

   The **Recovery Point Server** Information dialog opens.

d. Provide the RPS details and click Yes.

   The dialog **Recovery Point Server Information** closes. All the recovery points from the selected machine are displayed below the **Date Filter** option.

For RPS server

a. Select RPS server and click Add.

   The **Recovery Point Server Information** dialog opens.

b. Provide the RPS details and click Load

c. Select the data store from the drop-down list and click **Yes**.

   The **Recovery Point Server Information** dialog closes and you see the wizard.

d. Click **Connect**.

   All the machines are listed in the Machine drop-down list.

e. Select the machine from the drop-down list.

   All the recovery points from the selected machine are displayed below the **Date Filter** option.

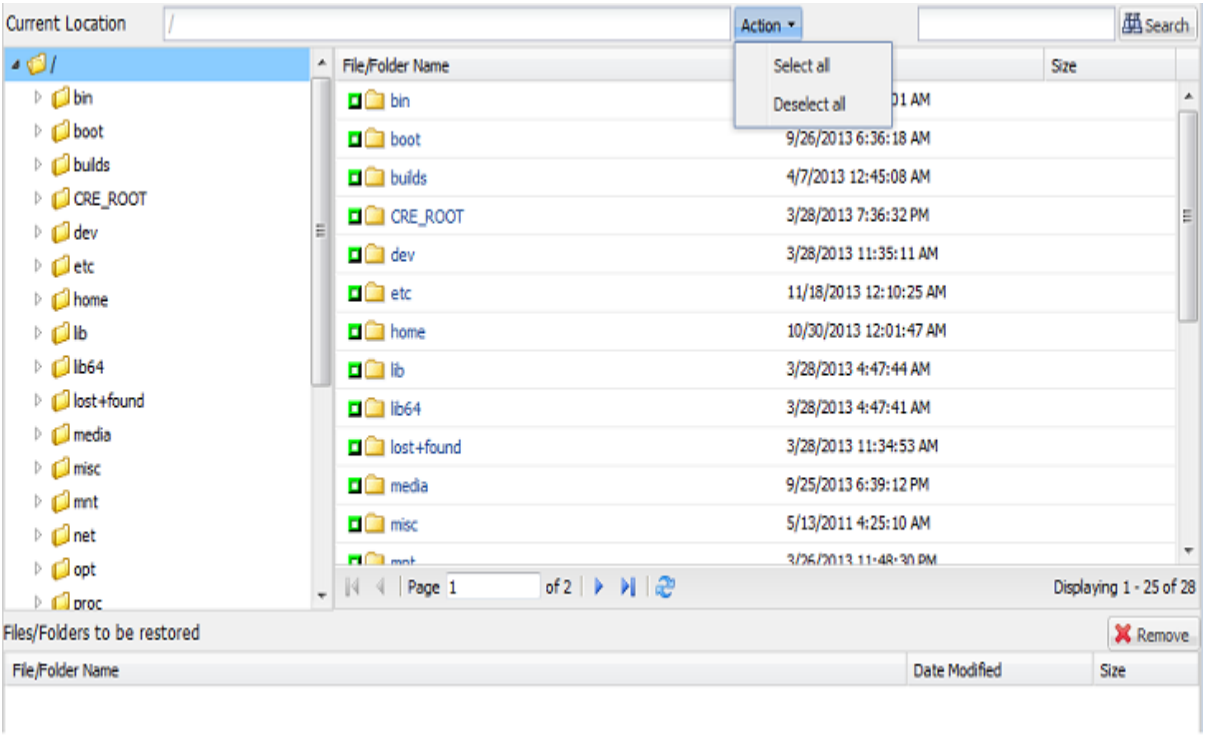
5. Apply the date filter to display the recovery points that are generated between the specified date and click **Search**.

   Default: Recent two weeks.

   All the recovery points available between the specified dates are displayed.

6. Select the recovery point that you want to restore and click **Add**. If the recovery point is encrypted, enter the encryption password to restore data.

   The **Browse-<node name>** dialog opens.



Important! If you see the warning message, "The files/folders are displayed under device file. Click for more information." on the Console, refer the following Note for resolution.

Note: For some complex disk layout, the file system is shown by the device file. The change in the file system display behavior does not affect the function of host-based Linux VM file-level restore. You can browse the file system under the device file. Also, you can use the search function to search specific file or directory.

7. Select the files and folders that you want to restore and click **OK**.

   Note: If you try to locate a file or folder using the **Search** field, ensure that you select the highest folder in the hierarchy. The search is conducted on all the child folders of the selected folder.

   The **Browse-<node name>** dialog closes and you return to the **Recovery Points** page. The selected files and folders are listed under **Files/Folders to be restored**.

8. Click **Next**.

   The **Target Machine** page opens.

   The recovery point is specified.

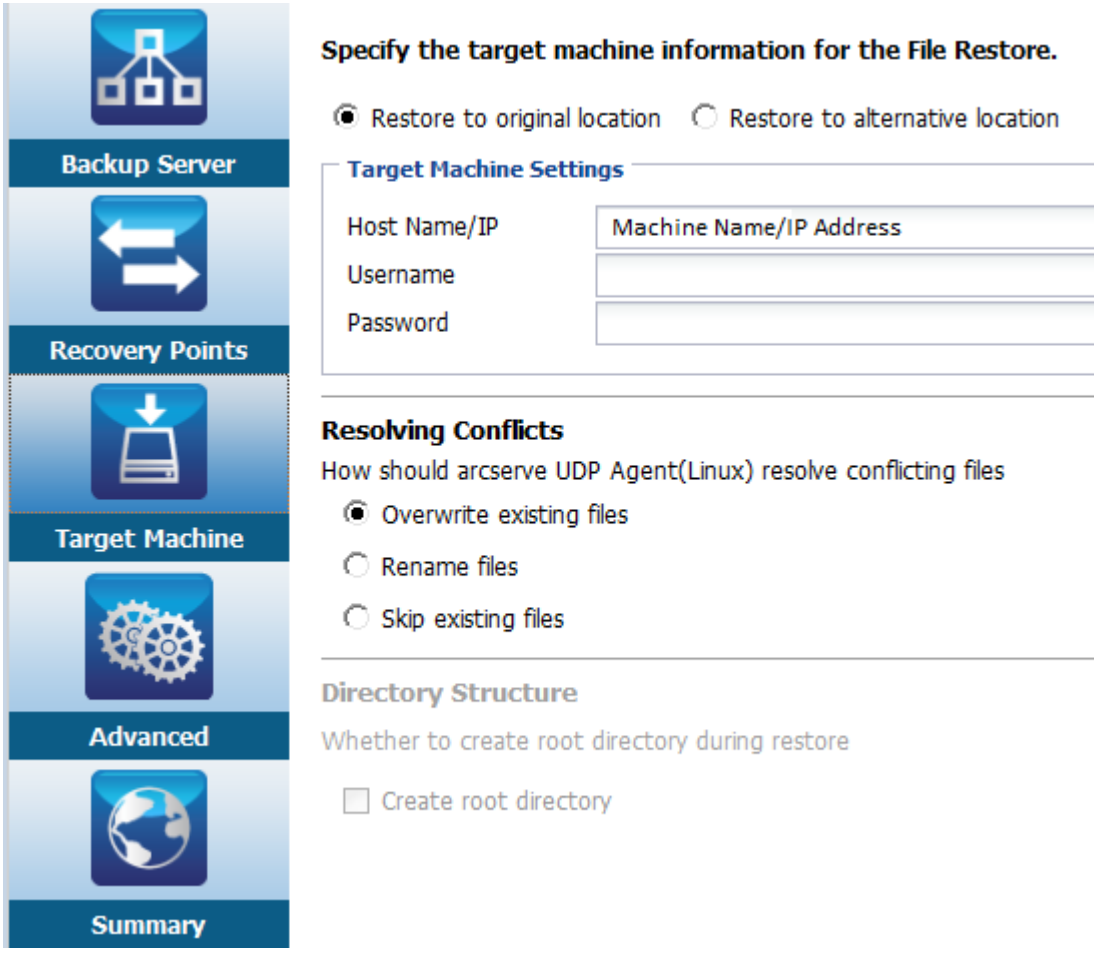
Specify the Target Machine Details

Specify the target node details so that data is restored to that node. You can restore the selected files or folders to the source node or to a new node.

To restore to the node from where the data was backed up, follow these steps:

1. Select **Restore to original location** on the **Target Machine** page.

   The **Host Name** field in **Target Machine Settings** gets populated with the name of the source node.



2. Enter the user name and the password of the node.
3. Select one of the following options to resolve conflicting files:

   Overwrite existing files

   Specifies that if the file exists in the target machine then the backup file from the recovery point replaces the existing file.

   Rename files

Specifies that if the file exists in the target machine, then a new file is created with the same file name and *.d2dduplicate<x>* file extension. <x> specifies the number of times the file is restored. All the data is restored to the new file.

Skip existing files

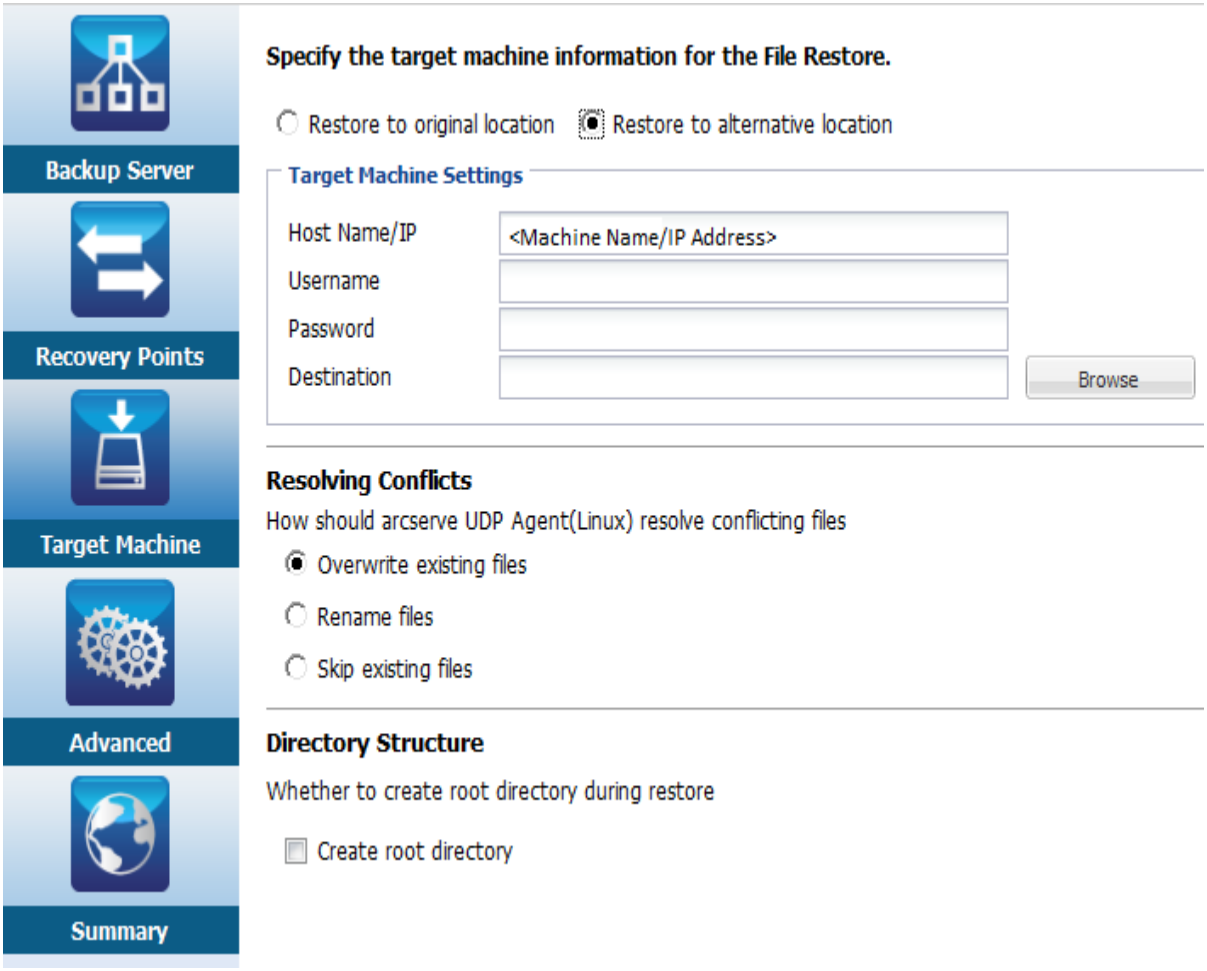
Specifies that if the same file exists in the target machine, then those files are not restored from the recovery point.

4. Click **Next**.

   The **Advanced** page opens.

To restore to a new node, follow these steps:

1. Select **Restore to alternative location** on the **Target Machine** page.



2. Enter the host name or the IP address of the target node.
3. Enter the user name and the password of the node.

4. Enter the path where the data is restored, or click **Browse** to select the folder where the data is restored and click **OK**.
5. Select one of the following options to resolve conflicting files:

   Overwrite existing files

   Specifies that if the file exists in the target machine then the backup file from the recovery point replaces the existing file.

   Rename files

   Specifies that if the file exists in the target machine, then a new file is created with the same file name and *.d2dduplicate<x>* file extension. <x> specifies the number of times the file is restored. All the data is restored to the new file.

   Skip existing files

   Specifies that if the same file exists in the target machine then those files are not restored from the recovery point.
6. (Optional) Select **Create root directory**.
7. Click **Next**.

   The **Advanced** page opens.

The target machine details are specified.

Specify the Advanced Settings

Specify the advanced settings to perform a scheduled recovery of your data. Scheduled recovery ensures that your data is recovered at the specified time even in your absence.

Follow these steps:

1. Set the start date and time by selecting one of the following options:

   Run Now

   Starts the file-level restore job as soon as you submit the job.

   Set Starting Date and Time

   Starts the file-level restore job at the specified date and time after submitting the job.

2. (Optional) Select **Estimate file size**.
3. (Optional) Select a script from the **Pre/Post Scripts Settings** option.

   These scripts run script commands for actions to take before the start of the job and/or upon the completion of the job.

   Note: The **Pre/Post Script Settings** fields are populated only if you already created a script file and placed it at the following location:

   /opt/Arcserve/d2dserver/usr/prepost

   Note: For more information about creating the pre/post scripts, see *Manage Pre/Post Scripts for Automation*.

4. Click **Next**.

   The **Summary** page opens.

   The advanced settings are specified.

(Optional) Manage Pre/Post Scripts for Automation

Pre/Post scripts let you run your own business logic at specific stages of a running job. You can specify when to run your scripts in **Pre/Post Script Settings** of the **Backup Wizard** and the **Restore Wizard** in the UI. The scripts can be run on the Backup Server depending on your setting.

Managing the pre/post script is a two part process, consisting of creating the pre/-post script and placing the script in the prepost folder.

Create Pre/Post Scripts

Follow these steps:

1. Log into the Backup Server as a root user.
2. Create a script file using the environment variables in your preferred scripting language.

Pre/Post Script Environment Variables

To create your script, use the following environment variables:

D2D_JOBNAME

Identifies the name of the job.

D2D_JOBID

Identifies the job ID. Job ID is a number provided to the job when you run the job. If you run the same job again, you get a new job number.

D2D_TARGETNODE

Identifies the node that is being backed up or restored.

D2D_JOBTYPE

Identifies the type of the running job. The following values identify the D2D_ JOBTYPE variable:

backup.full

Identifies the job as a full backup.

backup.incremental

Identifies the job as an incremental backup.

backup.verify

Identifies the job as a verify backup.

restore.bmr

Identifies the job as a bare-metal recovery (BMR). This is a restore job.

restore.file

Identifies the job as a file-level restore. This is a restore job.

D2D_SESSIONLOCATION

Identifies the location where the recovery points are stored.

D2D_PREPOST_OUTPUT

Identifies a temp file. The content of the first line of the temp file is displayed in the activity log.

D2D_JOBSTAGE

Identifies the stage of the job. The following values identify the D2D_JOBSTAGE variable:

pre-job-server

Identifies the script that runs on the Backup Server before the job starts.

post-job-server

Identifies the script that runs on the Backup Server after the job completes.

pre-job-target

Identifies the script that runs on the target machine before the job starts.

post-job-target

Identifies the script that runs on the target machine after the job completes.

pre-snapshot

Identifies the script that runs on the target machine before capturing the snapshot.

post-snapshot

Identifies the script that runs on the target machine after capturing the snapshot.

D2D_TARGETVOLUME

Identifies the volume that is backed up during a backup job. This variable is applicable for pre/post snapshot scripts for a backup job.

D2D_JOBRESULT

Identifies the result for a post job script. The following values identify the D2D_JOBRESULT variable:

success

Identifies the result as successful.

fail

Identifies the result as unsuccessful.

D2DSVR_HOME

Identifies the folder where Backup Server is installed. This variable is applicable for the scripts that run on the Backup Server.

The script is created.

Note: For all scripts, a return value of zero indicates success and a nonzero return value indicates failure.

Place the Script in the Prepost Folder and Verify

All the pre/post scripts for a Backup Server are centrally managed from the prepost folder at the following location:

/opt/Arcserve/d2dserver/usr/prepost

Follow these steps:

1. Place the file in the following location of the Backup Server:

   /opt/Arcserve/d2dserver/usr/prepost
2. Provide the execution permission to the script file.
3. Log into the Arcserve Unified Data Protection Agent for Linux web interface.
4. Open the **Backup Wizard** or the **Restore Wizard** and navigate to the **Advanced** tab.
5. Select the script file in the **Pre/Post Script Settings** drop-down list and then submit the job.
6. Click **Activity Log** and verify that the script is executed to the specified backup job.

   The script is executed.

   The pre/post scripts are successfully created and placed in the prepost folder.

Create and Run the Restore Job

Create and run the restore job so that you can initiate the file-level recovery. Verify the recovery point information before you restore the files. If needed, you can go back and can change the restore settings on the wizard.

Follow these steps:

1. Verify the restore details on the **Summary** page of the **Restore Wizard**.
2. (Optional) Click **Previous** to modify the information that you have entered on any page of the **Restore Wizard**.
3. Enter a job name and click **Submit**.

   The **Job Name** field has a default name initially. You can enter a new job name of your choice but you cannot leave this field empty.

   The **Restore Wizard** closes. You can see the status of the job in the **Job Status** tab.

   The restore job is successfully created and run.

Verify that Files are Restored

After the completion of restore job, verify that all the files are restored in the target node. Check the **Job History** and **Activity Log** tabs in the **Status** pane to monitor the progress of the restore process.

Follow these steps:

1. Navigate to the target machine where you restored data.
2. Verify that the required data from the recovery point is restored.

   The files are successfully verified.

   The file-level recovery is successfully performed.

How to Perform IVM migration (from Linux recovery point) from Cloud to Local

If you have Linux recovery point on Amazon S3, you can perform an IVM job on AWS, and then migrate the IVM from AWS to local.

Complete the following tasks to perform IVM Migration:

- Review the Prerequisites for IVM Migration
- Perform IVM Migration from Cloud to Local

Review the Prerequisites and Consideration for IVM Migration

Prerequisites:

- You have a valid recovery point and the encryption password, if any, for restore.
- You have Amazon EC2 and Amazon S3 account.
- Samba server is installed on the Linux Backup Server.

Consideration:

This feature is supported only on Linux Backup Server that is installed on RHEL/CentOS/Oracle Linux 7.x or SLES 12.x.

Perform IVM Migration from Cloud to Local

You can migrate IVM from Cloud to Local using three detailed procedures. Click links of each procedure and perform to complete migration.

- To create an instant virtual machine on Amazon EC2, run Instant Virtual Machine from Linux VM recovery point on Amazon EC2.
- Install Linux Backup Server inside this instant virtual machine you just created. Run backup inside the IVM, and back up to Amazon S3.
- Back up the linux machine to Amazon S3.
- Run BMR job from the recovery point in Amazon S3 that you backed up.

How to Perform a Bare Metal Recovery (BMR) for Linux Machines

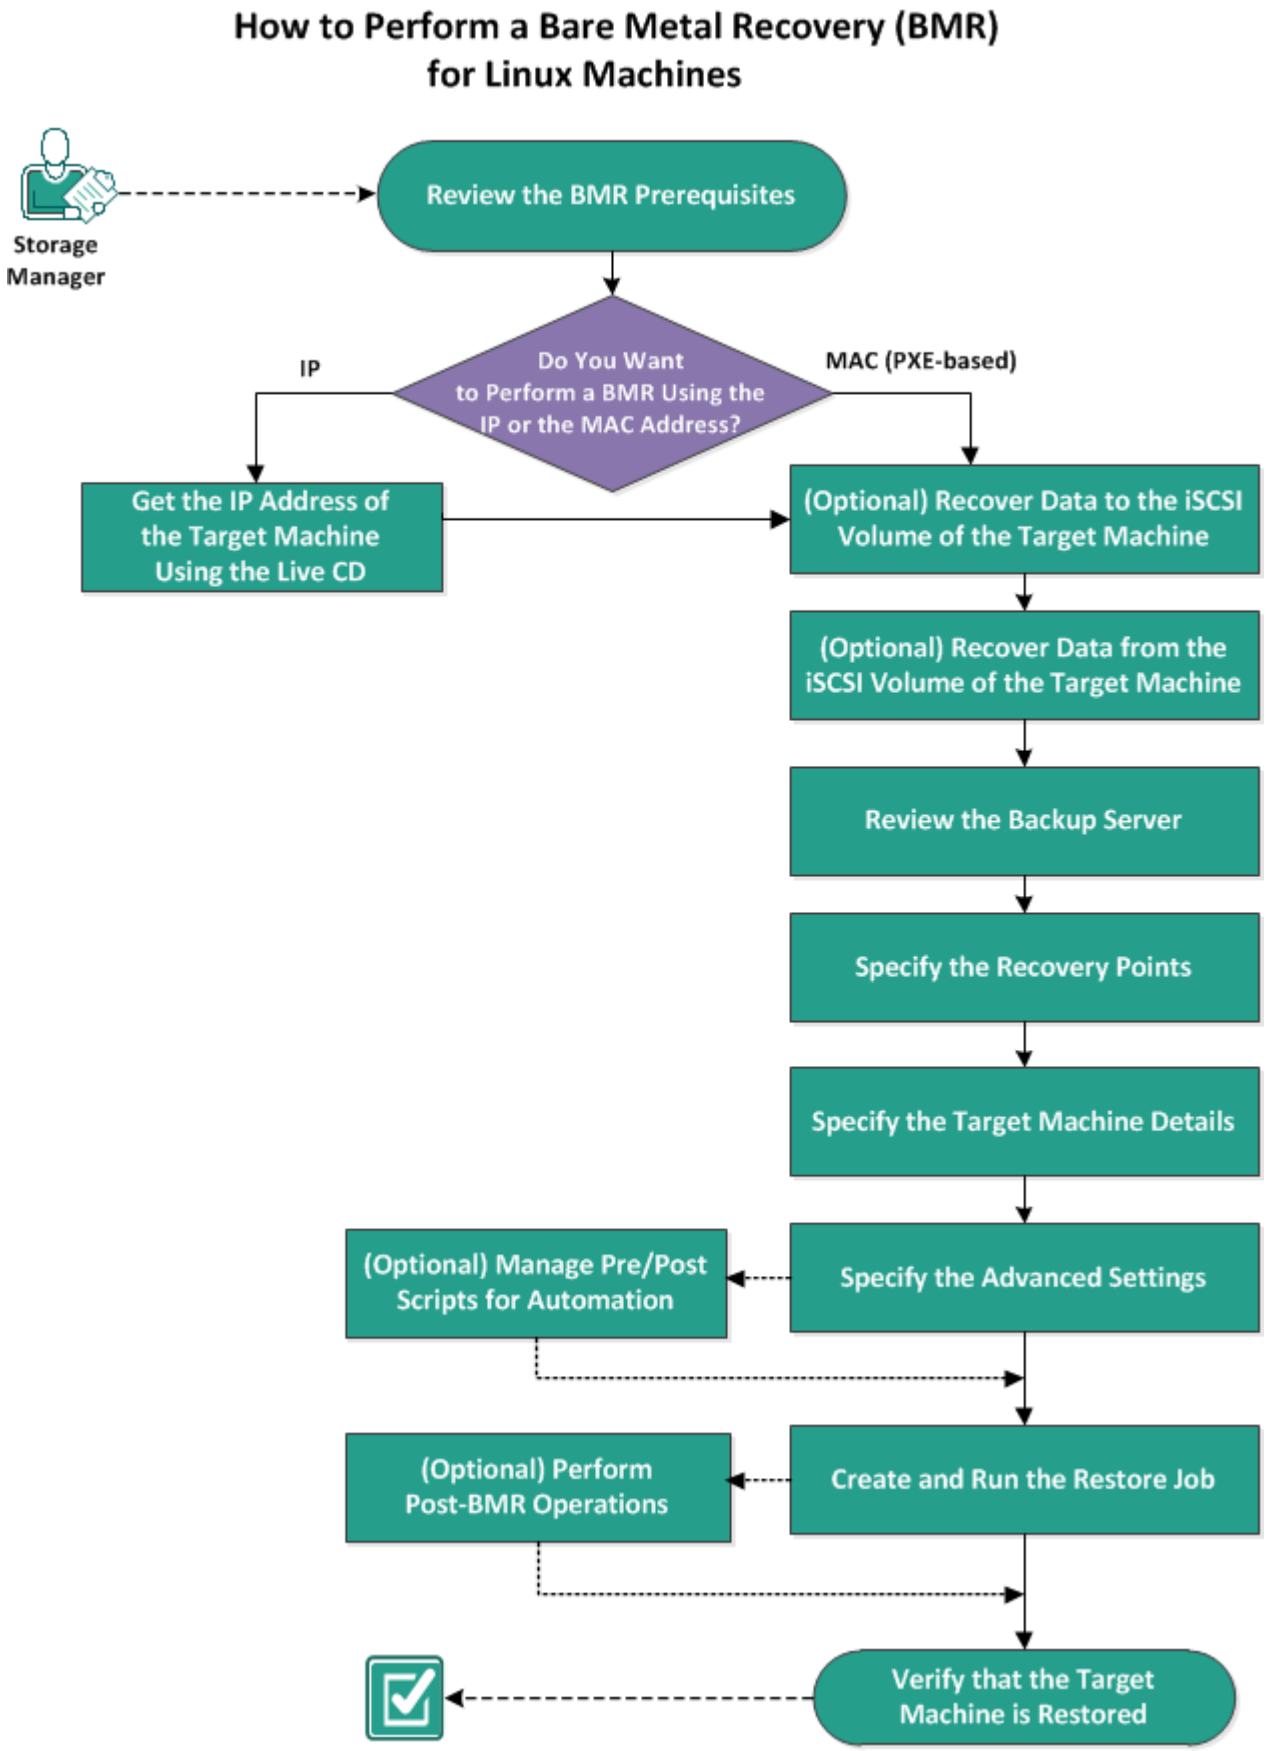
A BMR restores the operating system and software applications, and recovers all the backed-up data. BMR is the process of restoring a computer system from *bare metal*. Bare metal is a computer without any operating system, drivers, and software applications. After the restore is complete, the target machine automatically reboots in the same operating environment as the backup source node and all data is restored.

A complete BMR is possible because when you back up data, the backup also captures information related to the operating system, installed applications, drivers, and so on.

You can perform a BMR using the IP address or the Media Access Control (MAC) address of the target machine. If you boot the target machine using the Arcserve Unified Data Protection Agent for Linux Live CD, you can get the IP address of the target machine.

Note: Machine can boot up. Only one NIC is configured.

The following diagram displays the process to perform a BMR:



Complete the following tasks to perform a BMR:

- Review the BMR Prerequisites
- Get the IP Address of the Target Machine Using the Live CD
- (Optional) Recover Data to the iSCSI Volume of the Target Machine
- (Optional) Recover Data from the iSCSI Volume to the Target Machine
- Review the Backup Server
- Specify the Recovery Points
- Specify the Target Machine Details
- Specify the Advanced Settings
- (Optional) Manage Pre/Posts Scripts for Automation
- Create and Run the Restore Job
- (Optional) Perform Post-BMR Operations
- Verify that the Target Machine is Restored

Review the BMR Prerequisites

Consider the following options before performing a BMR:

- You have a valid recovery point and the encryption password, if any, for restore.
- You have a valid target machine for BMR.
- You have created the Arcserve Unified Data Protection Agent for Linux Live CD.
- If you want to perform a BMR using the IP address, you must get the IP address of the target machine using the Live CD.
- If you want to perform a PXE-based BMR using the MAC address, you must have the MAC address of the target machine.
- The recovery point must be from the Linux agent-based backup.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Get the IP Address of the Target Machine Using the Live CD

Before performing a BMR using the IP address, you need to get the IP address of the target machine. A bare-metal machine does not have any IP address initially. So, you have to boot the bare-metal machine using the default Live CD, which is Arcserve Unified Data Protection Agent for Linux Live CD, or the CentOS-based Live CD to get the IP address. After you get the IP address of the target machine, you can configure the static IP of the target machine.

Follow these steps:

1. Insert the Live CD or mount the .iso file of the Live CD into the CD-ROM drive of the target node.
2. Boot the target machine from CD-ROM.

   The target machine boots into the Arcserve Unified Data Protection Agent for Linux Live CD environment. On the screen, the IP address of the target machine is displayed.
3. To configure the static IP of the target machine using the default Live CD, follow these steps:
   a. On the target machine's screen, press Enter to enter the shell environment.
   b. Run the following command to configure the static IP:

      ifconfig <NIC name> <static IP address> netmask <netmask>

      route add default gw <gateway IP address> <NIC name>

   Note: The Network Interface Card (NIC) name depends on your hardware. For example, the typical NIC names are eth0 or em0.
4. To configure the static IP of the target machine using the CentOS-based Live CD, follow these steps:
   a. Open a terminal window on the target machine by clicking Applications, System Tools, Terminal.
   b. Run the following commands:

      sudo ifconfig <NIC name> <static IP address> netmask <netmask>

      sudo route add default gw <gateway IP address> <NIC name>

   The static IP is configured.

   The IP address of the target machine is acquired.

Important! Maintain a record of this IP address as it is used in the **Restore Wizard** when you have to specify the target machine details.

(Optional) Recover Data to the iSCSI Volume of the Target Machine

You can integrate the iSCSI volume to the target machine and make that volume a part of the target machine. Then you can restore data to the iSCSI volume of the target machine. By doing so, you can manage data and transfer data over a network.

Important! When you integrate the iSCSI volume with the target machine, you will lose all the existing data from the iSCSI volume.

Follow these steps:

1. Insert the Arcserve Unified Data Protection Agent for Linux Live CD or mount the iso file of the Arcserve Unified Data Protection Agent for Linux Live CD into the CD-ROM drive of the target machine.
2. Boot the target machine from the CD-ROM.

   The target machine boots into the Arcserve Unified Data Protection Agent for Linux Live CD environment. On the screen, the IP address of the target machine is displayed.
3. Enter the shell environment of the target machine.
4. Run the following command to start the iSCSI initiator daemon:

   /etc/init.d/iscsid start
5. Run a discovery script to discover the iSCSI target host.

   iscsiadm -m discovery -t sendtargets -p <ISCSI-SERVER-IP-ADDRESS>:<Port_Number>

   The default port value of iSCSI target host is 3260.
6. Make a note of the iSCSI Qualified Name (IQN) of the iSCSI target host found by the discovery script before you manually log into the discovered target.
7. List the available block device of the target node.

   #fdisk -l
8. Log in to the discovered target.

   iscsiadm -m node -T <iSCSI Target IQN name> -p <ISCSI-SERVER-IP-ADDRESS>:<Port_Number> -l

   You can see a block device in the /dev directory of the target node.
9. Run the following command to obtain the new device name:

   #fdisk -l

   You can see an additional device named /dev/sd<x> on the target node.

The iSCSI volume is integrated with the target volume.

(Optional) Recover Data from the iSCSI Volume to the Target Machine

If you have stored your data in an iSCSI target volume, you can connect to the iSCSI volume and recover data. The iSCSI volume lets you manage data and transfer data over a network.

Follow these steps:

1. Insert the Arcserve Unified Data Protection Agent for Linux Live CD or mount the iso file of the Arcserve Unified Data Protection Agent for Linux Live CD into the CD-ROM drive of the target machine.
2. Boot the target machine from the CD-ROM.

   The target machine boots into the Arcserve Unified Data Protection Agent for Linux Live CD environment. On the screen, the IP address of the target machine is displayed.
3. Enter the shell environment of the target machine.
4. Run the following command to start the iSCSI initiator daemon:

   /etc/init.d/iscsid start
5. Run a discovery script to discover the iSCSI target host.

   iscsiadm -m discovery -t sendtargets -p <ISCSI-SERVER-IP-ADDRESS>:<Port_Number>

   The default port value of iSCSI target host is 3260.
6. Make a note of the iSCSI Qualified Name (IQN) of the iSCSI target host found by the discovery script before you manually log into the discovered target.
7. List the available block device of the target node.

   #fdisk -l
8. Log in to the discovered target.

   iscsiadm -m node -T <iSCSI Target IQN name> -p <ISCSI-SERVER-IP-ADDRESS>:<Port_Number> -l

   You can see a block device in the /dev directory of the target node.
9. Run the following command to obtain the new device name:

   #fdisk -l

   You can see an additional device named /dev/sd<x> on the target node.

   For example, consider the name of the device is /dev/sdc. This device name is used to create a partition and a file system in the following steps.

10. Mount the iSCSI volume using the following commands:

    # mkdir /iscsi

    # mount /dev/sdc1 /iscsi

    Note: When you specify the session location in the Restore Wizard, you need to select Local and enter the path /iscsi.

    Example: <path>/iscsi

    The target machine can now connect to the iSCSI volume and can recover data from the iSCSI volume.

Review the Backup Server

When you open the **Restore Wizard**, review the Backup Server where you want to perform the restore operation.

Follow these steps:

1. Access the Restore Wizard in one of the following ways:

   From Arcserve UDP:

   a. Log in to Arcserve UDP.

   b. Click the **resources** tab.

   c. Select **All Nodes** in the left pane.

      All the added nodes are displayed in the center pane.

   d. In the center pane, select the node and click **Actions**.

   e. Click **Restore** from the **Actions** drop-down menu.

      The Arcserve Unified Data Protection Agent for Linux web interface opens. The restore type selection dialog is displayed in the Agent UI.

   f. Select the restore type and click **OK**.

      Note: You are automatically logged in to the agent node and the **Restore Wizard** opens from the agent node.

   From Arcserve Unified Data Protection Agent for Linux:

   a. Open the Arcserve Unified Data Protection Agent for Linux web interface.

      Note: During the installation of Arcserve Unified Data Protection Agent for Linux, you received the URL to access and manage the server. Log in to Arcserve Unified Data Protection Agent for Linux.

   b. Click **Restore** from the **Wizard** menu and select **Bare Metal Recovery (BMR)**.

      The **Backup Server** page of the **Restore Wizard - BMR** opens.

2. Verify the server from the **Backup Server** drop-down list in the **Backup Server** page.

   You cannot select any option from the **Backup Server** drop-down list.

3. Click **Next**.

   The **Recovery Points** page of the **Restore Wizard - BMR** opens.

   The Backup Server is specified.

Specify the Recovery Points

Each time that you perform a backup, a recovery point is created. Specify the recovery point information in the **Restore Wizard** so that you can recover the exact data that you want. You can restore specific files or all files depending on your requirement.

Important! To perform a BMR from a recovery point, the root volume and the boot volume must be present in the recovery point.

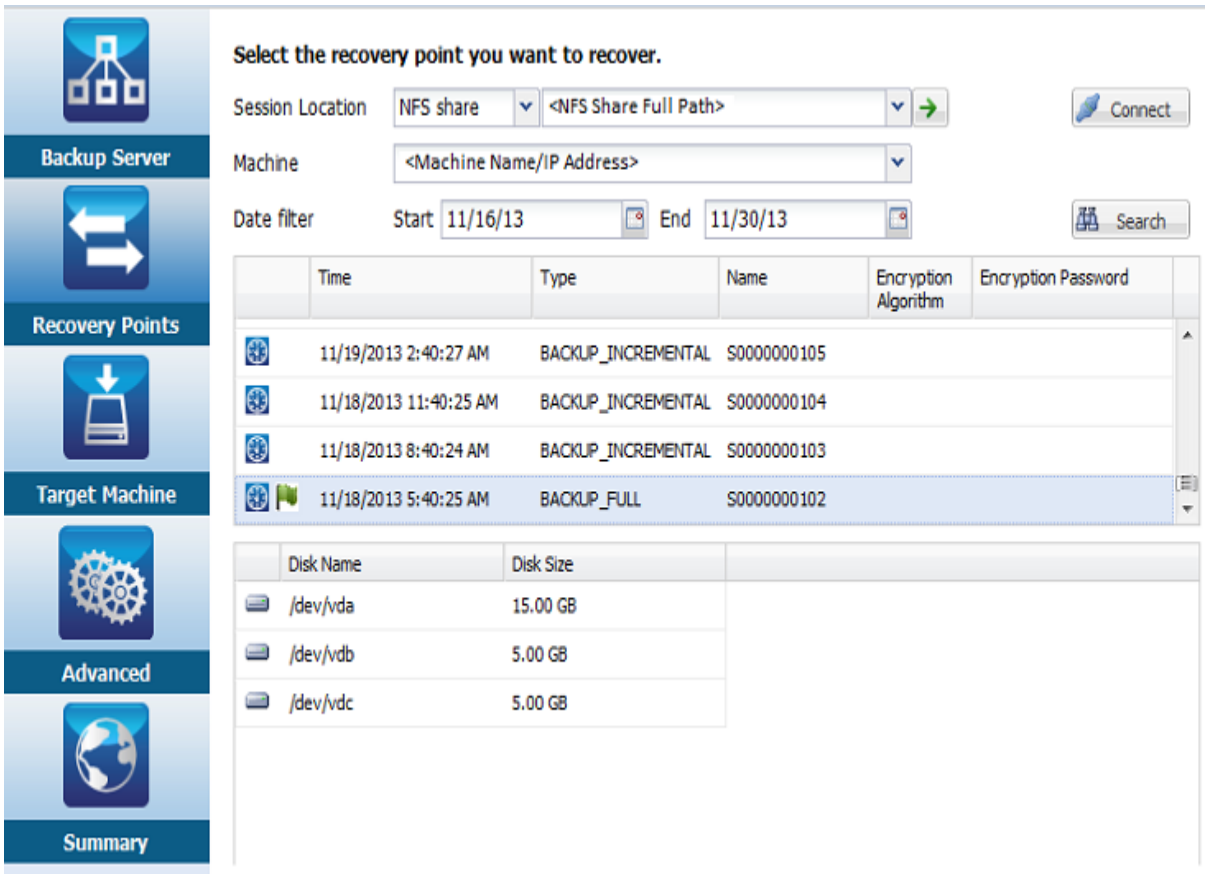
Follow these steps:

1. Perform one of the following steps depending on your backup storage.
   - Perform the following steps to access the recovery points if the recovery points are stored on a mobile device:
     a. Start the target machine using the Live CD.
     b. Log into the Arcserve Unified Data Protection Agent for Linux web interface from the Live CD.
     c. Open the **BMR Wizard**.
     d. Navigate to the **Recovery Points** page.
     e. Select **Local** as the **Session Location** on the **Recovery Points** page of the **BMR Wizard**.
   - Perform the following steps if the session location is NFS share or CIFS share:
     a. Select a session from the **Session Location** drop-down list and enter the full path of the share.

        For example, consider the Session Location as NFS share, xxx.xxx.xxx.xxx as the IP address of the NFS share, and the folder name is *Data.* You would enter xxx.xxx.xxx.xxx:/Data as the NFS share location.

        Note: If the backed up data is stored in Source local, then you must first convert the source node to an NFS server, and then share the ses-

sion location.



2. Click **Connect**.

   All the nodes that have been backed up to this location get listed in the **Machine** drop-down list.

3. Select the node that you want to restore from the **Machine** drop-down list.

   All the recovery points of the selected node get listed.

4. Apply the date filter to display the recovery points that are generated between the specified date and click **Search**.

   Default: Recent two weeks.

   All the recovery points available between the specified dates are displayed.

5. Select the recovery point that you want to restore and click **Next**.

   The **Target Machine** page opens.

   The recovery point is specified.

Specify the Target Machine Details

Specify the target machine details so that data is restored to that machine. A target machine is a bare metal machine where you will perform a BMR. If you restore using the IP address, you need the IP address of the target machine that you previously recorded at the beginning of this process. If you restore using the Media Access Control (MAC) address, you need the MAC address of the target machine.

Follow these steps:

1. Enter the MAC address or the IP address of the target machine in the **MAC/IP Address** field.
2. Enter a name in the **Host Name** field.

   The target machine will use this name as the host name after the restore process is complete.
3. Select one of the following options as the network:

   DHCP

   Automatically configures the IP address. This is the default option. Use this option if you have a Dynamic Host Configuration Protocol (DHCP) server to restore with the DHCP network.

   Static IP

   Manually configures the IP address. If you select this option, then enter the **IP Address**, **Subnet Mask**, and **Default Gateway** of the target machine.

   Important! Ensure that the Static IP is not used by any other machines on the network during the restore process.
4. (Optional) Select the **Enable instant BMR** option so that you can use the target machine instantly.

   When you enable this option, Arcserve Unified Data Protection Agent for Linux first recovers all the necessary data that is required to start the machine. The remaining data are recovered after the target machine is started. The network connection must be constantly available during instant BMR.

   Example: If you have 100-GB data and you want to perform a BMR and you *do not* select this option, first all 100-GB data will be recovered and then you can use the target machine. However, only around 1-GB data is required to start the machine. When you enable this option, first the required 1-GB data is recovered so that you can start and use the machine. After the machine is started, the remaining 99-GB data is automatically recovered.

Note: The necessary data that is required to start the machine depends on the operating system configuration. You can also pause or resume the auto recovery of data if the **Do not recover data automatically after machine is started** option is not selected.

5. (Optional) Select the **Do not recover data automatically when machine is started** option to stop the automatic recovery of data when the target machine is started.

   When you select the **Enable instant BMR** option, the default behavior is to recover the necessary data first and start the machine. After the machine starts, the remaining data gets recovered automatically. If you update any source data during the recovery, then by selecting this option, the data will be recovered until the point before they are updated.

6. Click **Next**.

   The **Advanced** page opens.

   The target machine details are specified.

Specify the Advanced Settings

Specify the advanced settings to perform a scheduled BMR of your data. Scheduled BMR ensures that your data is recovered at the specified time even in your absence.

Follow these steps:

1. Set the start date and time by selecting one of the following options:

   Run Now

   Starts the restore job as soon as you submit the job.

   Set Special Time

   Starts the restore job at the specified time after submitting the job.

2. (Optional) Select a script from the **Pre/Post Scripts Settings** option for the Backup Server and the target machine.

   These scripts run script commands for actions to take before the start of the job and/or upon the completion of the job.

   Note: The **Pre/Post Script Settings** fields are populated only if you already created a script file and placed it at the following location:

   /opt/Arcserve/d2dserver/usr/prepost

   Note: For more information about creating the pre/post scripts, see *Manage Pre/Post Scripts for Automation*.

3. (Optional) Click **Show More Settings** to display more settings for BMR.
4. (Optional) Reset the password for the specified user name for the recovered target machine.
5. (Optional) Enter the full path of the backup storage location of the recovery points in **Recover Point Local Access**.
6. (Optional) Enter the full name of the disk in the **Disks** field to exclude those disks on the target machine from participating in the recovery process.
7. (Optional) Select **Enable Wake-on-LAN** if you are performing Preboot Execution Environment (PXE) BMR.

   Note: The **Enable Wake-on-LAN** option is applicable only for physical machines. Ensure whether you have enabled the Wake-on-LAN settings in the BIOS settings of your physical machine.

8. (Optional) Select the **Reboot** option to automatically restart the target node after the BMR is complete.

9. Click **Next**.

   The **Summary** page opens.

   The advanced settings are specified.

(Optional) Manage Pre/Post Scripts for Automation

Pre/Post scripts let you run your own business logic at specific stages of a running job. You can specify when to run your scripts in **Pre/Post Script Settings** of the **Backup Wizard** and the **Restore Wizard** in the UI. The scripts can be run on the Backup Server depending on your setting.

Managing the pre/post script is a two part process, consisting of creating the pre/-post script and placing the script in the prepost folder.

Create Pre/Post Scripts

Follow these steps:

1. Log into the Backup Server as a root user.
2. Create a script file using the environment variables in your preferred scripting language.

Pre/Post Script Environment Variables

To create your script, use the following environment variables:

D2D_JOBNAME

Identifies the name of the job.

D2D_JOBID

Identifies the job ID. Job ID is a number provided to the job when you run the job. If you run the same job again, you get a new job number.

D2D_TARGETNODE

Identifies the node that is being backed up or restored.

D2D_JOBTYPE

Identifies the type of the running job. The following values identify the D2D_ JOBTYPE variable:

backup.full

Identifies the job as a full backup.

backup.incremental

Identifies the job as an incremental backup.

backup.verify

Identifies the job as a verify backup.

restore.bmr

Identifies the job as a bare-metal recovery (BMR). This is a restore job.

restore.file

Identifies the job as a file-level restore. This is a restore job.

D2D_SESSIONLOCATION

Identifies the location where the recovery points are stored.

D2D_PREPOST_OUTPUT

Identifies a temp file. The content of the first line of the temp file is displayed in the activity log.

D2D_JOBSTAGE

Identifies the stage of the job. The following values identify the D2D_JOBSTAGE variable:

pre-job-server

Identifies the script that runs on the Backup Server before the job starts.

post-job-server

Identifies the script that runs on the Backup Server after the job completes.

pre-job-target

Identifies the script that runs on the target machine before the job starts.

post-job-target

Identifies the script that runs on the target machine after the job completes.

pre-snapshot

Identifies the script that runs on the target machine before capturing the snapshot.

post-snapshot

Identifies the script that runs on the target machine after capturing the snapshot.

D2D_TARGETVOLUME

Identifies the volume that is backed up during a backup job. This variable is applicable for pre/post snapshot scripts for a backup job.

D2D_JOBRESULT

Identifies the result for a post job script. The following values identify the D2D_JOBRESULT variable:

success

Identifies the result as successful.

fail

Identifies the result as unsuccessful.

D2DSVR_HOME

Identifies the folder where Backup Server is installed. This variable is applicable for the scripts that run on the Backup Server.

The script is created.

Note: For all scripts, a return value of zero indicates success and a nonzero return value indicates failure.

Place the Script in the Prepost Folder and Verify

All the pre/post scripts for a Backup Server are centrally managed from the prepost folder at the following location:

/opt/Arcserve/d2dserver/usr/prepost

Follow these steps:

1. Place the file in the following location of the Backup Server:

   /opt/Arcserve/d2dserver/usr/prepost
2. Provide the execution permission to the script file.
3. Log into the Arcserve Unified Data Protection Agent for Linux web interface.
4. Open the **Backup Wizard** or the **Restore Wizard** and navigate to the **Advanced** tab.
5. Select the script file in the **Pre/Post Script Settings** drop-down list and then submit the job.
6. Click **Activity Log** and verify that the script is executed to the specified backup job.

   The script is executed.

   The pre/post scripts are successfully created and placed in the prepost folder.

Create and Run the Restore Job

Create and run the restore job so that you can initiate the process of BMR. Verify the recovery point information before you perform a BMR. If needed, you can go back and can change the restore settings.

Follow these steps:

1. Verify the restore details on the **Summary** page of the **Restore Wizard**.
2. (Optional) Click **Previous** to modify the restore settings on any of the **Restore Wizard** pages.
3. Enter a job name and click **Submit**.

   The **Job Name** field has a default name initially. You can enter a new job name of your choice but you cannot leave this field empty.

   The **Restore Wizard** closes. You can see the job in the **Job Status** tab. If you use the IP address for the BMR, the target machine automatically reboots to the same operating system as the backup source after the BMR process.

   If you use the MAC address for BMR, the status in the **Job Status** tab changes to *Waiting for target node startup*.
4. (Optional) For BMR using the MAC address, start the target machine when you see the *Waiting for target node startup* message in the **Job Status** tab.

   Note: If the target machine is already started before you submit the restore job, you must restart the target machine. Ensure that BIOS is configured to boot from the network.

   The status in the **Job Status** column changes to **Restoring volume**. This indicates the restore is in progress. After the restore job is complete, the target machine automatically reboots with the same operating system as the backup source.

   The restore job was successfully created and run.

(Optional) Perform Post-BMR Operations

The following topics are optional configuration settings that you may have to perform after a BMR:

Configure X Windows

When you perform a BMR across a dissimilar hardware, X Windows of the restored OS does not function properly and the target node displays an error dialog. The error dialog appears because the display configuration has changed. To resolve this error, follow the instructions in the error dialog to configure the graphic card. After that, you can see the X Windows and the desktop UI.

Configure the System Fully Qualified Domain Name (FQDN)

When you need an FQDN, then you must configure the FQDN. The BMR process does not automatically configure the FQDN.

Maximum character count for FQDN: 63

Follow these steps to configure the FQDN:

1. Edit the */etc/hosts* file and provide the IP Address, the FQDN name, and the server name.

   #vi /etc/hosts

   ip_of_system servername.domainname.com servername

2. Restart the network service.

   #/etc/init.d/network restart

3. Verify the host name and the FQDN name.

   #hostname

   servername

   #hostname -f

   servername.domainname.com

   The FQDN is configured.

Extend the Data Volume after a BMR on Dissimilar Disks

When you perform a BMR to a larger disk than the disk on the original node, some disk space is left unused. The BMR operation does not automatically process the unused disk space. You can format the disk space to a separate partition or resize the existed partition with the unused disk space. The volume that you want to resize must be unused, so you must avoid resizing a system

volume. In this section, we will focus on how to extend a data volume with the unused disk space.

Note: To avoid data loss, resize your volumes immediately after the BMR process. You can also back up the node before starting the volume resizing task.

When the target machine successfully restarts after the BMR, you can extend the data volume.

Raw partition volume

For example, a 2-GB disk in the session is restored to a 16-GB disk named */dev/sdb* with only one partition. The */dev/sdb1* raw partition is directly mounted on the */data* directory.

This example is used to explain the procedure of extending Raw partition volume.

Follow these steps:

1. Check the status of the /dev/sdb1 volume.

   # df –h /dev/sdb1

   /dev/sdb1 2.0G 40M 1.9G 3% /data

2. Umount the /dev/sdb1 volume.

   # umount /data

3. Resize /dev/sdb1 to occupy the entire disk space using the fdisk command.

   To perform this operation, first delete your existing partition and then recreate it with the same start sector number. The same start sector number is responsible for avoiding the data loss.

   # fdisk -u /dev/sdb

   Command (m for help): p

   Disk /dev/sdb: 17.1 GB, 17179869184 bytes

   255 heads, 63 sectors/track, 2088 cylinders, total 33554432 sectors

   *Units = sectors of 1 * 512 = 512 bytes*

   Device Boot Start End Blocks Id System

   /dev/sdb1 63 4192964 2096451 83 Linux

   Command (m for help): d

   Selected partition 1

   Command (m for help): n

Command action

e extended

p primary partition (1-4)

p

Partition number (1-4): 1

First sector (63-33554431, default 63):

Using default value 63

Last sector or +size or +sizeM or +sizeK (63-33554431, default 33554431):

Using default value 33554431

Command (m for help): p

Disk /dev/sdb: 17.1 GB, 17179869184 bytes

255 heads, 63 sectors/track, 2088 cylinders, total 33554432 sectors

*Units = sectors of 1 * 512 = 512 bytes*

Device Boot Start End Blocks Id System

/dev/sdb1 63 33554431 16777184+ 83 Linux

Command (m for help): w

The partition changes to the same start sector number as the original partition and the end sector number is 33554431.

4. Resize the volume using resize2fs command. If necessary, first run the e2fsck command.

   # e2fsck -f /dev/sdb1

   # resize2fs /dev/sdb1

5. Mount the volume to the mount point and check the volume status again.

   # mount /dev/sdb1 /data

   # df –h /dev/sdb1

   /dev/sdb1 16G 43M 16G 1% /data

   The volume is extended to 16 GB and is ready for use.

LVM volume:

For example, an 8-GB disk in the session is restored to a 16-GB disk named */dev/sdc* with only one partition. The */dev/sdc1* raw partition is used as the

only physical volume of the */dev/mapper/VGTest-LVTest* LVM logical volume whose mount point is */lvm*.

This example is used to explain the procedure of extending LVM volume.

Follow these steps:

1. Check the status of the /dev/mapper/VGTest-LVTest volume.

   # lvdisplay –m /dev/mapper/VGTest-LVTest

   --- Logical volume ---

   LV Name /dev/VGTest/LVTest

   VG Name VGTest

   LV UUID udoBIx-XKBS-1Wky-3FVQ-mxMf-FayO-tpfPl8

   LV Write Access read/write

   LV Status available

   # open 1

   LV Size 7.88 GB

   Current LE 2018

   Segments 1

   Allocation inherit

   Read ahead sectors 0

   Block device 253:2

   ---Segments---

   Logical extent 0 to 2017:

   Type linear

   Physical volume /dev/sdc1

   Physical extents 0 to 2017

   The physical volume is */dev/sdc1*, the volume group is *VGTest*, and the logical volume is */dev/VGTest/LVTest or /dev/mapper/VGTest-LVTest*.

2. Umount the /dev/mapper/VGTest-LVTest volume.

   # umount /lvm

3. Disable the volume group in which the /dev/sdc1 physical volume is located.

   # vgchange -a n VGTest

4. Create a partition to occupy the unused disk space using the fdisk command.

   # fdisk -u /dev/sdc

   Command (m for help): p

   Disk /dev/sdc: 17.1 GB, 17179869184 bytes

   255 heads, 63 sectors/track, 2088 cylinders, total 33554432 sectors

   *Units = sectors of 1 * 512 = 512 bytes*

   Device Boot Start End Blocks Id System

   /dev/sdc1 63 16777215 8388576+ 83 Linux

   Command (m for help): n

   Command action

   e extended

   p primary partition (1-4)

   p

   Partition number (1-4): 2

   First sector (16777216-33554431, default 16777216):

   Using default value 16777216

   Last sector or +size or +sizeM or +sizeK (16777216-33554431, default 33554431):

   Using default value 33554431

   Command (m for help): p

   Disk /dev/sdc: 17.1 GB, 17179869184 bytes

   255 heads, 63 sectors/track, 2088 cylinders, total 33554432 sectors

   *Units = sectors of 1 * 512 = 512 bytes*

   Device Boot Start End Blocks Id System

   /dev/sdc1 63 16777215 8388576+ 83 Linux

   /dev/sdc2 16777216 33554431 8388608 83 Linux

   Command (m for help): w

   The /dev/sdc2 partition is created.

5. Create a new physical volume.

   # pvcreate /dev/sdc2

6. Extend the volume group size.

 # vgextend VGTest /dev/sdc2

7. Enable the volume group that you have already disabled.

   # vgchange -a y VGTest

8. Extend the logical volume size using the lvextend command.

   # lvextend -L +8G /dev/VGTest/LVTest

9. Resize the volume using the resize2fs command. If necessary, first run the e2fsck command.

   # e2fsck -f /dev/mapper/VGTest-LVTest

   # resize2fs /dev/mapper/VGTest-LVTest

10. Mount the volume to the mount point and check the volume status again.

    # mount /dev/mapper/VGTest-LVTest /lvm

    # lvdisplay -m /dev/mapper/VGTest-LVTest

    ---Logical volume---

    LV Name /dev/VGTest/LVTest

    VG Name VGTest

    LV UUID GTP0a1-kUL7-WUL8-bpbM-9eTR-SVzl-WgA11h

    LV Write Access read/write

    LV Status available

    # open 0

    LV Size 15.88 GB

    Current LE 4066

    Segments 2

    Allocation inherit

    Read ahead sectors 0

    Block device 253:2

    --- Segments ---

    Logical extent 0 to 2046:

    Type linear

    Physical volume /dev/sdc1

    Physical extents 0 to 2046

    Logical extent 2047 to 4065:

Type linear

Physical volume /dev/sdc2

Physical extents 0 to 2018

The LVM volume extends to 16 GB and is ready for use.

Verify that the Target Node is Restored

After the completion of restore job, verify that the target node is restored with relevant data.

Follow these steps:

1. Navigate to the target machine that you restored.
2. Verify that the target machine has all the information that you backed up.

   The target machine is successfully verified.

   The BMR is successfully performed for Linux Machines.

How to Perform a Migration BMR for Linux Machines

A migration BMR is a two part process where the data is first restored to a temporary machine and then to the actual machine. A BMR with instant BMR option enabled lets you recover data to a temporary machine. You can use the temporary machine until the actual machine is ready. When you have the actual machine, a migration BMR lets you migrate data from the temporary machine to the actual machine. When performing a migration BMR, any data that you create on the temporary machine gets migrated to the actual machine.

Note: You can perform Migration BMR with an agent-based backup only. An agentless backup does not support Migration BMR.

You can perform a BMR using the IP address or the Media Access Control (MAC) address of the target machine. If you boot the target machine using the Arcserve Unified Data Protection Agent for Linux Live CD, you can get the IP address of the target machine.

Note: Machine can boot up. Only one NIC is configured.

Complete the following tasks to perform a Migration BMR:

- Review the Prerequisites for Migration BMR
- Perform a BMR to the Temporary Machine
- Perform a Migration BMR
- Verify that the Target Machine is Restored

Review the Prerequisites for Migration BMR

Consider the following options before performing a migration BMR:

- You have a valid recovery point and the encryption password, if any, for restore.
- You have a valid target machine for BMR.
- You have created the Arcserve Unified Data Protection Agent for Linux Live CD.
- If you want to perform a BMR using the IP address, you must get the IP address of the target machine using the Live CD.
- If you want to perform a PXE-based BMR using the MAC address, you must have the MAC address of the target machine.
- The recovery point must be from the Linux agent-based backup.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Perform a BMR to the Temporary Machine

Before you perform a migration BMR, you have to restore data from the source to a temporary machine. To restore the data temporarily, you can perform a BMR to the temporary machine. After the temporary machine is ready to use, you can continue working on the temporary machine.

When the actual machine is ready, you can perform a migration BMR from the temporary machine to the actual machine.

Note: For more information on performing a BMR, see How to Perform a Bare Metal Recovery (BMR) for Linux Machines.

Follow these steps:

1. Access the Restore Wizard in one of the following ways:

   From Arcserve UDP:

   a. Log into Arcserve UDP.

   b. Click the **resources** tab.

   c. Select **All Nodes** in the left pane.

      All the added nodes are displayed in the center pane.

   d. In the center pane, select the node and click **Actions**.

   e. Click **Restore** from the **Actions** drop-down menu.

      The Arcserve Unified Data Protection Agent for Linux web interface opens. The restore type selection dialog is displayed in the Agent UI.

   f. Select the restore type and click **OK**.

      Note: You are automatically logged in to the agent node and the **Restore Wizard** opens from the agent node.

   From Arcserve Unified Data Protection Agent for Linux:

   a. Open the Arcserve Unified Data Protection Agent for Linux web interface.

      Note: During the installation of Arcserve Unified Data Protection Agent for Linux, you received the URL to access and manage the server.

   b. Log in to Arcserve Unified Data Protection Agent for Linux.

2. Click **Restore** from the **Wizard** menu and select **Bare Metal Recovery (BMR)**.

   The **Backup Server** page of the **Restore Wizard - BMR** opens.

3. Provide all the details in the **Restore Wizard - BMR** and save the wizard.

4. Ensure that you select the **Enable instant BMR** check box on the **Target Machine** page of the wizard.
5. Ensure that you select the **Do not recover data automatically after machine is started** check box on the **Target Machine** page of the wizard.
6. Run the BMR job.

   The temporary machine is recovered using the BMR, with the instant BMR option enabled. You can use the temporary machine until the actual machine is ready.

Perform a Migration BMR

When the actual machine is ready, perform a migration BMR. Migration BMR restores the original data from the backup session and the new data from the temporary machine to the actual machine.

Follow these steps:

1. Click **Restore** from the **Wizard** menu and select **Migration BMR**.

   The **Backup Server** page of the **Restore Wizard - Migration BMR** opens.

2. Provide all the details in the **Restore Wizard - Migration BMR**.

   Note: For more information on performing a BMR, see How to Perform a Bare Metal Recovery (BMR) for Linux Machines.

3. Ensure that the following information is provided on the **Backup Server** page of the wizard.

   a. Select the instant VM recovery job or the Instant BMR job.

      Local Server

      Specifies that the Backup Server is locally managed. The BMR job for the temporary machine is run on the local server.

      Remote Server

      Specifies that the Backup Server is remotely managed. The BMR job for the temporary machine is run on the remote server. You have to provide the remote server details to connect to the remote server.

   b. Select the restore job from the Job Name drop-down list.

      The list displays the Instant VM recovery job or Instant BMR job, which is in the Ready to use job phase or Power off job phase, once it is ready to use.

4. Save the BMR job.

   In the home page, the **Job Phase** on the **Job Status** tab changes to **Click here to migrate data**.

5. (Optional) Boot the temporary machine using a Live CD when the selected job type is Instant BMR.

6. From the **Job Status** tab, use **Click here to migrate data**.

   The data migration begins.

   You have successfully performed a migration BMR.

Verify that the Target Node is Restored

After the completion of restore job, verify that the target node is restored with relevant data.

Follow these steps:

1. Navigate to the target machine that you restored.
2. Verify that the target machine has all the information from the temporary machine, including any new data that you created on the temporary machine.

   The target machine is successfully verified.

   The migration BMR is successfully performed for agent-based Linux machines.

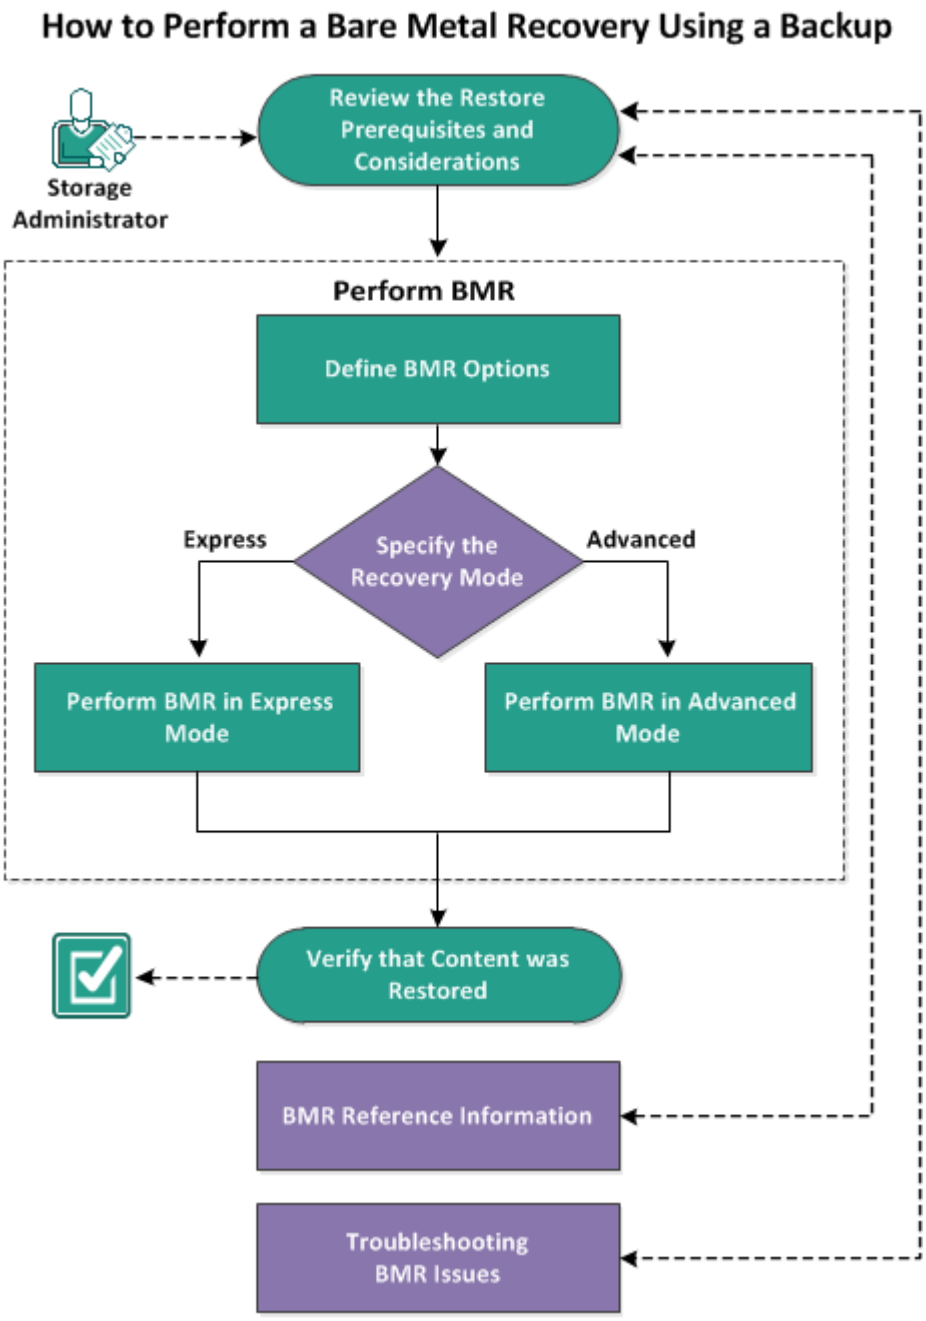
How to Perform a BMR Using a Backup

Bare Metal Recovery (BMR) is the process of restoring a computer system from "bare metal" including reinstalling the operating system and software applications, and then restoring the data and settings. The BMR process lets you restore a full computer with minimal effort, even to different hardware. BMR is possible because during the block-level backup process, Arcserve UDP Agent (Windows) not only captures the data, but also all information that is related to the following applications:

- Operating system
- Installed applications
- Configuration settings
- Necessary drivers

All relevant information that is necessary to perform a complete rebuild of the computer system from "bare metal" is backed up into a series of blocks and stored on the backup location.

The following diagram illustrates the process for how to perform a BMR using a backup:



Complete the following tasks to perform a BMR using a backup:

1. Review the BMR Prerequisites and Considerations
2. Define BMR Options
   - Perform BMR in Express Mode
   - Perform BMR in Advanced Mode
3. Verify that the BMR was Successful
4. BMR Reference Information
5. Troubleshooting BMR Issues

Review the BMR Prerequisites and Considerations

Verify that the following prerequisites exist before performing a BMR:

- You must have one of the following images:
  - A created BMR ISO image burned onto a CD/DVD
  - A created BMR ISO image burned onto a portable USB stick

  Note: Using Arcserve UDP Agent (Windows), you can utilize a Boot Kit Utility to combine a WinPE image and Arcserve UDP Agent (Windows) image to create a BMR ISO image. This ISO image is then burned onto a bootable media. You can then use either of these bootable media (CD/DVD or USB stick) to initialize the new computer system and allow the bare metal recovery process to begin. To ensure your saved image is always the most up-to-date version, create a new ISO image every time you update Arcserve UDP Agent (Windows).
- At least one full backup available.
- At least 2-GB RAM installed on the virtual machine and the source server that you are recovering.
- To recover VMware virtual machines to VMware virtual machines that are configured to behave as physical servers, verify the VMware Tools application is installed on the destination virtual machine.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Review the following restore considerations:

- Regardless of which method you used to create the Boot Kit image, the BMR process is basically the same.

  Note: The BMR process cannot create storage spaces. If the source machine had storage spaces, during BMR you cannot create storage spaces at the destination machine. You can either restore those volumes to regular disks / volumes or manually create storage spaces before performing the BMR, and then restore the data into those created storage spaces.
- Dynamic disks are restored at the disk level only. If your data is backed up to a local volume on a dynamic disk, you cannot to restore this dynamic disk during BMR. In this scenario, to restore during BMR you must perform one of the following tasks and then perform BMR from the copied Recovery Point:
  - Back up to a volume on another drive.
  - Back up to a remote share.

- Copy a recovery point to another location.

Note: If you perform BMR with multiple dynamic disks, the BMR may fail because of some unexpected errors (such as fail to boot, unrecognized dynamic volumes, and so on). If this occurs, you should restore only the system disk using BMR, and then after the machine reboot you can restore the other dynamic volumes on a normal environment.

- (Optional) Review the BMR Reference Information. For more information, see the following topics:
  - How Bare Metal Recovery Works
  - Operating Systems that Support UEFI/BIOS Conversion
  - Managing the BMR Operations Menu

Review the following considerations:

- If you upgrade to a newer version or update of Arcserve UDP, you must re-create the BMR ISO using the proper Windows AIK or ADK level to include support for latest features and bug fixes. However, once a BMR ISO is created, the ISO file can be used for the same OS level. The following OS levels can use the same ISO:
  - ISO created using Windows 7 WAIK – works for Windows 2003, Vista, 2008, 2008 R2
  - ISO create using Windows 8/8.1 ADK – works for Windows 8, 8.1, Server 2012, Server 2012 R2
  - ISO created using Windows 10 ADK – works for Windows 10, Windows Server 2016, and Windows Server 2019

Define BMR Options

Prior to initiating the BMR process, you must specify some preliminary BMR options.

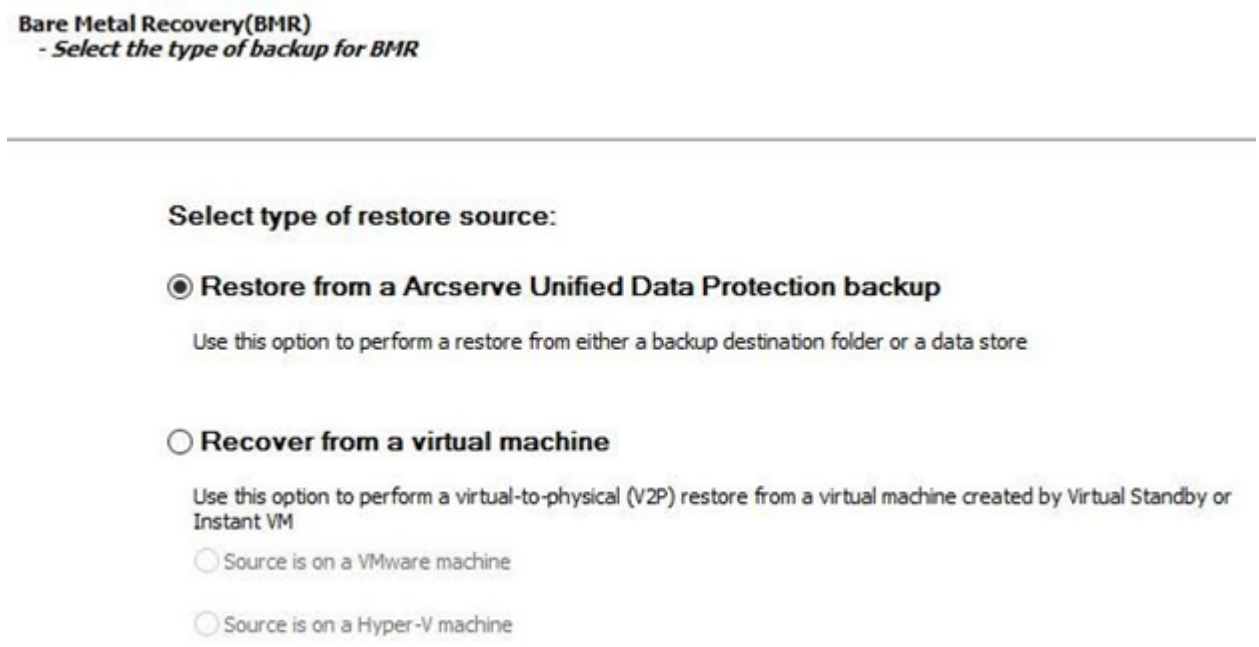
Follow these steps:

1. Insert the saved Boot Kit image media and boot the computer.
   - If you are using a BMR ISO image burned onto a CD/DVD, insert the saved CD/DVD.
   - If you are using a BMR ISO image burned onto a USB stick, insert the saved USB stick.

   The **BIOS Setup Utility** screen is displayed.
2. From the **BIOS Setup Utility** screen, select the CD-ROM Drive option or the USB option to launch the boot process. Select an architecture (x86/x64) and press **Enter** to continue.
3. The Arcserve UDP Agent (Windows) language select screen is displayed. Select a language and click **Next** to continue.



The Bare Metal Recovery process is initiated and the initial BMR wizard screen is displayed.



The BMR wizard screen allows you to select the type of BMR you want to perform:

- **Restore from an Arcserve Unified Data Protection backup**

  Use this option to perform a restore from either a backup destination folder or a data store.

  This option lets you recover data that was backed up using Arcserve UDP Agent (Windows). This option is used in connection with backup sessions performed with Arcserve UDP Agent (Windows) or with the Arcserve UDP host-based VM backup application.

  If you select this option, continue this procedure from here.

- **Recover from a virtual machine**

  Use this option to perform a virtual-to-physical (V2P) restore from a virtual standby VM. Virtual-to-physical (V2P) is a term that refers to the migration of an operating system (OS), application programs and data from a virtual machine or disk partition to a computer's main hard disk. The target can be a single computer or multiple computers.

  - **Source is on a VMware machine**

    Lets you recover data for a machine for which virtual conversion is done to a VMware virtual machine. This option is used in connection with the Arcserve Central Virtual Standby application.

Note: For this option, you can only recover data if the virtual conversion to a VMDK file (for VMware) was performed using Arcserve Central Virtual Standby.

If you select this option, see Recover using a VMware Virtual Standby VM to continue this procedure.

For more information, see Recover using a VMware Virtual Standby VM in the Agent for Windows online help.

- **Source is on a Hyper-V machine**

  Lets you recover data for a machine for which virtual conversion is performed to a Hyper-V virtual machine. This option is used in connection with the Arcserve Central Virtual Standby application.

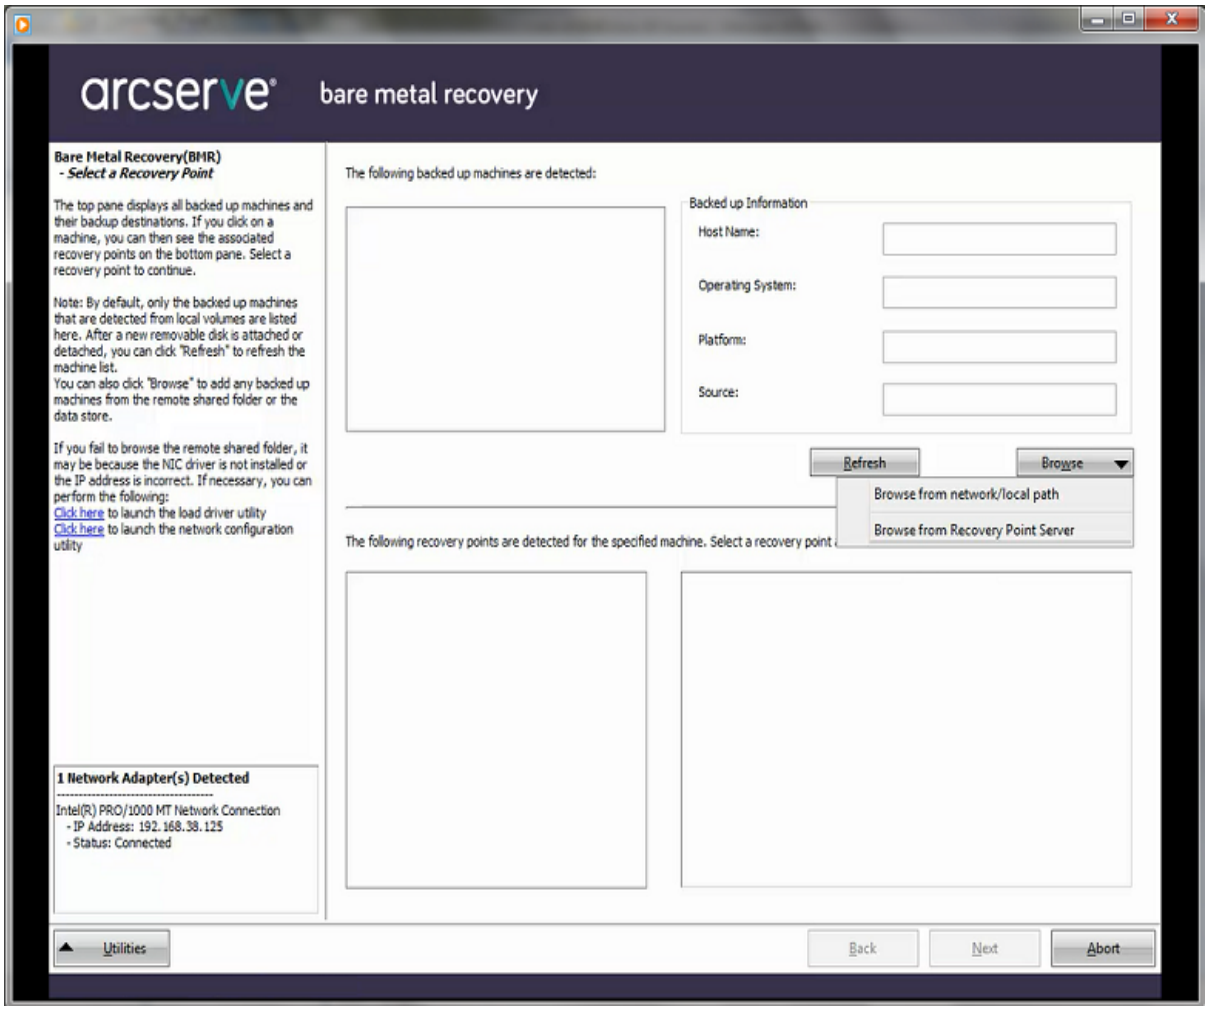
  Note: For this option, you can only recover data if the virtual conversion to a VHD file (for Hyper-V) was performed using Arcserve Central Virtual Standby.

  If you select this option, see Recover using a Hyper-V Virtual Standby VM to continue this procedure.

  For more information, see Recover using a Hyper-V Virtual Standby VM in the Agent for Windows online help.

4. Select **Restore from an Arcserve Unified Data Protection backup** and click **Next**.

   The **Select a Recovery Point** wizard screen is displayed.



5. From the **Select a Recovery Point** wizard screen, click **Browse** and select either **Browse from network/local path** or select **Browse from Recovery Point Server**.

    a. If you select **Browse** from network/local path, select the machine (or volume) which contains recovery points for your backup image.

    Arcserve UDP Agent (Windows) lets you recover from any local drive or from a network share.

    - If you recover from a local backup, the BMR wizard automatically detects and displays all volumes containing recovery points.
    - If you recover from a remote share, browse to the remote location where the recovery points are stored. If there are multiple machines containing recovery points, all machines are displayed.

    You may also need access information (User Name and Password) for the remote machine.

Note: The network must be up and running to browse to remote recovery points. If necessary, you can check/refresh your network configuration information or you can load any missing drivers from the Utilities menu.

- If the BMR module cannot detect any local destination volume, the **Select a Folder** dialog automatically displays. Provide the remote share where the backups are residing.
- If you are restoring from an iSCSI destination, the BMR module may not detect this destination and you need to perform the following:

  1. Click **Utilities**, select **Run** from the pop-up menu, type **cmd**, and then click **OK**.

  2. In the command prompt window, use the following Windows iSCSI commands to set up iSCSI connections:

  > net start msiscsi

  > iSCSICLI QAddTargetPortal <TargetPortalAddress>

  > iSCSICLI QLoginTarget <TargetName > [CHAP username] [CHAP password]

  Note: CHAP = Challenge-Handshake Authentication Protocol

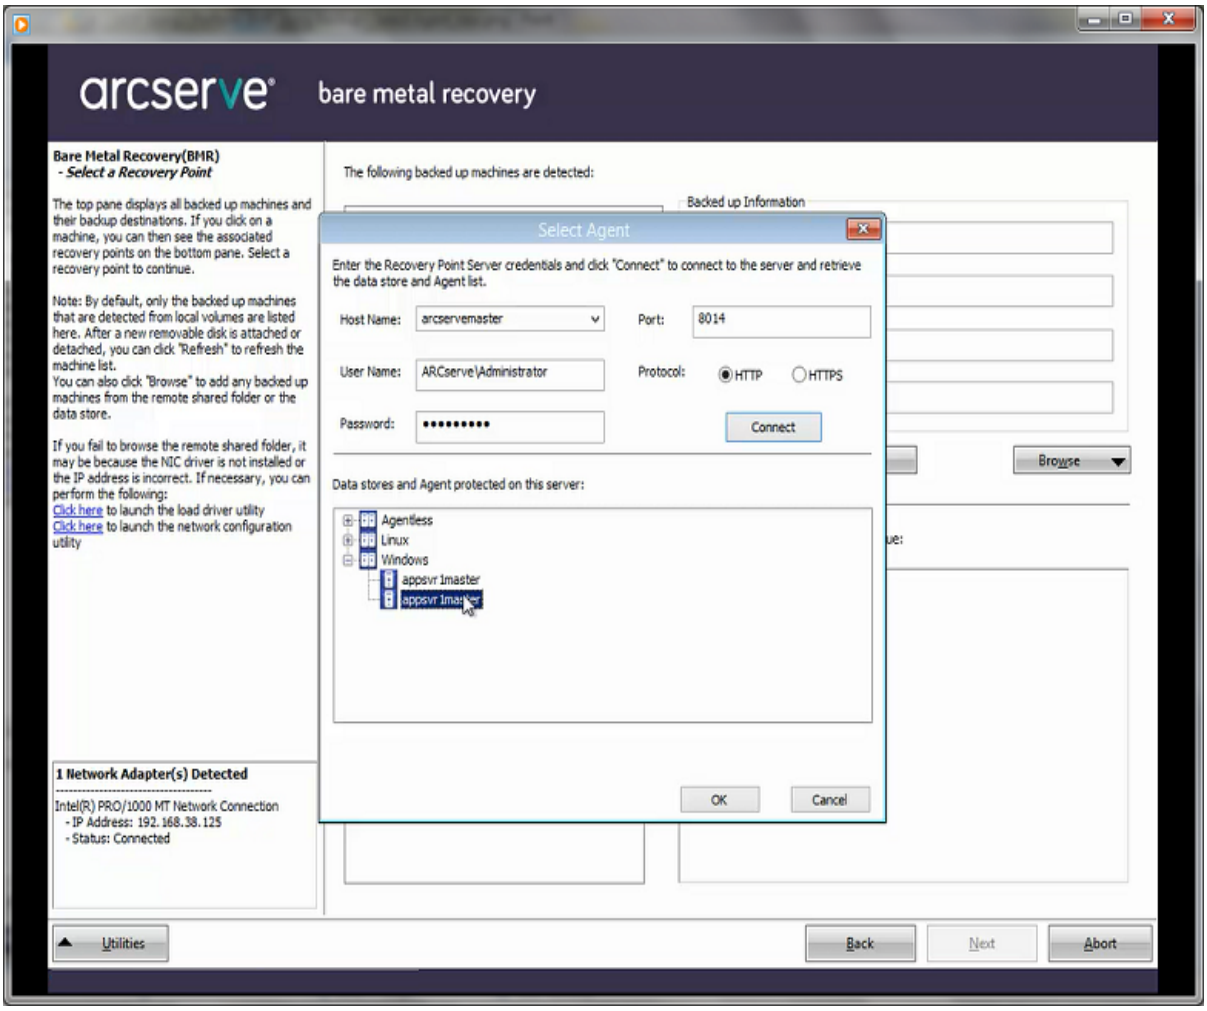
  For more information about Windows iSCSI command line options, see link.

  Note: Extra steps may be needed depending on the iSCSI target software being used. For more information, see the manual of the iSCSI target software.

  3. From the BMR screen the disks/volumes connected through the iSCSI disk should be displayed. The iSCSI disk can now be used as the source volume or the backup destination volume.

  Note: BMR does not support the case where the OS is installed on an iSCSI disk. Only data disks are supported.

b. If you select **Browse the Recovery Point Server**, the **Select Agent** dialog displays. Provide the **Recovery Point Server Host Name**, **User Name**, **Password**, **Port**, and **Protocol**. Click **Connect**.



6. Select the folder or Agent Name under Data Store where the recovery points for your backup are stored and click **OK**.

   The BMR wizard screen now displays the following information:

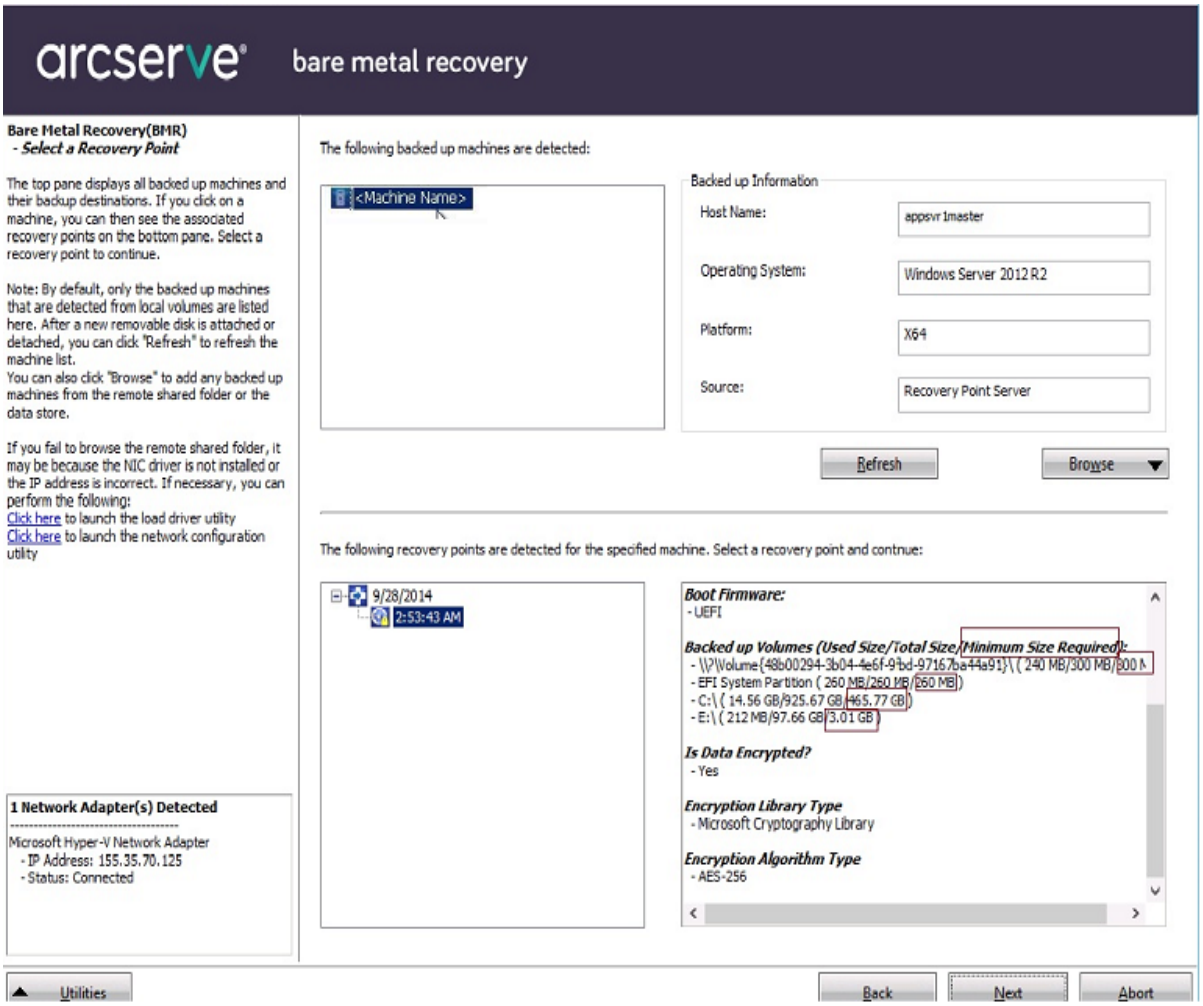
   - Machine name (in the upper left pane).
   - Related backup information (in the upper right pane).
   - All the corresponding recovery points (in the lower left pane).

     Note: For supported operating systems, you can perform a BMR from a backup performed on a UEFI machine to a BIOS-compatible machine and from a BIOS machine to a UEFI-compatible machine. See Operating Systems that Support UEFI/BIOS Conversion for a complete listing of firmware conversion supported systems.

   - For operating systems that do not support firmware conversion, to perform BMR for a UEFI system, you must boot the computer in UEFI mode. BMR

does not support restoring a computer with different firmware. To verify that the boot firmware is UEFI and not BIOS, click **Utilities**, **About**.

- For operating systems that do support firmware conversion, after you select a recovery point, if it is detected that the source machine is not the same firmware as your system, you will be asked if you want to convert UEFI to a BIOS-compatible system or BIOS to UEFI-compatible system.

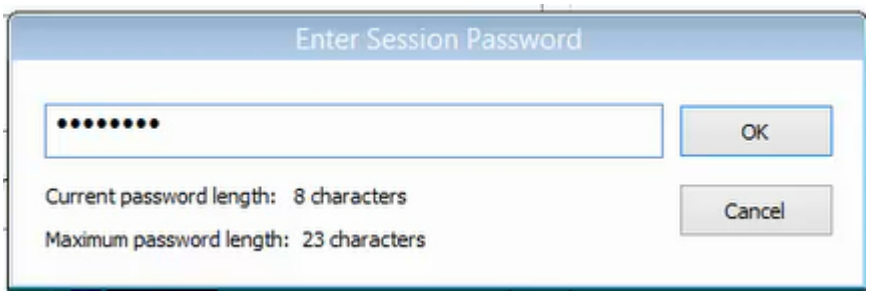


Note: The Arcserve UDP Version 5.0 Update 2 only supports BMR to a smaller disk when the sessions are backed up from Arcserve UDP Version 5.0 Update 2. See the field **Minimum Size Required** for the destination disk size. BMR to a smaller disk is only supported in **Advanced Mode**.

7. Select which recovery point to restore.

   The related information for the selected recovery point is displayed (in the lower right pane). This display includes such information as the type of backup that was performed (and saved), the backup destination, and the volumes that were backed up.

If the recovery point contains encrypted sessions (the recovery point clock icon includes a lock), a password required screen appears. Enter the session password and click **OK**.



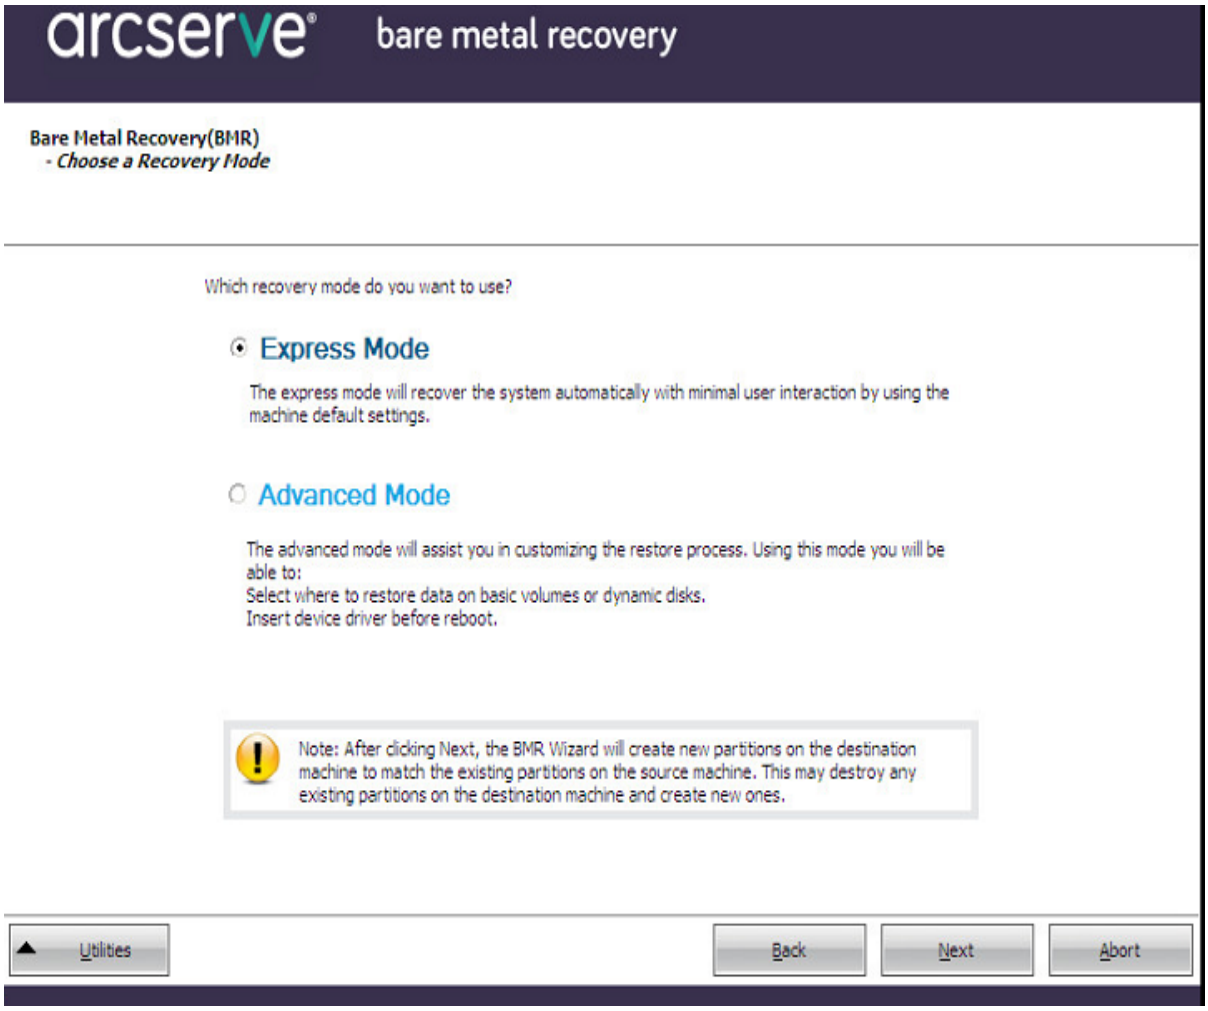
Notes:

If you are restoring from a Arcserve UDP Recovery Point Server, you are asked to provide a session password.

If your machine is a Domain Controller, Arcserve UDP Agent (Windows) supports a non-authoritative restore of the active directory (AD) database file during BMR. (It does not support restoring MSCS clusters).

8. Verify the recovery point that you want to restore and click **Next**.

   A BMR wizard screen is displayed with the available recovery mode options.



The available options are **Advanced Mode** and **Express Mode**.

- Select Express Mode if you want minimal interaction during the recovery process.
- Select Advanced Mode if you want to customize the recovery process.

  Default: Express Mode.

Perform BMR in Express Mode

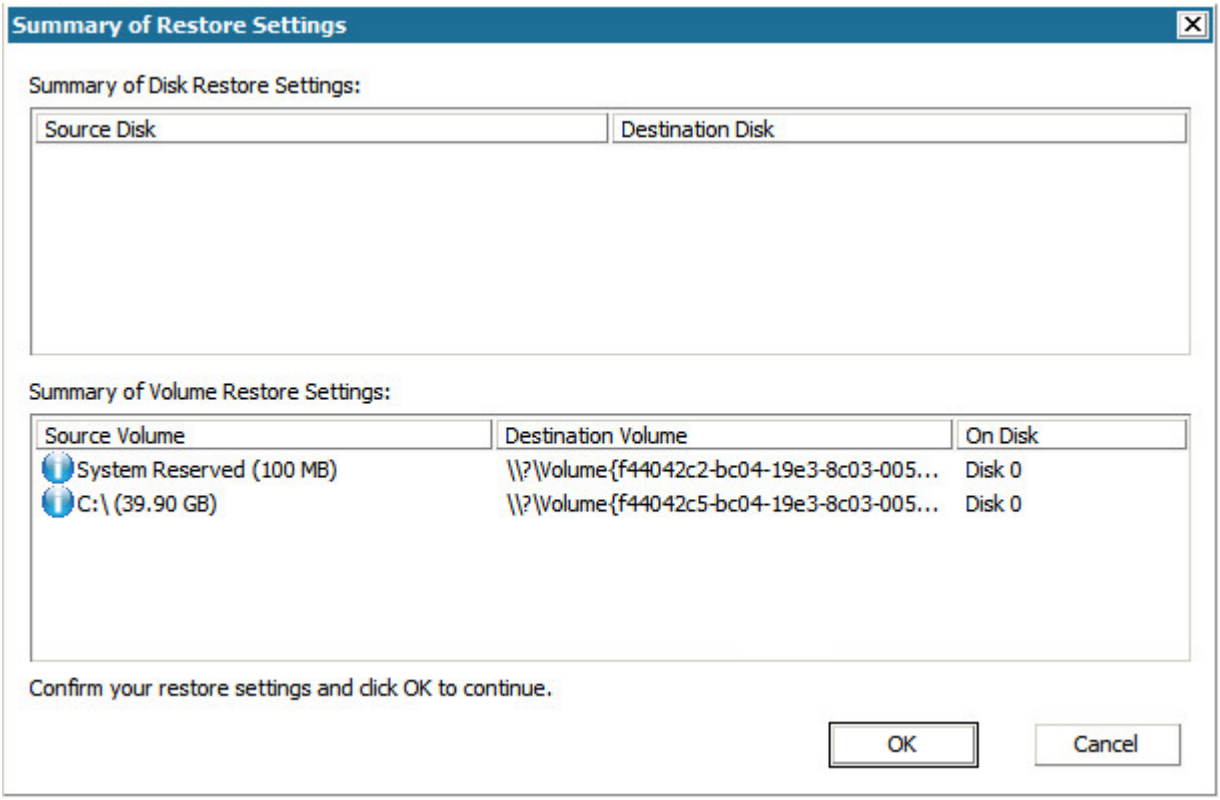
The Express Mode requires minimal interaction during the recovery process.

Follow these steps:

1. From the **Choose a Recovery Mode** dialog, select **Express Mode** and click **Next**.

   The **Summary of Disk Restore Settings** screen opens, displaying a summary of the volumes that are going to be restored.

   Note: On the bottom of restore summary window, the drive letters listed in **Destination Volume** column are automatically generated from the Windows Pre-installation Environment (WinPE). They can be different from the drive letters listed in **Source Volume** column. However, the data is still restored to proper volume even if drive letters are different.



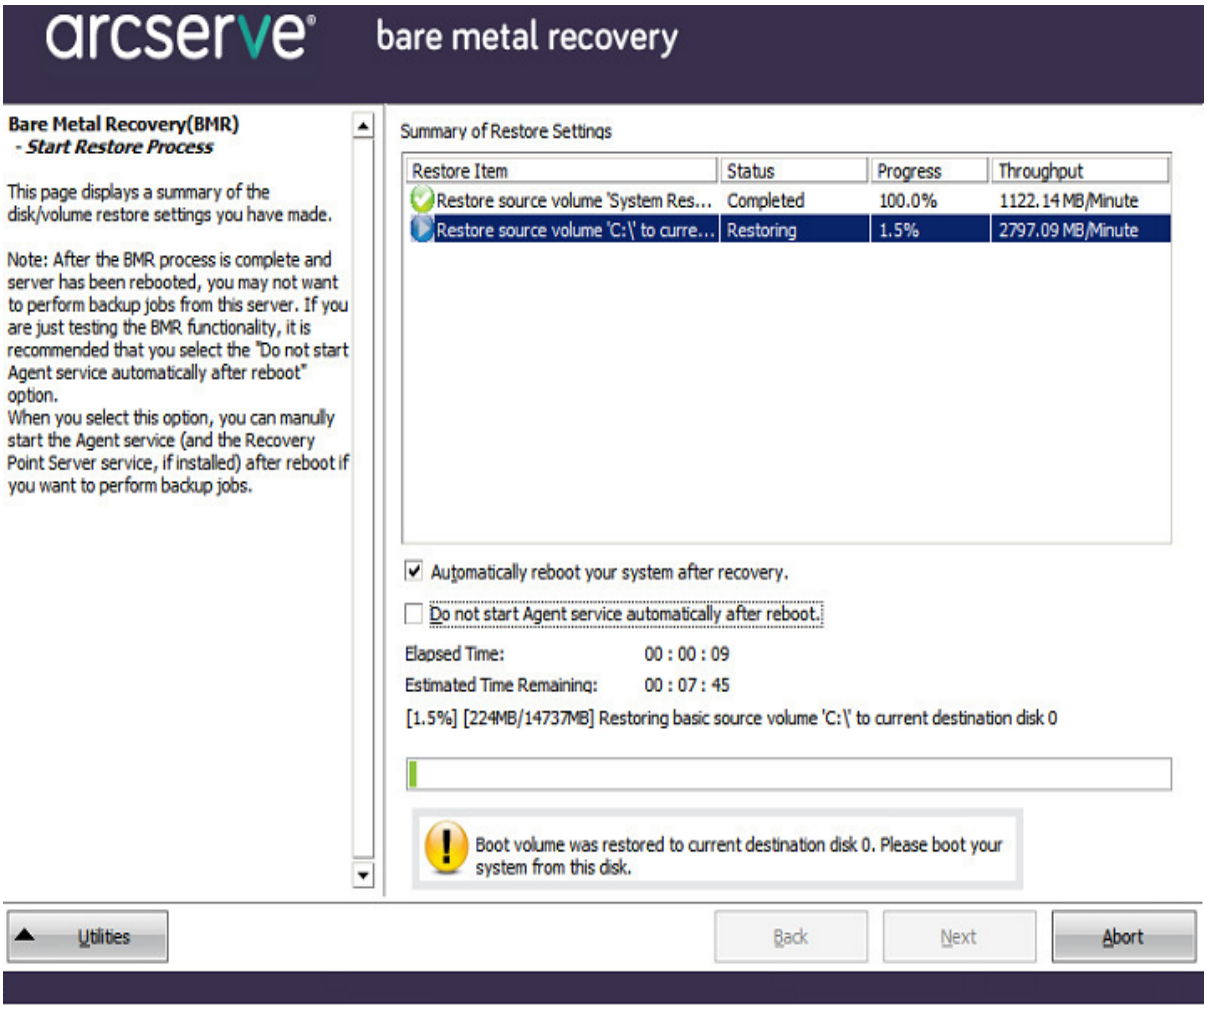
2. After you have verified that the summary information is correct, click **OK**.

   The restore process starts. The BMR wizard screen displays the restore status for each volume.

   - Depending upon the size of the volume being restored, this operation can take some time.

- During this process you are restoring, block-by-block whatever you had backed up for that recovery point and creating a replica of the source machine on the target machine.
- By default, the option to reboot your system automatically after recovery is selected. If necessary, you can clear this option and you can reboot manually at a later time.

  Important: If you are performing an authoritative restore of an active directory after a BMR, you must uncheck the option **Automatically reboot your system after recovery** and for more information, see How to Perform an Authoritative Restore of an Active Directory after a BMR.
- If necessary, you can select Do not start Agent service automatically after reboot.
- If necessary, you can cancel or abort the operation at any time.



3. From the **Utilities** menu, you can access the BMR **Activity Log** and you can use the **Save** option to save the Activity Log.

By default, the Activity Log is saved to the following location:

X:\windows\system32\dr\log

Note: To avoid getting a Windows-generated error, do not save the Activity Log on your desktop or create a folder on your desktop using the **Save As** option from the BMR Activity Log window.

4. If you are restoring to dissimilar hardware (the SCSI/FC adapter which used to connect hard drives could have been changed) and no compatible driver is detected in your original system, a "driver injection" page is displayed to allow you to provide drivers for these devices.

   You can browse and select drivers to inject to the recovered system so that even if you are recovering to a machine with dissimilar hardware, you can still bring back the machine after BMR.

   When the BMR process is completed, a confirmation notification is displayed.

Perform BMR in Advanced Mode

The **Advanced Mode** option lets you customize the recovery process.

Follow these steps:

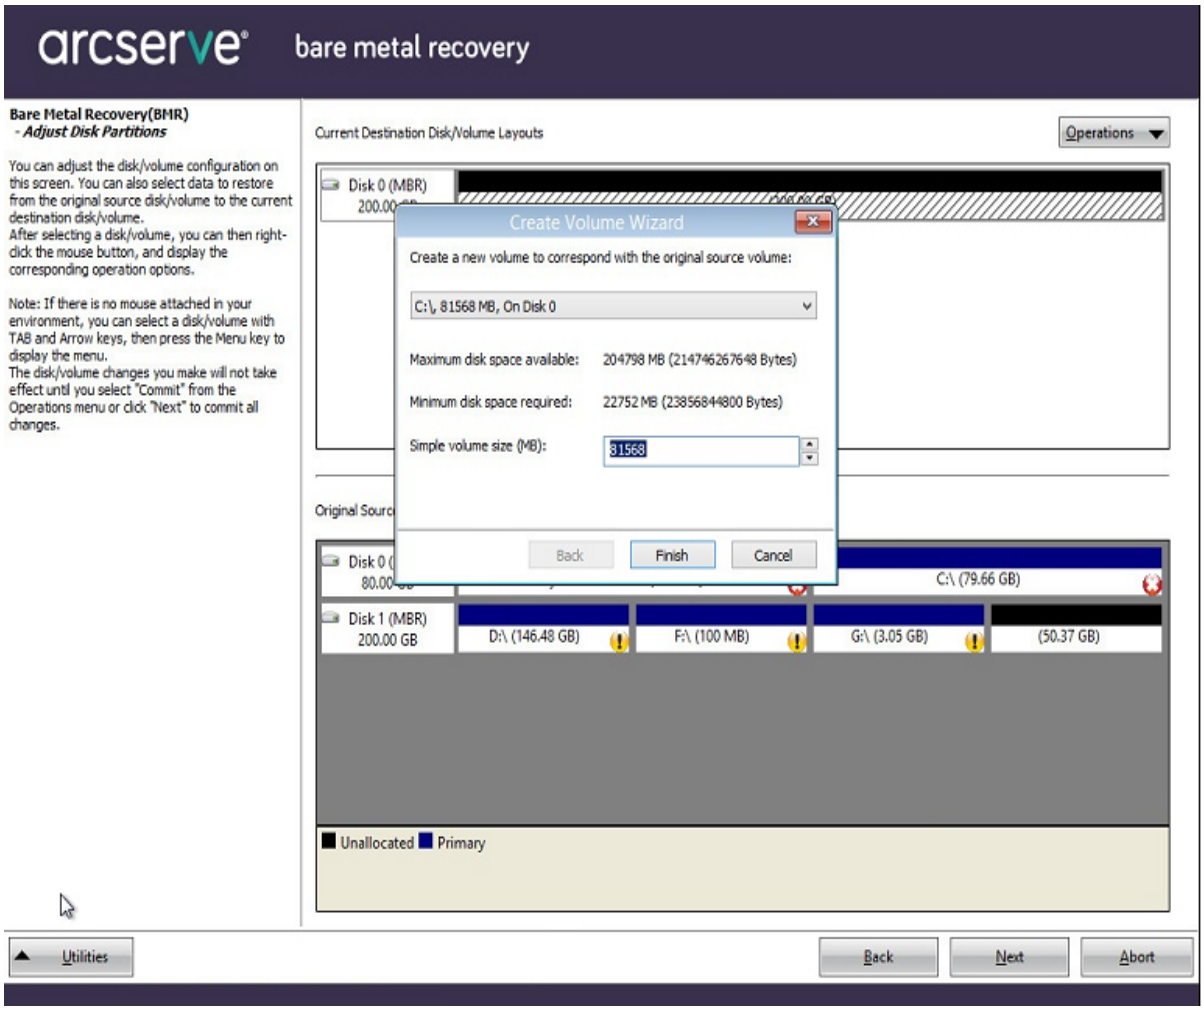
1. From the **Choose a Recovery Mode** dialog, select **Advanced Mode** and click **Next**.

   The BMR utility starts locating the machine that is going to be recovered and displays the corresponding disk partition information.

   The upper pane shows the disk configuration that you have on the current (target) machine and the lower pane shows the disk partition information that you had on the original (source) machine.

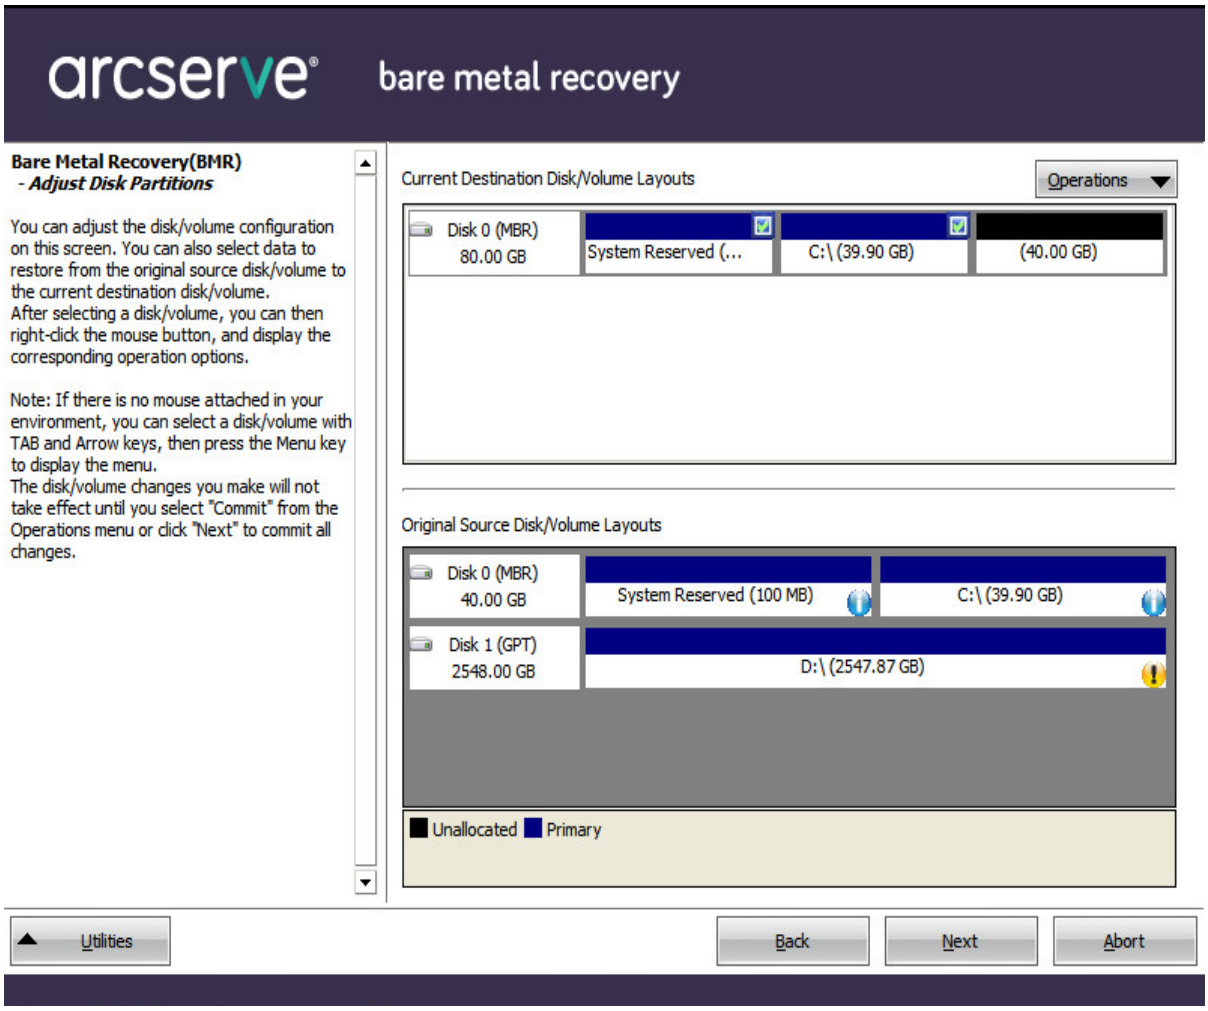
   Important! A red X icon displaying for a source volume in the lower pane indicates that this volume contains system information and has not been assigned (mapped) to the target volume. This system information volume from the source disk must be assigned to the target disk and restored during BMR or the reboot fails.

   You can create volumes to a smaller disk based on the suggested **Minimum disk space required**. In the example, the original size of the volume is 81568 MB. When you create the volume on the target disk, the suggested minimum size is 22752 MB. In this case, you can create the original volume with a size of 22752 MB.



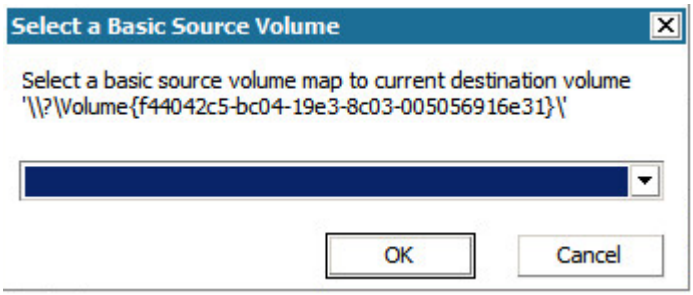
Note: If you perform BMR and you restore the system volume to a disk which is not configured as the boot disk, it will fail to boot the machine after BMR is completed. Ensure that you are restoring the system volume to a properly configured boot disk.

Note: When restoring to another disk/volume, the capacity of new disk/volume can be the same size, larger than original disk/volume, or smaller than the original disk/-volume. In addition, volume resizing is not for dynamic disks.



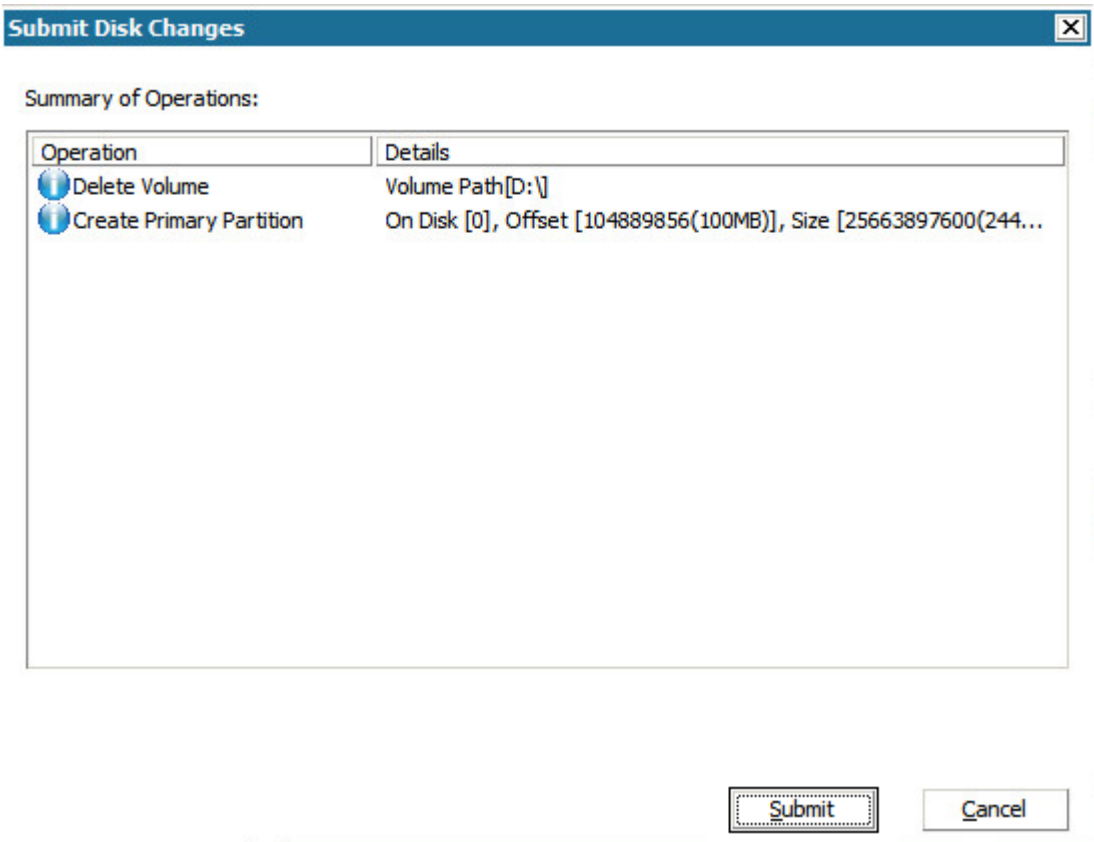
2. If the current disk information you are seeing does not appear correct, you can access the **Utilities** menu and check for missing drivers.
3. If necessary, on the target disk/volume pane you can click the **Operations** drop-down menu to display the available options. For more information about these options, see Managing the BMR Operations Menu.
4. Click on each target volume and from the pop-up menu, select the **Map Volume From** option to assign a source volume to this target volume.

   The **Select a Basic Source Volume** dialog opens.



5. From **Select a Basic Source Volume** dialog, click the drop-down menu and select the available source volume to assign to the selected target volume. Click **OK**.
   - On the target volume, a checkmark icon is displayed, indicating that this target volume has been mapped to.
   - On the source volume, the red X icon changes to a green icon, indicating that this source volume has been assigned to a target volume.
6. When you are sure all volumes that you want to restore and all volumes containing system information are assigned to a target volume, click **Next**.

   The **Submit Disk Changes** screen opens, displaying a summary of the selected operations. For each new volume being created, the corresponding information is displayed.



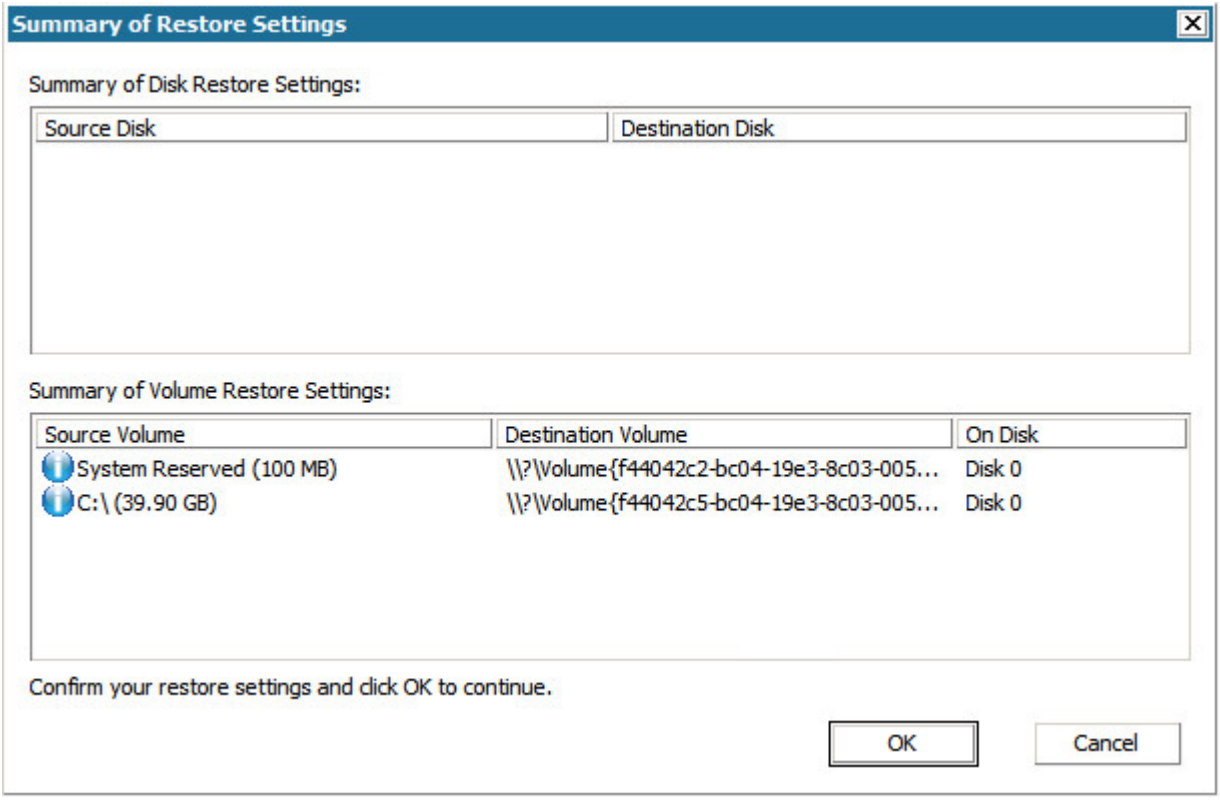
7. When you have verified the summary information is correct, click **Submit**. (If the information is not correct, click **Cancel**).

   Note: All operations to the hard drive do not take effect until you submit it.

   On the target machine, the new volumes are created and mapped to the corresponding source machine.
8. When the changes are completed, click **OK**.

The Summary of Disk Restore Settings screen opens, displaying a summary of the volumes that are going to be restored.

Note: On the bottom of restore summary window, the drive letters listed in "Destination Volume" column are automatically generated from the Windows Pre-installation Environment (WinPE). They can be different from the drive letters listed in "Source Volume" column. However, the data is still restored to proper volume even if drive letters are different.



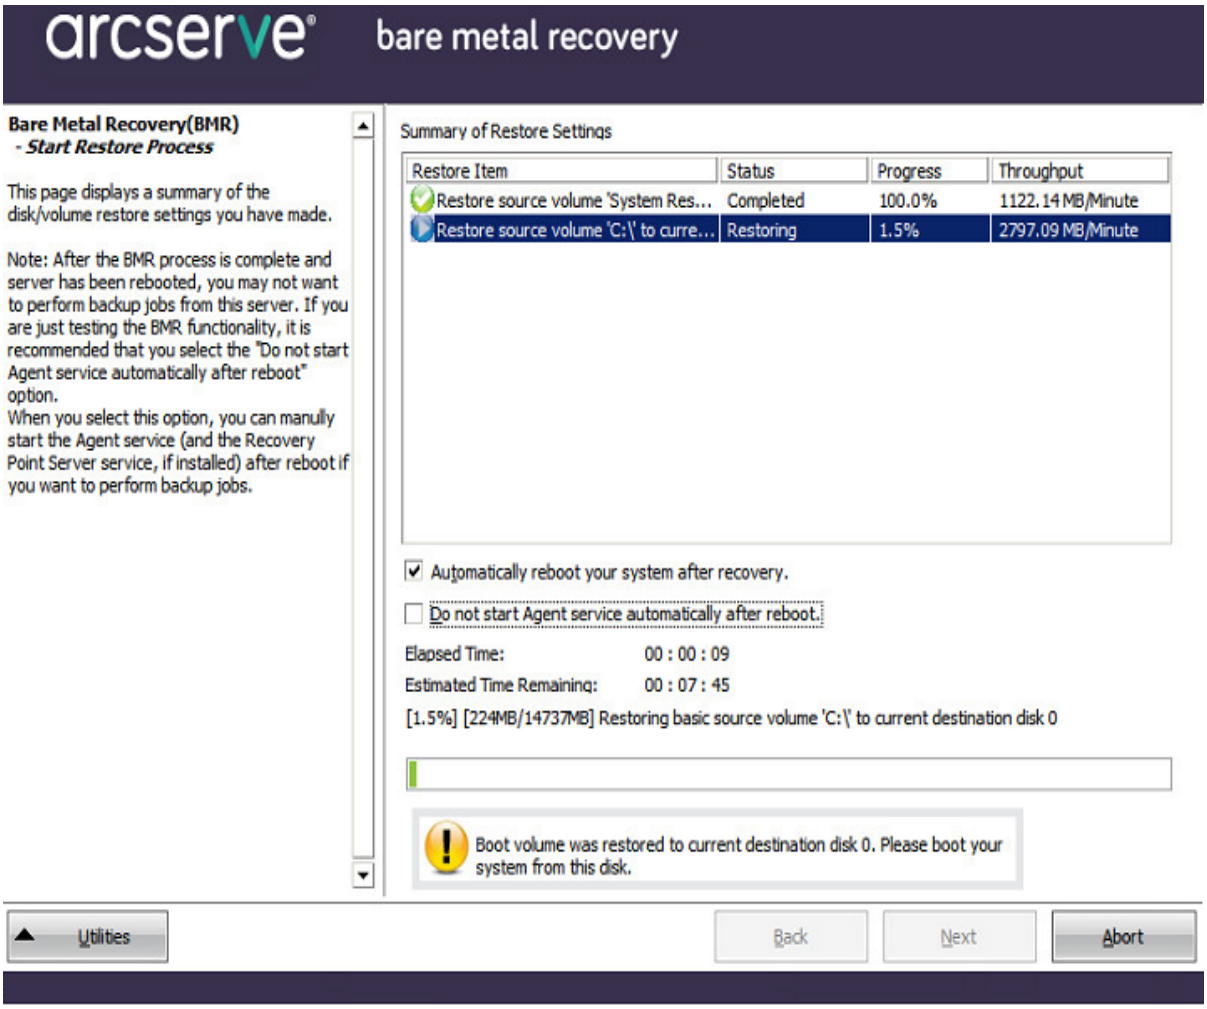
9. After you have verified that the summary information is correct, click **OK**.

   The restore process starts. The BMR wizard screen displays the restore status for each volume.

   - Depending upon the size of the volume being restored, this operation can take some time.
   - During this process you are restoring, block-by-block whatever you had backed up for that recovery point and creating a replica of the source machine on the target machine.
   - By default, the option to reboot your system automatically after recovery is selected. If necessary, you can clear this option and you can reboot manually at a later time.

Important: If you are performing an authoritative restore of an active directory after a BMR, you must uncheck the option **Automatically reboot your system after recovery** and for more information, see How to Perform an Authoritative Restore of an Active Directory after a BMR.

- If necessary, you can select Do not start Agent service automatically after reboot.
- If necessary, you can cancel or abort the operation at any time.



10. From the **Utilities** menu, you can access the BMR **Activity Log** and you can use the **Save** option to save the Activity Log.

    By default, the Activity Log is saved to the following location:

    X:\windows\system32\dr\log.

    Note: To avoid getting a Windows-generated error, do not save the Activity Log on your desktop or create a folder on your desktop using the **Save As** option from the BMR **Activity Log** window.

11. If you are restoring to dissimilar hardware (the SCSI/FC adapter which used to connect hard drives could have been changed) and no compatible driver is detected in your original system, a "driver injection" page is displayed to allow you to provide drivers for these devices.

    You can browse and select drivers to inject to the recovered system so that even if you are recovering to a machine with dissimilar hardware, you can still bring back the machine after BMR.

12. When the BMR process is completed, a confirmation notification is displayed.

Verify that the BMR was Successful

To verify that the BMR was successful, perform the following tasks:

- Reboot the operating system.
- Verify all systems and applications function correctly.
- Verify all network settings are properly configured.
- Verify the BIOS is configured to boot from the disk on which the boot volume was restored to.
- When the BMR is completed, be aware of the following conditions:
  - The first backup that is performed after the BMR is a Verify Backup.
  - When the machine has been rebooted, you may need to configure the network adapters manually if you restored to dissimilar hardware.

    Note: When the machine is rebooting, a Windows Error Recovery screen may be displayed indicating that Windows did not shut down successfully. If this occurs, you can safely ignore this warning and continue to start Windows normally.
  - For dynamic disks, if the status of the disk is offline, you can manually change it to online from the disk management UI (accessed by running the Diskmgmt.msc control utility).
  - For dynamic disks, if the dynamic volumes are in a failed redundancy status, you can manually resynchronize the volumes from the disk management UI (accessed by running the Diskmgmt.msc control utility).

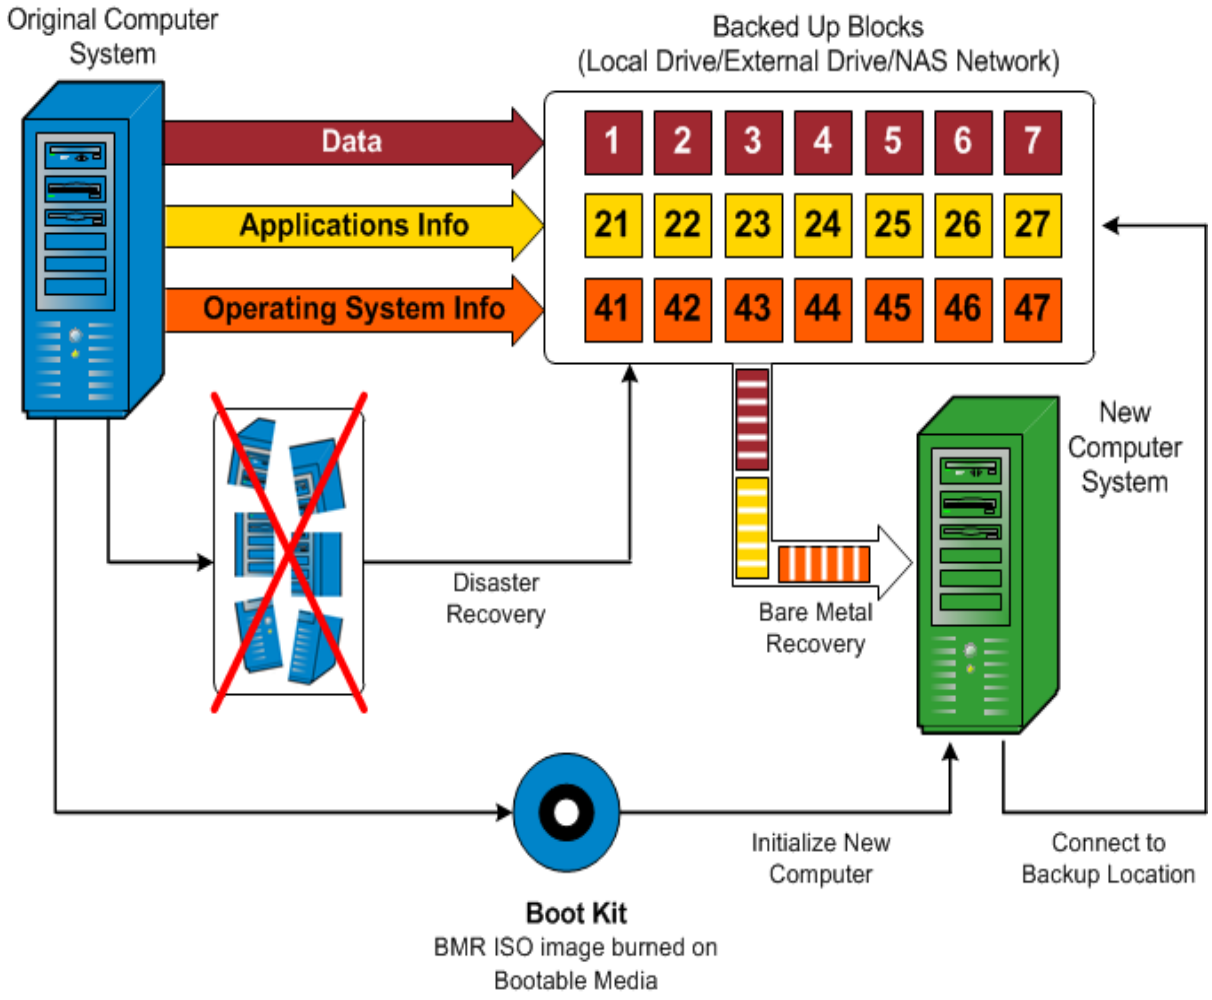
BMR Reference Information

- How Bare Metal Recovery Works
- Operating Systems that Support UEFI/BIOS Conversion
- Managing the BMR Operations Menu

How Bare Metal Recovery Works

Bare Metal Recovery is the process of restoring a computer system from "bare metal" by reinstalling the operating system and software applications, and then restoring the data and settings. The most common reasons for performing a bare metal recovery are because your hard drive either fails or becomes full and you want to upgrade (migrate) to a larger drive or migrate to newer hardware. Bare metal recovery is possible because during the block-level backup process, Arcserve UDP Agent (Windows) captures not only the data, but also all information related to the operating system, installed applications, configuration settings, necessary drivers, and so on. All relevant information that is necessary to perform a complete rebuild of the computer system from "bare metal" is backed up into a series of blocks and stored on the backup location.

Note: Dynamic disks are restored at disk level only. If your data is backed up to a volume on a dynamic disk, you will not be able to restore this dynamic disk (including all its volumes) during BMR.



When you perform a bare metal recovery, the Arcserve UDP Agent (Windows) boot disk is used to initialize the new computer system and allow the bare metal recovery process to begin. When the bare metal recovery is started, Arcserve UDP Agent (Windows) will prompt you to select or provide a valid location to retrieve these backed up blocks from, as well as the recovery point to be restored. You may also be prompted to provide valid drivers for the new computer system if needed. When this connection and configuration information is provided, Arcserve UDP Agent (Windows) begins to pull the specified backup image from the backup location and restore all backed up blocks to the new computer system (empty blocks will not be restored). After the bare metal recovery image is fully restored to the new computer system, the machine will be back to the state that it was in when the last backup was performed, and Arcserve UDP Agent (Windows) backups will be able to continue as scheduled. (After completion of the BMR, the first backup will be a Verify Backup).

Operating Systems that Support UEFI/BIOS Conversion

If it is detected that the operating system of your source machine is not the same firmware as your system, you will be asked if you want to convert UEFI to a BIOS-compatible system or BIOS to UEFI-compatible system. The following table lists each operating system and the type of conversion supported:

Operating System (OS)	CPU	uEFI to BIOS	BIOS to uEFI
Windows Vista (None SP)	x86	No	No
Windows Vista (None SP)	x64	No	No
Windows Vista SP1	x86	No	No
Windows Vista SP1	x64	Yes	Yes
Windows Server 2008	x86	No	No
Windows Server 2008	x64	Yes	Yes
Windows Server 2008 R2	x64	Yes	Yes
Windows 7	x86	No	No
Windows 7	x64	Yes	Yes
Windows 8	x86	No	No
Windows 8	x64	Yes	Yes
Windows Server 2012	x64	Yes	Yes
Windows 8.1	x86	No	No
Windows 8.1	x64	Yes	Yes
Windows 10	x86	No	No
Windows 10	x64	Yes	Yes
Windows Server 2012 R2	x64	Yes	Yes
Windows Server 2016	x64	Yes	Yes
Windows Server 2019	x64	Yes	Yes

Managing the BMR Operations Menu

The BMR Operations menu consists of the following three types of operations:

- Disk Specific Operations
- Volume/Partition Specific Operations
- BMR Specific Operation

Disk Specific Operations:

To perform disk specific operations, select the disk header and click **Operations**.

Clean Disk

This operation is used to clean all partitions of a disk and is:

- An alternate method to delete all volumes of a disk. With the **Clean Disk** operation, you do not have to delete each volume one by one.
- Used to delete the non-Windows partitions. Due to a VDS limitation, the non-Windows partition cannot be deleted from the UI, but you can use this operation to clean them all.

  Note: During BMR, when the destination disk has non-Windows partitions or OEM partitions, you cannot select this partition and delete it from the BMR UI. Usually this would occur if you ever installed Linux/Unix on the destination disk. To resolve this issue, perform one of the following tasks:
- Select the disk header on the BMR UI, click **Operations**, and use the **Clean Disk** operation to erase all partitions on the disk.
- Open a command prompt and type **Diskpart** to open the Diskpart command console. Then type "select disk x" , where 'x' is the disk number and "clean" to erase all partitions on the disk.

Convert to MBR

This operation is used to convert a disk to MBR (Master Boot Record). It is available only when the selected disk is a GPT (GUID Partition Table) disk and there are no volumes on this disk.

Convert to GPT

This operation is used to convert a disk to GPT. It is available only when the selected disk is an MBR disk and there are no volumes on this disk.

Convert to Basic

This operation is used to convert a disk to Basic. It is available only when the selected disk is a Dynamic disk and there are no volumes on this disk.

Convert to Dynamic

This operation is used to convert a disk to Dynamic Disk. It is available only when the selected disk is a Basic disk.

Online Disk

This operation is used to bring a disk online. It is available only when the selected disk is in the offline status.

Disk Properties

This operation is used to view detailed disk properties. It is always available and when you select this operation, a **Disk Properties** dialog appears.

Volume/Partition Specific Operations:

To perform volume/partition operations, select the disk body area and click **Operations**. From this menu, you can create new partitions to correspond to the disk partitions on the source volume.

Create Primary Partition

This operation is used to create a partition on a basic disk. It is available only when the selected area is an unallocated disk space.

Create Logical Partition

This operation is used to create a logical partition on a basic MBR disk. It is available only when the selected area is an extended partition.

Create Extended Partition

This operation is used to create an extended partition on a basic MBR disk. It is available only when the disk is an MBR disk and the selected area is an unallocated disk space.

Create System Reserved Partition

This operation is used to create the System Reserved Partition on a BIOS firmware system and builds a mapping relationship with the source EFI System Partition. It is only available when you restore a UEFI system to a BIOS system.

Note: If you previously converted from UEFI to a BIOS-compatible system, use the Create System Reserved Partition operation for destination disk resizing.

Create EFI System Partition

This operation is used to create the EFI System Partition on a basic GPT disk. It is available only when the target machine firmware is UEFI and the selected disk is a basic GPT disk.

Note: If you previously converted from BIOS to a UEFI-compatible system, use the Create EFI System Partition operation for destination disk resizing.

Note: Systems that support UEFI also require that the boot partition reside on a GPT (GUID Partition Table) disk. If you are using a MBR (Master Boot Record) disk, you must convert this disk to a GPT disk, and then use the Create EFI System Partition operation for disk resizing.

Resize Volume

This operation is used to resize a volume. It is an alternate method of Windows "Extend Volume/Shrink Volume". It is available only when the selected area is a valid disk partition.

Delete Volume

This operation is used to delete a volume. It is available only when the selected area is a valid volume.

Delete Extended Partition

This operation is used to delete the extended partition. It is available only when the selected area is the extended partition.

Volume Properties

This operation is used to view detailed volume properties. When you select this operation, a **Volume Properties** dialog appears.

BMR Specific Operations:

These operations are specific to BMR. To perform BMR operations, select the disk header or the disk body area and click **Operations**.

Map Disk From

This operation is used to build a mapping relationship between the source and target dynamic disks. It is available only when the selected disk is a Dynamic disk.

Note: When mapping to another disk, the capacity of each mapped target volume must be the same size or larger than the corresponding source volume.

Map Volume From

This operation is used to build a mapping relationship between the source and target basic volume. It is available only when the selected volume is a Basic volume.

Note: When mapping to another disk, the capacity of each mapped target volume must be the same size or larger than the corresponding source volume.

Commit

This operation is always available. All of the operations are cached in memory and they do not modify the target disks until you select the **Commit** operation.

Reset

This operation is always available. The **Reset** operation is used to relinquish your operations and restore the disk layout to the default status. This operation cleans all the cached operations. Reset means to reload the source and target disk layout information from the configure file and current OS, and discard any user changed disk layout information.

Troubleshooting BMR Issues

When a problem is detected, Arcserve UDP Agent (Windows) generates a message to help you identify and resolve the problem. These messages are contained in the Arcserve UDP Agent (Windows) **Activity Log**, which is accessed from the **View Logs** option on the home page UI. In addition, when an incorrect action is attempted, Arcserve UDP Agent (Windows) generally displays a pop-up message to help you identify and quickly resolve the problem.

- Slow throughput performance during BMR
- After BMR, dynamic volumes are not recognized by the operating system
- Unable to Reboot Hyper-V VM After BMR
- Unable to Reboot VMware VM After BMR
- Unable to boot the server after performing a BMR
- Failed to submit BMR job to Recovery Point Server

Slow throughput performance during BMR

This problem can be caused by SATA controllers with "AHCI" enabled.

During BMR, Arcserve UDP Agent (Windows) will install drivers for critical unknown devices. If the device already has a driver installed, Arcserve UDP Agent (Windows) will not update that driver again. For some devices, Windows 7PE may have the drivers for them, but these drivers may not be the best ones and this can cause the BMR to run too slow.

To resolve this problem, perform one of the following tasks:

- Check if the driver pool folder contains the newest disk drivers. If it does, and you are restoring to the original machine, please install the new driver from the driver pool folder. If you are restoring to alternate machine, download the latest disk drivers from the Internet, and load it before you start data recovery. To load the driver, you can use the "drvload.exe" utility, which is included in Windows PE.
- Change the device operating mode from "AHCI" (Advanced Host Controller Interface) to Compatibility mode. (Compatibility mode provides a better throughput).

If the problem persists, use Live Chat to contact Arcserve support. Live Chat lets you optimize intelligent conversation between you and the Technical Support team, allowing you to address your concerns and questions immediately, while still maintaining access to the product.

After BMR, Dynamic Volumes are Not Recognized by the Operating System

To keep dynamic disks in a consistent state, the Windows operating system automatically synchronizes the Logical Disk Manager (LDM) metadata on each dynamic disk. So when BMR restores one dynamic disk and brings it online, the LDM metadata on this disk is automatically updated by the operating system. This may result in a dynamic volume not being recognized by the operating system and missing after the reboot.

To resolve this problem, when you perform BMR with multiple dynamic disks, do not perform any pre-BMR disk operations such as cleaning, deleting volume, and so on.

If the problem persists, use Live Chat to contact Arcserve support. Live Chat lets you optimize intelligent conversation between you and the Technical Support team, allowing you to address your concerns and questions immediately, while still maintaining access to the product.

Unable to Reboot Hyper-V VM After BMR

If you performed BMR to a Hyper-V machine consisting of more than one disk connected to an Integrated Drive Electronics (IDE) controller and if the server does not reboot, perform the following troubleshooting procedure:

1. Verify that the disk that contains the system volume is the master disk.

   The Hyper-V BIOS searches for the system volume on the master disk (disk 1) which is connected to the master channel. If the system volume is not located on the master disk, the VM will not reboot.

   Note: Verify that the disk that contains the system volume is connected to an IDE controller. Hyper-V cannot boot from a SCSI disk.

2. If necessary, modify the Hyper-V settings, to connect the disk that contains the system volume to the IDE master channel and reboot the VM again.

   If the problem persists, use Live Chat to contact Arcserve support. Live Chat lets you optimize intelligent conversation between you and the Technical Support team, allowing you to address your concerns and questions immediately, while still maintaining access to the product.

Unable to Reboot VMware VM After BMR

If you performed BMR to a VMware machine consisting of more than one disk connected to an Integrated Drive Electronics (IDE) controller or a SCSI adapter and the server does not reboot, perform the following troubleshooting procedure:

1. Verify that the disk that contains the system volume is the master disk.

   The VMware BIOS searches for the system volume on the Master disk (disk 0) which is connected the master channel. If the system volume is not on the Master disk, the VM does not reboot.

2. If necessary, modify the VMware settings, to connect the disk that contains the system volume to the IDE master channel and reboot the VM again.

3. If the disk is a SCSI disk, verify the disk which contains boot volume is the first disk which connects to the SCSI adapter. If not, assign the boot disk from the VMware BIOS.

4. Verify the disk which contains boot volume is in the previous eight disks, because the VMware BIOS only detect eight disks during the boot. If there are more than seven disks ahead the disk which contains system volumes connected to the SCSI adapter, the VM cannot boot.

   If the problem persists, use Live Chat to contact Arcserve support. Live Chat lets you optimize intelligent conversation between you and the Technical Support team, allowing you to address your concerns and questions immediately, while still maintaining access to the product.

Unable to boot the server after performing a BMR

Symptom

When the source machine is an Active Directory server performing a BMR to a physical machine with different hardware or to a virtual machine on a Hyper-V server, the server does not boot and a blue screen displays with the following message:

STOP: c00002e2 Directory Services could not start because of the following error: a device attached to the system is not functioning. Error status: 0xc0000001.

Solution

Reboot the system to the BMR PE environment, rename all *.log files in the C:\Windows\NTDS folder, and restart the system. For example, rename the file edb.log to edb.log.old and restart the system.

If the problem persists, use Live Chat to contact Arcserve support. Live Chat lets you optimize intelligent conversation between you and the Technical Support team, allowing you to address your concerns and questions immediately, while still maintaining access to the product.

Failed to submit BMR job to Recovery Point Server

Only one BMR job is supported when restoring from same RPS server for the same node (Agent backup or Host-Based Backup). This is controlled by the job monitor on the RPS server.

If the machine where the BMR job is running is shut down or rebooted unexpectedly, the job monitor at the RPS server side will wait 10 minutes and then time out. During this time you cannot start another BMR for the same node from the same RPS server.

If you abort the BMR from the BMR UI, this problem does not exist.

If the problem persists, use Live Chat to contact Arcserve support. Live Chat lets you optimize intelligent conversation between you and the Technical Support team, allowing you to address your concerns and questions immediately, while still maintaining access to the product.

BMR Target Machine Boots into System Recovery Options Screen after Performing a BMR

Symptom

If you perform a BMR to a virtual machine on an ESX server, at times the BMR target machine boots into the System Recovery Options screen.

Solution

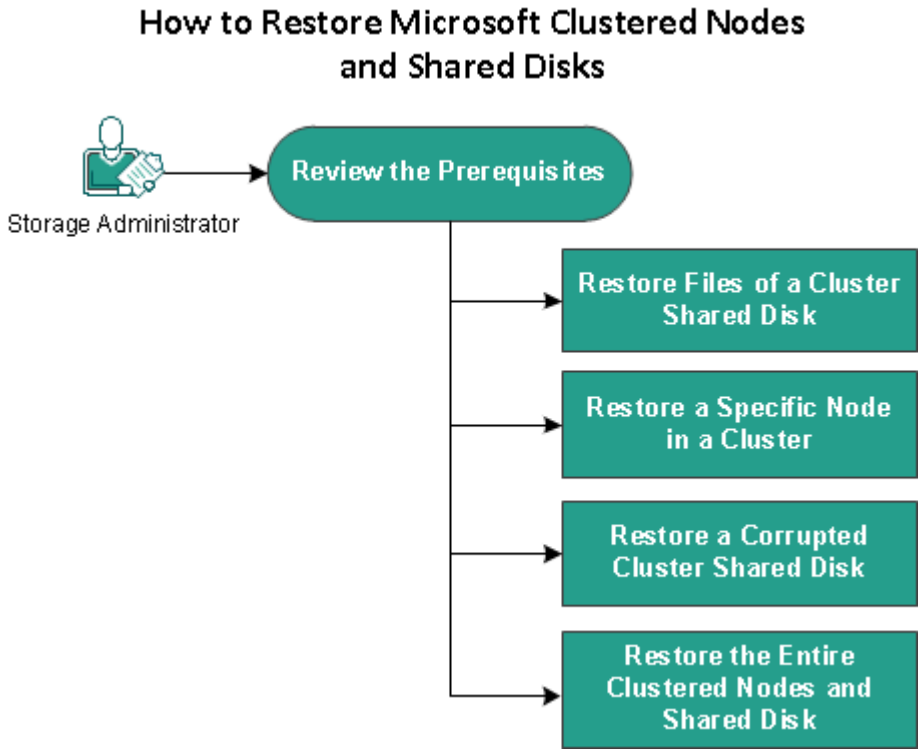
To resolve this problem, Disable specific VSS writers with VMware Tools on the Instant VM machine, and then perform a BMR again.

How to Restore Microsoft Clustered Nodes and Shared Disks

If you have a clustered environment and the clustered nodes and shared disk are not functioning properly, you can easily recover the nodes and disks. You can restore the following items:

- Individual files and folders in a shared disks
- Specific nodes in a cluster
- Entire shared disk
- Entire cluster setup (all clustered nodes and shared disk)

The following diagram illustrates the process to restore clustered nodes and shared disks:



Follow these steps to restore Microsoft clustered nodes and shared disks:

- Review the Prerequisites
- Restore Files of a Cluster Shared Disk
- Restore a Specific Node in a Cluster
- Restore a Corrupted Cluster Shared Disk
- Restore the Entire Clustered Nodes and Shared Disk

Review the Prerequisites

Verify that you have completed the following prerequisites:

- You have a valid recovery point for restore.
- You have a valid ISO image for a BMR.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Restore Files of a Cluster Shared Disk

The shared disk belongs to one of the nodes from the cluster. When you recover any files from the shared disk (not the cluster quorum disk), you need to find the parent node of the shared disk. After you identify the parent node, you can recover files to the parent node from the shared disk.

Note: After a failover happens, you have to browse the recovery point of a different agent to find out the desired recovery point.

Follow these steps:

1. Log in to the agent that owns the shared disk.
2. Open the Restore Wizard and select Find Files/Folders to Restore.

   Note: For more information on restoring the files and folders, see How to Restore Files/Folders.
3. Select all the files from the Restore Wizard that you want to restore to the original location.
4. Complete the Restore Wizard configurations and submit the job.

   The files are recovered.
5. Log in to the parent node of the shared disk and verify the files are recovered.

   The files of the shared disk are recovered.

Restore a Specific Node in a Cluster

If a specific node in a cluster is down, you can perform a BMR for only that node. Typically, in this scenario the shared disk is in a good state and does not need a recovery.

Follow these steps:

1. Prepare the BMR image (CD / DVD or USB stick).
2. Remove all the connections between the node that you want to recover and the shared disks.

   Example: Disconnect the fibre channel connection.
3. Perform a BMR for the cluster node.

   Note: For more information on performing a bare metal recovery, see How to Perform a BMR Using a Backup.

   The specific node in a cluster is recovered.
4. Check the status of the recovered node in the cluster management console and ensure that it acts as part of the cluster.

   The specific node in a cluster is recovered.

Restore a Corrupted Cluster Shared Disk

The shared disk belongs to one of the nodes from the cluster. If the shared disk is corrupted or broken, you can restore the specific files or folders of the shared disk, without recovering the clustered nodes. Typically, in this scenario the quorum disk and all the cluster nodes are in a good state.

Follow these steps:

1. Replace the corrupted disk manually and reconfigure the cluster shared disk.
2. Identify the agent that owns the shared disk and log in to that agent.
3. Open the Restore Wizard and select Find Files/Folders to Restore.

   Note: For more information on restoring the files and folders, see How to Restore Files/Folders.
4. Select all the files from the Restore Wizard that you want to restore to the original location.
5. Complete the Restore Wizard configurations and submit the job.

   The shared disk is recovered.
6. Check the status of the shared disk in the cluster management console and ensure that it acts as a part of the cluster.

   The shared disk is recovered.

Restore the Entire Clustered Nodes and Shared Disk

If the entire clustered setup is corrupted or not functioning, you can recover the entire cluster. Recovering the entire cluster is a two-part process. First you recover individual clustered nodes using BMR. Then you recover the files and folders of the shared disk.

Note: For quorum disks, rebuild the disk using the cluster management console instead of recovering it using the Restore Wizard in Arcserve UDP Agent (Windows).

Follow these steps:

1. Prepare the BMR image (CD / DVD or USB stick).
2. Remove all the connections between the node that you want to recover and the shared disks.

   Example: Disconnect the fibre channel connection.
3. Perform a BMR for the cluster node.

   Note: For more information on performing a bare metal recovery, see How to Perform a BMR Using a Backup.

   The specific node in a cluster is recovered.
4. Check the status of the recovered node in the cluster management console and ensure that it acts as part of the cluster.

   The specific node in a cluster is recovered.
5. Repeat the steps to recover all the clustered nodes.

   All the clustered nodes are recovered. Now recover the shared disk.
6. Replace the corrupted disk manually and reconfigure the cluster shared disk.
7. Identify the agent that owns the shared disk and log in to that agent.
8. Open the Restore Wizard and select Find Files/Folders to Restore.

   Note: For more information on restoring the files and folders, see How to Restore Files/Folders.
9. Select all the files from the Restore Wizard that you want to restore to the original location.
10. Complete the Restore Wizard configurations and submit the job.

    The shared disk is recovered.
11. Verify the files of the shared disk and ensure the files are recovered.

    The entire cluster is recovered.

How to Restore an Active Directory

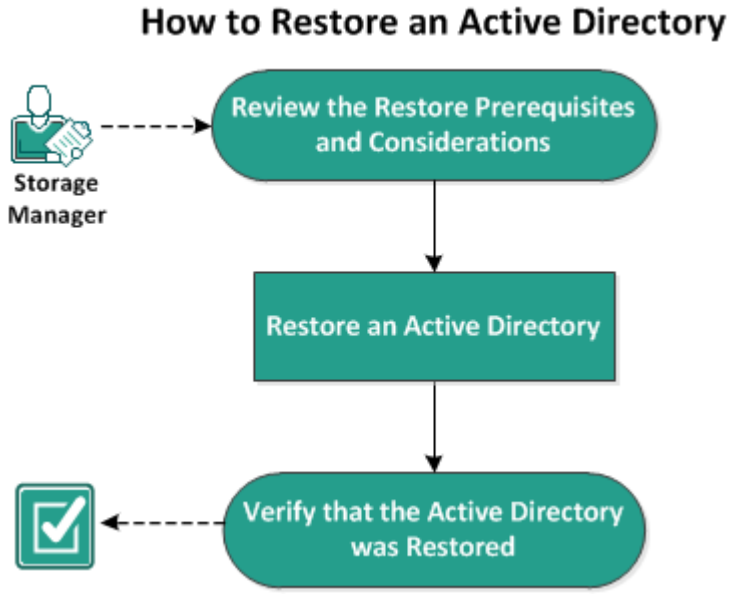
You need to restore a backed up Active Directory session if you have any of the following scenarios:

- You want to recover an attribute of the Active Directory object from any available backed up Active Directory session (not only the last backed up session).
- You want to recover the Active Directory object from any available backed up Active Directory session (not only the last backed up session).
- You want to recover multiple Active Directory attributes or objects from any available backed up Active Directory session (not only the last backed up session).

Important! To perform a granular recovery of an Active Directory, an agent-based backup is required.

List of Active Directory Objects Recovered by Object-level Restore		
Organizational Unit	Site	Lost and found class
User	Site Container	Build in Domain class
Group	Site Link	DNS Zone class
Computer	Site Link Bridge	Domain Class
Contact	Site Settings	Domain DNS Class
Connection	Subnet Container	DMD Class
Shared Folder	Trusted Domain	Organizational Unit Class
Printer	Configuration Class	Containerecifiers class

The scenario describes how you can restore an Active Directory.



Perform the following tasks to restore an Active Directory:

1. Review the Restore Prerequisites and Considerations
2. Restore an Active Directory
3. Verify that the Active Directory was Restored

Review the Restore Prerequisites and Considerations

Verify that the following prerequisites exist before performing a restore:

- You have already backed up the volumes that include the Active Directory database folder and Log files folder.
- You have the Arcserve UDP Agent (Windows) installed on Domain Controller.
- You have performed an agent-based backup.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Review the following restore considerations:

- For a recovery point without a file system catalog created, to ensure you can browse and select files/folders to restore, the account/group should be granted access permission to all the folders/files on all volumes with read/list access before the backup is taken.
- You can perform an Active Directory restore only on the Arcserve UDP Agent (Windows).

Restore an Active Directory

After you have installed the Active Directory in different volumes and have performed a backup for both volumes, you may want to restore the volumes with the Active Directory. This scenario describes how you can restore the backed up Active Directory volumes.

Note: Verify that you have completed the prerequisites and backed up Active Directory volumes.

Follow these steps:

1. Access the restore method selection dialog in one of the following ways:

   From Arcserve UDP:

   a. Log in to Arcserve UDP.

   b. Click the **resources** tab.

   c. Select **All Nodes** in the left pane.

      All the added nodes are displayed in the center pane.

   d. In the center pane, select the node and click **Actions**.

   e. Click **Restore** from the **Actions** drop-down menu.

      The restore method selection dialog opens.

      Note: You are automatically logged in to the agent node and the restore method selection dialog is opened from the agent node.

   From Arcserve UDP Agent (Windows):

   a. Log in to Arcserve UDP Agent (Windows).

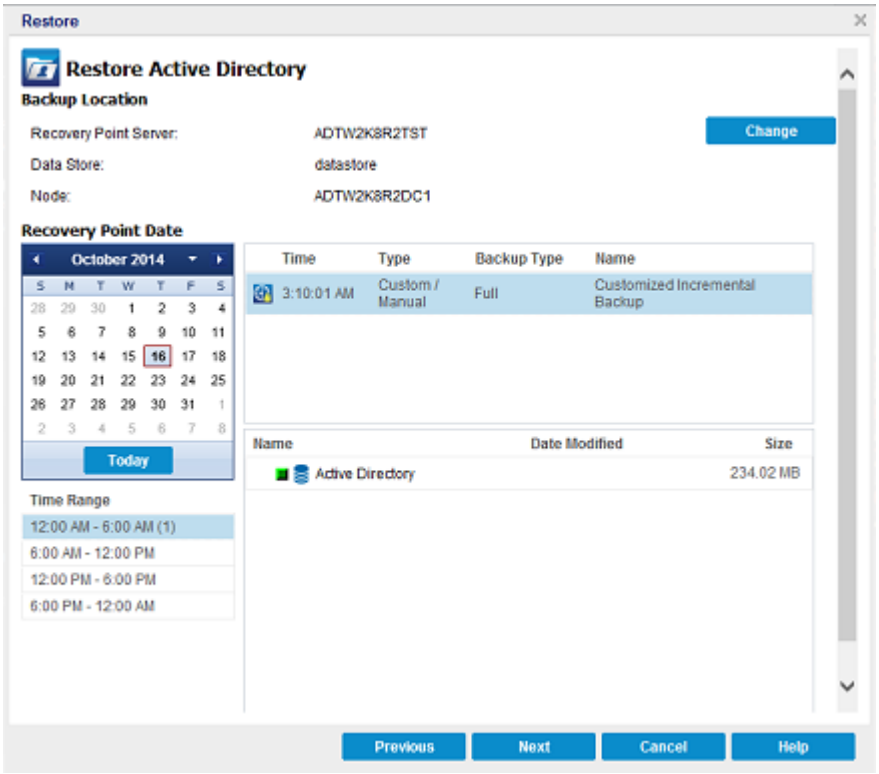
   b. From the home page, select **Restore**.

   The restore method selection dialog opens.

2. From the Restore screen, click Restore Active Directory.

   The Restore Active Directory dialog opens.

3. From the Restore Active Directory screen, perform the following steps:



a. From the calendar, select Backup date for the Active Directory that you want to restore.

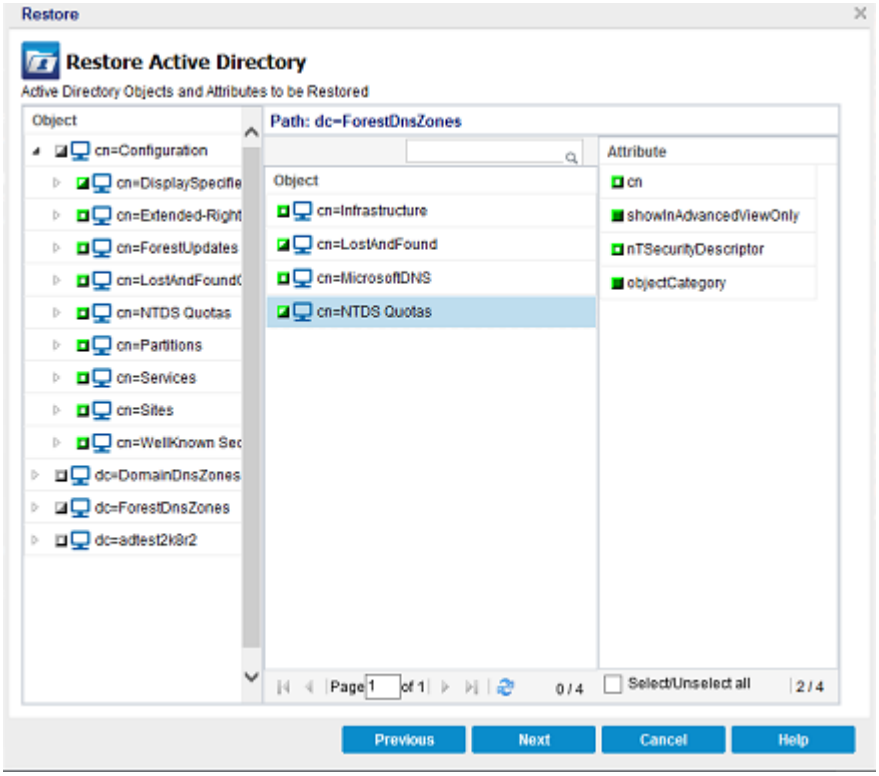
b. From the Time range, select Backup time.

c. From the Restore Active Directory screen, select a recovery point that is identified with Time, Type, Backup Type, and Name.

d. From the Name section, select an Active Directory backup session to restore.

4. Click Next.

5. Select the following options to further define the objects, path, and attributes to restore:



a. From the Object column, select the name of an object. The paths related to the selected object are displayed.

b. From the Path column, select a path. The attributes related to the selected path are displayed.

   Note: You can use the search icon to browse for the path.

c. From the Attribute column, select one or more attributes.

6. Click Next.

   The Restore Options screen is displayed.

7. From the Restore Options, select the following objects according to your requirement:

   a. If the selected object was renamed after backup, click the "Restore with original name of Renamed Objects" option to restore the renamed object.

      Note: If you do not select this option, the object will not be restored.

   b. If the selected object was moved to another container after backup, click the "Restore to original location of Moved Objects" option to restore the moved object.

      Note: If you do not select this option, the object will not be restored.

c. If the selected object was deleted permanently after backup, click the "Restore with the new object ID of Deleted Objects" option to restore the permanently deleted object.

   Note: Using this option helps you keep the restored object with the new object ID.

8. Click Next.

   The Restore Summary screen is displayed

9. Review the details and perform one of the following action:

   - Click Previous, if you want to modify the details.
   - Click Finish to run restore.

   A status message is displayed to inform you when the Restore job is completed. If the restore is unsuccessful, view the logs and try again.

Verify that the Active Directory was Restored

After the completion of the restore process, you can use the Active Directory Users and Computers utility to verify that the Active Directory (object and/or attribute) was restored to the specified destination.

Note: The Active Directory utility is installed automatically with the Active Directory.

How to Restore Active Directory Data Using Arcserve UDP Active Directory Object Level Restore Utility

After you have installed the Active Directory in different volumes and have performed a backup for both volumes using a Host-Based Agentless Backup task, you may want to restore objects and attributes from Active Directory granularly. This scenario describes how you can restore the backed up Active Directory objects and attributes from its containing volumes.

Before you perform a restore, make sure that the following prerequisites are available:

- The Active Directory Object Level Restore utility is available in the following location:

  <Arcserve UDP installation path>\Engine\BIN\

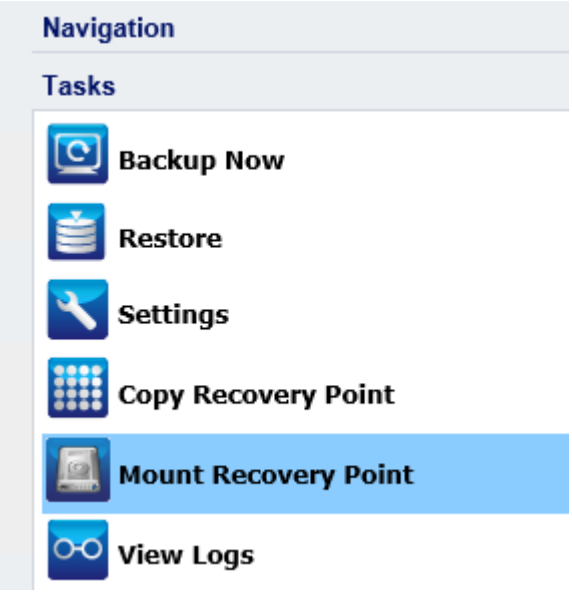
  Note: The tool is installed with Arcserve UDP Agent.
- The Restore job is set to run from the host-based backup proxy machine.

  Note: If you want to run the restore job on any other machine, search the recovery point from the backup destination.
- The path to Active Directory database (NTDS.dit) is identified to perform the restore job.

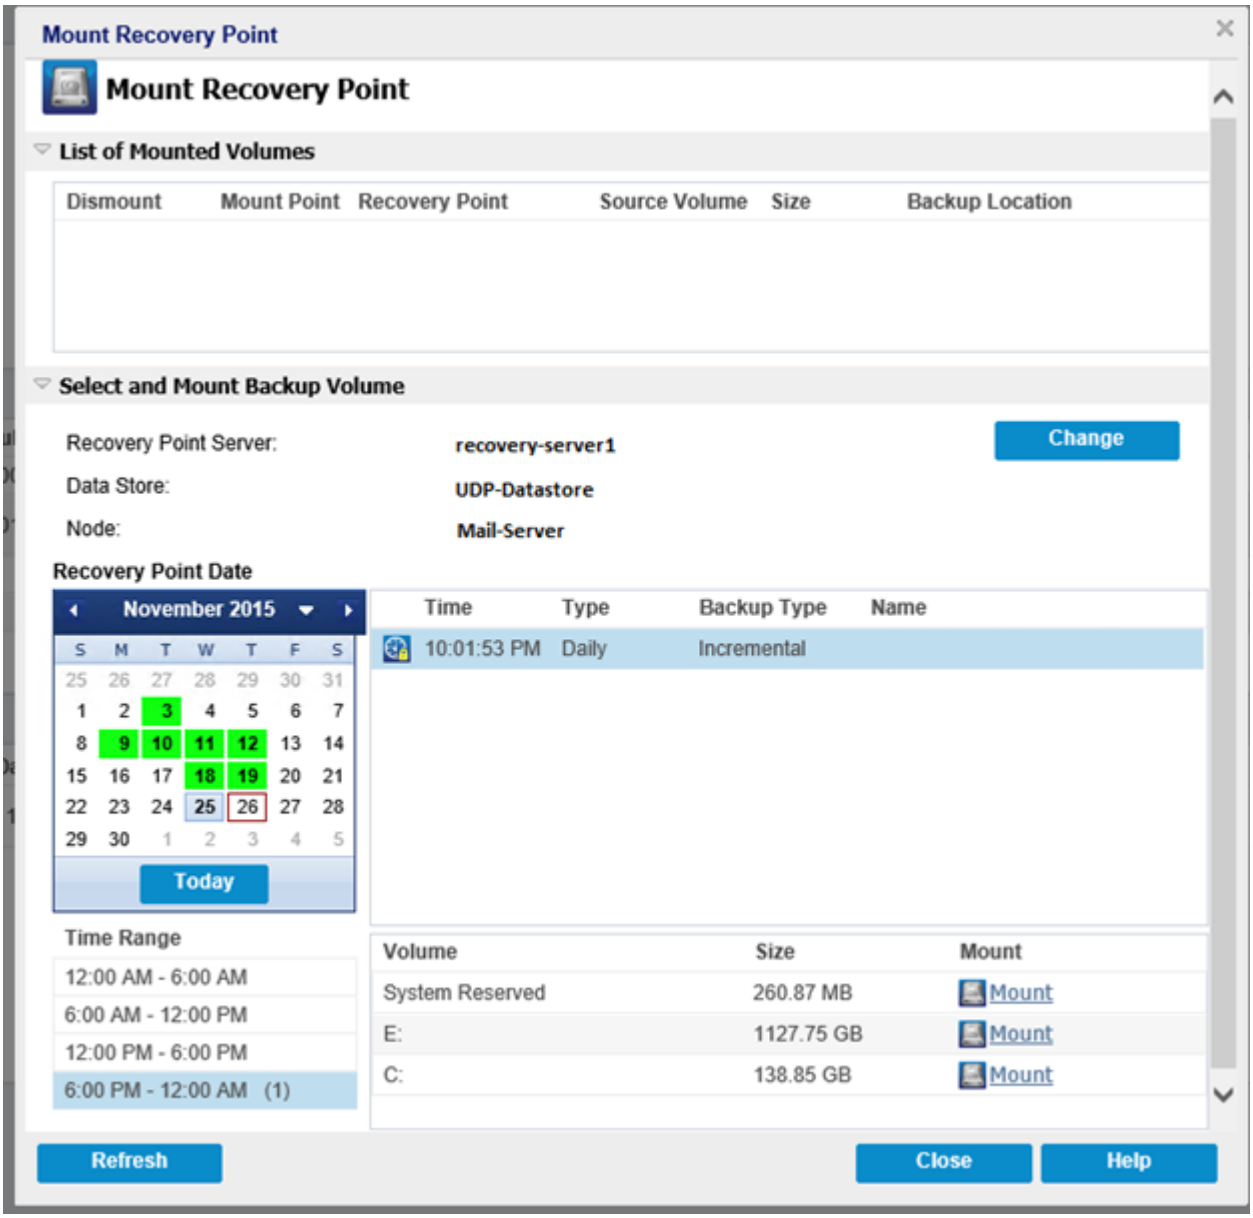
  Note: By default the path to NTDS.dit is C:\Windows\NTDS\NTDS.dit.

Follow these steps:

1. From the Arcserve UDP Agent console on the host-based backup proxy machine, select the Mount Recovery Point task. The Mount Recovery Point dialog opens.

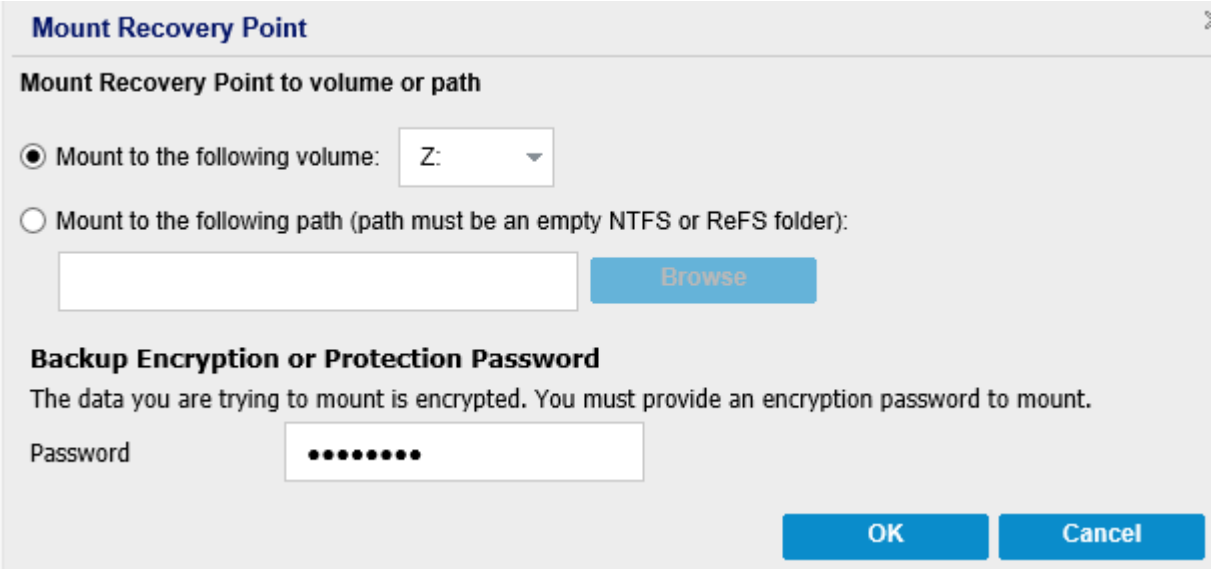


2. Select the recovery point date.
3. For the volume that contains the Active Directory database, click **Mount**.



Note: If the server that is running the restore job is not the HBBU proxy, click **Change** to select the appropriate Recovery Point Server, Data Store, and Active Directory Server.

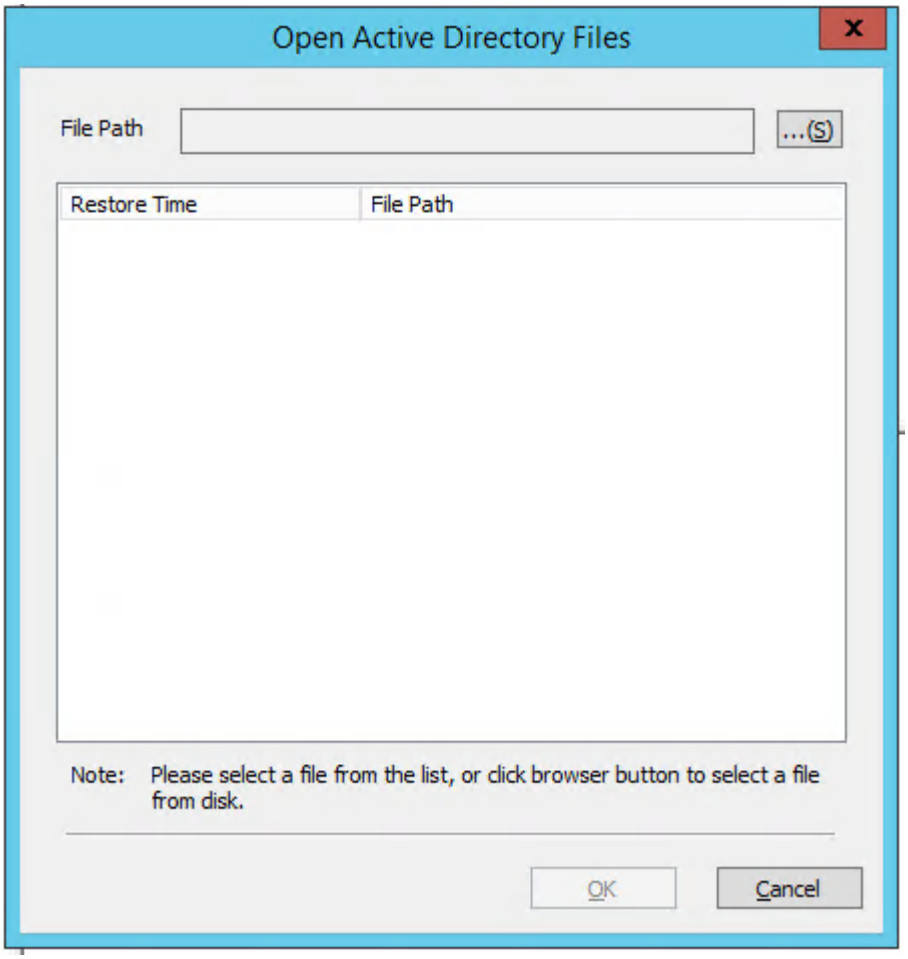
4. Select the drive letter to mount the volume and click **OK**.



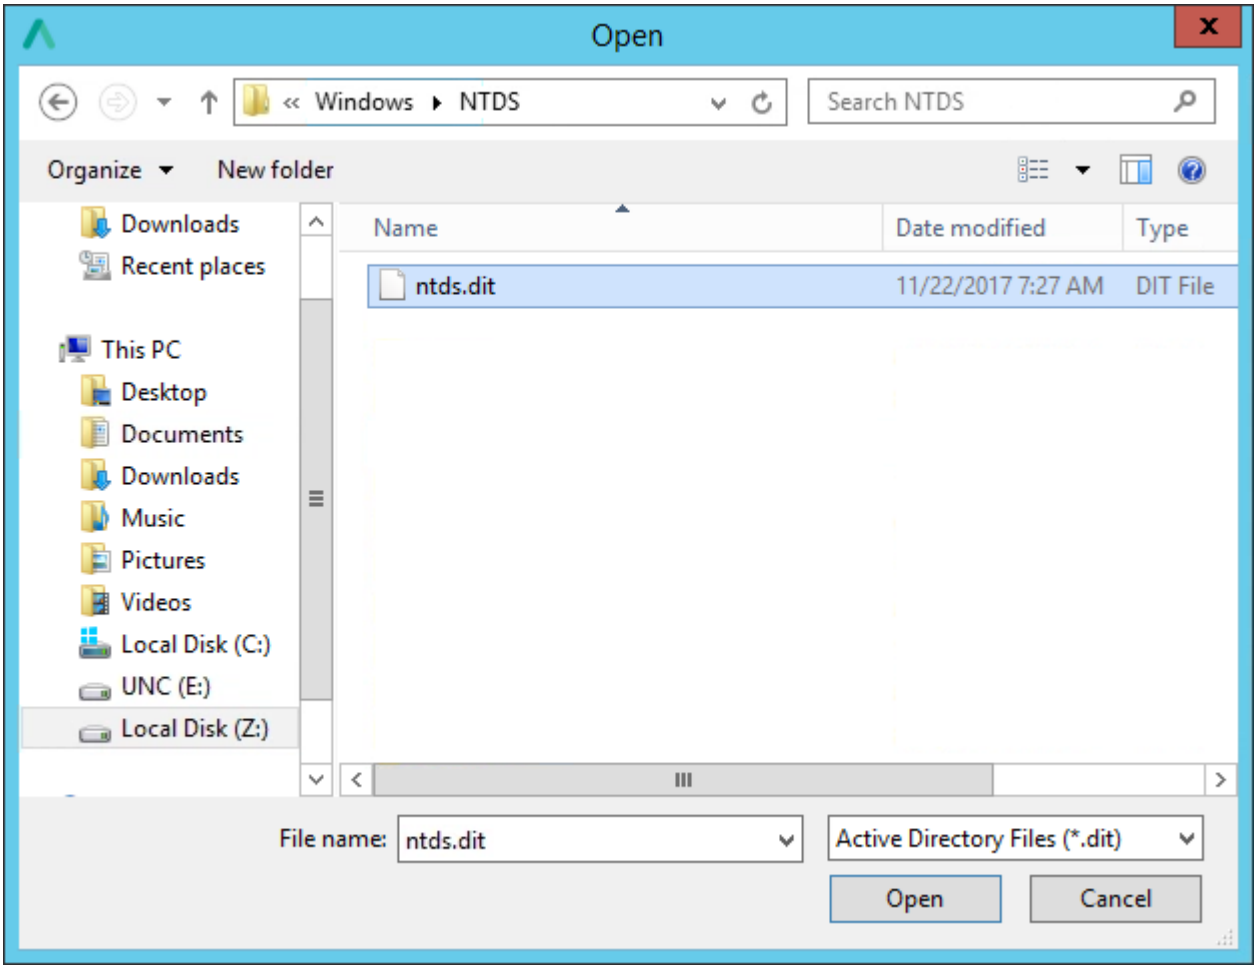
5. Launch the Active Directory Object Level Restore Utility from the following location:

   <Arcserve UDP installation path>\Engine\BIN\AD_restore.exe

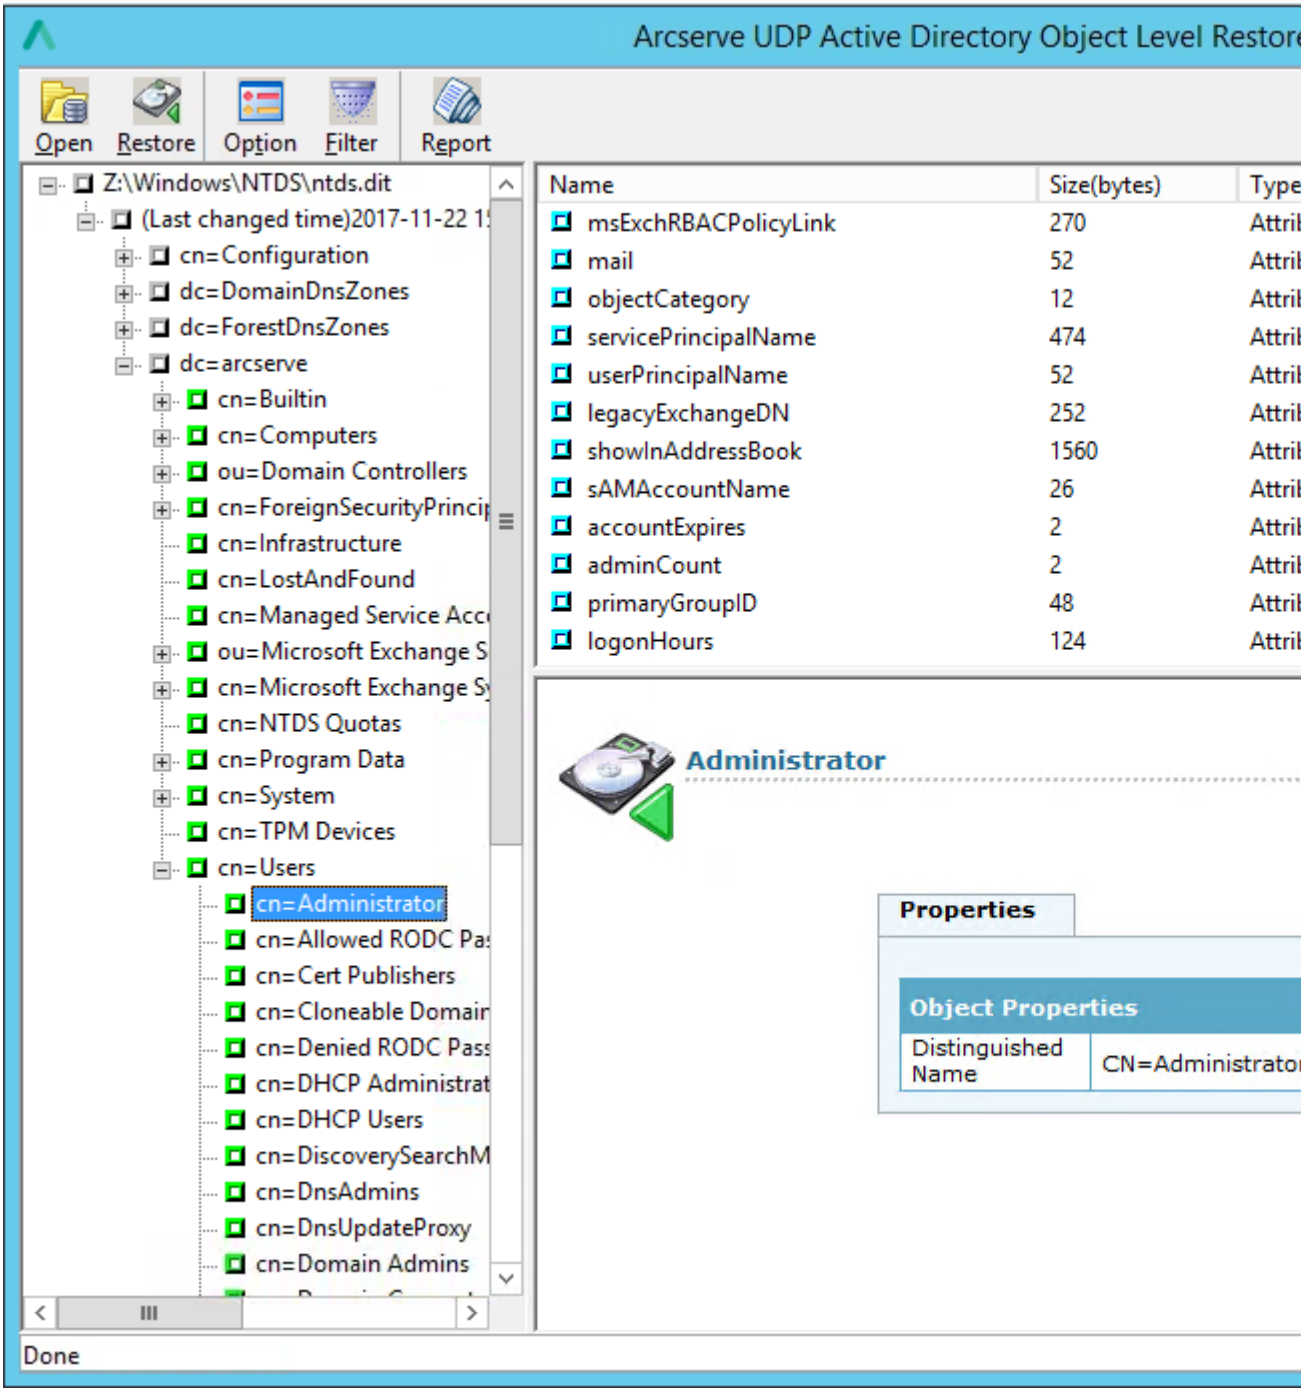
6. Click **Open** to open the selection window.



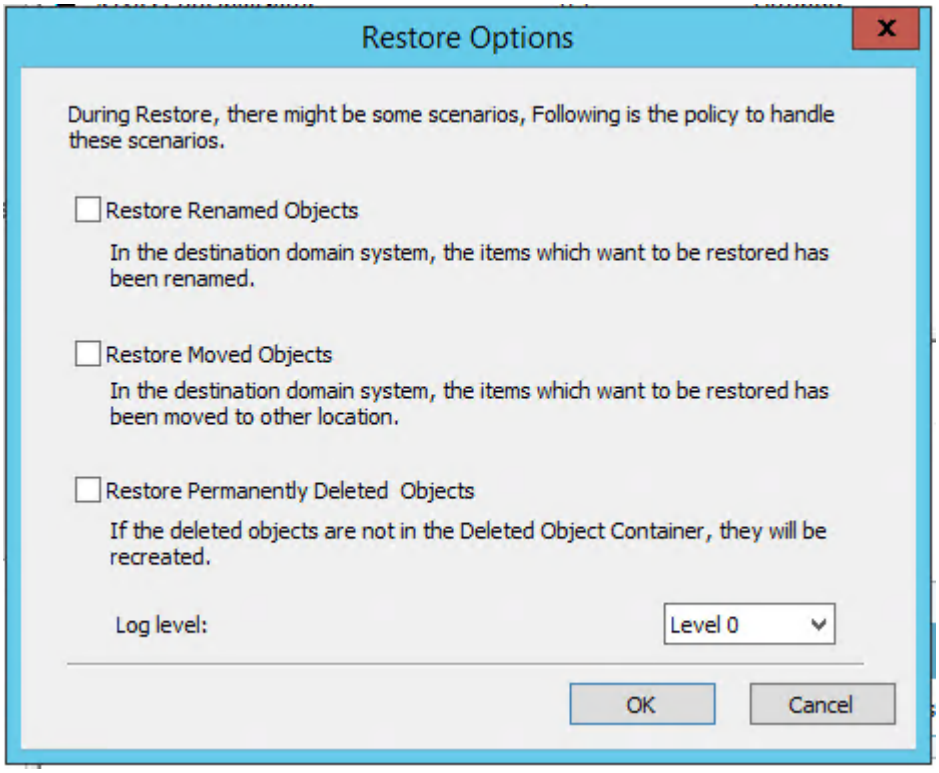
7. Click the ...(S) icon and browse the Active Directory database (NTDS.dit) on the mounted recovery point, click **Open**, and then click **OK**.



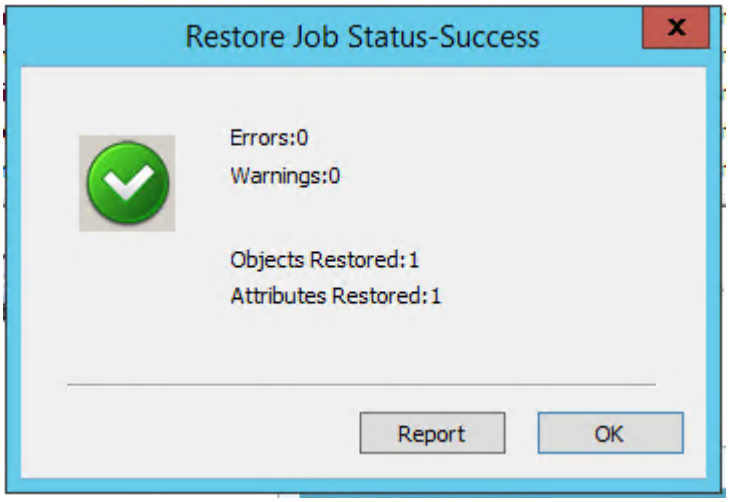
8. Browse and select the Active Directory objects or attributes you want to restore.



9. Click **Options** to make adjustments to the default behavior, as needed.



10. When ready, click **OK** to run the restore job.
11. After restore job is finished, a result window is displayed. Click **Report** if you wish to see the details or **OK** to close.



Notes:

- By default, the utility uses the current user who is logged into Windows to establish the connection.
- If an error is reported, the recommended action to take is to log into the machine with an account that has domain admin rights to perform the restore.

12. When the restore job completes, dismount the volume that was used for the recovery.
13. To dismount the volume, from the Arcserve UDP Agent console on the host-based backup machine, click **Mount Recovery Point**, and then click **Dismount**.



How to Restore Microsoft Office 365 Data

This section contains the following topics:

How to Restore Exchange Online Mailbox Data

You can restore Exchange Online mailbox data (emails, calendars, contacts, notes, tasks, and so on) from Microsoft cloud using any computer. You can restore data from recovery point to the original or an alternate location.

Perform the following tasks to restore Exchange Online mailbox data:

1. Select the Exchange Online Mailbox Items to Restore
2. Define the Restore Options
3. Restore the Recoverable Items
4. Restore the Recovery Point Content
5. Verify that Content is Restored

Select the Exchange Online Mail Items to Restore

You can restore Exchange Online mail data from a recovery point. When you select a recovery date, and then specify the time, all the associated recovery points for that duration are displayed. You can then browse and select the backup content (including applications) to restore.

Follow these steps:

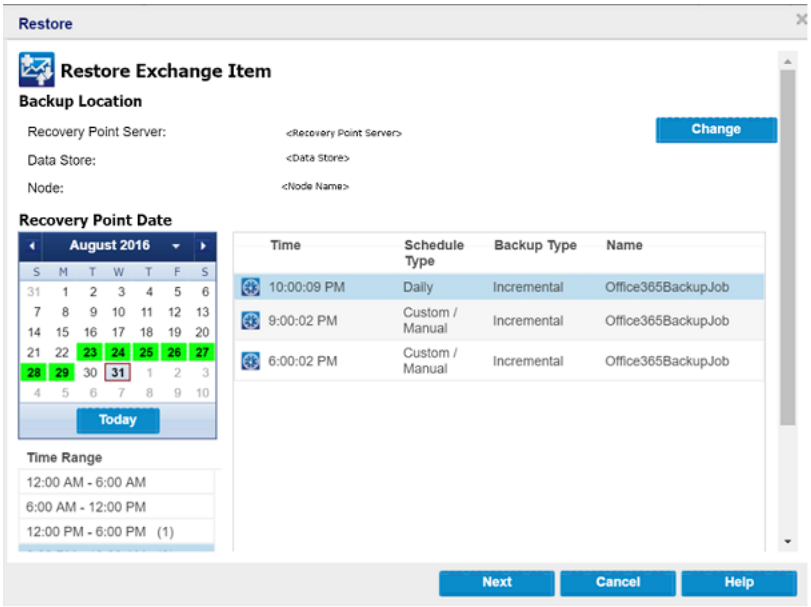
1. Log in to Arcserve UDP.
2. Click the **resources** tab.
3. Select **All Nodes** in the left pane.

   All the added nodes are displayed in the center pane.
4. In the center pane, select the Exchange Online node, and then click **Actions**.
5. Click **Restore** from the Actions drop-down menu.

   The Restore Exchange Item dialog opens.

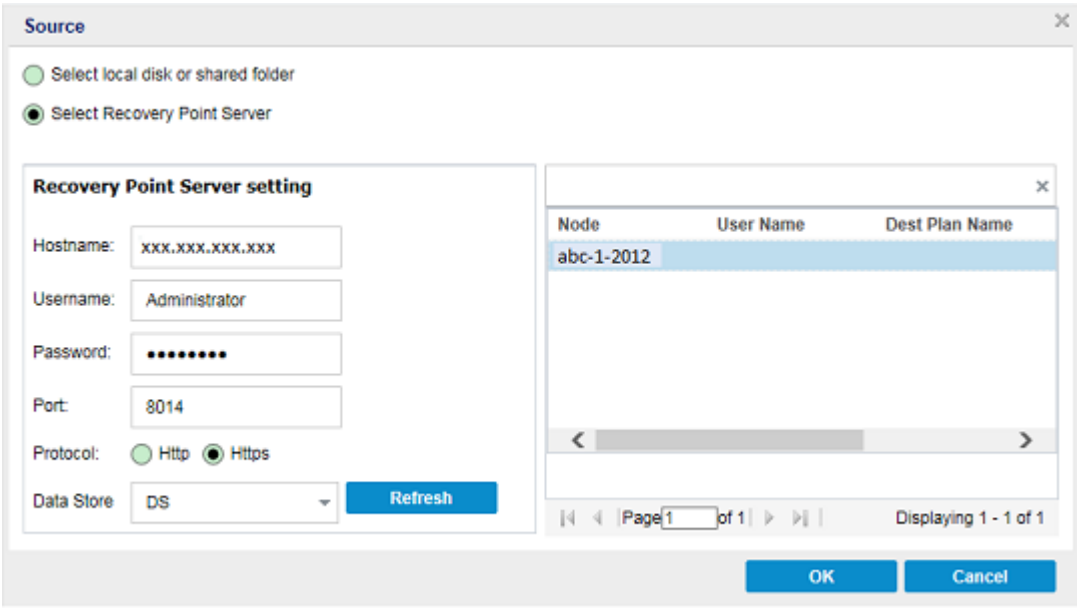
   Note: You are automatically logged in to the agent node and the Restore Exchange Item dialog opens.

   You can see the Recovery Point Server details in the Backup Location.



6. (Optional) Click **Change**, if you want to modify change the backup location.

   The Source dialog opens. You can select the backup location in this dialog.



7. To specify Source, select one of the following options, and then click **OK**:

   Select local disk or shared folder

   Note: In Arcserve UDP, we do not recommend to select the **Select local disk or shared folder** option.

   Select Recovery Point Server

   a. Specify the Recovery Point Server setting details, and then click **Refresh**.

      All the agents are listed in the Data Protection Agent column in the Source dialog.

   b. Select the agent from the displayed list and then click **OK**.

      The recovery points are listed in the Restore Exchange Item dialog.

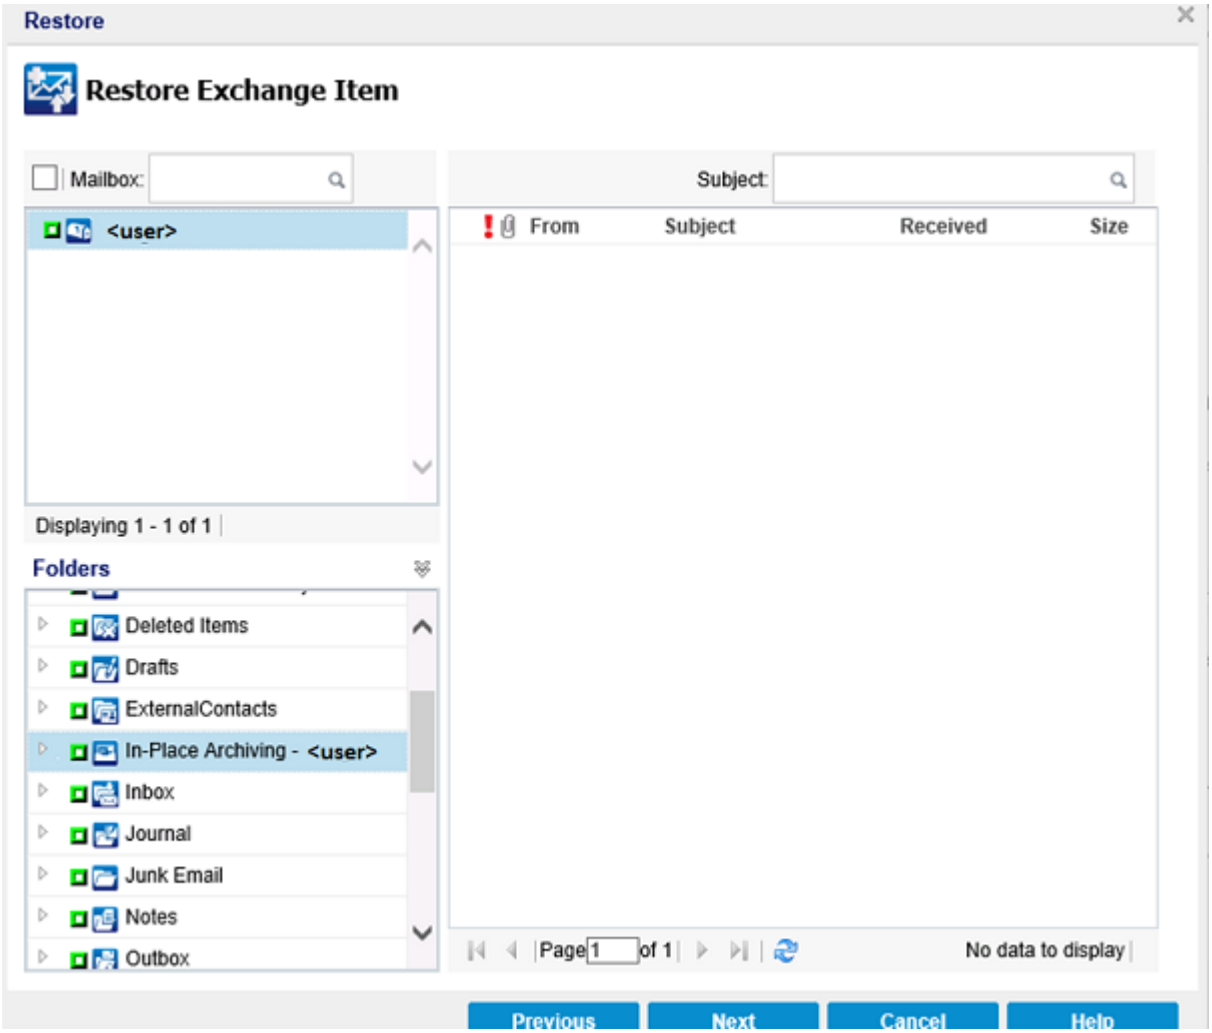
      Note: From the recovery point Folder, you may see many folders with the same name along with the exchange online nodes. This happens because the node GUID changes and a new recovery point folder is created when you delete a node and add again in the test plan.

8. Select the calendar date for the backup image to restore, and then click **Next**.

   All the dates containing recovery points for the specified backup source are high-lighted in green.

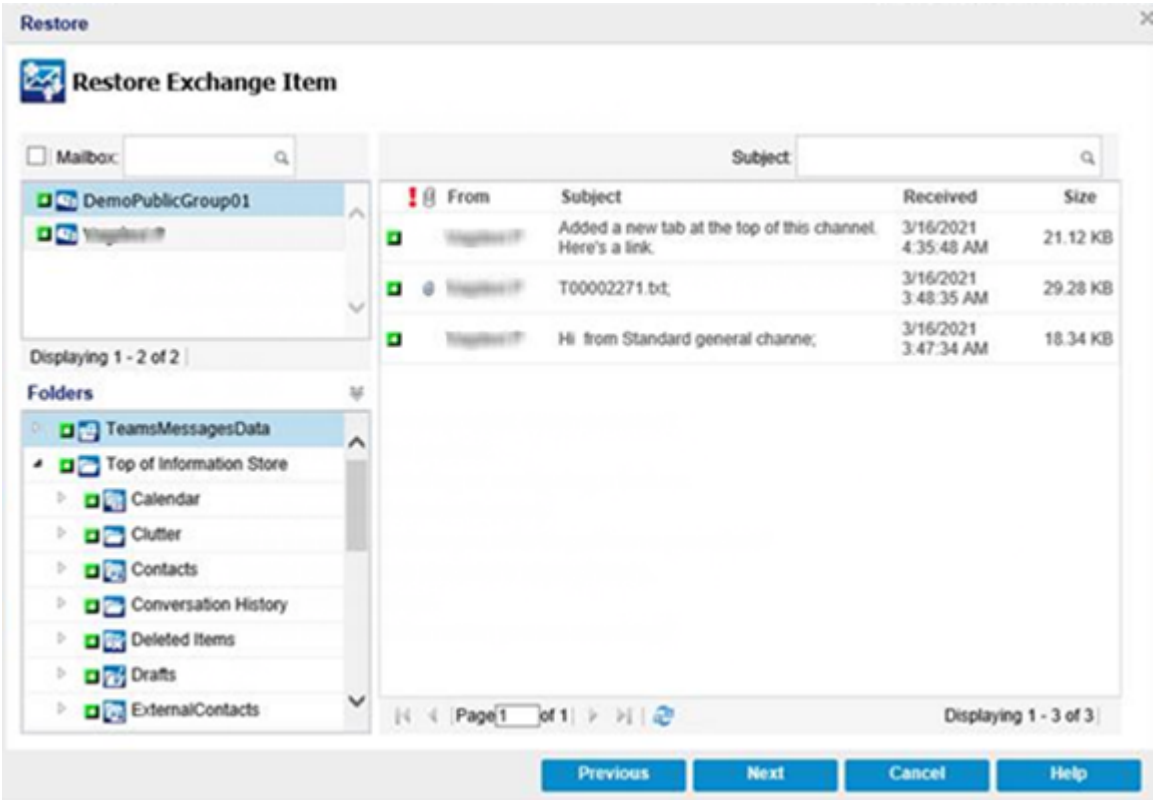
   The corresponding recovery points for that date are displayed, with the time of the backup, the type of backup that was performed (Full, Incremental, or Verify), and the name of the backup.

9. From the Mailbox pane, click the mailbox that you want to restore.



All the mail items related to the mailbox are displayed under the Folders pane.

Note: For Office 365 Exchange Online Groups backup, the folder structure for the group mailbox is different from the user mailbox. The group mailbox folders are shown under the Top of the Information Store folder. For example, when user backups AD Group mailboxes/distribution lists, all the folders that are backed up are shown under the Top of the Information Store folder.



10. Select the related mail items or folders (including the entire mailbox, emails, calendars, contacts, notes, tasks, and so on) that you want to restore from Folders, and then click **Next**.

    Notes:

    - You can select the entire content or partial content of the Exchange object to restore. To select partial content, expand the object, and click the check box of that content.
    - You can select multiple Exchange objects to restore.

    The Restore Options dialog opens.

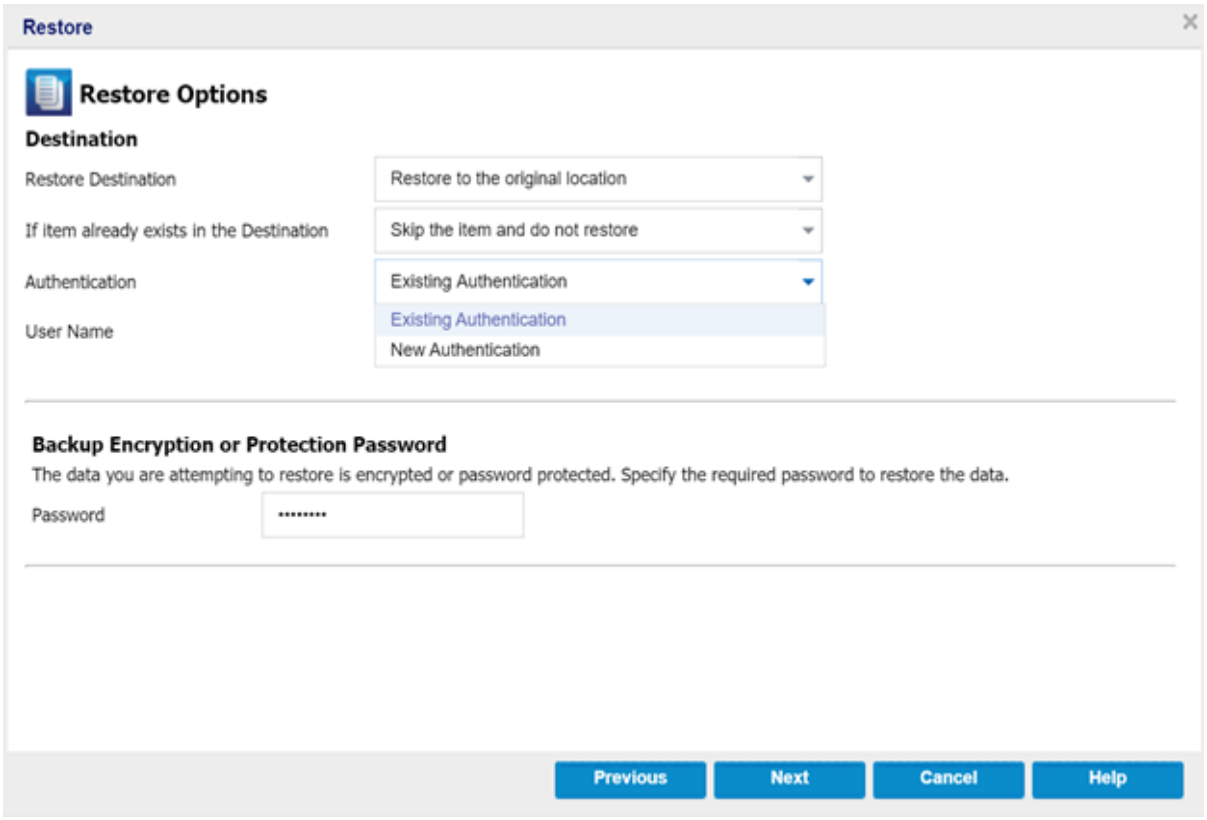
    The Exchange Online mail items to restore are selected. Now, you can define the restore options.

Define the Restore Options

After you specify the Exchange Online information to restore, define the restore options for the selected content.

Follow these steps:

1. On the **Restore Options** dialog, select the restore destination.



The available destination options are:

Restore to original location

Restores the mail data to the same location from where you took a back up.

Notes:

- If you restore a mail item to original location using the overwrite option twice, then after restoring second time, the first restore mail item is not overwritten. As a result, two similar mail items appear in original location.

- If you restore a mail item to original location using the skip option, and in the original folder an item similar to the restored item already exists, then the backup job displays incomplete result.

Restore to an alternate location

Restores the mail data to another mailbox or another folder in the original mailbox. When you select this option, you can browse and select the destination.

2. Specify one of the following options from the **If Item already exists in the Destination** drop-down:

Skip the item and do not restore

Skips over the items and does not restore.

Default: The Skip the item and do not restore.

Overwrite the item in the destination

Overwrites the item in the destination.

Note: The **If Item already exists in the Destination** drop-down list is available if you select the **Restore to original location** option from the **Restore Destination** drop-down list.

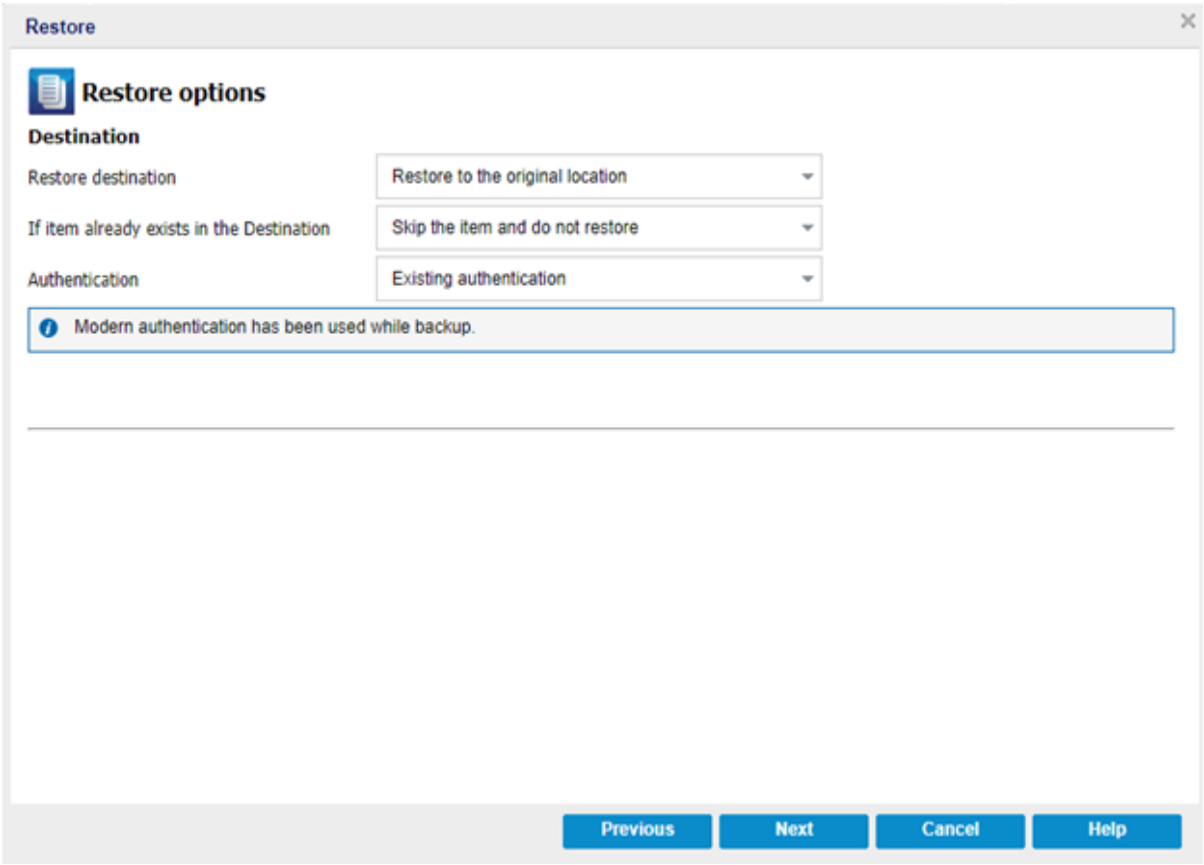
Authentication

Validates and identifies the user account. Select one of the following options:

- **Existing Authentication**

  When you select Existing Authentication, the authentication type that was selected during Node creation is used.

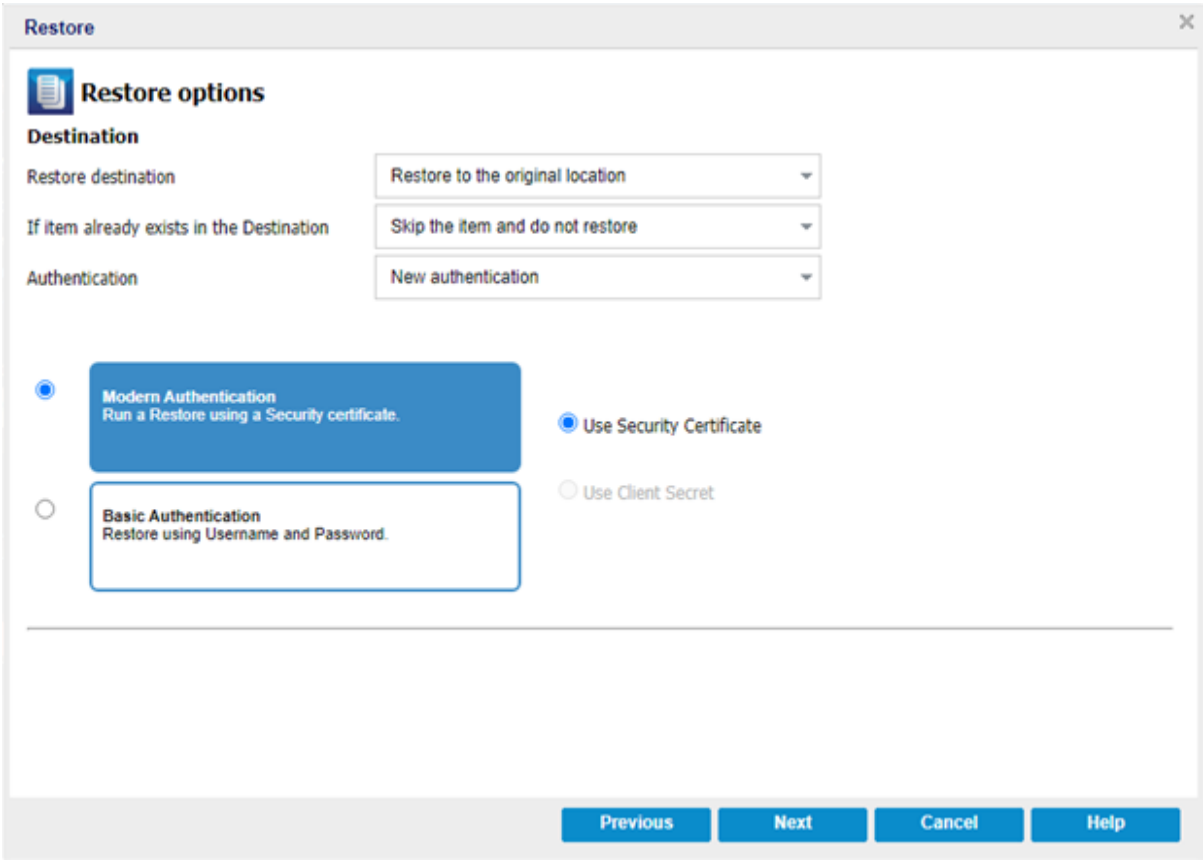
  - If the Node Authentication is basic, parameters such as Username, Password and Backup Encryption or Protection Password [If Encrypted] are loaded automatically to proceed with the restore.
  - If the Node Authentication is modern, parameters such as Username and Backup Encryption or Protection Password [If Encrypted] are loaded automatically to proceed with the restore.



- **New Authentication**

  During restore, you can select one of the following to change the authentication type:

  - Modern Authentication
  - Basic Authentication



3. To apply authentication, do one of the following:

   Modern Authentication

   For modern authentication, you can run a restore using security certificate.

   Use security certificate

   Note: By default, the backup plan selects **Use security certificate** for modern authentication.

   Follow these steps:

   a. To authenticate using a self-signed certificate, see How to Use Self-Signed Certificate.

   b. Enter the Application ID and Username.

   c. To apply the required roles and permissions, select the **Allow Arcserve UDP to set the required Azure roles and permissions for this application** check box, and then click **Next**.

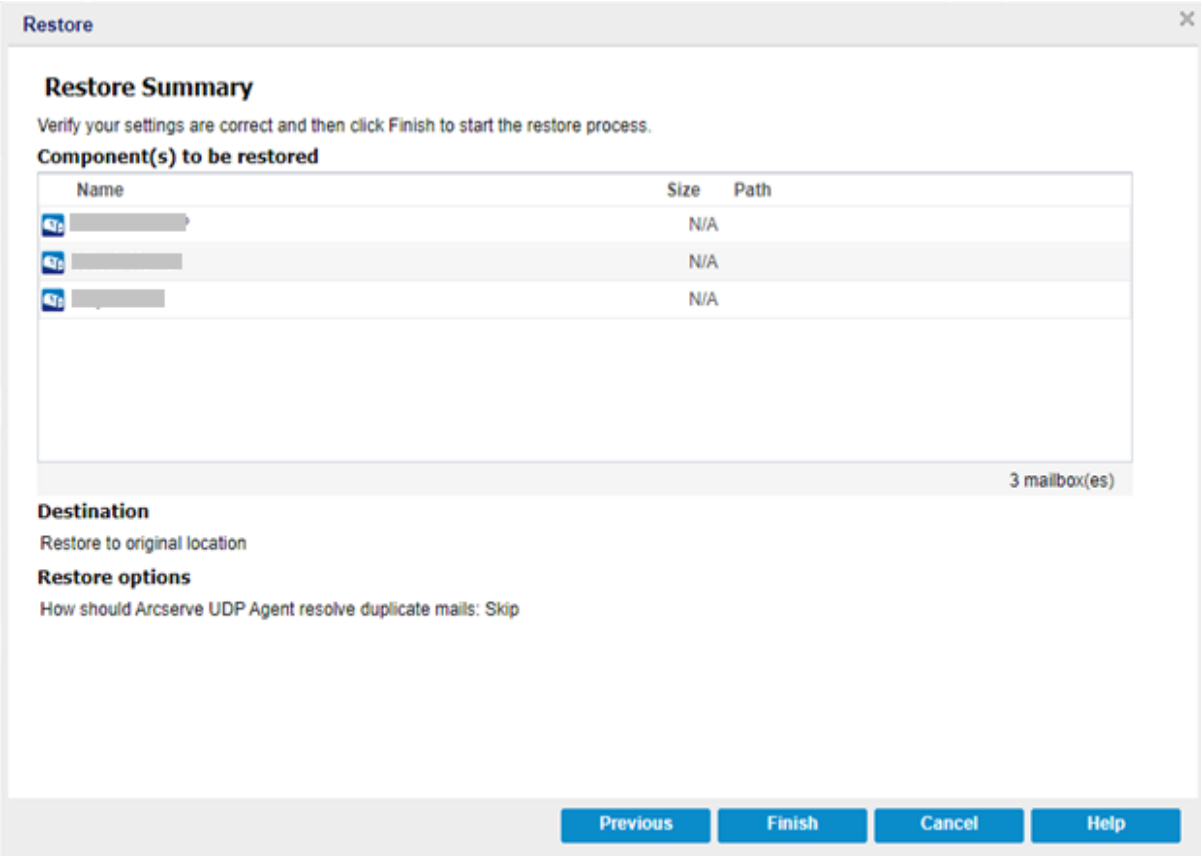
   Basic Authentication

   For Basic Authentication, to proceed with the restore, enter User Name and Password.

Note: This option fails if you have Modern Authentication enabled on the Tenant or Account.

4. Click **Next**.

   The **Restore Summary** dialog opens.



5. On the Restore Summary window, verify the settings, and then click **Finish**.

The restore options are defined to restore the Exchange Online information.

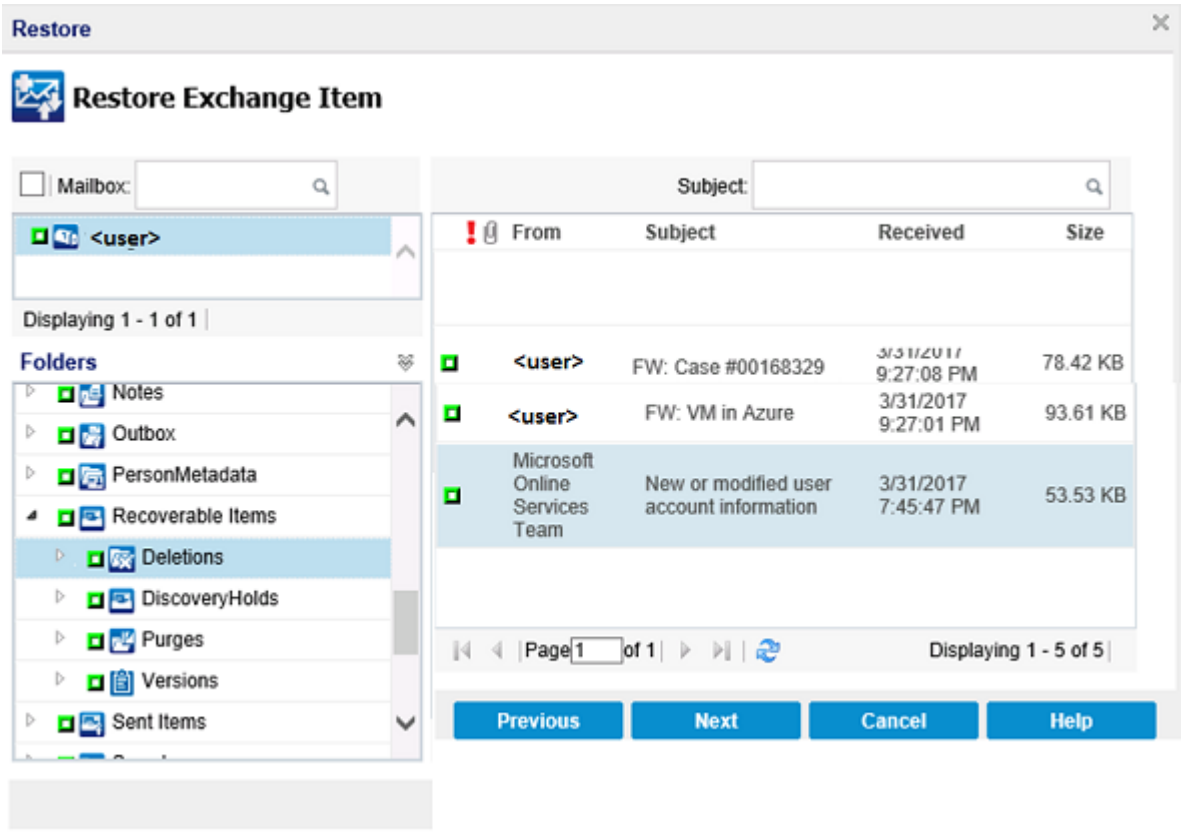
Restore the Recoverable Items

You can restore the recoverable items from the mailbox that enables the In-Place Hold or Litigation Hold feature , from Exchange online node UI to browse backup session. To enable this feature, see Specify the Source.

After you enable the feature, using the Restore wizard you can restore the recoverable items.

Follow these steps:

1. On the **Restore Exchange Item** dialog, select required folders under **Recoverable Items** and click **Next**.



The Restore Option screen is displayed.

2. From the Restore Option screen, perform the following details and click **Next**:
   - Select Restore Destination.
   - Skip or overwrite if the item already exists in the destination.

     You can perform either Original or Alternate restore.

     Original restore

The items under Recoverable Items are restored at the *UDP_RecoverableItems_yyyyMMdd_HHmmssfff* folder that is created on the target mailbox. Skip and Overwrite options for restore are not applicable to these items.

Alternate restore

As with user mailbox, restores the selected items to the destination path in a unique time stamped folder, such as *titled /restore_yyyyMMdd-mmssff*.

- Provide user id and password of the destination where you want to restore.

The selected recoverable items are restored.

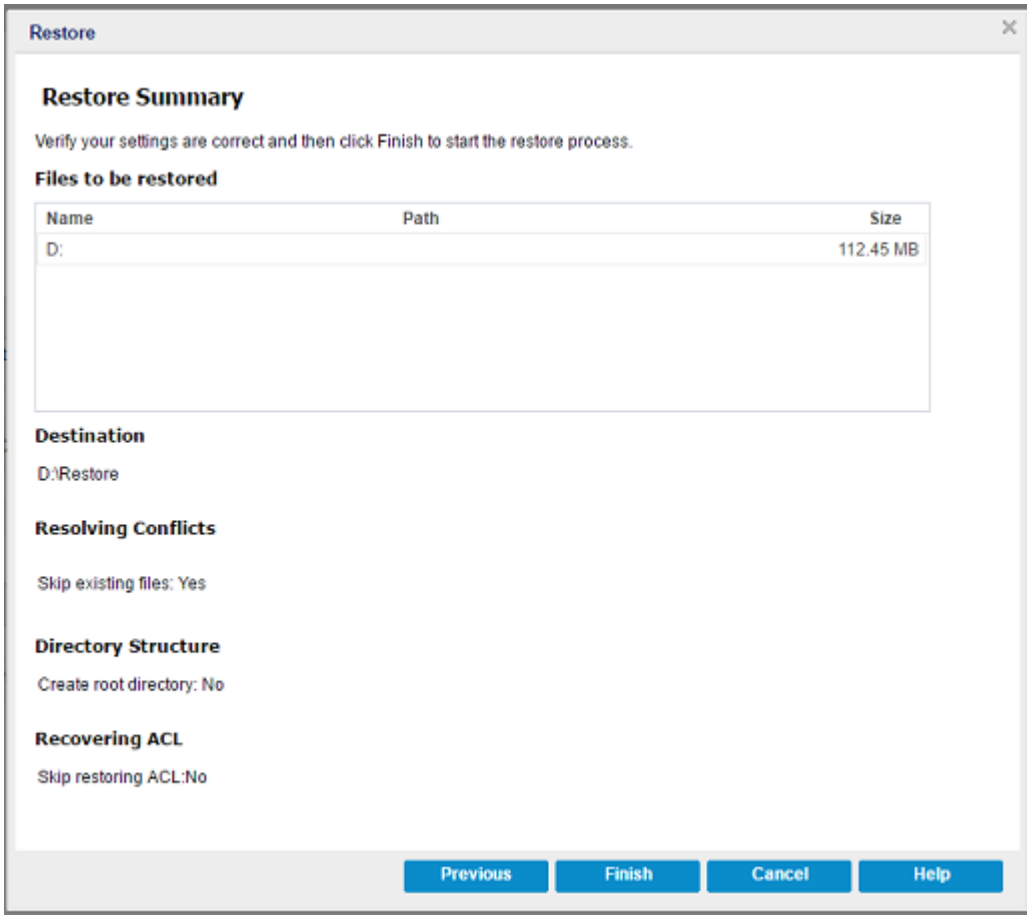
The recovery point content is stored.

Restore the Recovery Point Content

After you define the restore options, verify that your settings are correct and confirm the restore process. **Restore Summary** helps you to review all the restore options that you defined and modify them if necessary.

Follow these steps:

1. On the **Restore Summary** dialog, review the displayed information to verify that all the restore options and settings are correct.



- If the summary information is incorrect, click **Previous** and go back to the applicable dialog to change the incorrect setting.
- If the summary information is correct, click **Finish** to launch the restore job.

The recovery point content is stored.

Verify that Content is Restored

After the completion of the restore process, verify that content was restored to the specified destination.

Follow these steps:

1. Log into the destination mailbox.
2. Check the mailbox item that you restored.
3. Verify the restored content.

   The restored content is successfully verified.

How to Restore SharePoint Online Site Collection Data

You can restore SharePoint Online List/Library or List item in Site or SharePoint Parent Site. You can restore the data to original site with new name, restore the data to original location and export to disk from the recovery points.

Perform the following tasks to restore SharePoint Online List item:

1. Select the SharePoint Online site list Items to Restore
2. Define the Restore Options
3. Verify that Content is Restored

Select the SharePoint Online Site List Items to Restore

You can restore SharePoint Online list items data from a recovery point. When you select a recovery date, and then specify the time, all the associated recovery points for that duration are displayed. You can then browse and select the backup content (including applications) to restore.

Follow these steps:

1. Log into Arcserve UDP.
2. Click the **resources** tab.
3. Select **All Nodes** in the left pane.

   All the added nodes are displayed in the center pane.

   Or

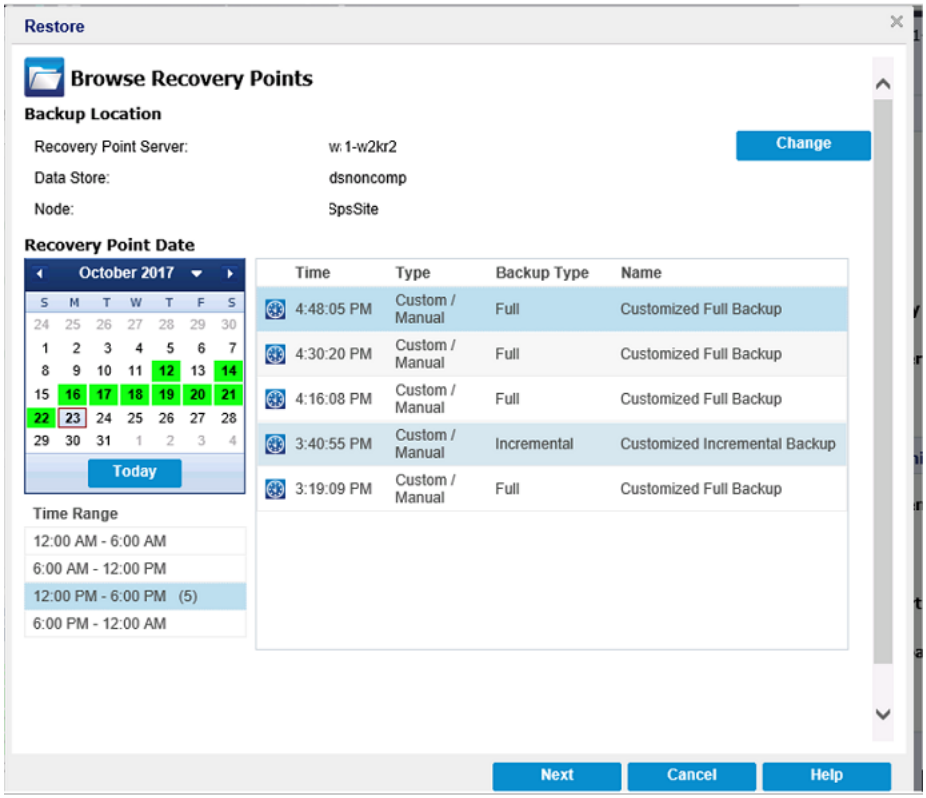
   Select **SharePoint Online Nodes** group.

   All the added SharePoint nodes are displayed in the center pane.
4. In the center pane, select the SharePoint Online node and click **Actions**.
5. Click **Restore** from the **Actions** drop-down menu.

   The **Restore SharePoint Item** dialog opens.

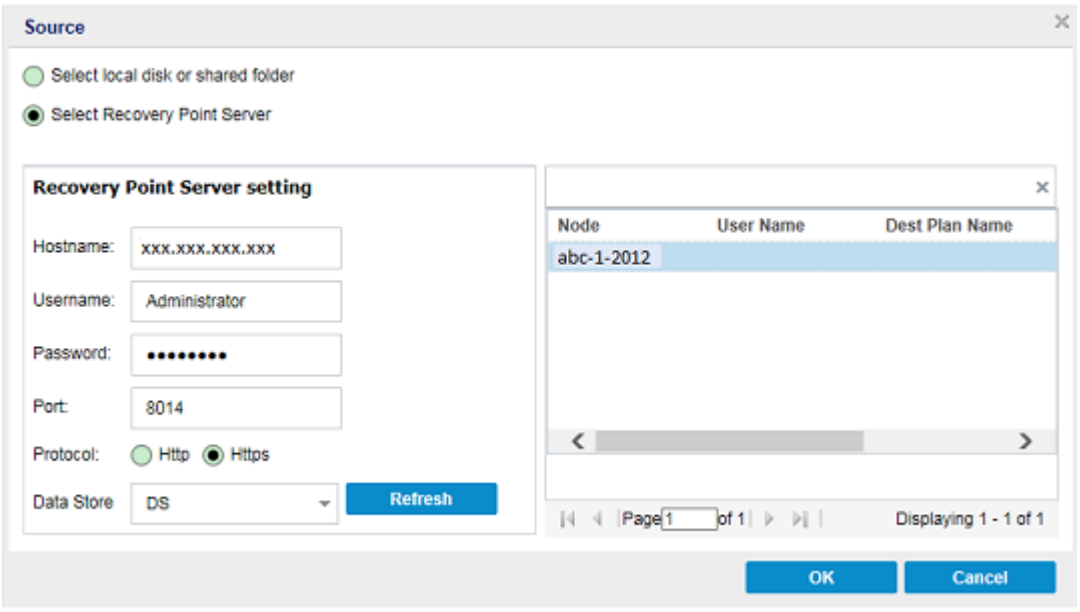
   Note: You are automatically logged into the agent node and the **Restore SharePoint Item** dialog opens.

   The **Backup Location** displays the **Recovery Point Server** details.



6. (Optional) Click **Change** to modify the backup location.

   The **Source** dialog opens. You can select the backup location in this dialog.



7. To specify Source, select one of the following options, and click **OK**:

   Select local disk or shared folder

Note: In Arcserve UDP, we do not recommend to select the **Select local disk or shared folder** option.

Select Recovery Point Server

a. Specify the Recovery Point Server setting details and click **Refresh**.

   All the agents are listed in the Data Protection Agent column in the **Source** dialog.

b. Select the agent from the displayed list/library and click **OK**.

   The recovery points are listed in the **Restore SharePoint Item** dialog.

8. Select the calendar date for the backup image to restore and click **Next**.

   All the dates containing recovery points for the specified backup source are high-lighted in green.

   The corresponding recovery points for that date are displayed, with the time of the backup, the type of backup that was performed (Full or Incremental), and the name of the backup.

9. From the **Restore SharePoint Online items** dialog, expand the site collection.

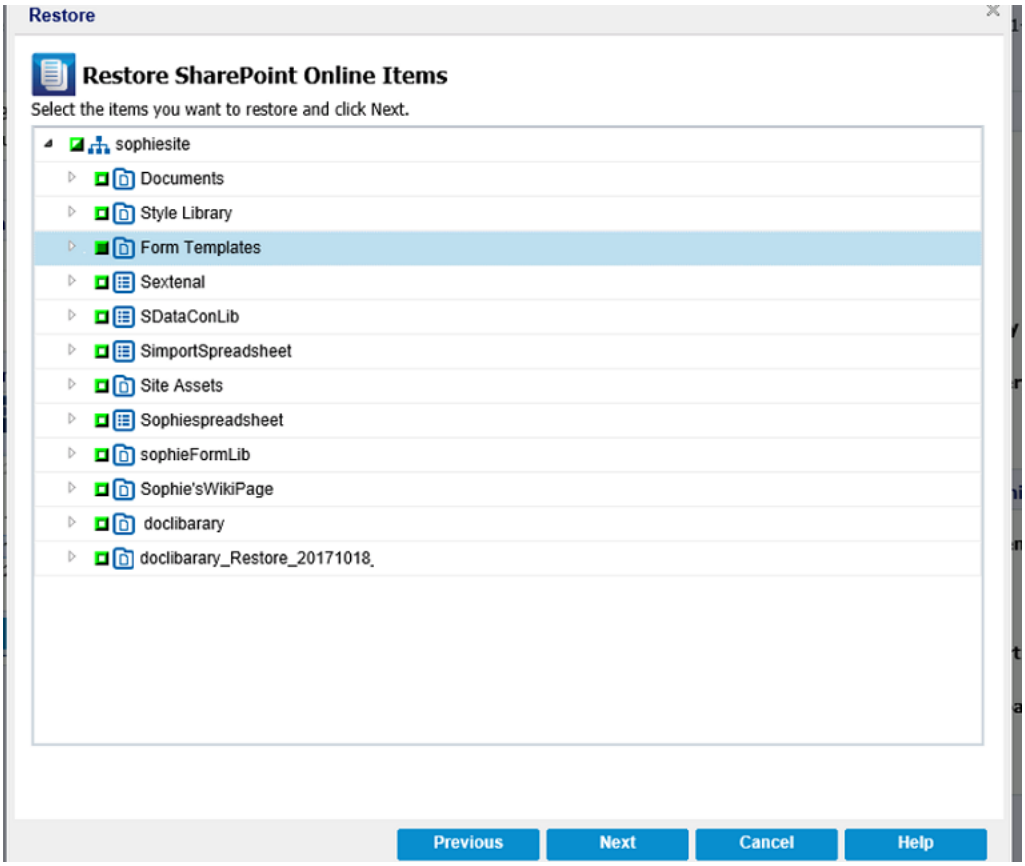
   All lists/Libraries and list are displayed.

10. Select the List/Library or item available in the Site collection that you want to restore from SharePoint site collection and click **Next**.

    Notes:

    - You can select the entire content or partial content of the SharePoint object to restore. To select partial content, expand the object, and click the check box of that content.
    - You can select multiple SharePoint lists/Libraries or list items to restore.

    The **Restore Options** dialog opens.



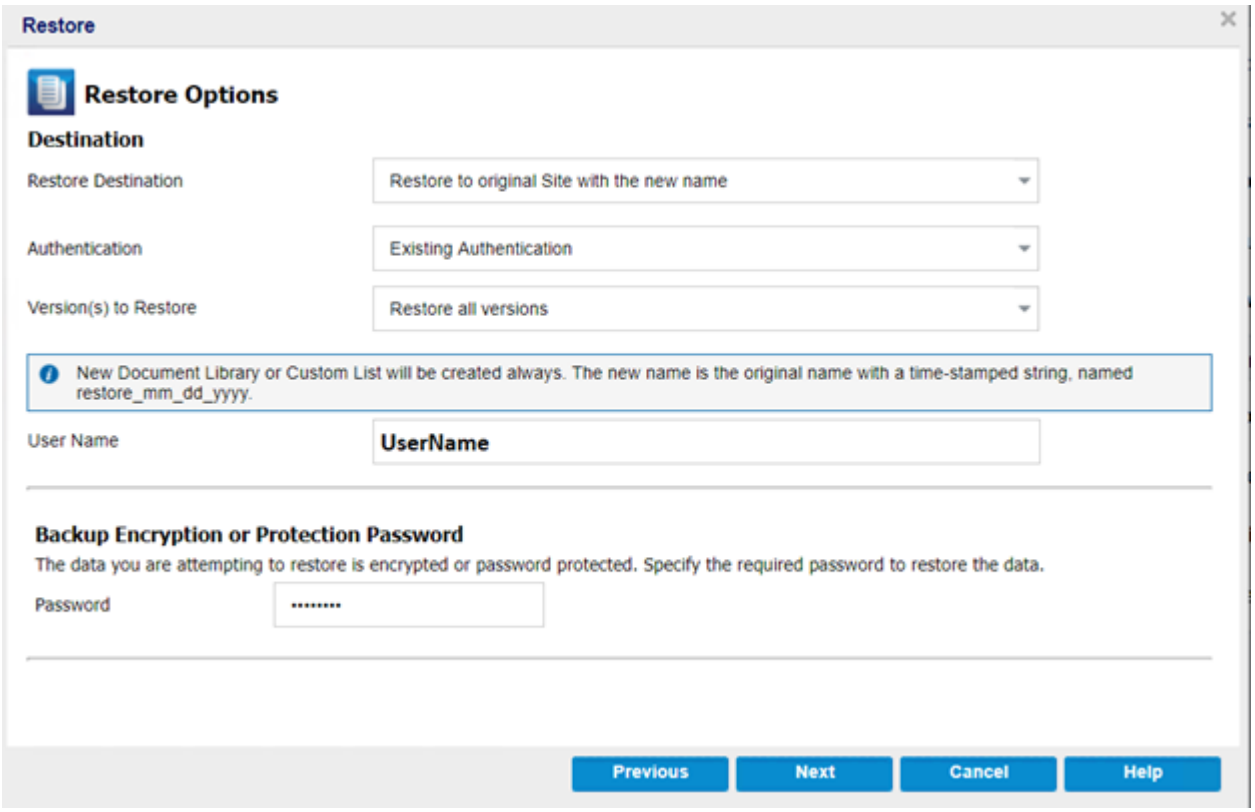
The SharePoint Online lists/Libraries or list items to restore are selected. Now, you can define the restore options.

Define the Restore Options

After you specify the SharePoint Online information to restore, define the restore options for the selected content.

Follow these steps:

1. On the **Restore Options** dialog, select the restore destination.



The available destination options are:

Restore to original Site with the new name

Restores the list/Library and list items data to the same site with new list name from where you took a backup.

Restore to the original location

Restores the list/Library and list items data to the same location from where you took a backup.

2. Specify one of the following options from the **If Item already exists in the Destination** drop-down list:

Append as new version if versioning is enable

This option works when your version setting is enabled from library setting in SharePoint Site. Once you choose this option, new version is appended to the current versions of the list items if the list items are existing.

Skip the item and do not restore

Skips over the items and does not restore.

Overwrite the item in the destination

Overwrites the item in the destination.

Export to disk

Restores the lists/Libraries or list items in the site collection to folder or share folder in disk.

Note: For lists restore when selecting restore option of Export to disk, only export attachment(s) of list to disk.

Authentication

Validates and identifies the user account. Select one of the following options:

- **Existing Authentication**

  When you select Existing Authentication, the authentication type that was selected during Node creation is used.

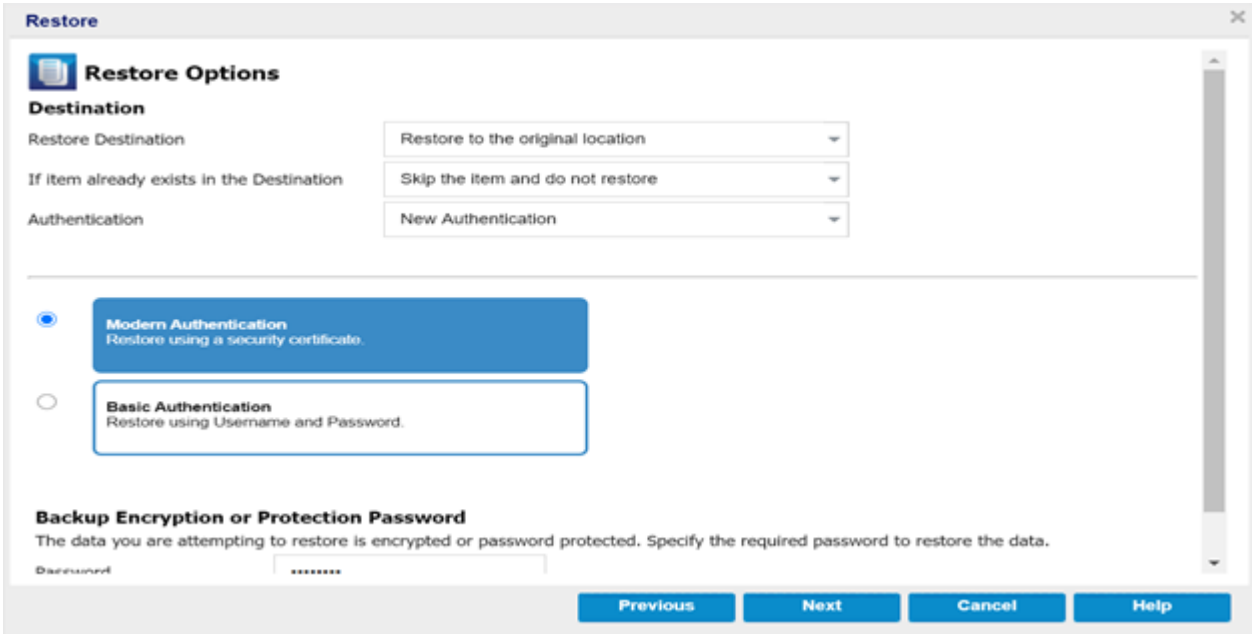
  - If the Node Authentication is basic, parameters such as Username, Password and Backup Encryption or Protection Password [If Encrypted] are loaded automatically to proceed with the restore.
  - If the Node Authentication is modern, parameters such as Username and Backup Encryption or Protection Password [If Encrypted] are loaded automatically to proceed with the restore.

- **New Authentication**

  During restore, you can select one of the following to change the authentication type:

  - Modern Authentication

- Basic Authentication



3. To apply authentication, do one of the following:

   Modern Authentication

   For modern authentication, you can run a restore using security certificate.

   Use security certificate

   Note: By default, the backup plan selects **Use security certificate** for modern authentication.

   Follow these steps:

   a. To authenticate using a self-signed certificate, see How to Use Self-Signed Certificate.

   b. Enter the Application ID and Username.

   c. To apply the required roles and permissions, select the **Allow Arcserve UDP to set the required Azure roles and permissions for this application** check box, and then click **Next**.

   Basic Authentication

   For Basic Authentication, to proceed with the restore, enter UserName and Password.

   Note: This option fails if you have Modern Authentication enabled on the Tenant or Account.

4. Specify the list item Versions that you want to restore when your version setting is enabled from library setting in SharePoint Site.

Restore all versions

Restores all versions on backup.

Restore only the latest version

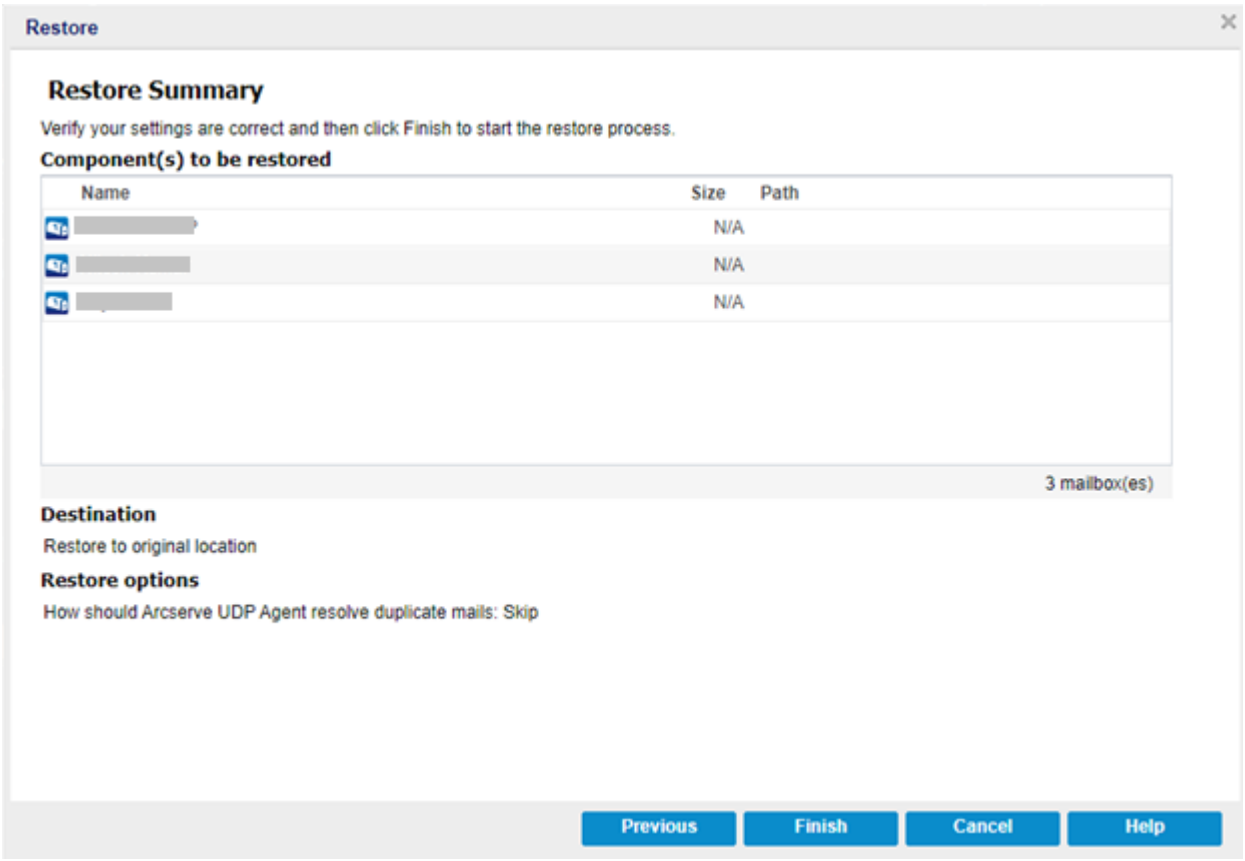
Restores only the latest version on backup.

Restore only the latest major version

Restores only the latest major version on backup.

5. Specify the Session password if have.
6. Click **Next**.

   The **Restore Summary** dialog opens.



   You can verify the restore information form Restore Summary dialog.

7. Click **Finish** to submit the restore job.

   The restore options are defined to restore the SharePoint Online information.

Verify that Content is Restored

After the completion of the restore process, verify that content was restored to the specified destination.

Follow these steps:

1. Log into the SharePoint site collection.
2. Verify the list/Library and list items.
3. Verify the restored content.

   The restored content is successfully verified.

How to Restore OneDrive Data

You can restore OneDrive data (files, folders, and so on) using any computer. You can restore using the restore option or also by using Mount Volume option from agent user interface.

Perform the following tasks to restore OneDrive data using the restore option:

1. Select the OneDrive Items to Restore
2. Define the Restore Options
3. View Restore Summary

Or else

Restore OneDrive data using the Mount Volume option.

Select OneDrive Items to Restore

You can restore OneDrive data from a recovery point. When you select a recovery date, and then specify the time, all the associated recovery points for that duration are displayed. You can then browse and select the backup content (including applications) to restore.

Follow these steps:

1. Log into Arcserve UDP.
2. Click the **resources** tab.
3. Select **All Nodes** in the left pane.

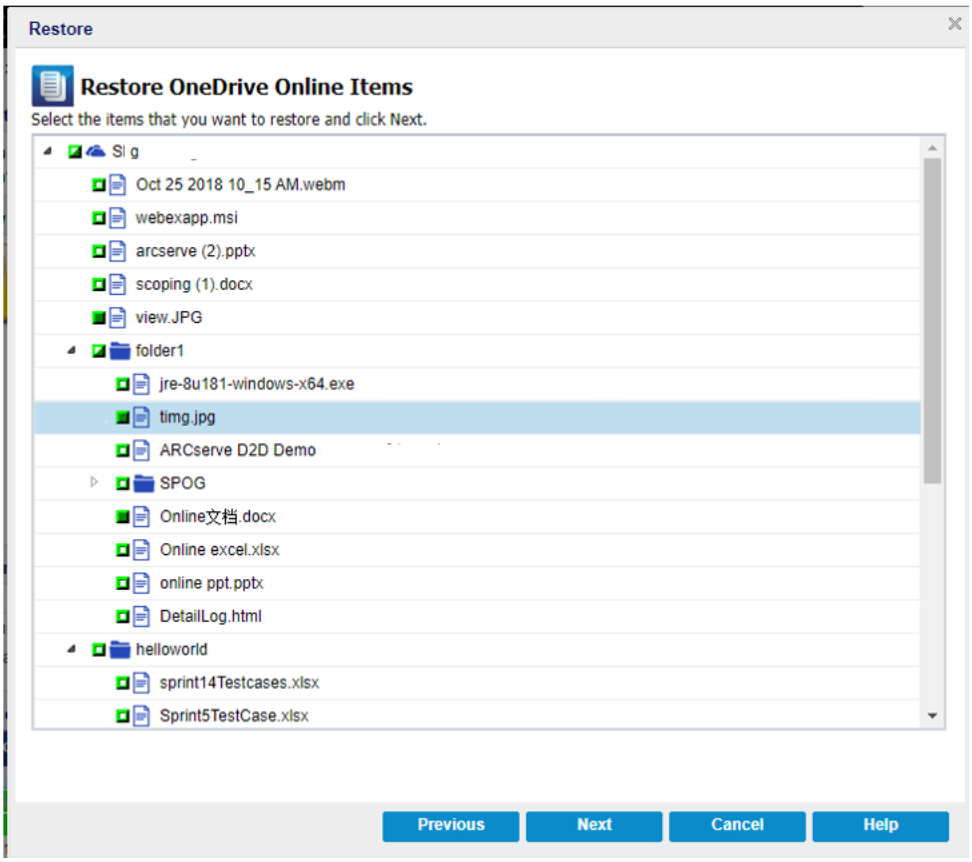
   All the added nodes are displayed in the center pane.
4. In the center pane, select the OneDrive node and click **Actions**.
5. Click **Restore** from the **Actions** drop-down menu.

   Note: You are automatically logged in to the agent node and the **Node** dialog opens.

   You can view the **Browse Recovery Points** details in the Backup Location. The name of selected *Recovery Point Server* is displayed. If desired, click **Change** and modify **Recovery Point Server setting** from the **Source** pop-up.
6. Select the calendar date for the backup image to restore and click **Next**.

   All the dates containing recovery points for the specified backup source are highlighted in green.

   The corresponding recovery points for that date are displayed, with the time of the backup, the type of backup that was performed (Full or Incremental), and the name of the backup.
7. From the **Restore OneDrive Node Items** pane, select the check boxes of related items or folders that you want to restore from Folders, and click **Next**.



Notes:

- You can select the entire content or partial content of the OneDrive object to restore. To select partial content, expand the object, and click the check box of that content.
- You can select multiple OneDrive objects to restore.

The **Restore Options** dialog opens.

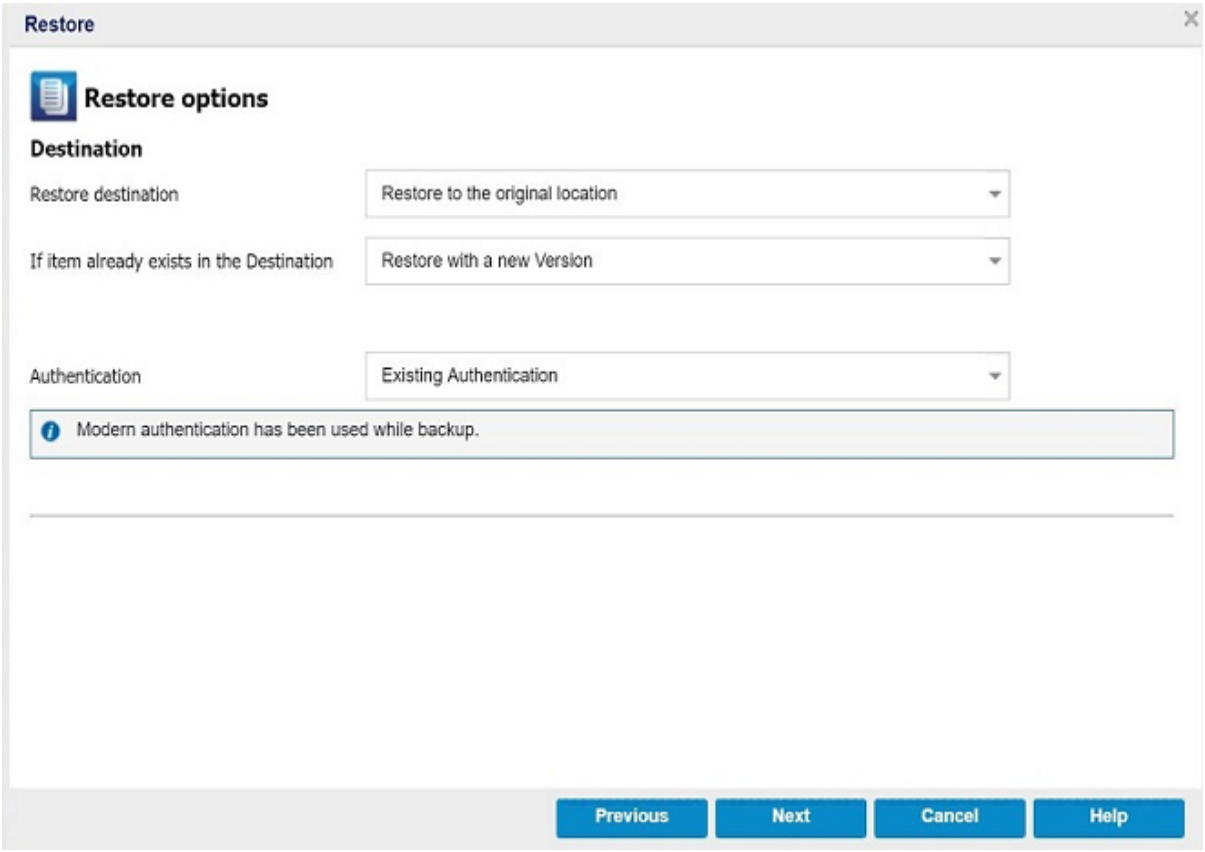
The OneDrive Node items to restore are selected. Now, you can define the restore options.

Define the Restore Options

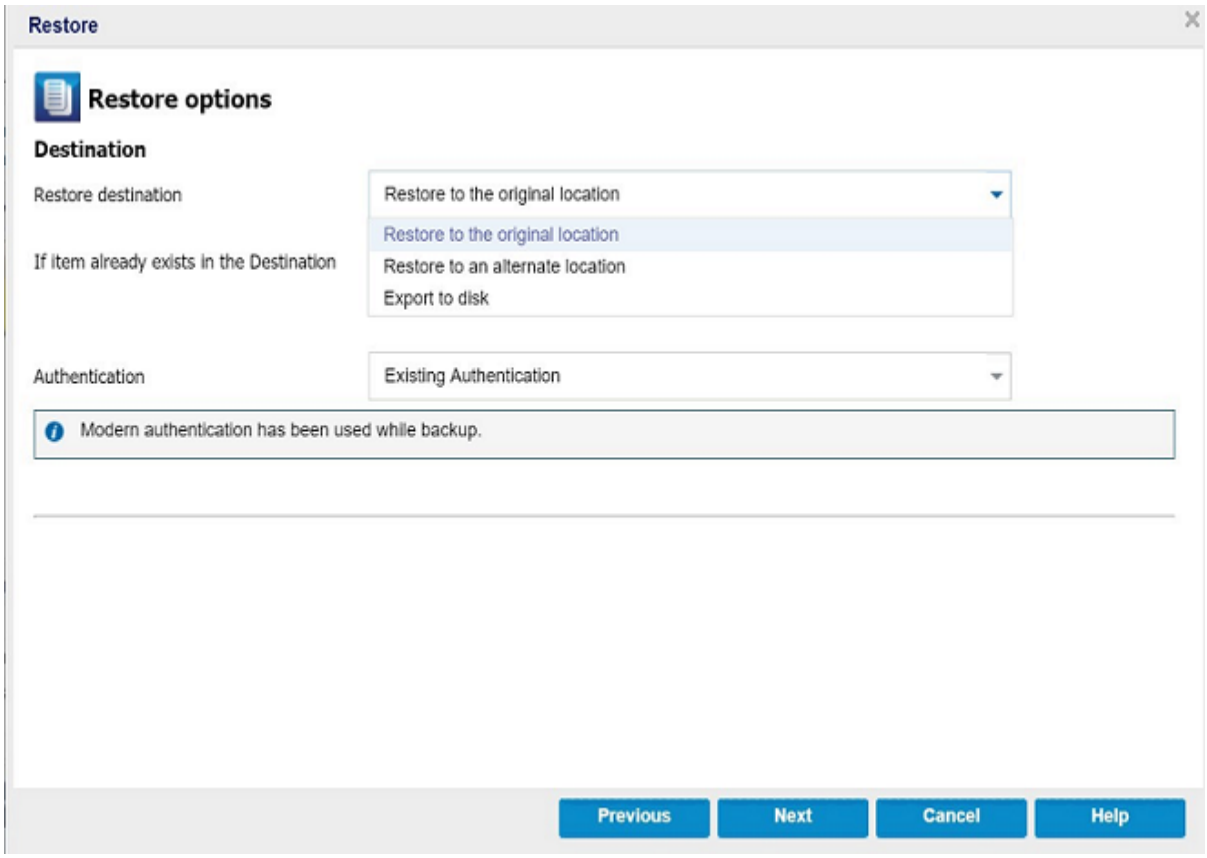
After you specify the OneDrive information to restore, define the restore options for the selected content.

Follow these steps:

1. On the **Restore Options** dialog, select the restore destination.



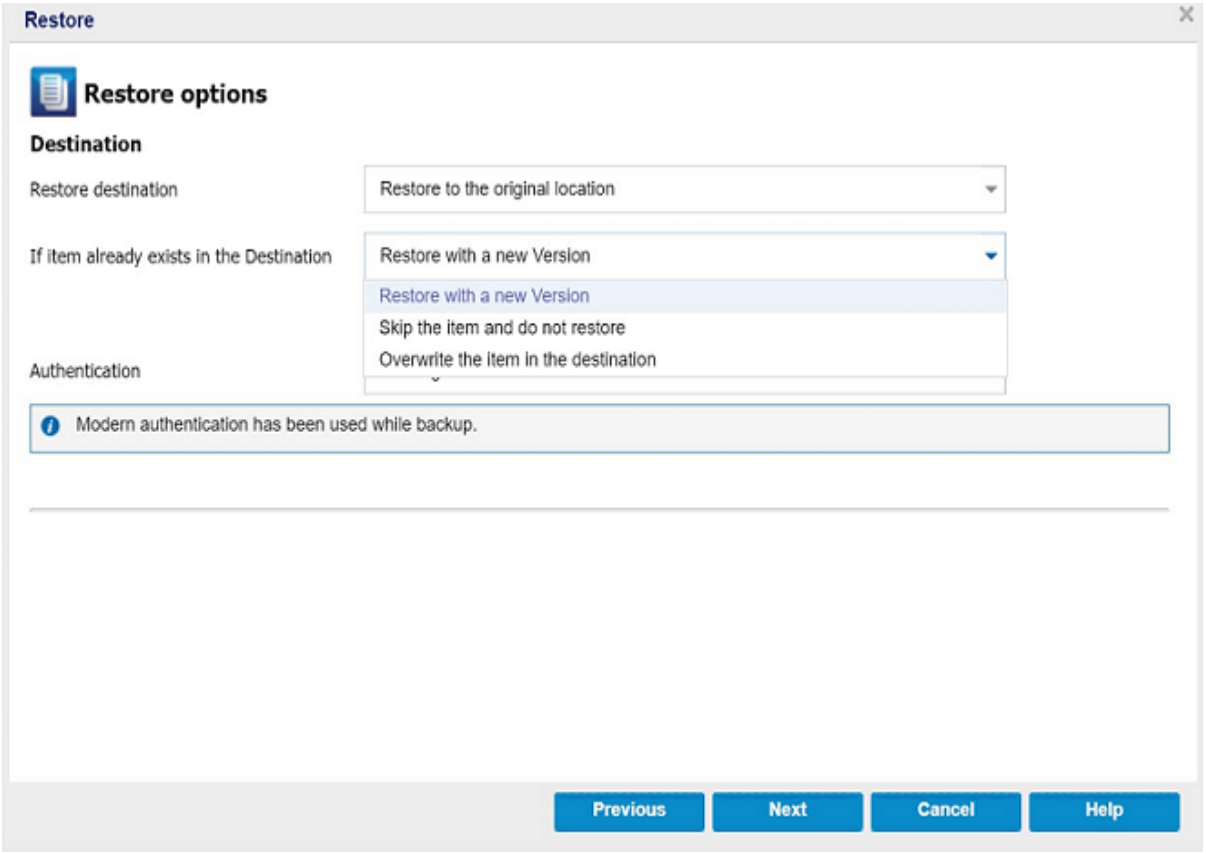
The available destination options are:



Restore to the original location

Restores the files or folders to the same location from where you took a backup.

If you select the **Restore to the Original location** option, from the **If item already exists in the Destination** drop-down, select one of the following options:



Restore with a new version

Appends new version to the existing file(s).

This option is selected by default.

Skip the item and do not restore

Skips over the items and does not restore.

Overwrite the item in the destination

Overwrites the item in the destination.

Restore to an alternate location

Restores the files or folders or accounts to an account in the alternate location. When you select this option, you can browse and select the destination.

Hierarchy of the source

Retains the hierarchy tree as specified in the source while creating backup plans.

Note: This option displays only when Restore to an alternate location is selected.

Export to disk

Restores to a folder or share folder in disk.

2. To validate and identify the user account, from the **Authentication** drop-down, select one of the following:

   Existing Authentication

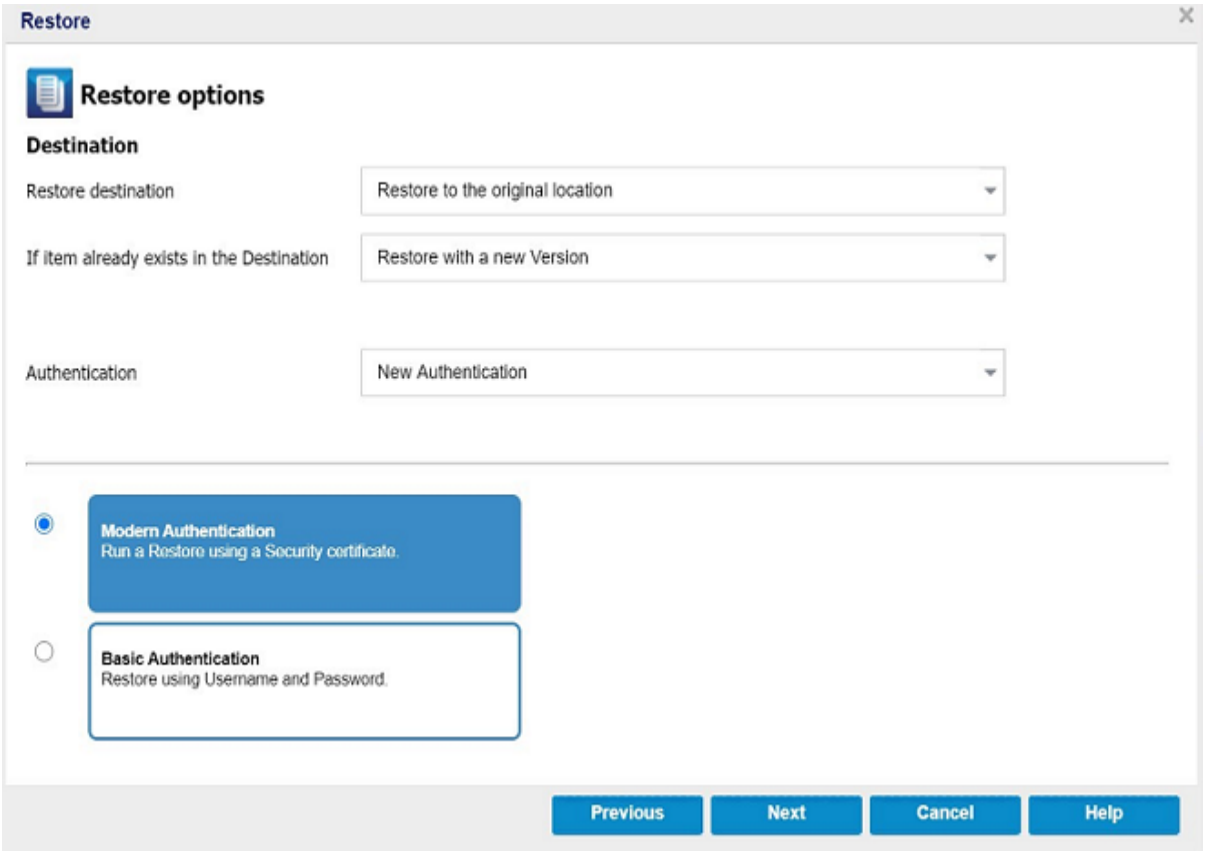
   When you select Existing Authentication, the authentication type that was selected during Node creation is used.

   - If the Node Authentication is basic, parameters such as Username, Password and Backup Encryption or Protection Password [If Encrypted] are loaded automatically to proceed with the restore.
   - If the Node Authentication is modern, parameters such as Username and Backup Encryption or Protection Password [If Encrypted] are loaded automatically to proceed with the restore.

   New Authentication

   During restore, you can select one of the following to change the authentication type:

   - Modern Authentication
   - Basic Authentication



3. To apply the new authentication, do one of the following:

Modern Authentication

For modern authentication, you can run a restore using security certificate.

Follow these steps:

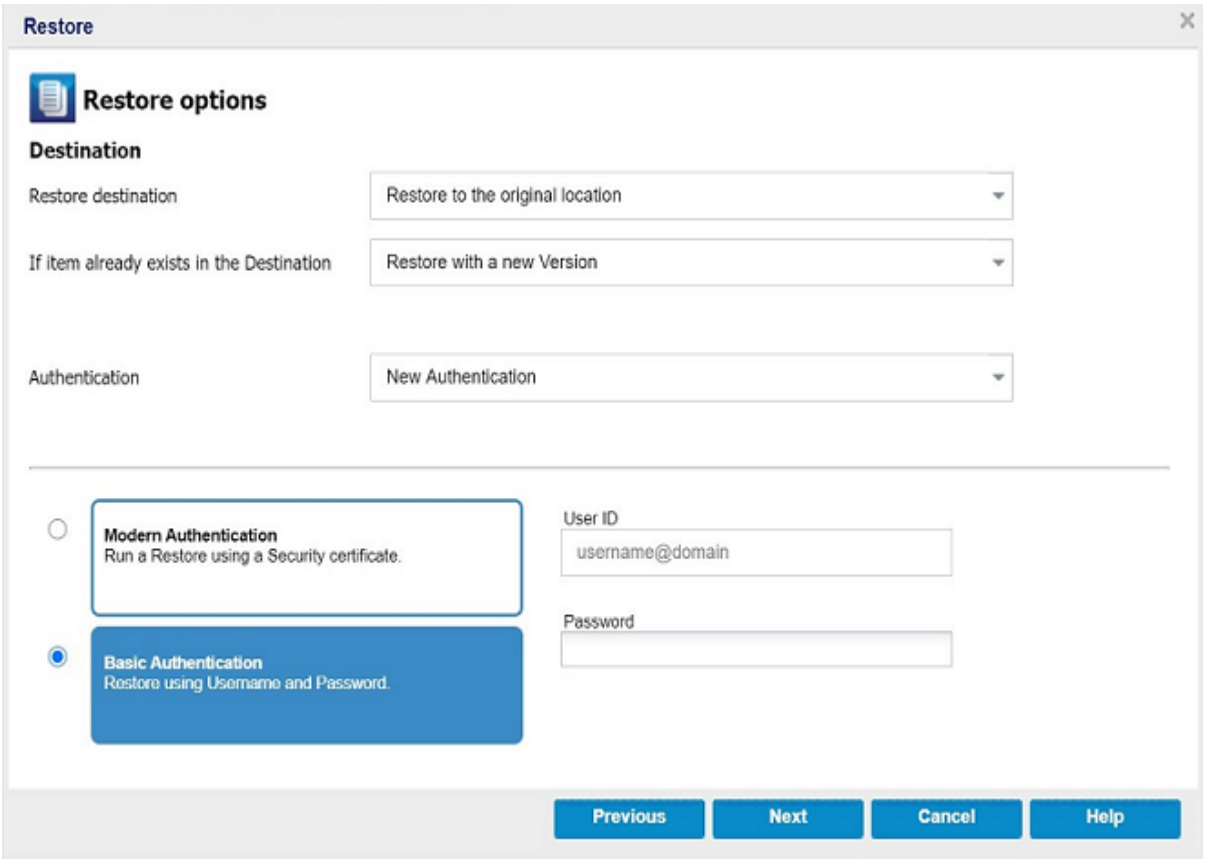
a. To authenticate using the self-signed certificate, see How to Use Self-Signed Certificate.

b. Enter the Application ID and Username.

c. To apply the required roles and permissions, select the **Allow Arcserve UDP to set the required Azure roles and permissions for this application** check box, and then click **Next**.

Basic Authentication

For Basic Authentication, to proceed with the restore, enter User ID and Password.



4. Click **Next**.

   The **Restore Summary** dialog opens.

5. On the Restore Summary dialog, verify your settings, and then click **Finish**.

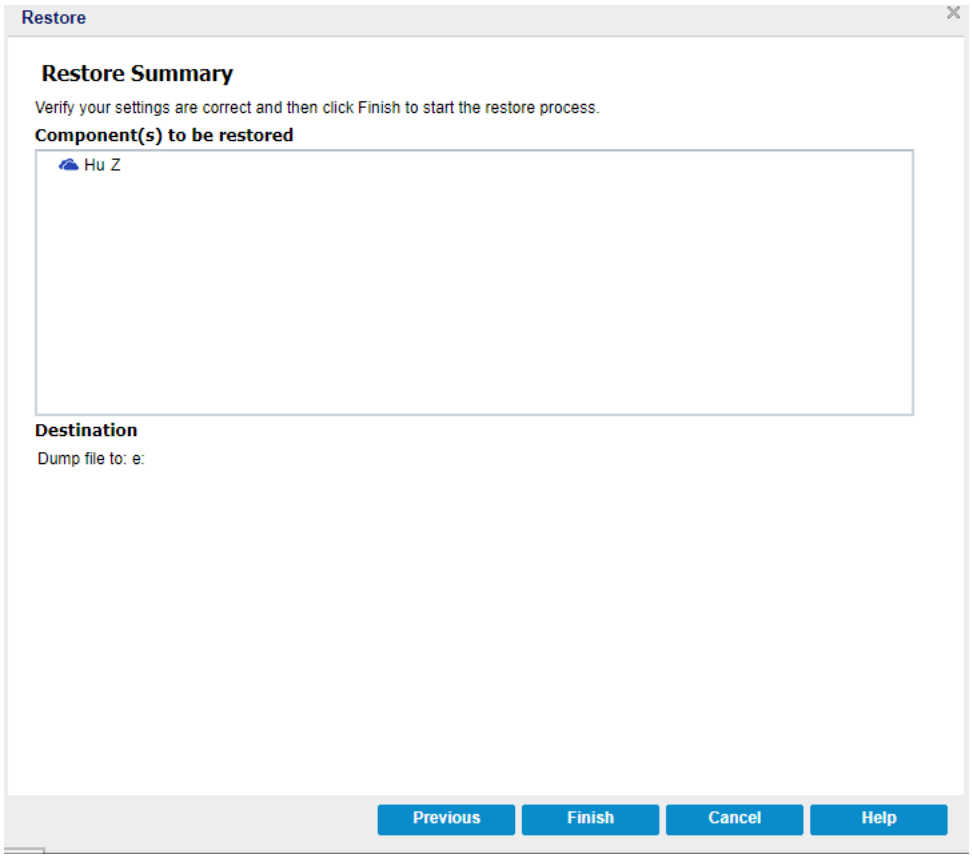
   The restore options are defined to restore the OneDrive information. Now, you can view restore summary.

View Restore Summary

After you define the restore options, verify that your settings are correct and confirm the restore process. **Restore Summary** helps you to review all the restore options and destination path that you defined. If you need to modify, click **Previous**.

Follow these steps:

1. On the **Restore Summary** dialog, review the displayed information to verify that all the restore options and settings are correct.



- If the summary information is incorrect, click **Previous** and go back to the applicable dialog to modify the settings.
- If the summary information is correct, click **Finish** to launch the restore job.

The recovery point content is restored.

Restore OneDrive Data Using Mount Volume Option

OneDrive Protection uses the universal backup session format. You can mount the backup session as a drive letter and then copy the file/folder from the mounted volume.

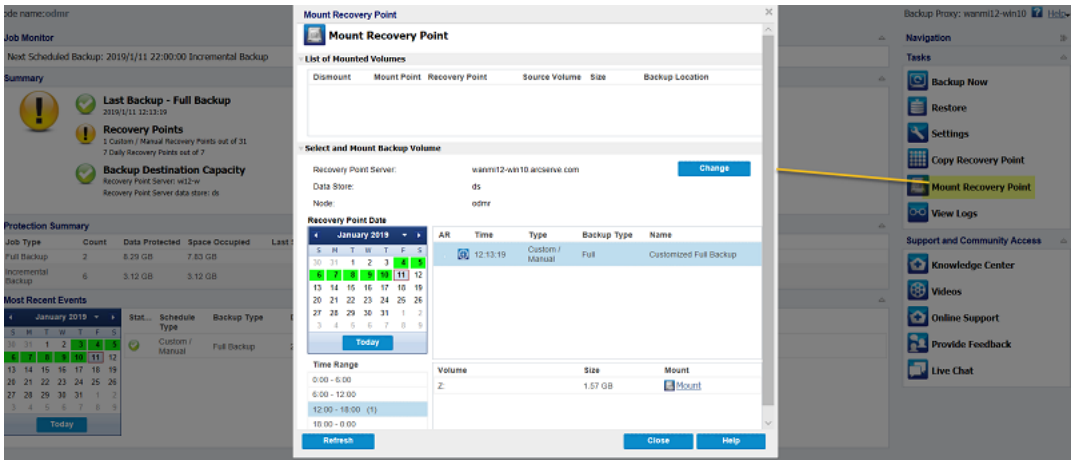
Follow these steps:

1. Log into Arcserve UDP.
2. Click the **resources** tab.
3. Select **All Nodes** in the left pane.

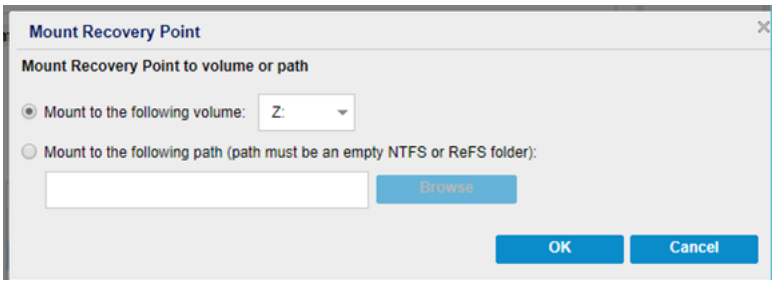
   All the added nodes are displayed in the center pane.
4. In the center pane, select the OneDrive node and click **Actions**.
5. Click **Log into Agent** from the **Actions** drop-down menu.

   Note: You are automatically logged in to the agent node. You can view complete details about the node, also on the right pane, view the list of tasks that you can perform.
6. From the right pane under **Tasks**, click **Mount Recovery Point**.

   The **Mount Recovery Point** dialog opens. You can select the backup session in this dialog.



7. Select the destination and mount the session as a Drive Letter or Mount into an empty NTFS folder.



You can browse the volume. In the Volume, Arcserve UDP saved all the meta data of OneDrive. Each Account will have a related folder in the root volume. The folder is named according to the account name.

You can open the folder created by that account name and verify if the OneDrive data is backed up.

How to Perform Point-In-Time Restore

Point-In-Time Restore supports restoring SQL Database to any specific time period between N and N+1 recovery points. Point-In-Time helps the administrators to restore the transactions happened in SQL Database between two recovery points. For example, consider that you have a recovery point at 03/16/2019 12:14:04:177 and subsequent recovery point at 03/29/2019 22:03:14:177. Using Point-In-Time, you can restore the transactions happened between the two recovery points. This helps the administrators to restore only the required transactions from a large size of backed-up data.

What To Do Next?

1. Review Prerequisites
2. Review Considerations
3. Perform Point-In-Time Restore

Prerequisites

Review the following prerequisites before performing the Point-In-Time restore:

- MS SQL Database Server must have installed on the Windows Agent machine.
- You must have at least two recovery points to perform Point-In-Time restore.
- You must manually enable Point-In-Time every time you perform restore.

Considerations

Review the following considerations before performing the Point-In-Time restore:

- You can perform Point-In-Time restore using PIT.EXE command line utility only.
- You must have two recovery points working together to perform Point-In-Time restore.
- Point-In-Time restore supports only the Agent Based backup. Host Based backup is not supported.
- Point-In-Time supports restore to original location only. Restore to alternate location and dump to disk are not supported.
- The backup of SQL Database Transaction Log is saved to the same directory of SQL Log File before taking the snapshot.
- The first recovery point provides the SQL Database level restore. After Point-In-Time restore, SQL Database is set to recovering state. You can back-up the recovery point in Point-In-Time disabled mode.
- To perform Point-In-Time restore, you must select the time point between N and N+1 recovery points. Any time point prior to N recovery point and after N+1 recovery point is not allowed. If you set the time point that is not between N and N+1 recovery point, the restore job fails and causes the database corruption.
- The Truncate Log option is disabled by default in Point-In-Time backup and restore plans. If enabled, after the backup transaction logs are truncated and the next Point-In-Time backup misses some transactions.

Perform Point-In-Time Restore

1. Run the file **PIT.EXE**.

   Point-In-Time command line utility opens and displays PIT$ command prompt.

   For more information, view Understanding Point-In-Time Command Line Utility.

2. From the command prompt, run **set pitbackup=1**.

   Point-In-Time backup is enabled.

3. From **Arcserve UDP Console**, perform a backup of SQL database.

   The N recovery point is created.

4. Verify that the following files are generated under **Catalog** folder:
   - If the SQL Server Instance name is MSSQLSERVER (default Instance Name):
     - ***_<DatabaseName>.idx.pit***

       Contains fixed size of time point summary.
     - ***_<DatabaseName>.cat.pit***

       Contains variable size of time point details.
     - ***_<DatabaseName>.map.pit***

       Contains mapping between the internal string and internal string identity.
   - If the SQL Server Instance name is other than MSSQLSERVER (not the default Instance Name):
     - ***<SQLInstanceName>_<DatabaseName>.idx.pit***

       Contains fixed size of time point summary.
     - ***<SQLInstanceName>_<DatabaseName>.cat.pit***

       Contains variable size of time point details.
     - ***<SQLInstanceName>_<DatabaseName>.map.pit***

       Contains mapping between the internal string and internal string identity.
5. From **Arcserve UDPConsole**, perform another backup of SQL database.

   The N+1 recovery point is created.
6. From the command prompt, run ***set pitrestore=1***.

   Point-In-Time restore is enabled.
7. From command prompt, run the below SQL queries to view the time point information of N+1 backup session:
   - *Query <SQLInstanceName>\<DatabaseName> <BackupDestination>\Catalog\<N+1 backup session folder name>*

     Displays all time points summary in backup
8. From command prompt, run ***set pittime="<MM/DD/YYYY HH:MM:SS:ss>"***
9. From **Arcserve UDP Console**, perform database level restore to original location using the N recovery point.

   The Point-In-Time restore is successfully completed. Verify that the expected data is restored to database.

Understanding Point-In-Time Command Line Utility

Point-In-Time command line utility is available under the path *<Arcserve UDP Installed path>\BIN* with the file name *PIT.exe*.

The command prompt in the command line utility appears as **PIT$**.

Following are the available options that you can run in command line utility:

HELP

Displays all the options available in PIT.exe.

COMMAND /?

Displays the usage of command.

CONFIG

Uses setting in registry.

OPTION

Specifies Global Option. VALUE is always a hexadecimal number.

SEQ

Specifies the sequence number in DB/TLog backup.

TIMEOUT

Specifies the connection timeout, in seconds.

DEVICE

Specifies the Backup device type that is always a disk.

DATABASE

Specifies the database name. Enter the value in *[server]\[instance]\<Database>* format.

LOGBACKUP

Specifies the path of TLog backup.

LOGRESTORE

Specifies the path of TLog restore.

DBBACKUP

Specifies the path of TLog backup.

DBRESTORE

Specifies the path of TLog restore.

CATALOG

Specifies the path of catalog.

STOPAT

Specifies the time point to restore.

NORECOVERY

Specifies disable(1) or enable(0) the auto recovery after UDP DB restore.

PITBACKUP

Specifies enable(1) or disable(0) point in time backup for UDP.

Example: set pitbackup=1

PITRESTORE

Specifies enable(1) or disable(0) point in time restore for UDP.

Example: *set pitrestore=1*

PITTEMP

Temp folder for point-in-time restore

PITTIME

Specifies time point for point-in-time to restore.

Example: set pittime="<MM/DD/YYYY HH:MM:SS:ss>

PITNOSTAGE

Specifies the use of log in the mounted volume instead of restore to temp folder.

BACKUPOPT

Specifies the backup option to control VSS.

RESTOREOPT

Specifies the restore option to control VSS.

Following are the SQL queries that you can run in command line utility:

- *Query <SQLInstanceName>\<DatabaseName> <BackupDestination>\Catalog\<N+1 backup session folder name>*

  Displays all time points summary in backup.

- *Query /d <SQLInstanceName>\<DatabaseName> <BackupDestination>\Catalog\<N+1 backup session folder name>*

  Displays all time point details in backup.

- *Query /i N <SQLInstanceName>\<DatabaseName> <BackupDestination>\Catalog\<N+1 backup session folder name>*

  Displays time point summary of N in backup.

- *Query /d /i N <SQLInstanceName>\<DatabaseName> <BackupDestination>\Catalog\<N+1 backup session folder name>*

  Displays time point details of N in backup.

How to Restore Cluster Shared Volume

Each time Arcserve UDP performs a successful backup, a point-in-time snapshot image of your backup is created (recovery point). This collection of recovery points allows you to locate and specify exactly which backup image you want to restore. If at some later time, you suspect any of the backed up information is missing, corrupted, or not reliable, you can locate and restore from a previous known good version.

What To Do Next?

1. Review Prerequisites and Considerations
2. Specify CSV Information to Restore
   a. Specify CSV and Content to Restore
   b. Define the Restore Options
3. Restore CSV Content

Review Prerequisites and Considerations

Verify that the following prerequisites exist before performing a restore:

- You have at least one recovery point available to restore.
- You have a valid and accessible recovery point destination to restore the recovery point content from.
- You have a valid and accessible target location to restore the recovery point content to.
- Review the Compatibility Matrix that provides the supported operating systems, databases, and browsers.

Review the following restore considerations:

- If the restore is to a remote destination and if all the drive letters (A - Z) are occupied, the restore to a remote path will not succeed. Arcserve UDP Agent (Windows) needs to use one drive letter to mount the remote destination path.
- (Optional) Understand how the restore process works. For more information, see How File Level Restores Work.

- (Optional) Review the files skipped during restore. For more information, see Files Skipped During Restore.
- When you attempt to restore an optimized backup session to a non-empty volume (unoptimized restore), the restore job may take more time than the estimated time displayed in the job monitor. The amount of data that is processed and the elapsed time may increase based on the data that is optimized on the volume.

Example:

The backup volume size is 100 GB and after optimization the volume size is reduced to 50 GB.

When you perform an unoptimized restore of this volume the restore job monitor displays 100% after restoring 50 GB, but it will take more time to restore the entire 100 GB.

- The following Activity log message will be displayed when restoring the system files:

  "System files were skipped. If necessary, you can use the Bare Metal Recovery (BMR) option to restore them."

How File Level Restores Work

During a block-level backup, each backed up file is made up of a collection of blocks that define that particular file. A catalog file is created containing a list of the backed up files, along with the individual blocks that were used for each file and the available recovery points for these files. When you need to restore a particular file, you can search your backup and select the file you want to restore and the recovery point you want to restore from. Then, Arcserve UDP collects the version of the blocks that were used for the recovery point of the specified file, and reassembles and restores the file.

Note: You can also perform a restore without a catalog file from a catalog-less backup recovery point.

The following flow diagram shows the process of how Arcserve UDP restores a specific file:

Files Skipped During Restore

While performing a restore by Arcserve UDP Agent (Windows) some files may be skipped intentionally.

The files and folders in the following table are skipped during a restore if the following two conditions exist:

- Files are skipped when such files exist before the restore and the conflict option is "skip existing files".
- Files and folders listed in the following table are skipped because they are an important component for Windows or Arcserve UDP Agent (Windows).

OS	Folder or Location	File or Folder Name	Remarks

All	Root folder of each volume	CAVolTrc.dat	Used by the Arcserve UDP tracking Driver.
		cavoltrcsnapshot.dat	
		System Volume Information*	Used to save files/folders by a Windows system, for example, volume shadow copy files.
		RECYCLER*	Used only on NTFS partitions. It contains a Recycle Bin for each user that logs on to the computer, sorted by their security identifier (SID).
		$Recycle.Bin*	When you delete a file in Windows NT Explorer or My Computer, the file is stored in the Recycle Bin until you empty the Recycle Bin or restore the file.
	Any folder contain pic-ture files	Thumbs.db	Stores thumbnail images for Windows Explorer thumbnail view.
	Root folder of volume	PageFile.Sys	Windows virtual memory swap file.
		Hiberfil.sys	Hibernate file, used to save the system data when a computer goes into hibernate mode.

The following files and folders are skipped only when you restore to the original or alternate location:

OS	Folder or Location	File or Folder Name	Remark
All	Folder specified in value record under: HKLM\Software\Microsoft\Windows NT\Cur-rentVersion\WinLogon\SfcDllCache	All files/folders (recursively)	Folder con-tains a cached dll file which is used for System File Checker (SFC) and contents of the system dll cache directory are rebuilt by using SFC.
	%SystemRoot%\SYSTEM32\dllCache		
	Root folder of quorum_device	MSCS*	Used for Microsoft Cluster Server.

	%SystemRoot%\SYSTEM32\	perf?00?.dat perf?00?.bak CATROOT*	Performance data used by the Windows Performance Counter. Used for Windows File Protection (WFP) records digital signatures of the operating system installs (such as DLL, EXE, SYS, OCX, and so on) to protect them from deletion or from replacement by older versions.
	%SystemRoot%\inetsrv\	metabase.bin	Metabase binary file of earlier IIS versions before 6.0.
	File or folder specified in value except "SIS Common Store" under HKLM\SYSTEM\CurrentControlSet\Control\BackupRestore\FilesNotToBackup	All files/folders (recursively)	Files and folders should not be backed up and restored. For more information, see link.
XP W2003	System volume	NTLDR	The main boot loader.

		BOOT.INI	Contains boot configuration (if missing, NTLDR will default to \Windows on the first partition of the first hard drive).
		NTDETECT.COM	Required for booting an NT-based OS. Detects basic hardware information needed for a successful boot.
Vista and later	Root folder of system volume	boot*	Boot folder for Windows.
		bootmgr	Windows boot manager file.
		EFI\Microsoft\Boot*	Used for EFI boot.
	%SystemRoot%\SYSTEM32\	LogFiles\WMI\RTBackup*	Stores ETW trace files (extension .etl) for real time event trace sessions.
		config\RegBack*	Backup of current registry table.

Win-8 and later	System volume	swapfile.sys	System controller file, normally around 256 MB. It is used by Metro style applications that do not fit the traditional paging characteristics (such as usage pattern, growth, space reservation) of pagefile.sys.
		BOOTNXT	Used to boot from OS, other than Windows 8. Created when enabling the startup options, and updated by Windows.

The Activity log provides the following information:

- Date Time Information: jobxxxx System Files skipped. You can use Bare-Metal Recovery Option (BMR) to restore them.
- Date Time Information: jobxxxx Files or Directories skipped. Skipped files or directories are available at: C:\Program Files\Arcserve\Unified Data Protection\Engine\Logs\Restore-<YYYYMMDD>-<hhmmss>-<Process ID>-<Job ID>.log.

Specify CSV Information to Restore

Arcserve UDP provides you with an option to restore data from a recovery point. The goal of running a successful restore job is to identify quickly the data you need and to retrieve it from the appropriate backup location. Each restore job requires a source and destination.

What To Do Next?

1. Specify CSV and Content to Restore
2. Define the Restore Options

Specify CSV and Content to Restore

Use the **Browse Recovery Points** option to restore from a recovery point. When you select a recovery date, and then specify the time, all the associated recovery points for that duration are displayed. You can then browse and select the backup content (including applications) to be restored.

Follow these steps:

1. Access the restore method selection dialog in one of the following ways:

   From Arcserve UDP Console:

   a. Log into **Arcserve UDP**.

   b. Click the **Resources** tab.

   c. Select **All Nodes** in the left pane.

      All the added nodes are displayed in the center pane.

   d. In the center pane, select the node and click **Actions**.

   e. Click **Restore** from the **Actions** drop-down menu.

      The restore method selection dialog opens.

      Note: You are automatically logged into the agent node and the restore method selection dialog is opened from the agent node.

   From Arcserve UDP Agent (Windows):

   a. Log into Arcserve UDP Agent (Windows).

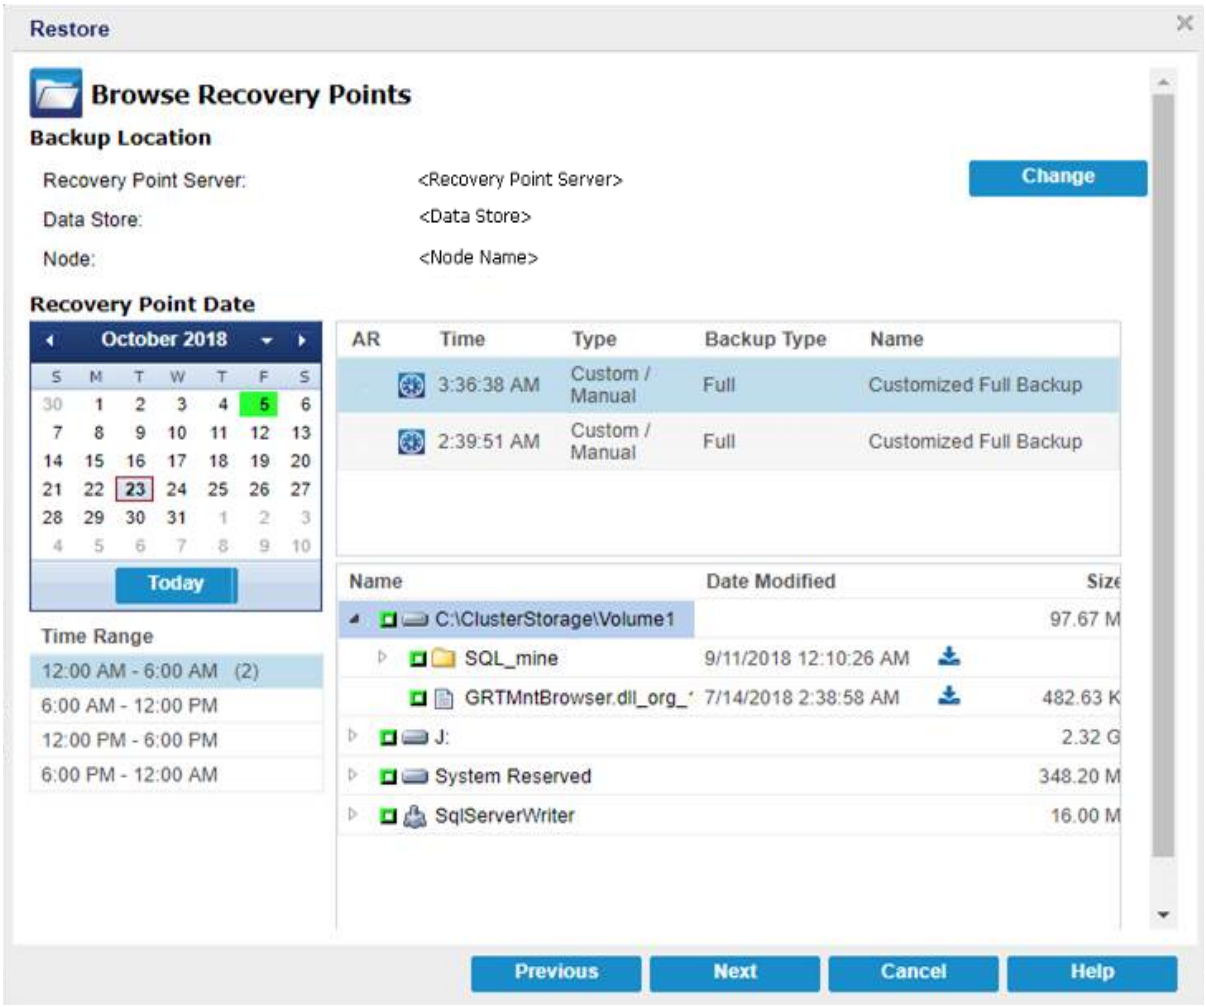
   b. From the home page, select **Restore**.

   The restore method selection dialog opens.

2. Click the **Browse Recovery Points** option.

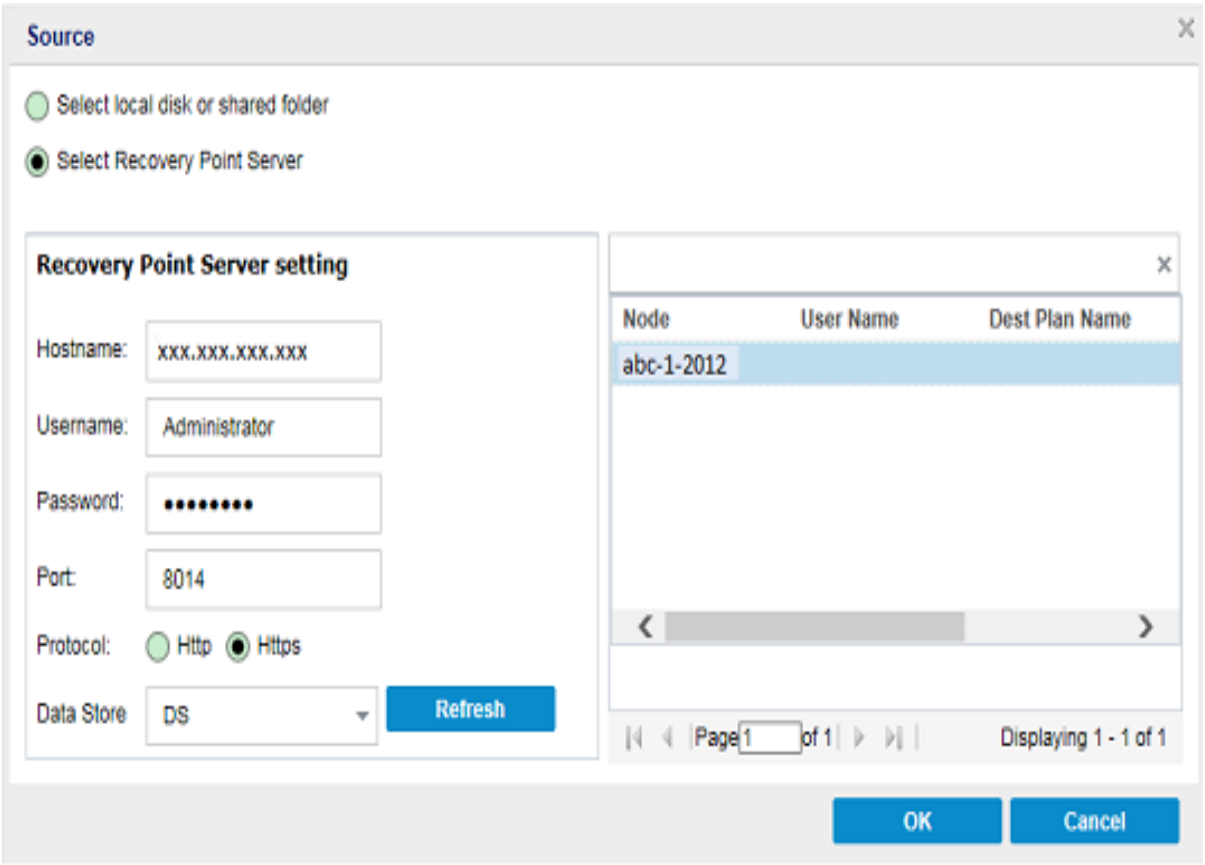
The **Browse Recovery Points** dialog opens. You can see the **Recovery Point Server** details in the **Backup Location**.

AR indicates the run result if Assured Recovery ran for the session.



3. Click **Change** to update the backup location.

   The **Source** dialog opens where you can select the backup location.



4. Select one of the following sources:

   Select local disk or shared folder

   a. Specify or browse to the location where your backup images are stored and select the appropriate backup source.

      You can click the green arrow button to verify the connection to the specified location. If necessary, enter the **Username** and **Password** credentials to gain access to that source location.

      The **Select backup location** dialog opens.

   b. Select the folder where the recovery points are stored and click **OK**.

      The **Select backup location** dialog closes and you can see the backup location in the **Source** dialog.

   c. Click **OK**.

      The recovery points are listed in the **Browse Recovery Points** dialog.

      Select Recovery Point Server

   d. Specify the Recovery Point Server setting details and click **Refresh**.

All the agents are listed in the Data Protection Agent column in the Source dialog.

e. Select the agent from the displayed list and click **OK**.

The recovery points are listed in the **Browse Recovery Points** dialog.

5. Select the calendar date for the backup image to restore.

   All the dates containing recovery points for the specified backup source are highlighted in green.

   The corresponding recovery points for that date are displayed, with the time of the backup, the type of backup that was performed (Full, Incremental, or Verify), and the name of the backup.

6. Select a recovery point to restore.

   The backup content (including any applications) for the selected recovery point displays.

   Note: If a clock icon appears with a lock symbol, it indicates that the recovery point contains encrypted information and may require a password for restore.

7. Select the CSV volume to restore.

   - For a volume-level restore, you can specify to restore the entire volume or selected files/folders within the volume.
   - For an application-level restore, you can specify to restore the entire application or selected components, databases, instances, and so on, within the application.

8. Click **Next**.

   The **Restore Options** dialog opens.

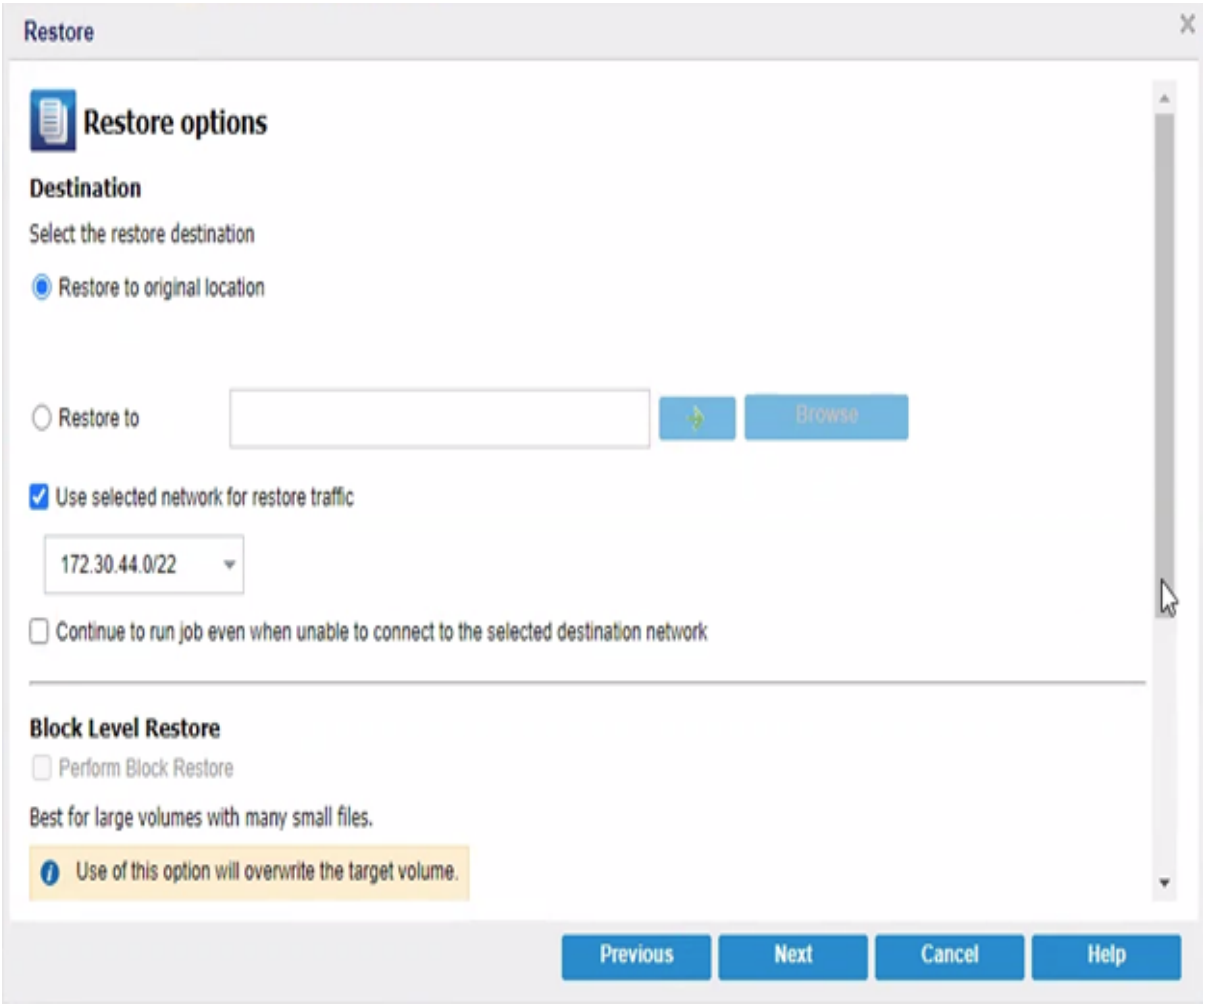
   The recovery point and content to restore is specified.

Define the Restore Options

After you specify a recovery point and content to restore, define the copy options for the selected recovery point.

Follow these steps:

1. On the **Restore Options** dialog, select the restore destination.



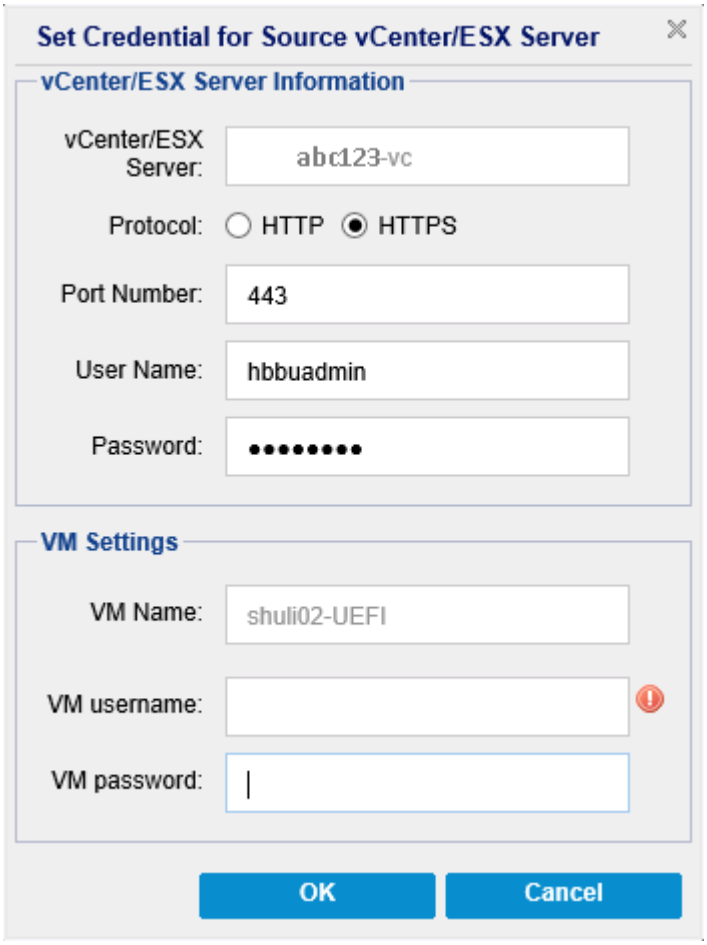
The available destination options are:

Restore to Original Location

Restores to the original location from where the backup image was captured.

Note: If you performed the recovery point backup using host-based agentless backup, restoring to original location is to restore the file back in to the virtual machine. In this case, a dialog box opens. You may enter the credentials of the hypervisor, and the operating system of the virtual machine.

For VMware VM:

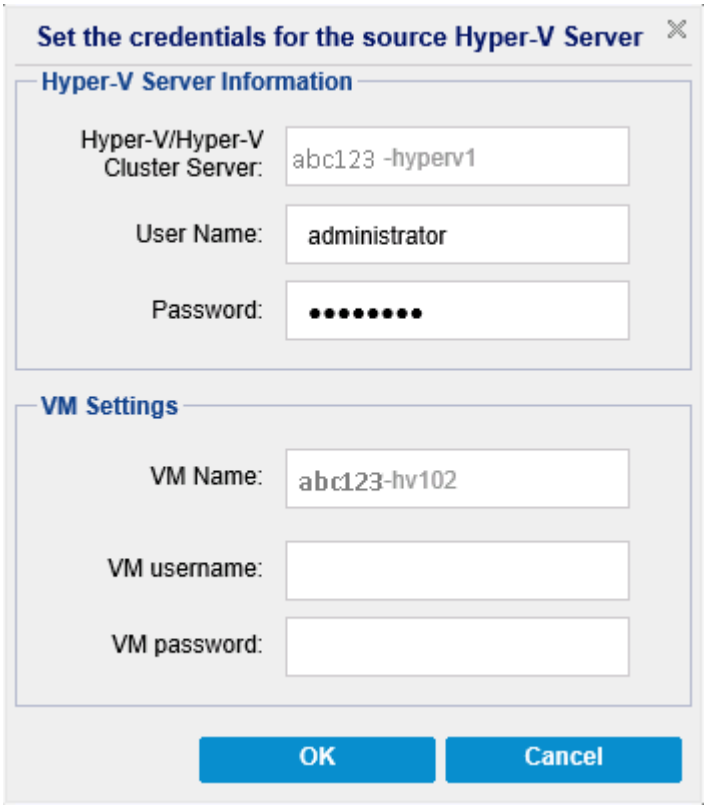


Note: To be able to create or write files inside the VM, consider the following requirements for the settings and account permission of virtual machine:

- VMware Tools is installed and running.
- Firewall must allow File and Printer Sharing.
- The account is built-in local administrator, built-in domain administrator, or domain account that is member of the local Administrators group. If other accounts are used, then:
  - Disable the UAC remote access. To disable UAC remote access, see Import Virtual Machine Using Additional Administrative Account.
  - Disable UAC in the Local Security Policy by disabling the setting Run all administrator in Admin Approval Mode at secpol.msc -> Local Policies -> Security Options. (Secpol.msc is Microsoft's security policy editor).

Important: Do not attempt to disable the UAC in the User Account Control Settings dialog box that opens from the control panel.

For Hyper-V VM:



Note: To be able to create or write files inside the VM, consider the following requirements for the settings and account permission of virtual machine:

- Hyper-V integration services are installed and running.
- Firewall must allow File and Printer Sharing.
- The account is built-in local administrator, built-in domain administrator, or domain account that is member of the local Administrators group. If other accounts are used:

  Disable the UAC remote access. To disable UAC remote access, see Import Virtual Machine Using Additional Administrative Account.
- If virtual machine guest OS is Client version Windows (such as Windows 10), you need to manually configure firewall to allow Windows Management Instrumentation (WMI).

Restore to

Restores to the specified location. You can click the green arrow button to verify the connection to the specified location. If necessary, enter the User-name and Password credentials to gain access to that location.

2. Specify the **Resolving Conflicts** option that Arcserve UDP performs if conflicts are encountered during the restore process.

   The available options are:

   Overwrite existing files

   Overwrites (replaces) any existing files, which are at the restore destination. All objects are restored from the backup files regardless of their current presence on your computer.

   Replace active files

   Replaces any active files upon reboot. If during the restore attempt Arcserve UDP Agent (Windows) discovers that the existing file is currently in use or being accessed, it will not immediately replace that file, but instead to avoid any problems will delay the replacement of the active files until the next time the machine is rebooted. The restore occurs immediately, but the replacement of any active files is performed during the next reboot.

   This option is only available if you select the **Overwrite existing files** option.

   Note: If you do not select this option, any active file is skipped from the restore.

   Rename files

   Creates a new file if the file name already exists. Selecting this option copies the source file to the destination with the same filename but a different extension. Data is then restored to the new file.

   Skip existing files

   Skips over and not overwrite (replace) any existing files that are located at the restore destination. Only objects that are not currently existing on your machine are restored from the backup files.

   Default: Skip existing files.

3. Specify the **Directory Structure** to create a root directory during restore.

   Create root directory

   Specifies that if a root directory structure exists in the captured backup image, Arcserve UDP recreates that same root directory structure on the restore destination path.

With this option not selected, the file or folder is restored directly to the destination folder.

For example, if during the backup you captured the files "C:\Folder-1\SubFolder2\A.txt" and "C:\Folder1\SubFolder2\B.txt" and during the restore you specified to the restore destination as "D:\Restore".

- If you select to restore the "A.txt" and "B.txt" files individually, the destination for the restored files will be "D:\Restore\A.txt" and "D:\Restore\B.txt" (the root directory above the specified file level will not be recreated).
- If you select to restore from the "SubFolder2" level, the destination for the restored files will be "D:\Restore\SubFolder2\A.txt" and "D:\Restore\SubFolder2\B.txt" (the root directory above the specified folder level will not be recreated).

With this option selected, the entire root directory path for the files/folders (including the volume name) is recreated to the destination folder. If the files/-folders to be restored are from the same volume name, then the destination root directory path does not include that volume name. However, if the files/-folders to be restored are from different volume names, then the destination root directory path does include the volume name.

For example, if during the backup you captured the files "C:\Folder-1\SubFolder2\A.txt", "C:\Folder1\SubFolder2\B.txt", and also E:\Folder-3\SubFolder4\C.txt" and during the restore you specified to the restore destination as "D:\Restore".

- If you select to restore just the "A.txt" file, the destination for the restored file will be "D:\Restore\ Folder1\SubFolder2\A.txt" (the entire root directory without the volume name will be recreated).
- If you select to restore both the "A.txt" and "C.txt" files, the destination for the restored files will be "D:\Restore\C\Folder1\SubFolder2\A.txt" and "D:\Restore\E\Folder3\SubFolder4\C.txt" (the entire root directory with the volume name will be recreated).

4. From **Recovering ACL**, select the **skip recovering ACL of files / folders** option to skip the original permission for the restored files/folders. Selecting the option lets you inherit the permissions of target folder instead. If you do not select the option, the original permissions are kept.
5. If necessary, specify the **Backup Encryption Password**, when the data you are trying to restore is encrypted.

A password is not required if you are attempting to restore from the same Arcserve UDP Agent (Windows) computer from where the encrypted backup was performed. However, if you are attempting to restore from a different Arcserve UDP Agent (Windows) computer, a password is required.

Note: If a clock icon appears with a lock symbol, it indicates that the recovery point contains encrypted information and may require a password for restore.

6. Click **Next**.

   The **Restore Summary** dialog opens.

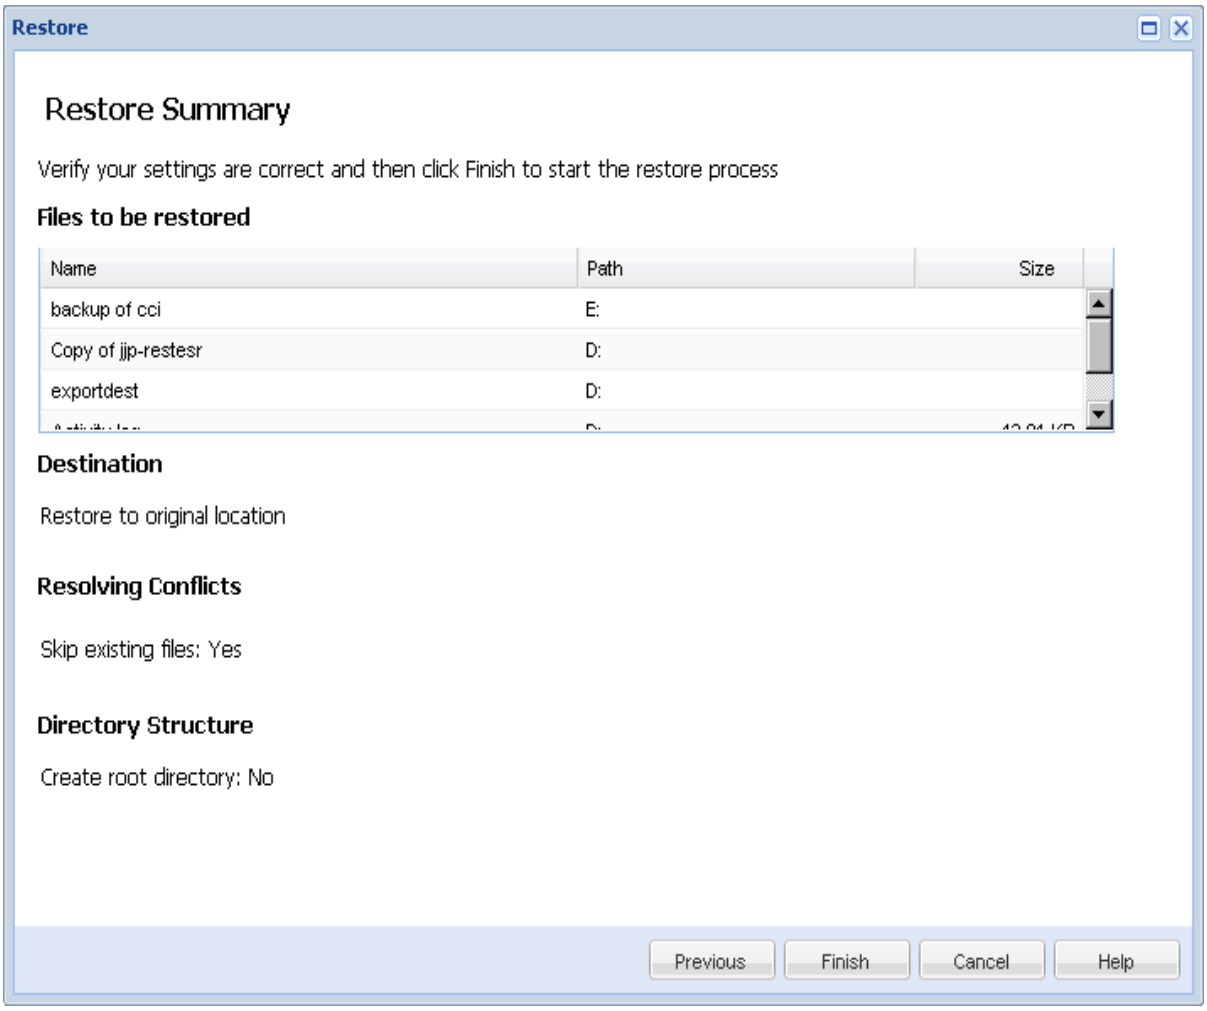
   The restore options are defined to restore from a recovery point.

Restore CSV Content

After you define the restore options, verify that your settings are correct and confirm the restore process. **Restore Summary** helps you to review all the restore options that you defined and modify them if necessary.

Follow these steps:

1. On the **Restore Summary** dialog, review the displayed information to verify that all the restore options and settings are correct.



- If the summary information is incorrect, click **Previous** and go back to the applicable dialog to change the incorrect setting.
- If the summary information is correct, click **Finish** to launch the restore process.

The recovery point content is restored.

Chapter 14: Managing Tape Backup and Restore

Arcserve UDP lets you back up data to a tape and restore the backed up data from the tape to a node.

This section contains the following topics:

How to Back Up a Deduplication Data Store to a Tape

To backup up a deduplication data store to a tape, use the following methods:

- Backup UDP nodes to tape. For more information, see the Method 1: Backing up the latest UDP backup session to tape section in the Arcserve Backup Administration Guide.
- Backup Arcserve UDP data from Datastores. For more information, see the Method 2: Backing up Arcserve UDP data from Datastores section in the Arcserve Backup Administration Guide.

How to Restore a Deduplication Data Store From a Tape

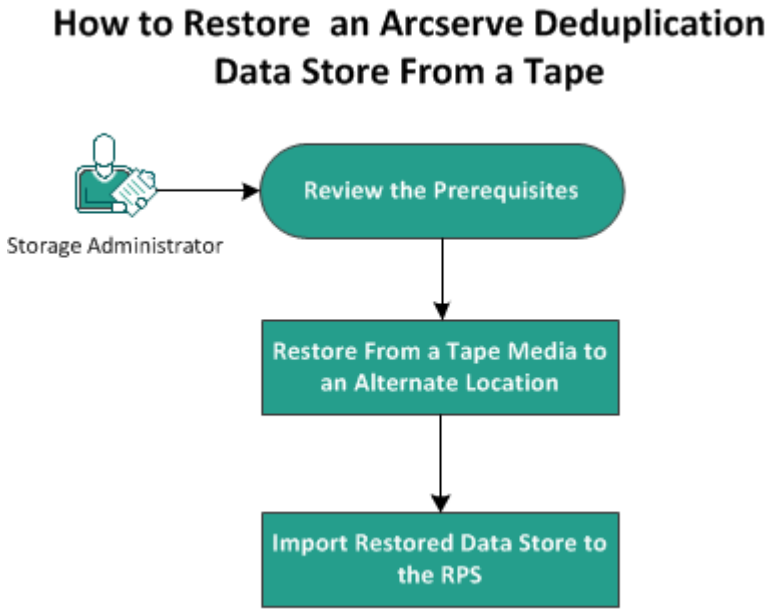
If you have previously backed up a deduplication data store from a recovery point server (RPS) to a tape device, you can restore the data store. For this procedure, Arcserve Backup and Arcserve UDP are used in conjunction with each other to restore a deduplication data store from a tape. Arcserve Backup is used to restore it from the tape to a specified destination, and then Arcserve UDP is used to import it to an RPS.

The following two processes are involved in the restore procedure:

1. The first process uses Arcserve Backup to restore the sessions from the tape media to a volume. It is recommended to restore the sessions to an alternate location.
2. The second process uses Arcserve UDP to import the restored data store to the RPS.

Note: You will need to provide the **Backup Destination Folder** path of the deduplication data store when you browse the location during import.

The following diagram illustrates how to restore AN Arcserve deduplication data store from a tape:



What To Do Next?

1. Review the Prerequisites
2. Restore From a Tape Media to an Alternate Location
3. Import Restored Data Store to the RPS

Review the Prerequisites

Review the following prerequisites before you begin the restore:

- You must have backed up an RPS data store to the tape.
- You will need to provide the session password, if necessary.
- You will need to provide the user name and password for the restore destination.

Restore From a Tape Media to an Alternate Location

To restore the session from the tape media to an alternate location, you will need to use the Arcserve Backup Manager.

After the restore is successful, you can then import the restored data store to the RPS using Arcserve UDP.

Follow these steps:

1. From Arcserve Backup, log in to the Arcserve Backup Manager.
2. From the **Quick Start** navigation pane, click **Restore** and then from the center pane, select the **Source** tab.
3. From the drop-down menu, select **Restore by Session** and select the session that you want to restore.
4. Click the **Destination** tab.
5. Clear (uncheck) the **Restore files to their original location(s)** check box.
6. Expand **Windows Systems** and browse the location where you want to restore to.
7. Click the **Schedule** tab and select **Once** for the **Repeat Method** option.
8. Click **Submit**.

   The **Restore Media** dialog opens.
9. Verify the restore media and click **OK**.

   The **Session User Name and Password** dialog opens.
10. Provide the user name and password for the restore location and the session password for the recovery points, if necessary.
11. Click **OK**.

    The **Submit Job** dialog opens.
12. Provide the required information on the **Submit Job** dialog and click **OK**.

    The restore job is submitted.

    After the restore job is complete, the Arcserve UDP data store files will be displayed at the location that you specified.

Import Restored Data Store to the RPS

To import the restored data store to the RPS, you will need to use the Arcserve UDP Console. The **Import Data Store** feature lets you add a data store to the recovery point server. You can import any existing data store to a recovery point server.

Follow these steps:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Recovery Point Servers**.

   The **Destinations: Recovery Point Servers** page is displayed.
3. Perform one of the following actions:
   - Right click a recovery point server.
   - Select a recovery point server, and from the center menu click the **Actions** drop-down list.
4. Click **Import Data Store**.

   The **Import a Data Store** page is displayed.
5. Perform the following actions, and click **Next**:
   a. Click **Browse** and select the **Backup Destination Folder** from where you want to import the data store.
   b. If necessary, enter the **Encryption Password**.

      Note: If the data store is not encrypted, you can leave this field empty.

   After authenticating the **Backup Destination folder**, the **Import a Data Store** page displays the details of the data store.

   Note: Importing of restored datastore is not supported if the original datastore is present in the same RPS. You need to delete the original datastore before importing the restored datastore or alternately import the datastore to another RPS.
6. If necessary, modify the data store details and click **Save**.

   When the restored data store is imported, Arcserve UDP will continue to read the index, hash, and data destinations from the data store configuration settings and show the actual paths where data store originally resided. After the import is completed, these path destinations need to be changed to the new restored paths.

   Note: You cannot enable or disable the encryption option for an existing data store.

   The data store is added to the recovery point server and displayed at the **Destinations: Recovery Point Servers** dialog.

Upon completion of a successful import, a green check mark is displayed next to the corresponding data store name.

You have successfully restored a deduplication data store from a tape.

Chapter 15: Using the PowerShell Interface

This section contains the following topics:

How to use the PowerShell Interface

Arcserve UDP provides PowerShell capabilities that allows you to submit a backup job, perform a restore, and recover VM from the command line. The PowerShell interface is named as UDPPowerCLI.ps1.

Review the Prerequisite

Review the following prerequisites before using the PowerShell interface:

- You must have Windows 2008 R2 Server or higher versions.
- You must have PowerShell 3 or higher version installed on your server.

Using the PowerShell Interface for Arcserve UDP

The PowerShell utility is bundled with the Arcserve UDP installation file.

When you install Arcserve UDP, typically the file gets installed at the following location:

C:\Program Files\Arcserve\Unified Data Protection

In such cases, on Console, UDPPowerCLI.ps1 gets installed at the following location:

C:\Program Files\Arcserve\Unified Data Protection\Management\PowerCLI

On RPS or Agent, UDPPowerCLI.ps1 gets installed at the following location:

C:\Program Files\Arcserve\Unified Data Protection\Engine\PowerCLI

Refer the following options to help you use the PowerShell interface:

- Update the PowerShell execution policy to allow the scripts to run. For example, update the execution policy to **Set-ExecutionPolicy RemoteSigned**.

  Note: For more information about changing the execution policy, see the Microsoft website.

- Run the following PowerShell command to get the detailed help messages and examples for the scripts:

  ```
  On Console:
  ```

  Get-Help 'C:\Program Files\Arcserve\Unified Data Protection\Management\PowerCLI\UDPPowerCLI.ps1' -full

  On RPS or Agent:

  Get-Help 'C:\Program Files\Arcserve\Unified Data Protection\Engine\PowerCLI\UDPPowerCLI.ps1' -full

PowerShell Syntax and Parameters

SYNTAX 1

```
UDPPowerCLI.ps1 -Command <CreatePswFile> -Password
<System.Security.SecureString> -PasswordFile <string>
[<CommonParameters>]
```

SYNTAX 2

```
UDPPowerCLI.ps1 -Command <Backup> [-UDPCon-
soleServerName <String>] [-UDPConsoleProtocol <
{http|https}>] [-UDPConsolePort <int>] [-UDPCon-
soleUserName [<String>]] [-UDPConsolePassword <Sys-
tem.Security.SecureString>] [-UDPConsolePasswordFile
<String>] [-UDPConsoleDomainName <String>] -planName
<String> -nodeName <String> [-backupJobType <String>]
[-jobDescription <String>] [-waitJobFinish <String
String>] [-timeOut <int>] [-jobType <{ agent-
base|agentless|uncpath} String>] [-backupScheduleType
<String>] [<CommonParameters>]
```

SYNTAX 3

```
UDPPowerCLI.ps1 -Command <Restore> [-UDPCon-
soleServerName <String>] [-UDPConsoleProtocol
<String>] [-UDPConsolePort <int>] [-UDPConsoleUserName
<String>] [-UDPConsolePassword <Sys-
tem.Security.SecureString>] [-UDPConsolePasswordFile
<String>] [-UDPConsoleDomainName <String>] [-UDPA-
gentServerName <String>] [-UDPAgentProtocol <String>]
[-UDPAgentPort <int>] [-UDPAgentUserName <String>] [-
UDPAgentPassword <System.Security.SecureString>] [-
UDPAgentPasswordFile <String>] [-UDPAgentDomainName
<String>] [-RestoreDirectoryPath <String>] [-
RestoreFilePath <String>] [-BackupSessionNumber <int>]
[-VmName <String>] -RestoreDestination <String> [-
RestoreDestinationUserName <String>] [-RestoreDestin-
ationPassword <System.Security.SecureString>] [-
CreateRootFolder <String>] [-ChangeFileName <String>]
[-ReplaceActiveFilesFlag <String>] [-Over-
writeExistFiles <String>] [<CommonParameters>]
```

SYNTAX 4

```
UDPPowerCLI.ps1 -command <RecoverVM> [-UDPCon-
soleServerName <String>] [-UDPConsoleProtocol
<String>] [-UDPConsolePort <int>] [-UDPConsoleUserName
<String>] [-UDPConsolePassword <Sys-
tem.Security.SecureString>] [-UDPConsolePasswordFile
```

```
<String>] [-UDPConsoleDomainName <String>] [-UDPA-
gentServerName <String>] [-UDPAgentProtocol <String>]
[-UDPAgentPort <int>] [-UDPAgentUserName <String>] [-
UDPAgentPassword <System.Security.SecureString>] [-
UDPAgentDomainName <String>] [-UDPAgentPasswordFile
<String>] [-BackupSessionNumber <int>] -RecoverVmName
<String> [-OverwriteExistingVM <String>] [-PoweronVM
<String>] [<CommonParameters>]
```

DESCRIPTION

A utility to connect to the Arcserve UDP Console service, and submit backup and restore jobs.

PARAMETERS

-Command <String>

Specifies the command that is used. Currently, the following strings are supported:

- CreatePswFile
- Backup
- Restore
- RecoverVM

Required? **true**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-Password <System.Security.SecureString>

Specifies the password you want to use for creating the password file.

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPConsoleServerName <String>

Specifies the DNS name of the UDP server (the server where you have installed the Console) to which you want to connect. If this value is not specified, then the cmdlet uses the default value, the DNS name of the local machine.

Required? **false**

Position? **named**

Default value **$env:COMPUTERNAME**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPConsolePort <int>

Specifies the port number you want to use for the connection. If this value is not specified, then the cmdlet uses the default value, 8015.

Required? **false**

Position? **named**

Default value **8015**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPConsoleProtocol <String>

Specifies the protocol on the server that you want to use for the connection. The protocol can be either http or https. If this value is not specified, then the cmdlet uses the default value, http.

Required? **false**

Position? **named**

Default value **http**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPConsoleUserName <String>

Specifies the user name you want to use for connecting to the UDP server. If the user name is not specified, then the cmdlet uses the user name currently used to log into the system.

Required? **false**

Position? **named**

Default value **$env:UserName**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPConsolePassword <System.Security.SecureString>

Specifies the password you want to use for connecting to the UDP server.

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-PasswordFile <String>

Specifies to generate the password file.

Required? **true**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPConsolePasswordFile <String>

Specifies the UDP password file you want to use for connecting to the UDP server.

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPAgentServerName <String>

Specifies the DNS name of the UDP agent server to which you want to connect for restore.

Required? **false**

Position? **named**

Default value **$env:COMPUTERNAME**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPAgentProtocol <String>

Specifies the internet protocol that you want to use to connect to the UDP agent server. It can be either http or https. If this value is not specified, then the cmdlet uses the default value, http.

Required? **false**

Position? **named**

Default value **http**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPAgentPort <int>

Specifies the port number that you want to use to connect to the UDP agent server. If this value is not specified, then the cmdlet uses the default value, 8014.

Required? **false**

Position? **named**

Default value **8014**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPAgentUserName <String>

Specifies the user name that you want to use to connect to the UDP agent server. If the user name is not specified, then the cmdlet uses the user name currently used to log into the system.

Required? **false**

Position? **named**

Default value **$env:UserName**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPAgentPassword <System.Security.SecureString>

Specifies the password that you want to use to connect to the UDP agent server.

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPAgentPasswordFile <String>

Specifies the UDP agent password file that you want to use to connect to the UDP agent server.

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPAgentDomainName <String>

Specifies the domain name where the specified UDP agent user is located.

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-NodeName <String>

Specifies the name of node that you want to back up.

Required? **true**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-RestoreFilePath <String>

Specifies the file that you want to restore.

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-RestoreDirectoryPath <String>

Specifies the directory that you want to restore.

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-BackupSessionNumber <int>

Specifies the session number to use for the restore job.

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-VmName <String>

Specifies the host name of a virtual machine for restoring file or directory from its backup session.

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-RestoreDestination <String>

Specifies the directory path where the files will be restored.

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-RestoreDestinationUserName <String>

Specifies the user name of the destination machine where you want to restore data. The user name belong to the user who can log in to the destination machine.

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-RestoreDestinationPassword <System.Security.SecureString>

Specifies the password that you will use to log into the destination machine.

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-CreateRootFolder <String>

Specifies that if a root directory structure exists in the captured backup image, Arcserve UDP recreates that same root directory structure on the restore destination path. When this option is not selected, the file or folder is restored directly to the destination folder. You can use any one of the following strings:

- True
- False

Required? **false**

Position? **named**

Default value **False**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-ChangeFileName <String>

Creates a new file if the file name already exists. Selecting this option copies the source file to the destination with the same filename but a different extension. Data is then restored to the new file. You can use any one of the following strings:

- True
- False

Required? **false**

Position? **named**

Default value **False**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-ReplaceActiveFilesFlag <String>

Replaces any active files after reboot. If during the restore attempt Arcserve Unified Data Protection Agent for Windows discovers that the existing file is currently in use or being accessed, it will not immediately replace that file, but instead, to avoid any problems, will delay the replacement of the active files until the next time the machine is rebooted. (The restore occurs immediately, but the replacement of any active files is done during the next reboot). This option is only available when the **OverwriteExistingFiles** parameter is True. You can use any one of the following strings:

- True
- False

Required? **false**

Position? **named**

Default value **False**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-OverwriteExistingFiles <String>

Overwrites (replaces) any existing files, which are at the restore destination. All objects are restored from the backup files regardless of their current presence on your computer. You can use any one of the following strings:

- True
- False

Required? **false**

Position? **named**

Default value **False**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPConsoleDomainName <String>

Specifies name of domain where the specified user is located. If this value is not specified, then the cmdlet uses the domain name of local machine; or the DNS name of local machine if it is not in a domain.

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-PlanName <String>

Specifies the plan name that defines the backup job setting.

Required? **true**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-BackupJobType <String>

Specifies the type of the backup job. One of the following values can be used: Full (indicates a Full backup), Incr (indicates an Incremental backup), or Rsyn (indicates a Resync backup). The following strings are supported:

- Full
- Incr
- Rsyn

Required? **true**

Position? **named**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-JobDescription <String>

Specifies the description for the backup job.

Required? **true**

Position? **named**

Default value **PowerCLIJo**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-RecoverVmName <String>

Specifies the host name of the virtual machine that you want to recover.

Required? **true**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-OverwriteExistingVM <String>

Specifies that if the value is true, the restore job overwrites the existing virtual machine. The default value is false. You can use any one of the following strings:

- True
- False

Required? **false**

Position? **named**

Default value **False**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-PoweronVM <String>

Specifies that if the value is true, the virtual machine is powered on after it is recovered. The default value is false. You can use any one of the following strings:

- True
- False

Required? **false**

Position? **named**

Default value **False**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-waitJobFinish <{true|false} String>

Specifies that if the value is true, the command waits for further instructions until the backup job is complete. The default value is false. You can use any one of the following strings:

- True
- False

Required? **false**

Position? **named**

Default value **False**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-timeOut <int>

Specifies the maximum waiting time (in seconds) for the backup job to complete.

Required? **false**

Position? **named**

Default value **600**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-jobType <String>

Specifies the backup job type for different node type. The default value is agentbase. Use any one of the following strings:

- agentbase
- agentless
- uncpath

Required? **false**

Position? **named**

Default value **agentbase**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-backupScheduleType <String>

Specifies schedule backup job, submits the specified schedule backup job immediately, and runs only once. The following strings are supported:

- Daily
- Weekly
- Monthly

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

<CommonParameters>

This cmdlet supports the common parameters such as **Verbose**, **Debug**, **ErrorAction**, **ErrorVariable**, **WarningAction**, **WarningVariable**, **OutBuffer**, and **OutVariable**. For more information, see about_CommonParameters.

INPUTS

OUTPUTS

- 0 or 1

If job is successfully submitted successfully, the command returns 0; otherwise returns 1.

PowerShell Examples

Create System.Security.SecureString type of password

Note: If you have a configured PowerShell script in Arcserve UDP 6.5 that uses a plain text password, modify the same to use $SecurePassword or use the password file.

Description

The command is used to create the password with System.Security.SecureString type and is also used by all other commands. The System.Security.SecureString type is system predefined type. There are many ways to generate it. The following two types are commonly used for different purposes:

- This command asks for inputting the password on the PowerShell Console. It requires interaction with the end users. Usually after executing this command, we could execute the command *CreatePswFile* to write the password to a password file. And then use the password file for running PowerShell scripts automatically.

  ```
  $SecurePassword = Read-Host -AsSecureString
  ```

- This command requires the password with plain text. It could use used directly for running PowerShell scripts automatically.

  ```
  $SecurePassword = ConvertTo-SecureString
  "<PlainPassword>" -AsPlainText -Force
  ```

- (Optional) After executing one of the previous commands to assign a password, execute the CreatePswFile command to write the password to a password file. Next, use the password file for running the PowerShell scripts automatically.

  The command encrypts secure password and saves it to the password file.

  ```
  C:\PS>UDPPowerCLI.ps1 -Command CreatePswFile -
  Password $SecurePassword -PasswordFile
  myUDPPasswordFile
  ```

Example 1

Description

The command encrypts secure password and save it to the password file.

```
C:\PS>UDPPowerCLI.ps1 -Command CreatePswFile -Password
$SecurePassword -PasswordFile myUDPPasswordFile
```

Example 2

Description

On the local server, the command connects to the UDP Console service with HTTP protocol over port 8015, and then submits an Incremental backup job for the plan named *myplan*.

```
C:\PS>UDPPowerCLI.ps1 -Command Backup -UDPCon-
soleUserName myUsr -UDPConsolePassword $SecurePassword
-PlanName myPlan -BackupJobType Incr
```

Example 3

Description

On the local server, the command connects to the UDP Console service with HTTP protocol over port 8015, and then submits an Incremental backup job for the node named *myNodeName*.

```
C:\PS>UDPPowerCLI.ps1 -Command Backup -UDPCon-
soleUserName myUsr -UDPConsolePasswordFile myUDPPass-
wordFile -NodeName myNodeName -BackupJobType Incr
```

Example 4

Description

The command connects to the UDP Console service on the server named *myServer* with HTTPS protocol over port 8018, and then submits a Full backup job for the plan named *myPlan*, and set the job description as *myJob*.

```
C:\PS>UDPPowerCLI.ps1 -Command Backup -UDPCon-
soleServerName myServer -UDPConsoleProtocol https -
UDPConsolePort 8018 -UDPConsoleUserName myUsr -UDPCon-
solePassword $SecurePassword -UDPConsoleDomainName
myDomain -PlanName myPlan -BackupJobType Full -JobDe-
scription myJob
```

Example 5

Description

The command connects to the UDP agent service on the server named *yourUDPAgentServer* with HTTP protocol over port 8014, and then submits an Incremental backup job for *yourUDPAgentServer*.

```
C:\PS>UDPPowerCLI.ps1 -Command Backup -UDPA-
gentServerName yourUDPAgentServer -UDPA-
gentPasswordFile myUDPAgentPasswordFile -BackupJobType
Incr
```

Example 6

Description

The command shortens the Parameter name.

```
C:\PS>UDPPowerCLI.ps1 -Cmd Backup -Svr myServer -Ptc
https -Prt 8018 -Usr myUsr -Psw $SecurePassword -Dmn
myDomain -Pln myPlan -Jbt Full -Jbd myJob
```

Example 7

Description

The command connects to the server named *yourUDPAgentServer* using the user name of the environment, the default HTTP protocol, and port 8014. It verifies the backup session number is 1 from the *yourUDPAgentServer* backup configuration and then restores the directory to the original location, with the restore option selected as Overwrite existing files.

```
C:\PS>UDPPowerCLI.ps1 -Command restore -UDPA-
gentServerName yourUDPAgentServer -UDPA-
gentPasswordFile myUDPAgentPasswordFile -
RestoreDirectoryPath 'c:\Test' -BackupSessionNumber 1
```

Example 8

Description

The command connects to the server named *yourUDPAgentServer* using the HTTPS protocol and port 8018. It verifies the backup session number is 1 from the *yourUDPAgentServer* backup configuration and then restores the 1.txt file to an alternate location, with the restore option selected as Overwrite existing file and create root directory.

```
C:\PS>UDPPowerCLI.ps1 -Command restore -UDPA-
gentServerName yourUDPAgentServer -UDPAgentUserName
UDPAgentUsername -UDPAgentPasswordFile myUDPA-
gentPasswordFile -UDPAgentProtocol 'https' -UDPA-
gentPort 8018 -UDPAgentDomainName UDPAgentdomainName -
BackupSessionNumber 1 -RestoreFilePath 'C:\1.txt' -
RestoreDestination 'C:\restore' -RestoreDestin-
ationUserName remoteAccessUser -RestoreDestin-
ationPassword remoteAccessPsw -CreateBaseFolder 'true'
```

Example 9

Description

The command connects to the server named *yourUDPAgentServer* using the user name of the environment, the default HTTP protocol, and port 8014. Then, it connects to the UDP server using the default port 8015 and protocol HTTP to check the backup session number is 1. Lastly, it restores the directory to an alternate location, with the restore option selected as Overwrite existing file and create root directory.

```
C:\PS>UDPPowerCLI.ps1 -Command restore -UDPA-
gentServerName yourUDPAgentServer -UDPA-
gentPasswordFile myUDPAgentPasswordFile -
RestoreDirectoryPath 'c:\Test' -BackupSessionNumber 1
-RestoreDestination 'C:\restore' -RestoreDestin-
ationUserName remoteAccessUser -RestoreDestin-
ationPassword remoteAccessPsw -UDPConsoleServerName
yourUDPServer -vmname sourceVMName -UDPCon-
solePasswordFile myUDPPasswordFile -domainname
yourUDPDomainName -OverwriteExistFiles 'true' -
CreateRootFolder 'true'
```

Example 10

Description

The command connects to the server named *yourUDPAgentServer* using the user name of the environment, the default HTTP protocol, and port 8014. Then, it connects to the UDP server using the default port 8015 and protocol HTTP to check the backup session number is 1. Lastly, it recovers the VM to the original location, with the recover VM option selected as Overwrite existing vm and power on vm after recovered.

```
C:\PS>UDPPowerCLI.ps1 -Command RecoverVM -UDPA-
gentServerName yourUDPAgentServer -UDPA-
gentPasswordFile myUDPAgentPasswordFile -
BackupSessionNumber 1 -UDPConsoleServerName
yourUDPServer -recovervmname sourceVMName -UDPCon-
solePasswordFile myUDPPasswordFile -UDPCon-
soleDomainName yourUDPDomainName -OverwriteExistingVM
'true' -PoweronVM 'true'
```

Example 11

Description

The command submits weekly backup job on the UDP Agent immediately and runs only one time.

```
C:\PS>UDPPowerCLI.ps1 -Command Backup -UDPA-
gentServerName myServer -UDPAgentPassword $Se-
curePassword -UDPAgentDomainName myDomainName -
UDPAgentUserName UDPAgentUsername -BackupJobType Incr
-backupScheduleType 'weekly' -jobDescription 'Power-
CLIJob'
```

Example 12

Description

The command submits the backup job and sets the timeout in seconds to wait for the job to complete.

```
C:\PS>UDPPowerCLI.ps1 -Command Backup -UDPCon-
soleServerName myServer -UDPConsolePasswordFile
myUDPPasswordFile -UDPConsoleDomainName myDomainName -
nodeName myNodeName -UDPConsoleUserName myAdmin -
BackupJobType Incr -jobDescription 'PowerCLIJob'
waitJobFinish 'true' -timeout 600 -jobType 'agentbase'
```

Configuration file and properties of PowerCLI

You can customize the following UDP PowerCLI configuration file and properties:

Configuration File Name: *PowerCLIConfiguration.properties*

Configuration File Installation Location:

- *C:\Program Files\Arcserve\Unified Data Protection\Management\Configuration*
- *C:\Program Files\Arcserve\Unified Data Protection\Engine\Configuration*

Configuration file format:

Enable24HoursSetting=yes

LogPath=D:\PowerCLILogPath

RotationSizeInMB=2

RetentionPeriodInDays=5

RetainNoOfFiles=6

Configuration Properties:

Property Name	Description	Default Value	Accepted Value
Enable24HoursSetting	▪ If the value is set to Yes, the timestamp of log file gets written in 24-hour format. ▪ If the value is set to No, the timestamp of log file gets written in 12-hour format.	Yes	Yes or No
LogPath	▪ The folder location of the PowerCLI log file(s). ▪ If the location is not provided, the log file gets written to its default loc-	-	-

	ation. ▪ LogPath accepts Relative and absolute paths. ▪ Symbols except \ / : * ? < > | are acceptable for file or folder name.		
RotationSizeInMB	▪ If the specified limit provided in RotationSizeInMB is reached, the PowerCLI log file gets renamed (Rotated) with a date time stamp.	0	Minimum = 1 Maximum = 1024
RetentionPeriodInDays	▪ If the rotated files older than the value provided in RetentinPeriodInDays, the old log files get deleted.	0	Minimum = 5 Maximum = 365
RetainNoOfFiles	▪ If there are more rotated log files than the value provided in RetainNoOfFiles, the extra files get deleted.	0	Minimum = 5 Maximum = 1000

Notes:

- If you provide values for configuration properties other than permitted value, the PowerCLI displays error and the execution gets stopped.
- If you provide default values for configuration properties, the execution of PowerCLI continues and runs with default behavior.

Chapter 16: Protecting Oracle System using Recovery Manager (RMAN)

Recovery Manager (RMAN) is an Oracle utility that performs backup job, restore and recovery tasks on databases and automates the administration of backup strategies. Using this feature, you can easily backup, restore, and recover the database files.

This section contains the following topics:

Review the Prerequisites and Considerations

For more information, see Review the Prerequisites for Oracle Database for RMAN (Windows and Linux).

How to Create a Backup Plan for Oracle Database (Windows and Linux)

To protect your Windows and Linux nodes or clustered nodes, you need to create a plan. The plan consists of a backup task. This backup task lets you specify the nodes you want to protect, the backup destination, and the backup schedule. The backup destination is a recovery point server where you want to store your backup data. The destination can also be a local destination or a remote share folder.

This section contains the following topics:

Create a Backup Plan with a Backup Task

A backup plan includes a backup task that performs a backup of node and stores data to a specified destination. Each task consists of parameters that define the source, destination, schedule, and other backup details. This section provides information about how to create a backup plan for Oracle database.

Follow these steps:

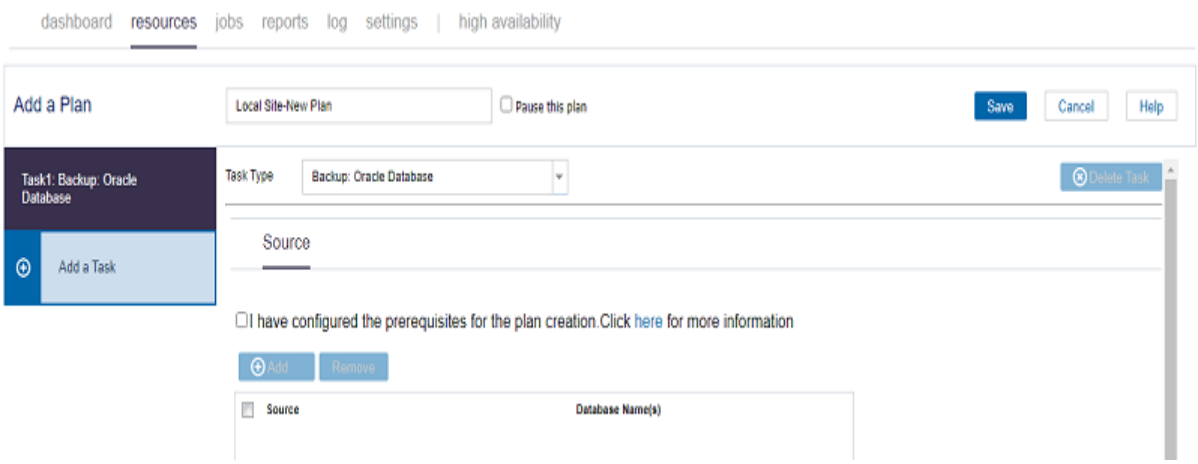
1. Log into Arcserve UDP Console.
2. Navigate to **resources** > **Plan** > **All Plans**.

   If you have created plans earlier, those plans are displayed on the center pane.
3. On the center pane, click **Add a Plan**.

   The Add a Plan page opens.
4. Enter a plan name.
5. (Optional) Select the **Pause this plan** check box.

   The plan will not run until you clear the check box to resume the plan.

   Note: If a plan is paused, then any in-progress job is not paused but all corresponding scheduled jobs associated with that plan are paused. However, you can manually run a job. For example, you can manually run backup job and replication job for a node even if the respective plan is paused. In such case, the following task to the on-demand (manual) job does not run. For example, there is a replication task after an on-demand backup job, the replication job does not run for the on-demand backup job. You need to manually run the replication job. When you resume the plan, the pending jobs do not resume immediately. After you resume the plan, the pending jobs run from the next scheduled time.
6. From the Task Type drop-down list, select **Backup: Oracle Database**.



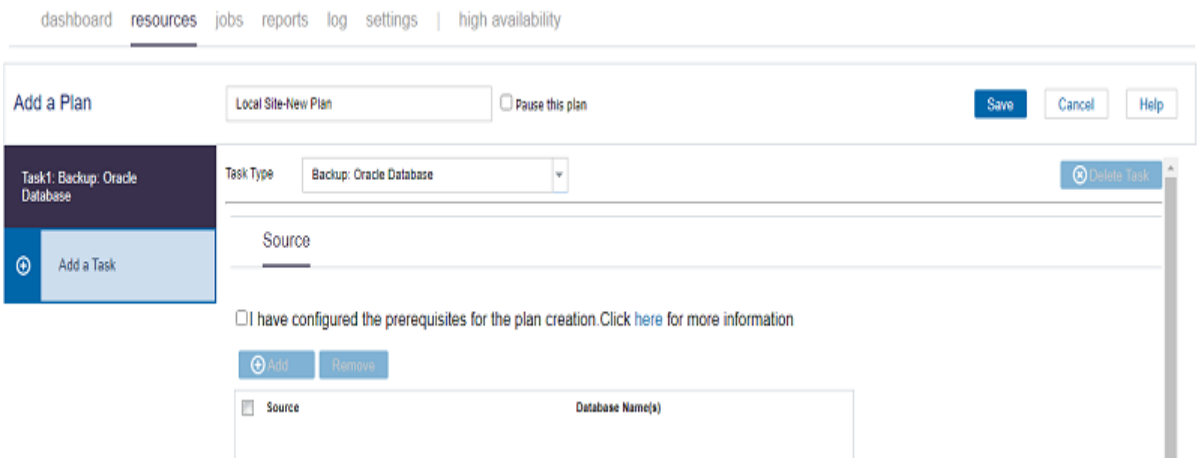
Now specify the Source, Destination, Schedule, and Advanced details.

Specify the Source

The Source page lets you specify the source nodes that you want to protect. You can select more than one node in a plan. If you have not added any nodes to the Console, you can add nodes when you create or modify a plan from the Source page. You can also save a plan without adding any source nodes. The plan gets deployed only after you add source nodes.

Follow these steps:

1. On the Source tab, to confirm the configurations of prerequisites, select the **I have configured the prerequisites for the plan creation** check box. For more information, see Review the Prerequisites and Considerations.



2. Click **+ Add**.
3. Select one of the following options:

   Select Sources to Protect in Arcserve UDP

   The Select Nodes to Protect dialog opens. Select the nodes from the displayed list. Select this option if you have already added the nodes to the Console.

   Add Windows Node

   The Add windows oracle node dialog opens. Select this option if you have not added the nodes and you want to manually add the nodes to protect.

   Add Linux Node

   The Add linux oracle node dialog opens. Select this option if you have not added the nodes and you want to manually add the nodes to protect.

4. If you select the **Select Sources to Protect in Arcserve UDP** option, on the Select Nodes to Protect dialog, do the following:

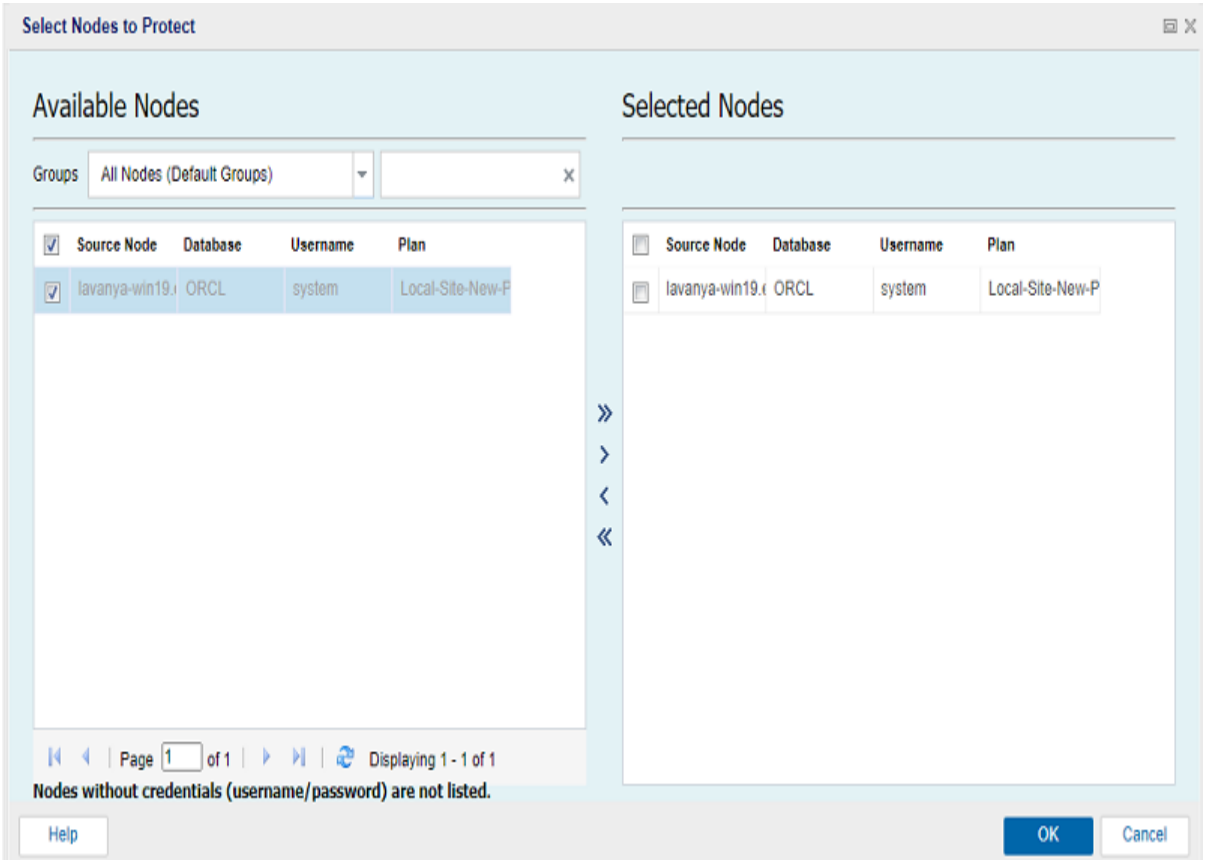
a. (Optional) Select a filter from the Groups drop-down list to filter nodes. You can enter keywords to further filter your nodes.

   The nodes are displayed on the Available Nodes column.

b. Select the nodes from the Available Nodes column, and then click the single right arrow (>) or double right arrow (>>) to move the selected nodes or all the nodes respectively to the Selected Nodes column.

   The nodes are displayed in the Selected Nodes column.

c. Click **OK** to close the dialog.



5. If you select the **Add Windows Node** option, on the Add windows oracle node dialog, do the following:

   a. Enter the node name or IP address of the Oracle node that you want to add.

   b. Type a description for the node.

   c. Enter the admin credentials.

   d. Click **Add**.

      Wait until the validation of username and password completes.

      The node is added on the center pane.

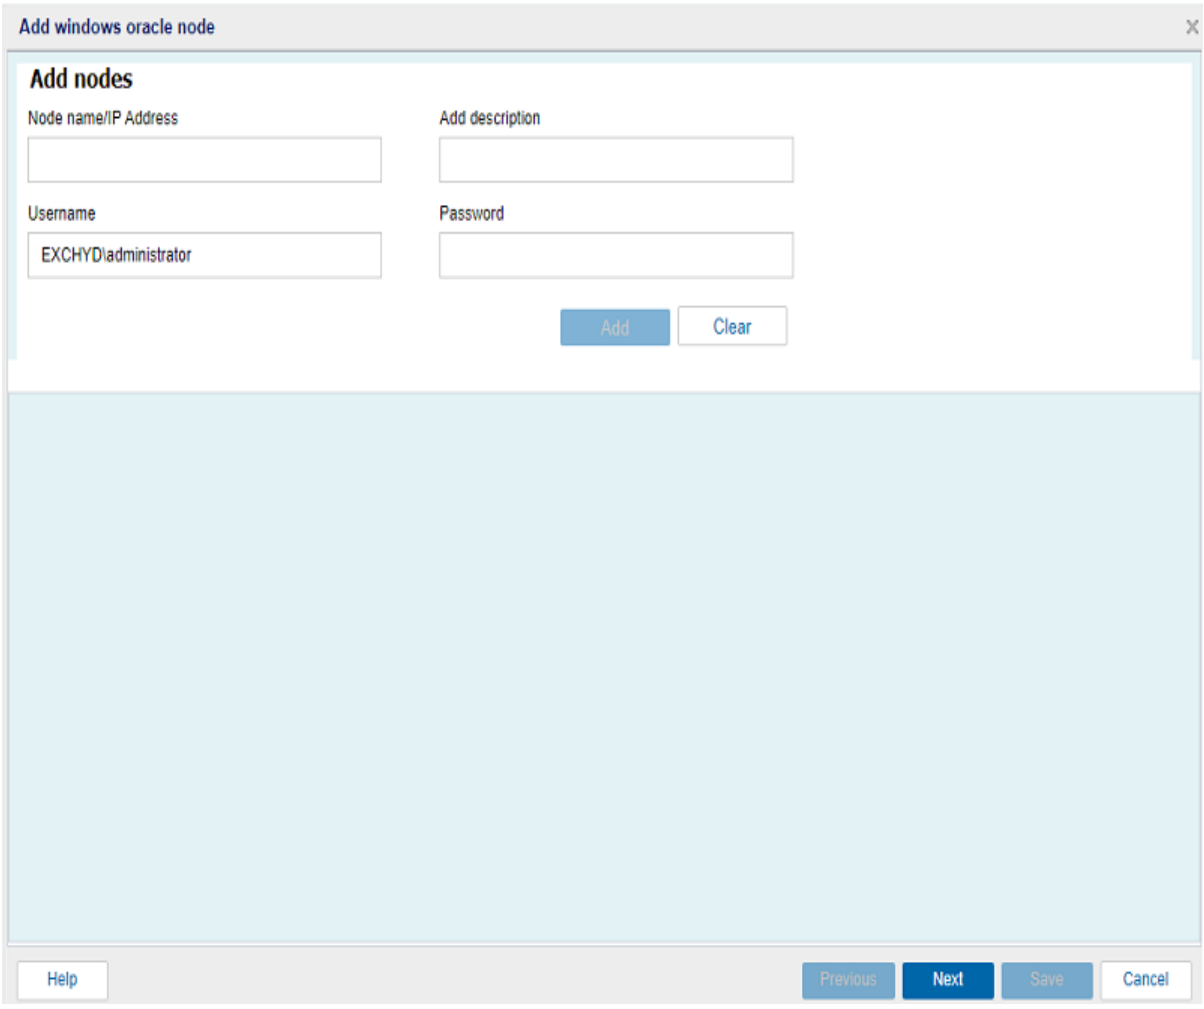
e. To add multiple nodes, follow the steps again. All the added nodes get listed on the center pane.

f. Click the expand symbol next to the IP address, and then verify the details that are added.

Notes:

- To edit the added nodes from the list, select the node that you want to edit, and then click **Edit**.
- To remove the added nodes from the list, select the nodes that you want to remove, and then click **Remove**.
- To remove all the added nodes, click **Remove all**.

g. Click **Next**.



6. If you select the **Add Linux Node** option, on the Add linux oracle node dialog, do the following:

a. Enter the node name or IP address of the Oracle node that you want to add.

b. Type a description for the node.

c. Enter the user name and password.

d. Enter the non-root user name and password.

e. Click **Add**.

   Wait until the validation of username and password completes.

   The node is added on the center pane.

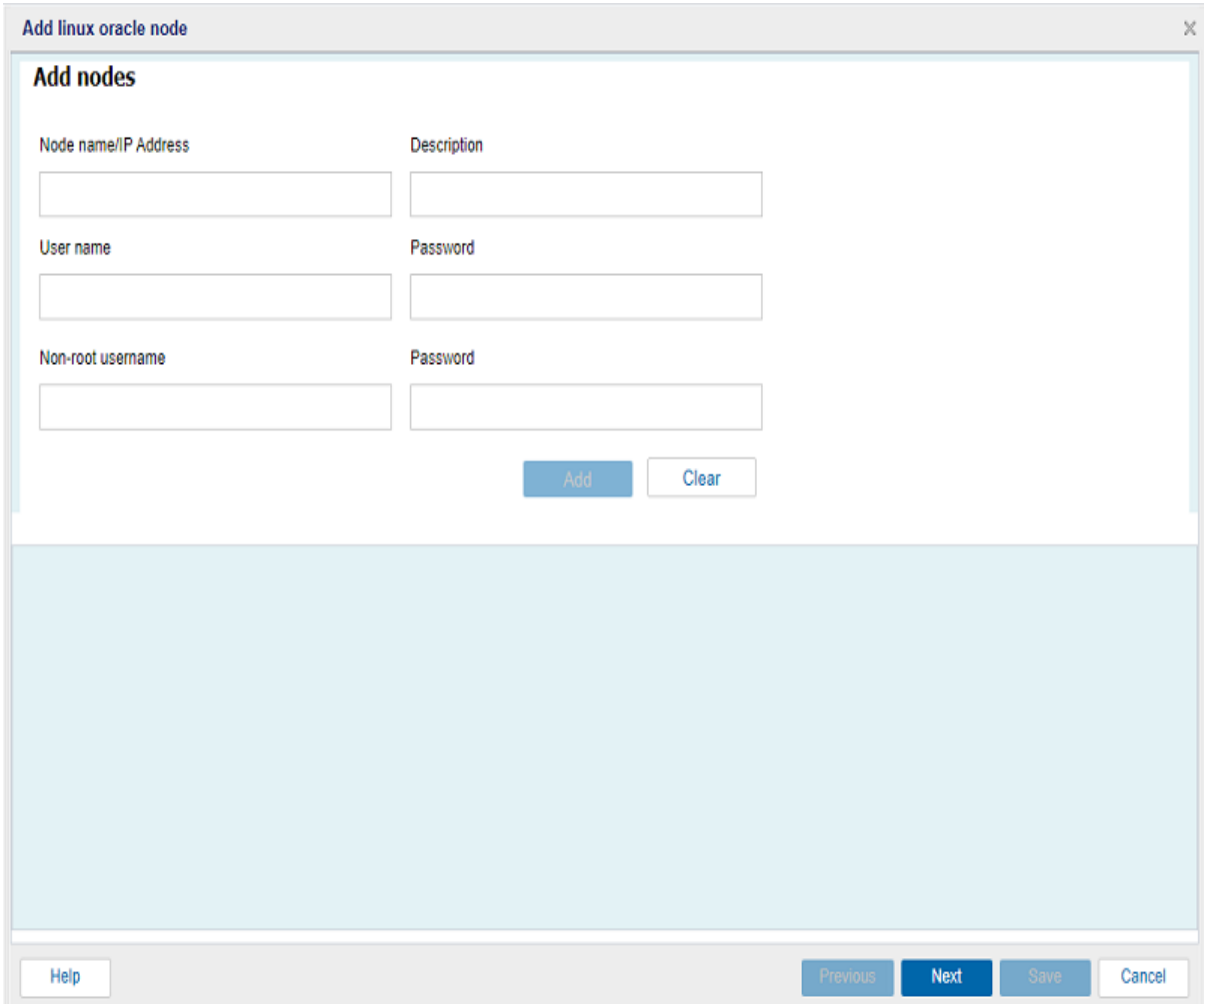
f. To add multiple nodes, follow the steps again. All the added nodes get listed on the center pane.

g. Click the expand symbol next to the IP address, and then verify the details that are added.

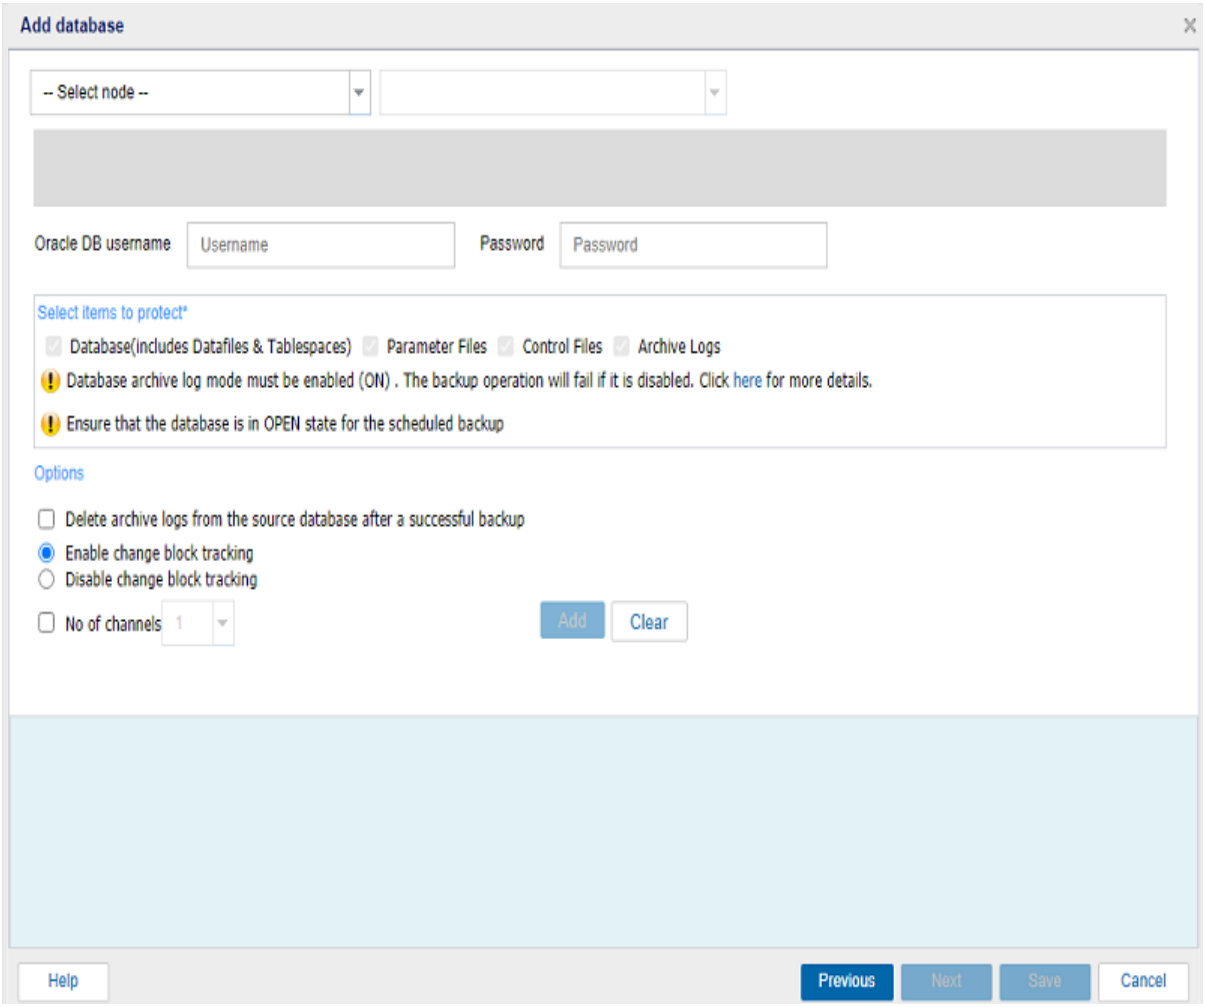
   Notes:

   - To edit the added nodes from the list, select the node that you want to edit, and then click **Edit**.
   - To remove the added nodes from the list, select the nodes that you want to remove, and then click **Remove**.
   - To remove all the added nodes, click **Remove all**.

h. Click **Next**.



7. On the Add database dialog, do the following:



a. Select the node from the drop-down list.

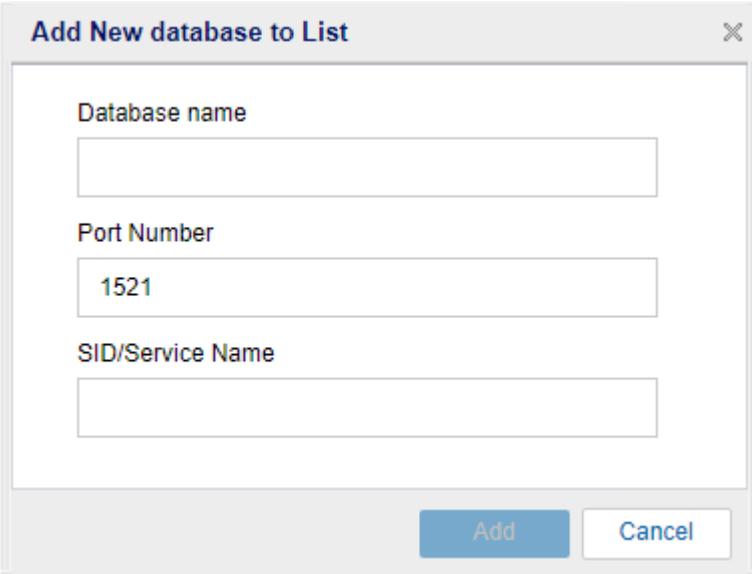
   Wait until the list of databases available on the selected host appears and the drop-down list gets enabled.

b. From the drop-down list, do one of the following:

   Add New Database

   To add a new database, follow these steps:

   1. Select the **Add New Database** option.
   2. On the Add New database to List dialog, do the following, and then click **Add**:
      - **Database Name:** Type the database name.
      - **Port Number:** Enter the port number. The default port value is 1521.
      - **SID/Service Name:** Enter the SID or service name.



The new database gets added in the drop-down list.

Select the existing database

Select the database from the drop-down list as needed. After the database is selected, it gets moved out of the list and appears in the center pane.

Notes:

- To avoid multiple inputs for each database, group the databases that you have configured with the same user name and password. For example, if you have databases configured with the same credentials and want to add the same file types to those databases, group the databases, and then add.
- To add different file types to the databases having different credentials, select the individual databases, and then add.
- To remove the database from the group (in the center pane), click the cross mark on the database name. After the database is removed, it gets added back to the drop-down list.

c. Enter the system user credentials for Oracle DB.

Notes:

- If you select the **Database** option, the Datafiles and Tablespace options are enabled by default.
- As you cannot skip the backing up of Control Files, the Control Files option is enabled by default.
- Make sure to enable the database ARCHIVELOG before the backup operation is triggered. To verify if the ARCHIVELOG log mode is

enabled, see Review the Prerequisites for Oracle Database. If the ARCHIVELOG is not enabled, the backup job fails.

d. For Options, specify the following:

Delete archive logs from the source database after a successful backup

Lets you delete the archive logs from the source database after a successful backup.

Enable change block tracking

Lets you enable the change block tracking feature. Use this option to enhance the performance of the RMAN. It helps the RMAN utility or the binary to find the differential data faster and perform the backup job.

Disable change block tracking

Lets you disable the change block tracking feature.

No of channels

Lets you select the number of channels or threads that you want to create for the backup job.

Note: RMAN supports a maximum of 255 channels or threads.

e. Click **Add**.

All the added databases get listed.

Note: After you click the Add button, all the database credentials and the firewall parameters are verified.

f. To add multiple databases, follow the steps again. All the added databases get listed.

g. Click the expand symbol next to the IP address, and then verify the details that are added.

Notes:

- To edit the added databases from the list, select the database that you want to edit, and then click **Edit**.
- To remove the added databases from the list, select the databases that you want to remove, and then click **Remove**.

8. Click **Save**.

Notes:

- You can add multiple nodes and multiple databases.
- To remove the source, select the check box next to the added source, and then click **Remove**.

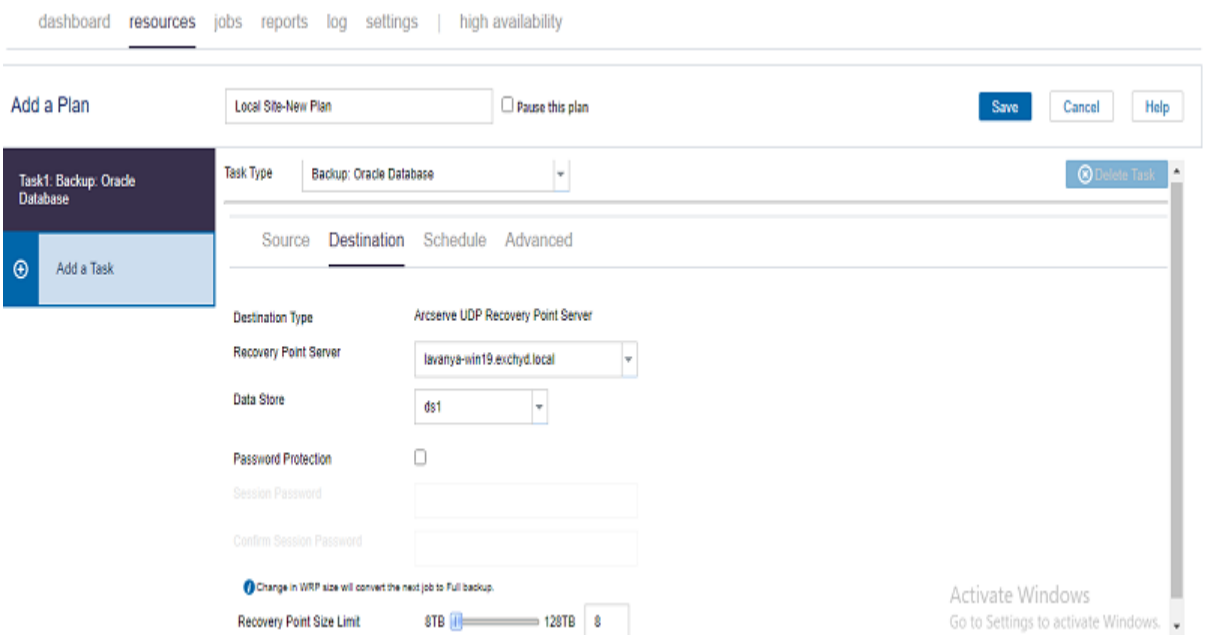
The source is specified.

Specify the Destination

The destination is a location where you store the backup data. You must at least specify the destination to save the plan.

Follow these steps:

1. Click the **Destination** tab.



2. Do the following:

   - **Recovery Point Server:** From the drop-down list, select the destination Recovery Point Server (RPS).
   - **Data Store:** From the drop-down list, select the destination data store.
   - **Password Protection:** Select the **Password Protection** check box, provide the session password, and then retype the session password to confirm.

     Note: The session password is optional when the backup destination is an unencrypted RPS data store.
   - **Recovery Point Size Limit:** Specifies the size of recovery point server for backup. Set the size of recovery point server on the scale.

     Range: 8 TB to 128 TB

   The destination is specified.

Specify the Schedule

The Schedule page lets you define a schedule for Backup, Merge, and Throttle functions to repeat at specific intervals. After you define a schedule, the jobs run automatically as per the schedule. You can add multiple schedules and can provide retention settings.

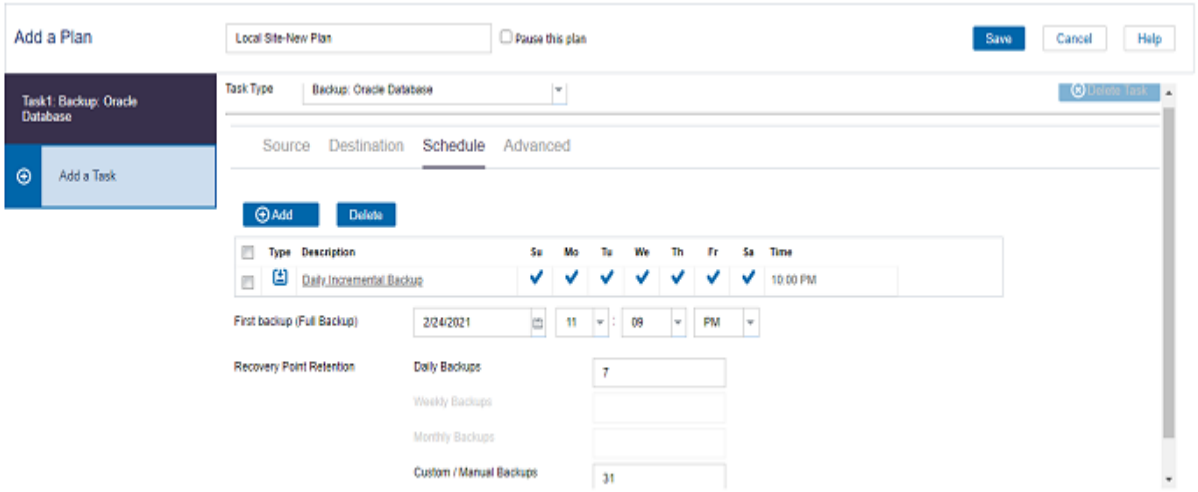
A Backup Schedule refers to regular schedule that is repeated multiple times a day based on the number of hours or minutes you select. Besides the regular schedule, a backup schedule also provides options to add daily, weekly, and monthly schedules.

If RPS server is the backup destination, the default value is a daily Incremental backup at every 10:00 PM.

You can edit or delete a backup job schedule.

Follow these steps:

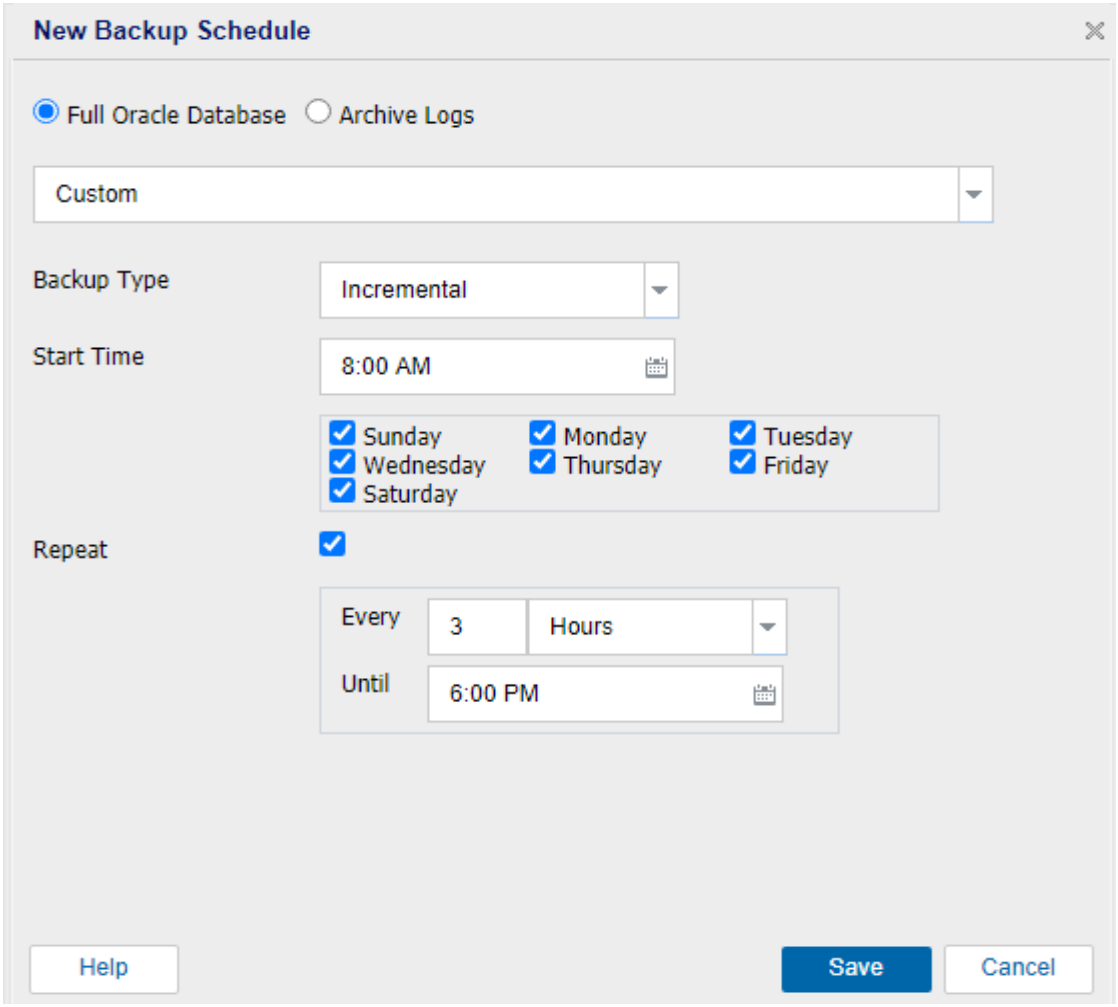
1. Click the **Schedule** tab, and then add backup, merge, and Disk Read throttle schedules.



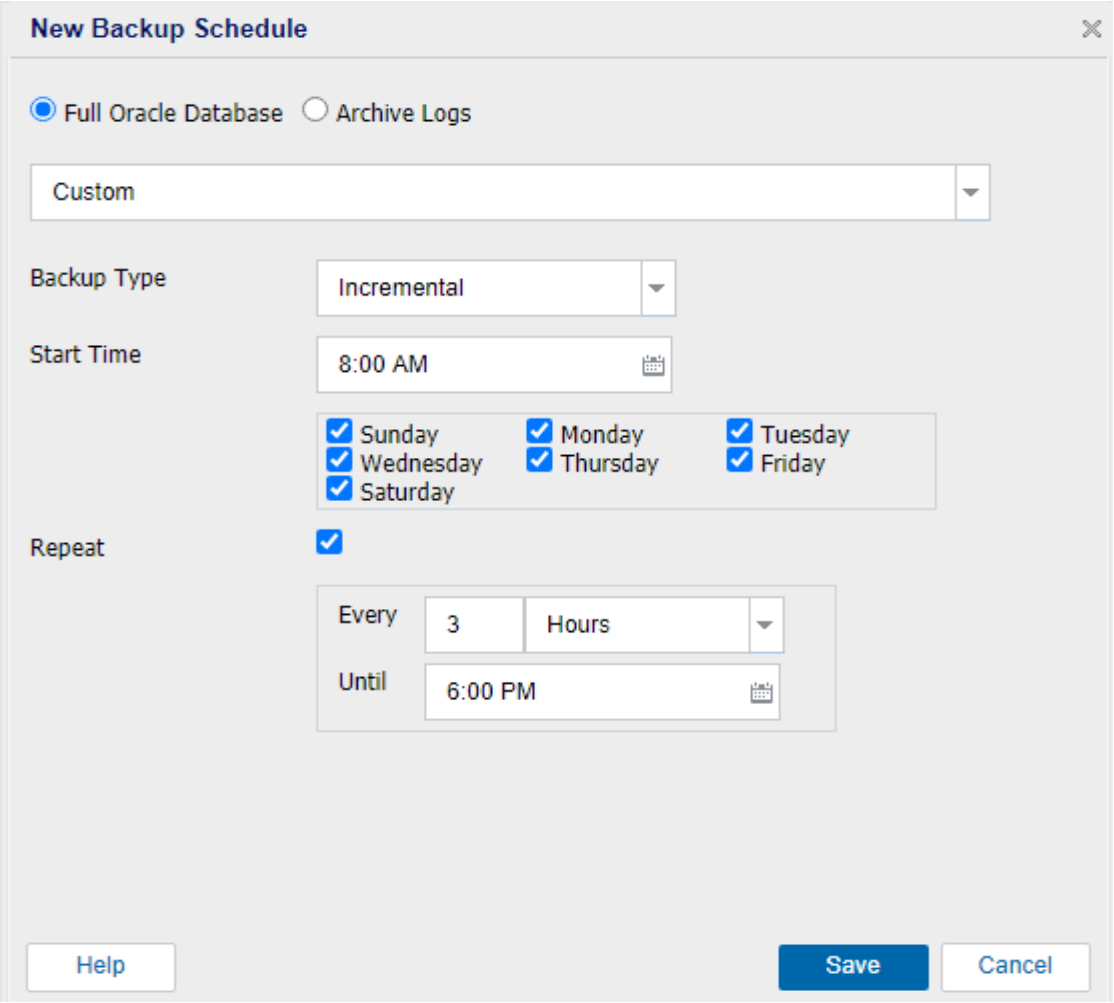
2. To add a backup schedule, perform the following steps:
   a. Click **Add**, and then select the **Add Backup Schedule** option.

      The New Backup Schedule dialog opens.

   b. On the New Backup Schedule dialog, select one of the following:



Full Oracle Database



Do the following:

a. From the drop-down list, select one of the following options:

Custom

Specifies the backup schedule that repeats multiple times a day.

Daily

Specifies the backup schedule that occurs once a day. By default, all the days of the week are selected for Daily backup. If you do not want to run the backup job on a specific day, clear the check box for that day of the week.

Weekly

Specifies the backup schedule that occurs once a week.

Monthly

Specifies the backup schedule that occurs once a month.

Hourly

Specifies the backup schedule that occurs once per hour.

Minute

Specifies the backup schedule that occurs once in a minute.

b. From the Backup Type drop-down list, select the backup type.

Full

Determines the backup schedule for Full Backups. As scheduled, Arcserve UDP performs a full backup of all used blocks from the source machine. A full backup typically consumes time depending on the backup size.

Incremental

Determines the backup schedule for Incremental Backups.

As scheduled, Arcserve UDP incrementally backs up only those blocks that have changed since the last successful backup. The advantages of Incremental Backups are that it is a fast backup and it produces a small backup image. This is the most optimal way to perform a backup.

c. Specify the backup start time.

d. (Optional) Select the **Repeat** check box and specify the repeat schedule.

Note: Make sure the time to repeat the schedule must not be less than 15 minutes.

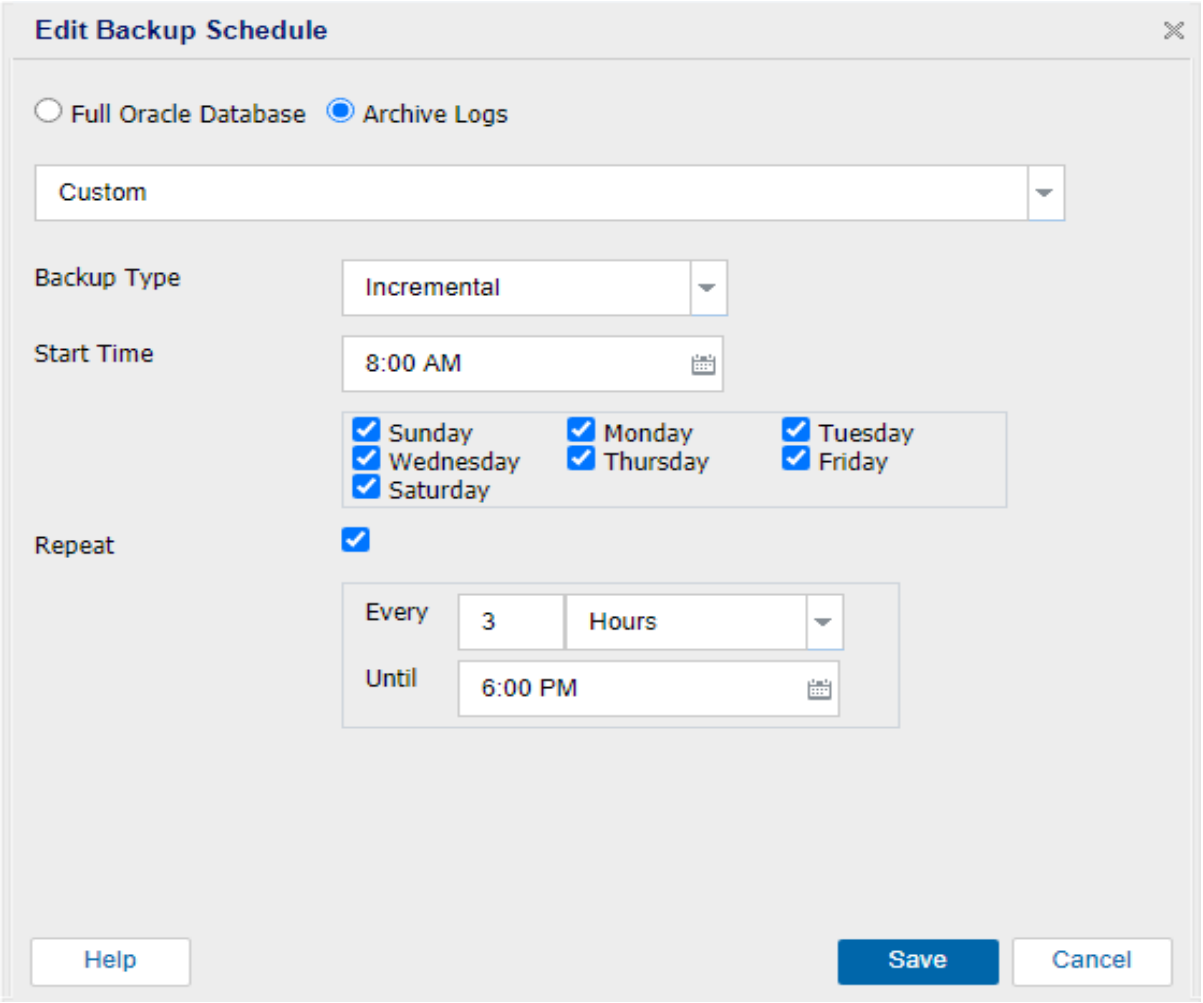
e. If you have selected the **Hourly** option from the drop-down list, do the following:

1. Select the backup type from the Backup Type drop-down list.
2. Specify the backup start time.
3. Specify the backup end time.
4. Select the backup repeat time from the drop-down list.

   Note: If you want to customize the backup repeat time, select Custom from the drop-down list, and then enter the time in hours.

f. Click **Save**.

Archive Logs



Do the following:

a. From the drop-down list, select one of the following options:

Custom

Specifies the backup schedule that repeats multiple times a day.

Daily

Specifies the backup schedule that occurs once a day. By default, all the days of the week are selected for Daily backup. If you do not want to run the backup job on a specific day, clear the check box for that day of the week.

Weekly

Specifies the backup schedule that occurs once a week.

Monthly

Specifies the backup schedule that occurs once a month.

Hourly

Specifies the backup schedule that occurs once per hour.

Minute

Specifies the backup schedule that occurs once in a minute.

b. From the Backup Type drop-down list, select the backup type.

Full

Determines the backup schedule for Full Backups. As scheduled, Arcserve UDP performs a full backup of all used blocks from the source machine. A full backup typically consumes time depending on the backup size.

Incremental

Determines the backup schedule for Incremental Backups.

As scheduled, Arcserve UDP incrementally backs up only those blocks that have changed since the last successful backup. The advantages of Incremental Backups are that it is a fast backup and it produces a small backup image. This is the most optimal way to perform a backup.

c. Specify the backup start time.

d. (Optional) Select the **Repeat** check box and specify the repeat schedule.

Note: Make sure the time to repeat the schedule must not be less than 15 minutes.

e. If you have selected the **Hourly** option from the drop-down list, do the following:

1. Select the backup type from the Backup Type drop-down list.
2. Specify the backup start time.
3. Specify the backup end time.
4. Select the backup repeat time from the drop-down list.

   Note: If you want to customize the backup repeat time, select **Custom** from the drop-down list, and then enter the time in hours.

f. Click **Save**.

Note: We recommend using the following schedules when you configure the Archive logs and Full database under Daily/Weekly/Monthly/Custom:

- **Daily Incremental:** Performs the first backup as full followed by infinite daily incremental backups of the full database.

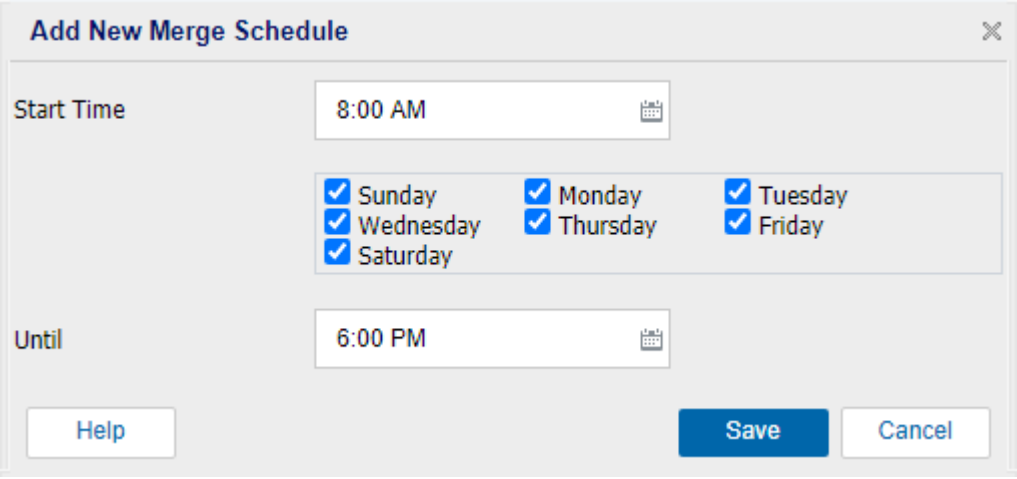
- **Daily Incremental with Custom Schedule for Archive Log (in hours/minutes):** Performs the first backup as full followed by an infinite daily incremental backup of full database and custom archive log backups based on a specified interval.
- **Daily Incremental with Weekly Full and Custom Schedule for Archive Log (in hours/minutes):** Performs the first backup as full followed by daily incremental backup, a weekly full backup of the full database, and custom archive log backups based on a specified interval.
- **Daily Incremental with Monthly Full and Custom Schedule for Archive Log (in hours/minutes):** Performs the first backup as full followed by daily incremental backup, a monthly full backup of the full database, and custom archive log backups based on a specified interval.

The Backup Schedule is specified and appears on the Schedule page.

3. To add a merge schedule, do the following:
   a. Click **Add**, and then select the **Add Merge Schedule** option.

      The Add New Merge Schedule dialog opens.

   b. On the Add New Merge Schedule dialog, do the following:
      - **Start Time:** Specify the start time to start the merge job.
      - **Until:** Specify the end time for the merge job.



   c. Click **Save**.

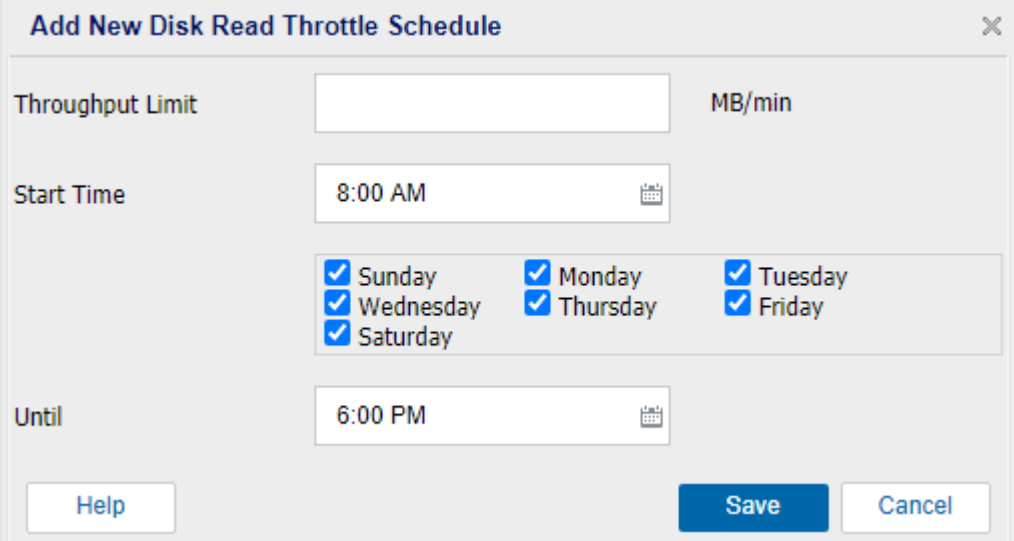
The Merge Schedule is specified and appears on the Schedule page.

4. To add a disk read throttle schedule, do the following:
   a. Click **Add**, and then select the **Add Disk Read Throttle Schedule** option.

      The Add New Disk Read Throttle Schedule dialog opens.

b. On the Add New Disk Read Throttle Schedule dialog, do the following:

- **Throughput limit:** Specify the throughput limit in MB per minute unit.
- **Start Time:** Specify the start time to start the backup throughput job.
- **Until:** Specify the end time for the backup throughput job.



c. Click **Save**.

The Disk Read Throttle Schedule is specified and appears on the Schedule page.

5. Specify the start time for the scheduled backup.



6. Specify the recovery points retention settings for Custom, Daily, Weekly, and Monthly schedule.

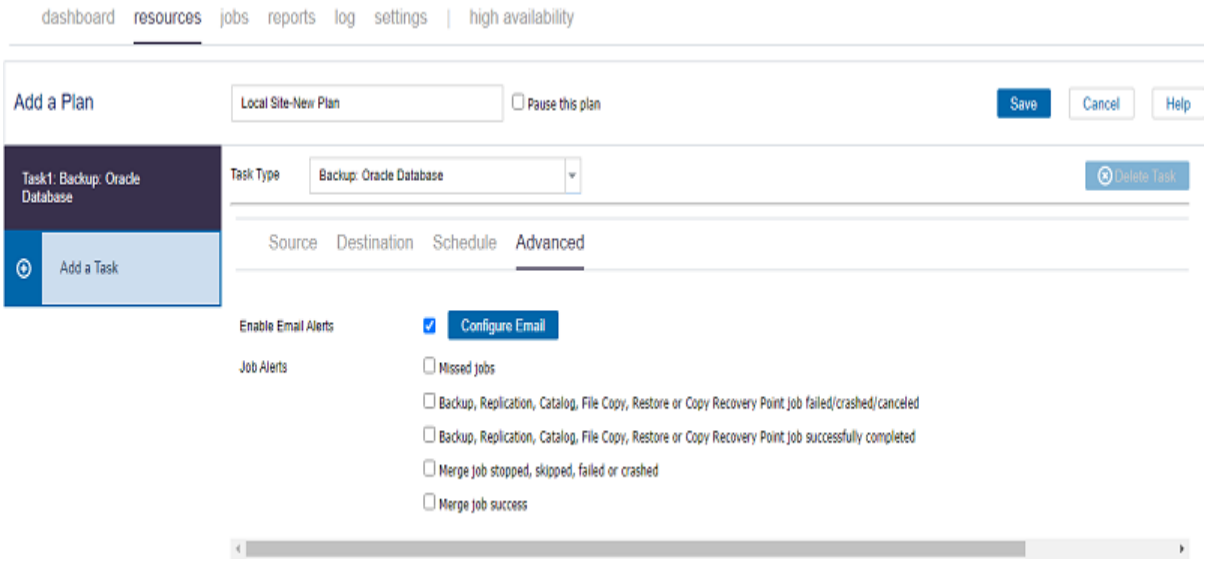
   These options are enabled if you have added the corresponding backup schedule. If you modify the retention settings on this page, the changes are reflected on the Backup Schedule dialog.



The schedule is specified.

Specify the Advanced Settings

The Advanced tab lets you specify some advanced settings for the backup job.



Follow these steps:

1. Specify the following details:

   Enable Email Alerts

   Lets you enable email alerts. You can configure email settings and can specify the types of alerts that you want to receive in an email. When you select this option, the following options are enabled for your selection.

   Configure Email

   Lets you configure the email settings.

   To configure the email server and proxy server details, click **Configure Email**.

   Job Alerts

   Lets you specify the types of job emails that you want to receive.

2. Click **Save**.

The Advanced settings are specified.

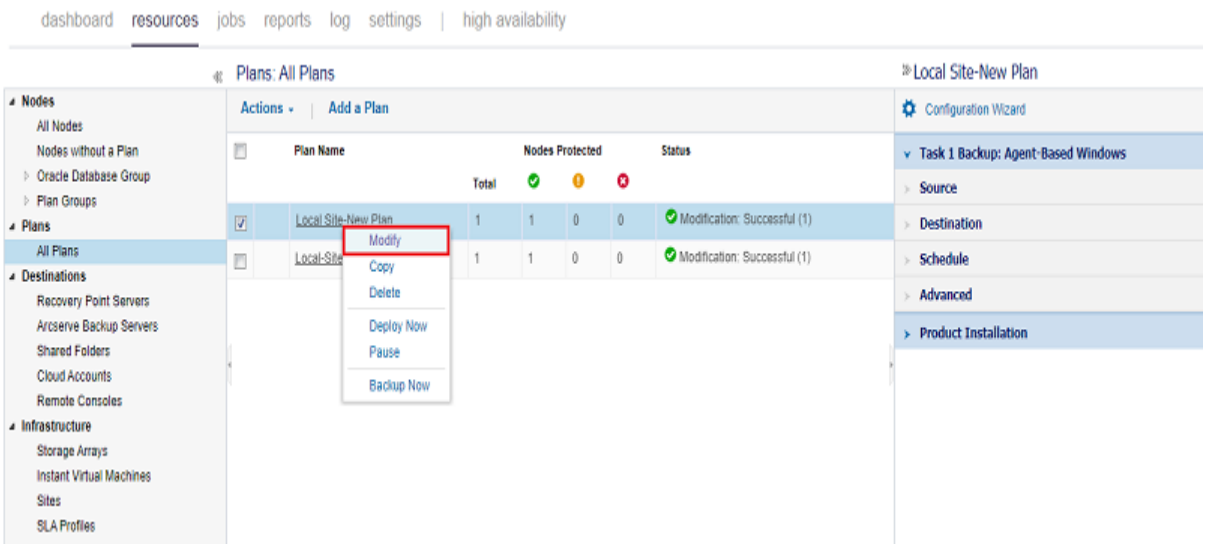
Add an Assured Recovery Task to the Plan

An assured recovery task includes an assured recovery task to an existing backup/replicate plan. Each task consists of parameters that define the source, test settings, schedule, and advanced settings. Create an Assured Recovery task based on backup/replication task to verify accessibility and assured recovery of the data and provide the data integrity check.

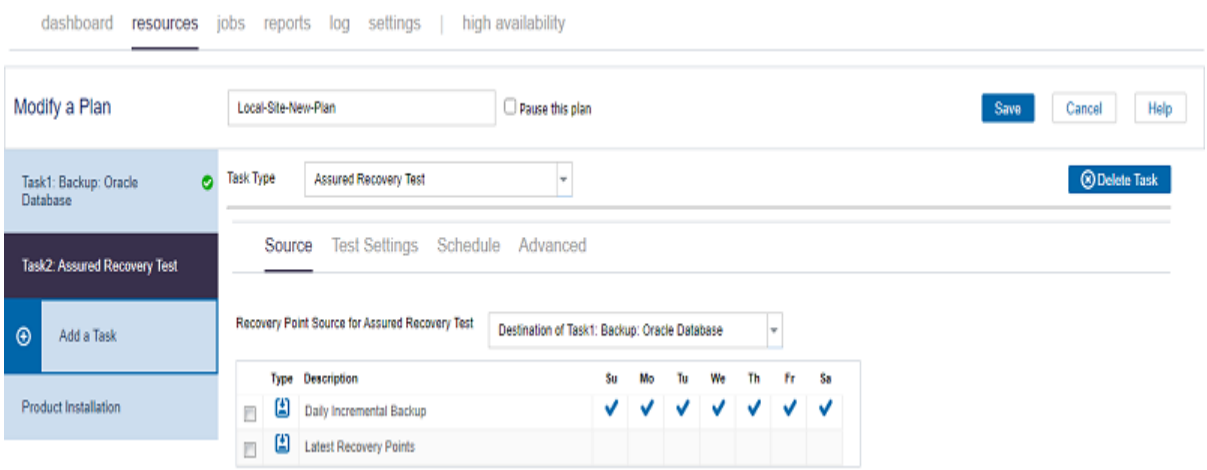
Follow these steps:

1. Click the **resources** tab on the Console.
2. From the left pane, navigate to **Plans**, and then click **All Plans**.

   If you have created plans earlier, those plans are displayed on the center pane.
3. From the center pane, right-click the existing backup plan, and then click **Modify**.



4. On the Modify a Plan page, click **Add a Task** from the left pane.



5. (Optional) Select the **Pause this plan** check box.

   The plan will not run until you clear the check box to resume the plan.

   Note: If you pause the plan, the Assured Recovery job will not start. When you resume the plan again, the Assured Recovery job is not resumed automatically. You have to manually run another backup job to start the Assured Recovery job.

6. From the Task Type drop-down list, select the **Assured Recovery Test** option.
7. Now specify the Source, Test settings, Schedule, and Advanced details.
8. Click **Save**

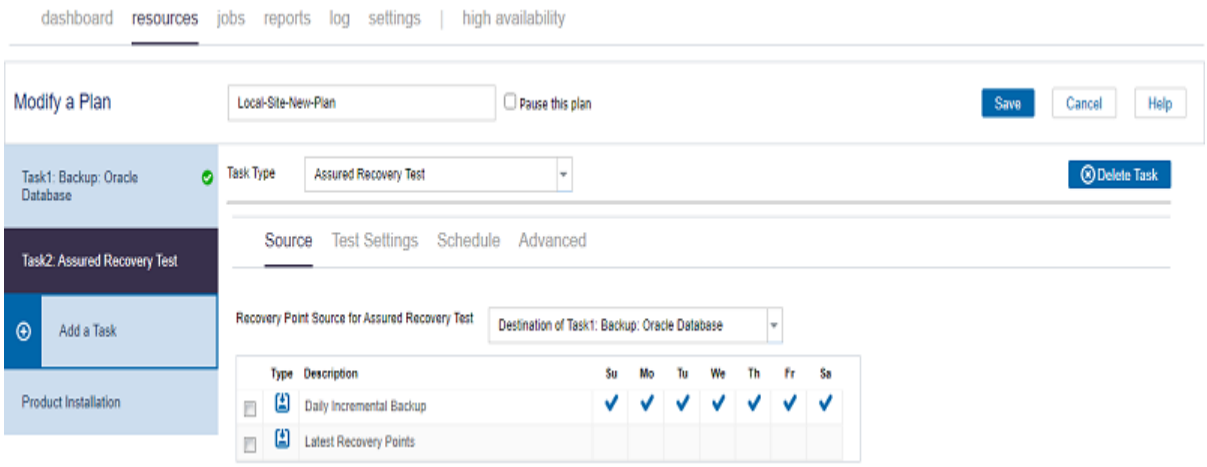
   The changes are saved and the Assured Recovery task is automatically deployed.

Specify the Source

The Source tab lets you specify the source recovery point that you want to protect.

Follow these steps:

1. On the Source tab, from the Recovery Point Source for Assured Recovery Test drop-down list, select the desired option.



2. Select Types of Recovery Points for Assured Recovery test. You can select any of the following two options:

   - Assured Recovery Test from selected backup types(s)

     You can select either Daily Backups, Weekly Backups, Monthly Backups. The daily, weekly, or monthly options are enabled depending on the schedule of the source backup. For example, if the source backup has only monthly backup scheduled, then the Assured Recovery test from selected backup types(s) option has only Monthly Backups enabled.

   - Assured Recovery Test from the Latest Recovery Points

   The source is specified

   Next Steps:

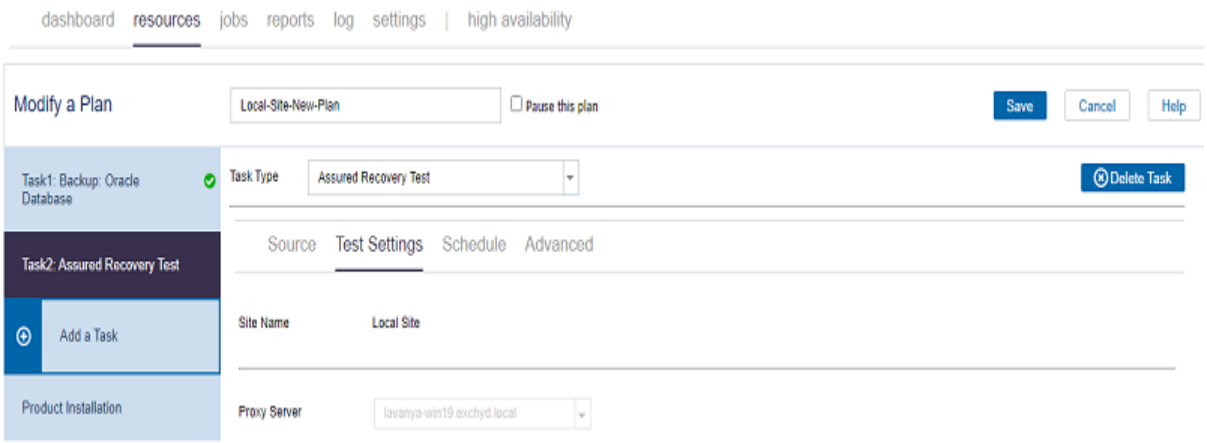
   - Specify the Test Settings
   - Specify Assured Recovery Test Job Schedule
   - Specify the Advanced Settings

Specify the Test Settings

The Test Settings tab lets you specify the proxy server.

Follow these steps:

1. Click the **Test Settings** tab.



2. From the Proxy Server drop-down list, select the proxy node.

   Note: For Oracle, Recovery Point Server itself acts as the proxy server in the Assured Recovery test.

   The Test Settings are specified.

Next Steps:

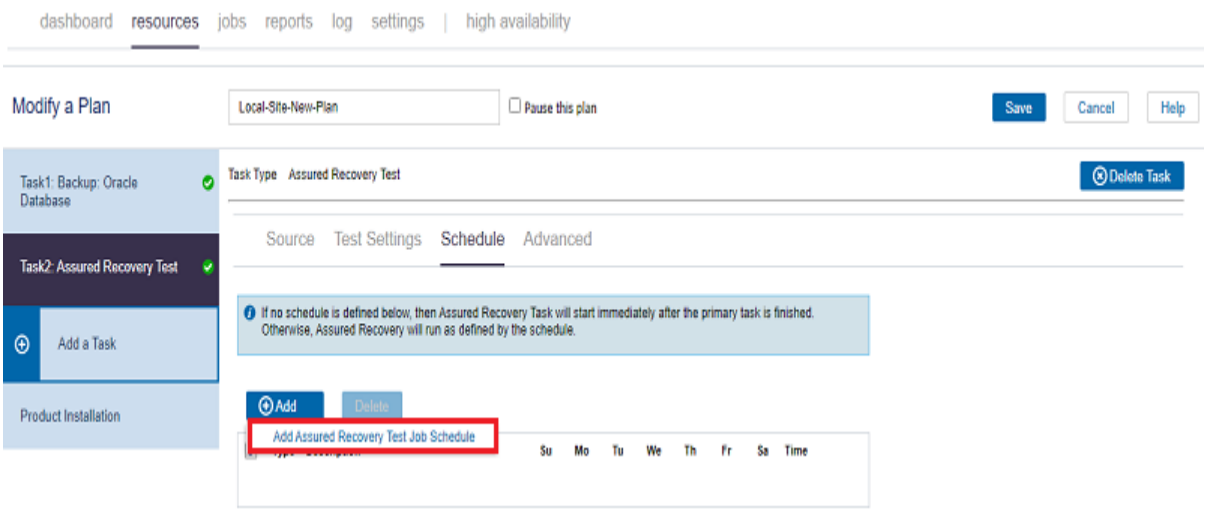
- Specify Assured Recovery Test Job Schedule
- Specify the Advanced Settings

Specify Assured Recovery Test Job Schedule

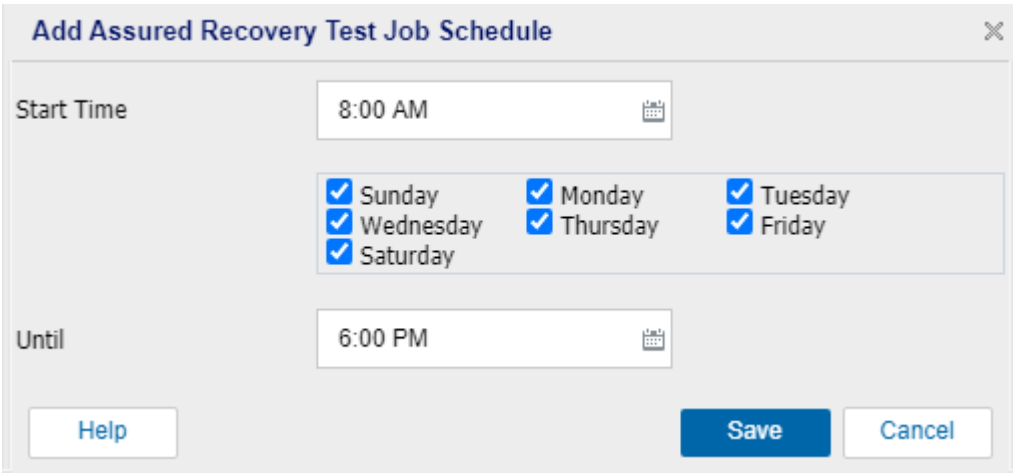
The schedule tab lets you specify the assured recovery test job schedule. If you do not specify a schedule, then the task starts immediately when you complete the primary task.

Follow these steps:

1. Click the **Schedule** tab.



2. Click **Add**, and then select **Add Assured Recovery Test Job Schedule**.
3. On the Add Assured Recovery Test Job Schedule dialog, do the following, and then click **Save**:
   - **Start Time:** Specify the start time to start the Assured Recovery test job.
   - **Until:** Specify the end time for the Assured Recovery test job.



The Assured Recovery test job is scheduled and appears on the Schedule page.

Notes:

- If Assured Recovery source is Latest Recovery Points + Enable Assured Recovery schedule, the default value is 1. As a result, by default the latest one is verified.

  For value <=0, all unverified recovery points are tested.

- If Assured Recovery source is Daily/Weekly/Monthly + Enable Assured Recovery schedule, the default value is 9 (7 daily + 1 weekly + 1 monthly).

  For value <=0, all unverified recovery points are tested.

- If you do not want to use the default value, you can add a DWORD 32bit registry key at the following location:

  HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AssuredRecovery\MaxNumberOfRecoveryPoint

Next Step:

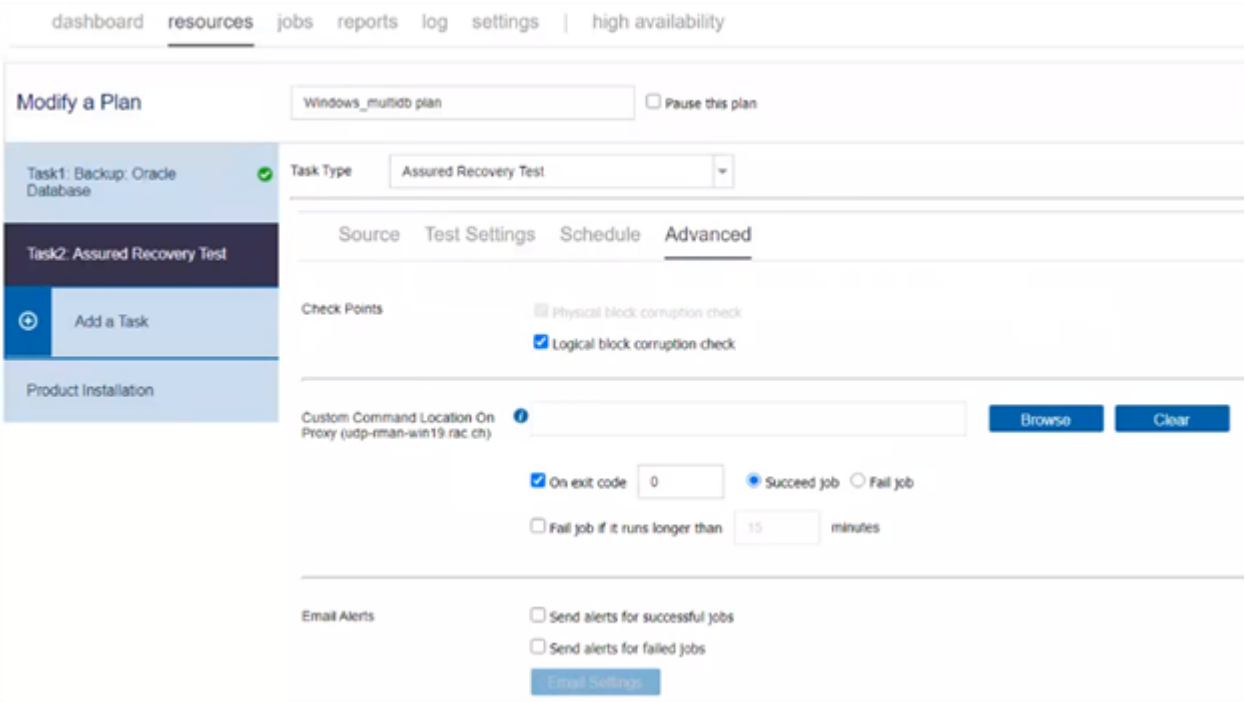
- Specify the Advanced Settings

Specify the Advanced Settings

The Advanced tab helps you apply other settings required to complete the plan. This is optional, if you do not need to add further settings.

Follow these steps:

1. Click the **Advanced** tab.



2. Do the following:

   Check Points

   Physical block corruption check

   RMAN detects only physically corrupt blocks with every backup it takes and stores the corruption information in the Control File and alert.log file. By default, the **Physical block corruption check** option is enabled.

   Logical block corruption check

   RMAN does not automatically detect logically corrupt blocks by default. Select the **Logical block corruption check** check box to detect the logical corruption along with the physical corruption. If logical corruption is found, RMAN stores the corruption information in the alert.log file.

   Custom Command Location on Proxy

Note: This options is not visible if you have selected Agent-Based Linux backup/replication as Recovery Point Source of Assured Recovery Test.

Specifies the custom script that is stored on the proxy server. You can browse the volume information of the proxy server. If IVM assured recovery type is selected, the script is copied into VM and executed when VM boots. If IVHD assured recovery type is selected, the script is executed on the proxy server.

Note: The Assured Recovery job can only support the executable Windows bat script or executable binary. But, this limitation does not restrict your usage of Assured Recovery job. You can use a Windows batch script written to start your powershell script, such as the name test.bat with the following content:

Powershell.exe -NoProfile -ExecutionPolicy ByPass -Command "& '%ScriptPath%\AR-check.ps1'"

Additionally, the following arguments are passed to the script as per your requirement.

- set PlanName=%1%

- set NodeName=%2%

- set ProxyServer=%3%

- set RecoveryPointName=%3%

- set MountPointRootPathName=%5% (Applicable only to Instant Virtual Disk test type)

As a result, the bat is invoked and run as below:

For Instant Virtual Machine test type:

"%ScriptPath%\test.bat PlanName NodeName ProxyServer RecoveryPointName

For Instant Virtual Disk test type:

%ScriptPath%\test.bat PlanName NodeName ProxyServer RecoveryPointName MountPointRootPathName

On exit code

Specifies the exit code for Succeed Job or Fail Job.

Succeed Job

Specifies that the assured recovery job is set to succeed when the script returns the exit code.

Fail Job

Specifies that the assured recovery job is set to fail when the script returns the exit code.

Fail Job if it runs longer than

Specifies if the script runs longer than the specified times.

Default value: 15 Minutes.

Email Alerts

Lets you enable email alerts. You can configure email settings and specify the type of alerts that you want to receive in an email.

- If you want to receive email alerts for assured recovery job success, select the **Send alerts for successful jobs** check box.
- If you want to receive email alerts for assured recovery job failure, select the **Send alerts for failed jobs** check box.

  Note: Unavailable for jobs running with Linux Backup Server.

Email Settings

Lets you configure the email settings.

To configure the email server and proxy server details, click **Email Settings**. For more information about how to configure Email Settings, see Configure Email and Alert.

The Advanced settings are specified.

(Optional) Perform a Manual Backup

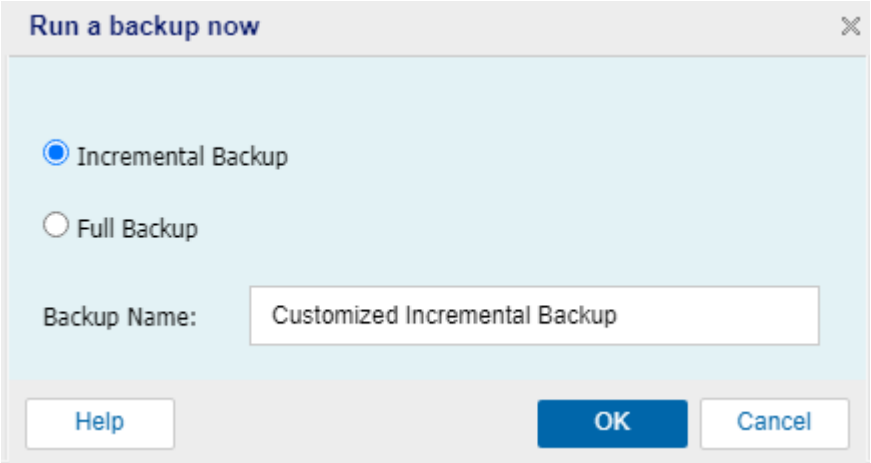
Typically, backups are performed automatically and are controlled by the schedule settings. In addition to the scheduled backup, a manual backup provides you the option to back up your nodes on a need basis. For example, if you have a repeat schedule for Full and Incremental backups and you want to make major changes to your machine, you must perform an immediate manual backup without waiting for the next scheduled backup to occur.

Follow these steps:

1. Log into Arcserve UDP Console.
2. Navigate to **resources** > **Plan** > **All Plans**.

   The list of all plans is displayed on the center pane.
3. Right-click the plan, and then select **Backup Now**.

   The Run a backup now dialog opens.
4. On the Run a backup now dialog, select the type of backup you want to perform, and then specify the backup name if required.



The available options are:

Incremental Backup

Initiates an Incremental Backup of your machine. An Incremental Backup backs up only those blocks that have changed after the previous backup.

Note: If you are performing a backup for the first time, the full backup gets initiated automatically even though you selected the Incremental Backup option.

The advantages of Incremental Backups are that it is a fast backup and produces a small backup image. This is the most optimal way to perform backups.

Full Backup

Initiates a Full Backup of your entire machine or the selected volumes.

Notes:

- If you add a new volume to the backup source, the newly added volume is fully backed up regardless of the overall backup method selected.
- If no backup name is specified, by default, it is automatically named as Customized Full or Incremental Backup.

5. Click **OK**.

   The Progress Information dialog appears.

6. Wait until the backup job completes, and then click **OK**.

7. To monitor the backup jobs, navigate to **jobs** > **Jobs in Progress**.

   Wait until the backup job is finished.

   Note: Only one job can be run at a time. If you manually attempt to launch a backup job when another job is currently running, an alert message informs you that another job is running and requests that you try again at a later time.

   The manual backup is successfully performed.

Verify the Backup Job

To verify your backup job, confirm that you have successfully created the backup plan. After you verify that the plan is created successfully, confirm whether the backup job is running as scheduled. You can verify the status of backup jobs from the **Jobs** tab.

Follow these steps to verify plans:

1. Click the **resources** tab.
2. From the left pane, navigate to **Nodes**, and then click **All Nodes**.

   A list of all nodes is displayed on the center pane.
3. Verify that plans are mapped with nodes.

Follow one of these methods to verify backup jobs:

Method 1

1. Click the **jobs** tab.
2. From the left pane, click **All Jobs Completed**.

   The status of each job is listed on the center pane.
3. Verify that the backup job is successful.

   The backup job is verified.

Method 2

1. Click the **jobs** tab.
2. From the left pane, click **Jobs Successfully Completed**.

   The list of backup jobs completed appear on the center pane.
3. Click the job you want to verify.

   The right pane refreshes.
4. Under Job Details, click the **View Logs** hyperlink.
5. Verify that the backup job is successful.

   The backup job is verified.

How to Restore an Oracle Database using RMAN

You can restore the entire Oracle Database or individual files such as Tablespaces, Data files, Tables, Parameter files, Archive logs, and Control files from agent nodes using RMAN. You can also restore database from recovery point to the Oracle base location or an alternative location.

To restore Oracle Database, follow these steps:

1. Select the Oracle Database Items to Restore
2. Define Restore Options

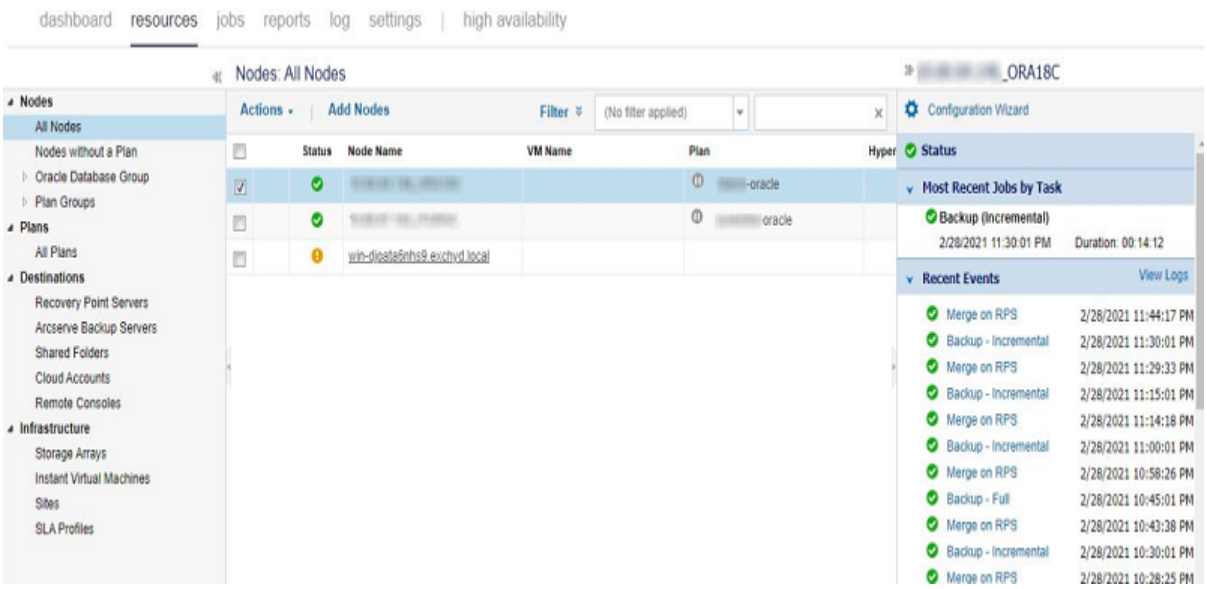
Select the Oracle Database Items to Restore

This section provides information about how to browse and select Oracle Database items to restore.

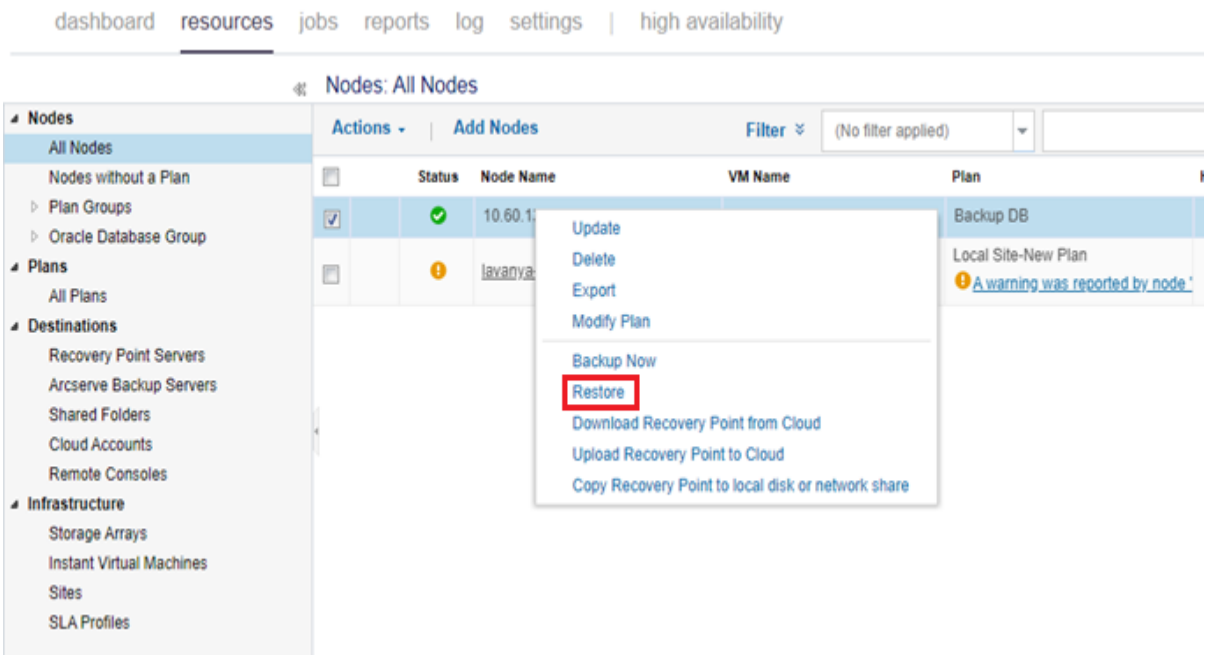
Follow these steps:

1. Log into the UDP console account.
2. On the UDP console page, click the **resources** tab.
3. From the left pane, navigate to **Nodes** > **All Nodes**.

   The list of existing nodes appears on the center pane.



4. Right-click the Oracle node, and then select the **Restore** option.

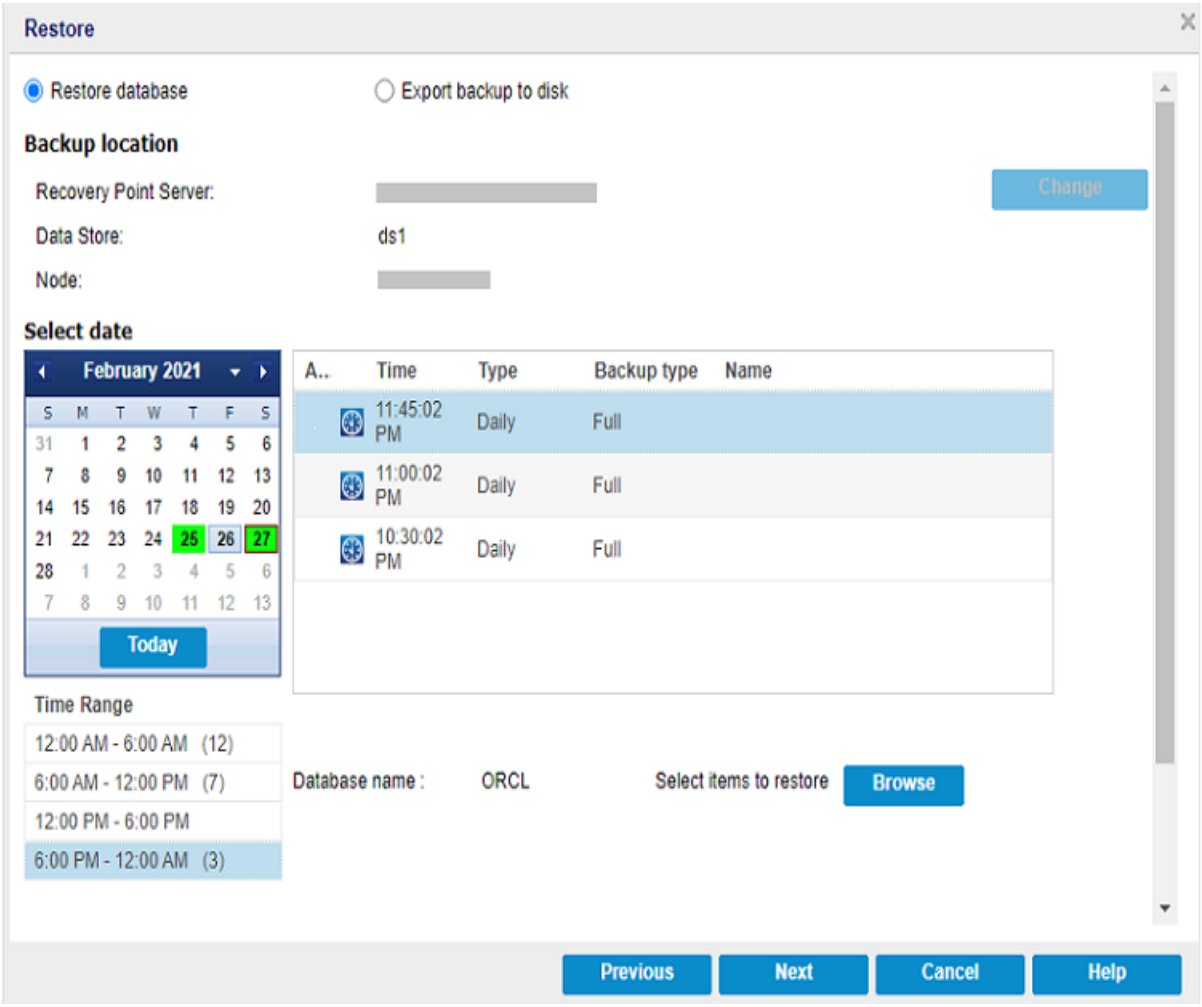


You are automatically logged into the agent node and the Restore wizard appears.

5. On the Restore wizard, select one of the following:

   Restore database

   Restore database includes steps for both Full database restore and Granular restore.



To select the Oracle database items, follow these steps:

a. (Optional) To modify the details of the backup location, click **Change**.

b. From the Select date section, select a date.

   Based on selection, the list of backup recovery points is displayed with the time of the backup, the type of backup that was performed (Full or Incremental), and the name of the backup.

c. From the list of backup recovery points, select a backup recovery point you want to restore.

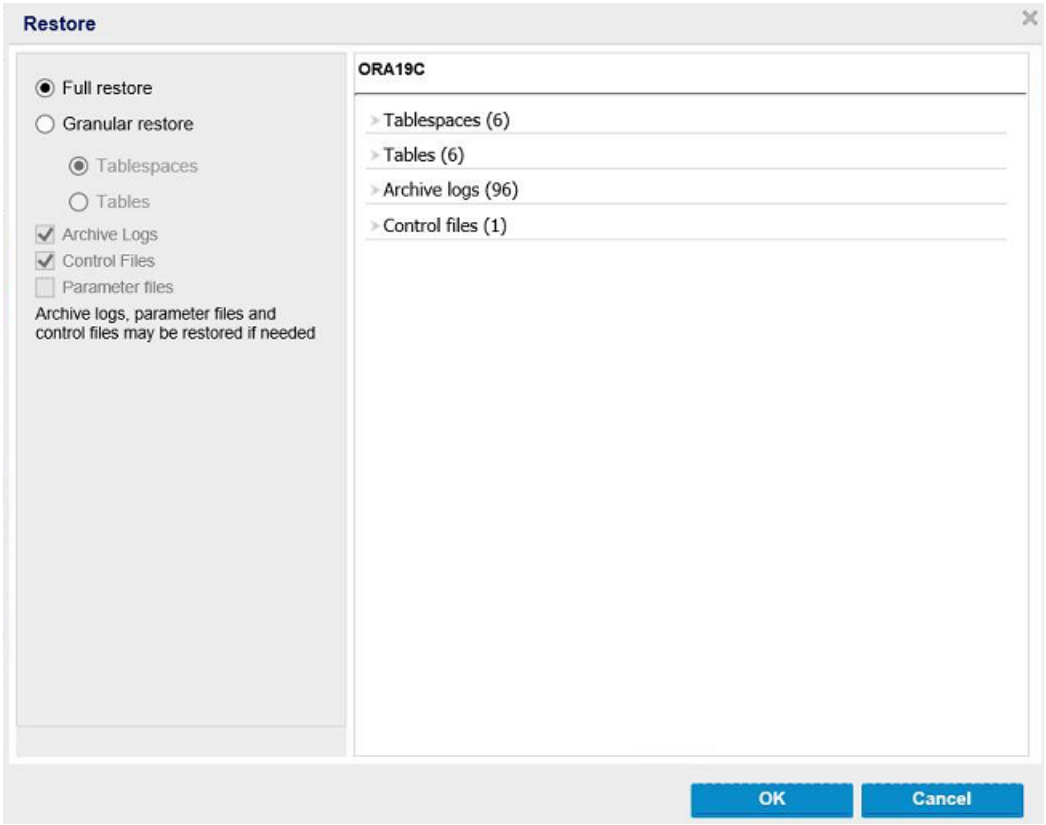
   Note: By default, the latest backup recovery point is selected.

d. To select the database items, follow these steps:

   - Click the **Browse** button.

     The Restore dialog appears.

   - On the Restore dialog, do the following:

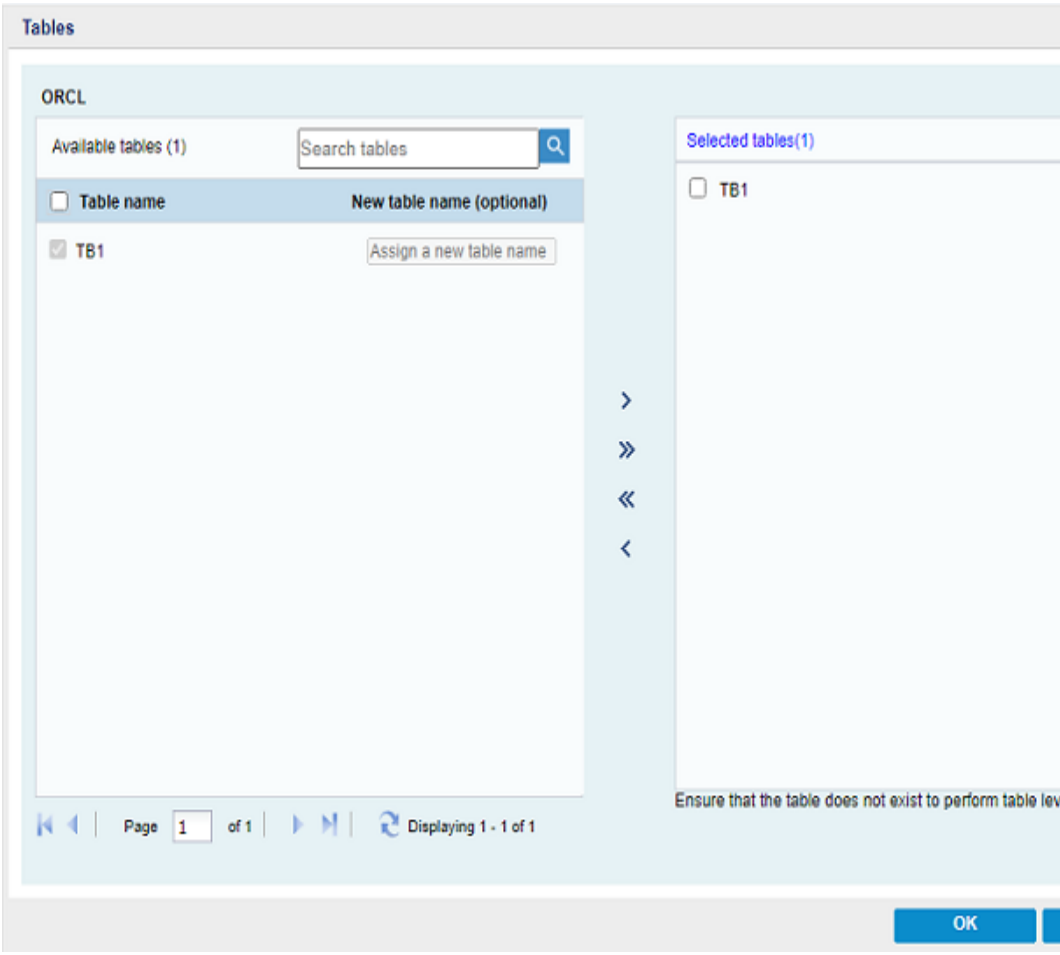


- From the left-pane, perform the following:
  - To perform full restore of entire database items, select **Full restore**.
  - To perform restore of Tablespaces and Tables, select **Granular restore**.

  Note: The Oracle database items such as Archive Logs, Control Files, and Parameter files are selected by default when the Restore dialog appears. Additionally, these items get restored if needed.
- From the right-pane, click each file to expand, and then view the details.
- (Optional) To select the specific tables, from the right-pane, click the **Tables** option.

  A table selection dialog appears.

  On the table selection dialog, do the following, and then click **OK**:



- From the left-pane, perform the following:
  - To search the available tables, in the search box, select the name of the table and then click the search icon.
  - Select the tables you want to restore.
  - (Optional) Type the new name for a table you want to restore.

  After the selection of tables, click the right arrow (>) to move them to the selected tables() list.
- From the right-pane, deselect the tables you do not want to restore, and then click the left arrow (<) to move them to the Available tables() list.

– Click **OK**.

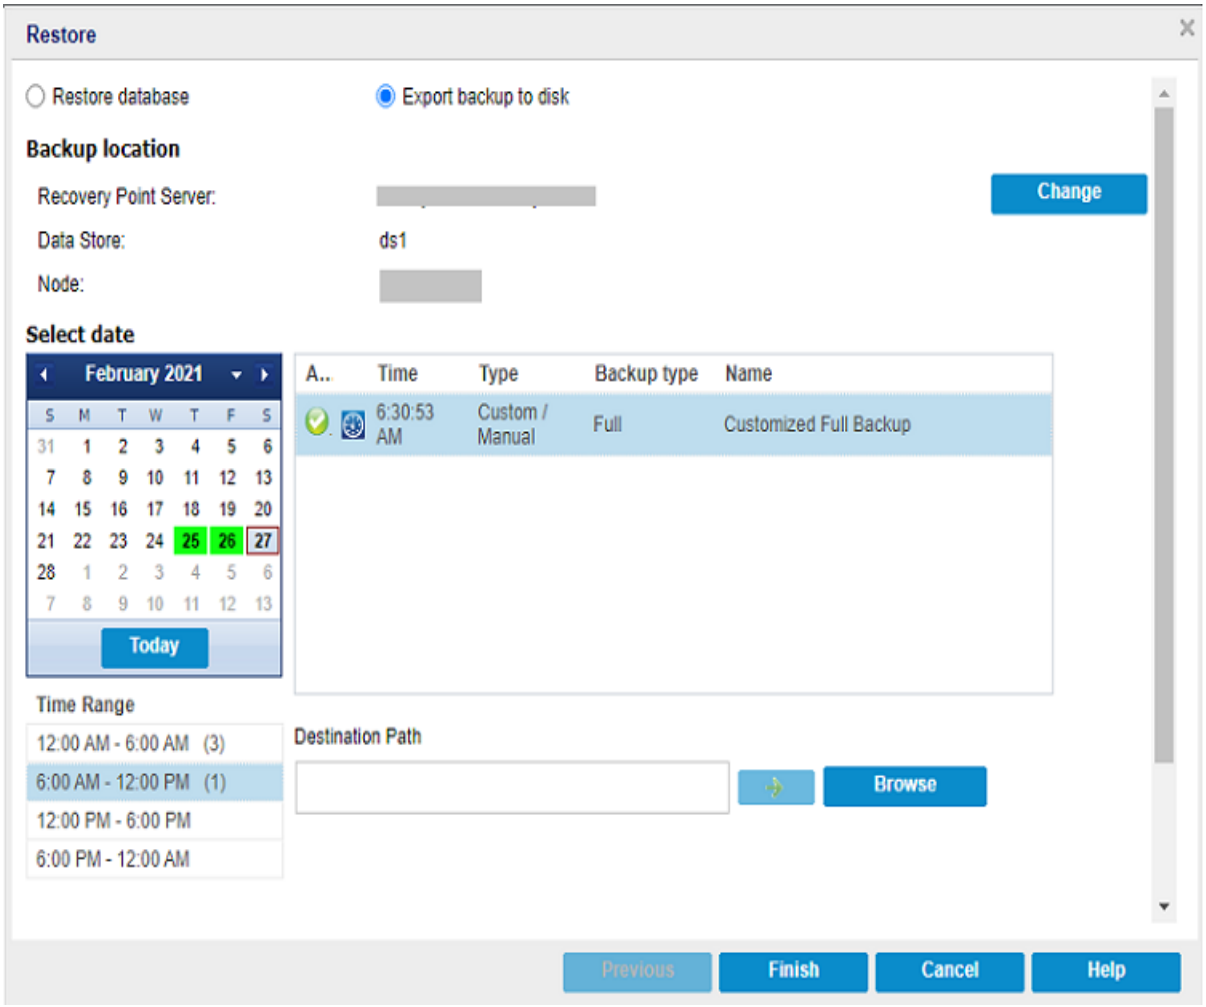
e. Click **Next**.

The Restore Options dialog appears.

The Oracle database items to restore are selected. Now, you can define the restore options.

Export backup to disk

Restores the Oracle database items to folder or share folder in disk.



To export backup to disk, follow these steps:

a. (Optional) To modify the details of the backup location, click **Change**.

b. From the Select date section, select a date.

   Based on selection, the list of backup recovery points is displayed with the time of the backup, the type of backup that was performed (Full or Incremental), and the name of the backup.

c. From the list of backup recovery points, select a backup recovery point you want to restore.

   Note: By default, the latest backup recovery point is selected.

d. In the Destination Path field, provide the full path of the remote network share.

e. To validate the connection, click the green arrow icon next to the Destination Path field, enter the username and password as needed, and then click **OK**.

f. On successful connection to the remote network share, click the **Browse** button.

   The Select backup location dialog appears.

g. On the Select backup location dialog, select a folder to restore, and then click **OK**.

h. Click **Finish**.

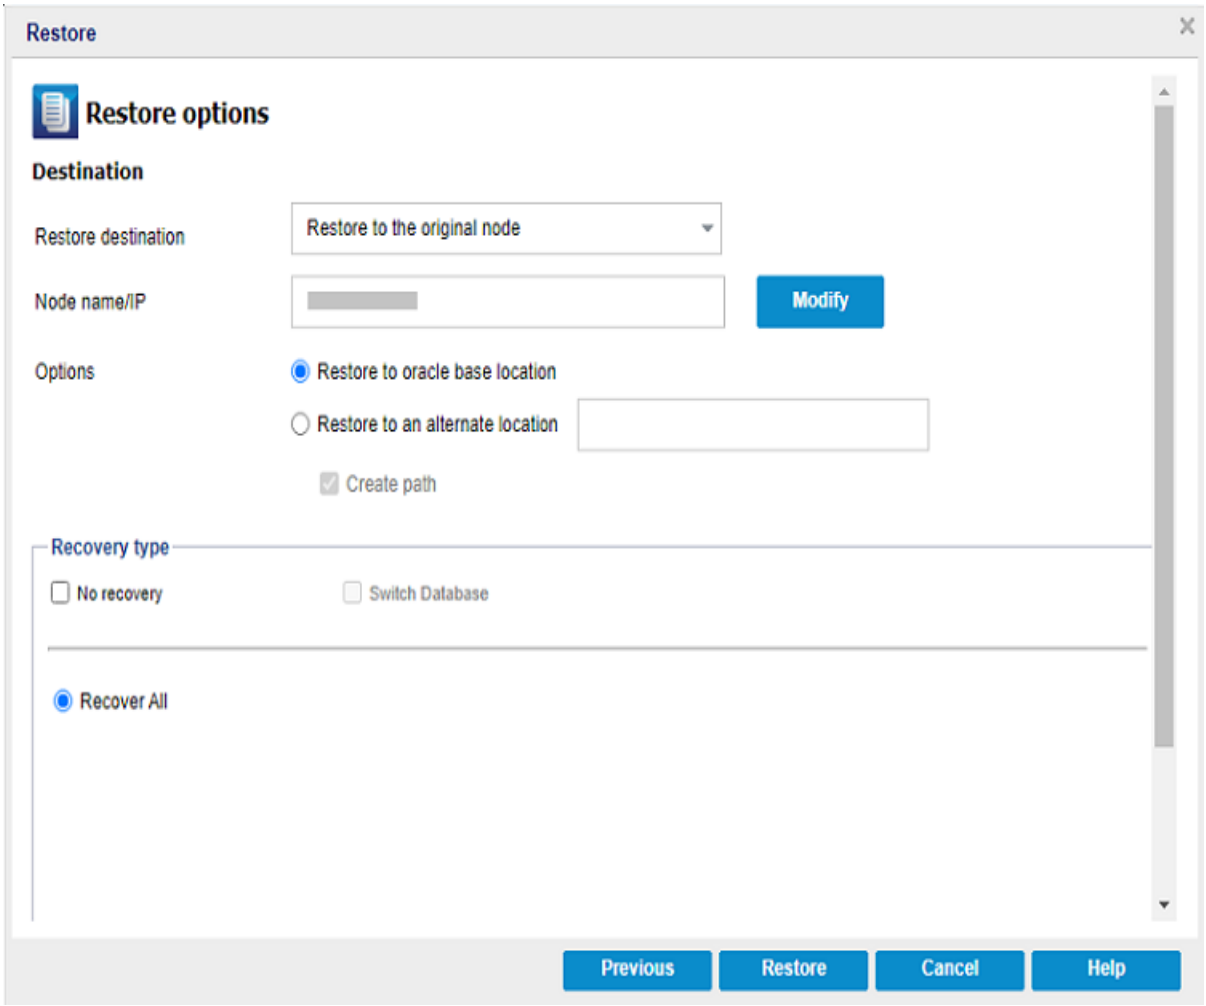
   The Oracle database items get exported to the disk.

Define the Restore Options

This section provides information about how to define restore options for the selected Oracle database items.

Follow these steps:

1. On the Restore Options dialog, select the restore destination.



Notes:

- During restore, for any Full Backup recovery point, only the Recover All option is available.
- Recovery points associated to Incremental backups include all options such as Recover All, Point in time, Until SCN, and Until log sequence number for restore.

The available destination options are:

Restore to the original node

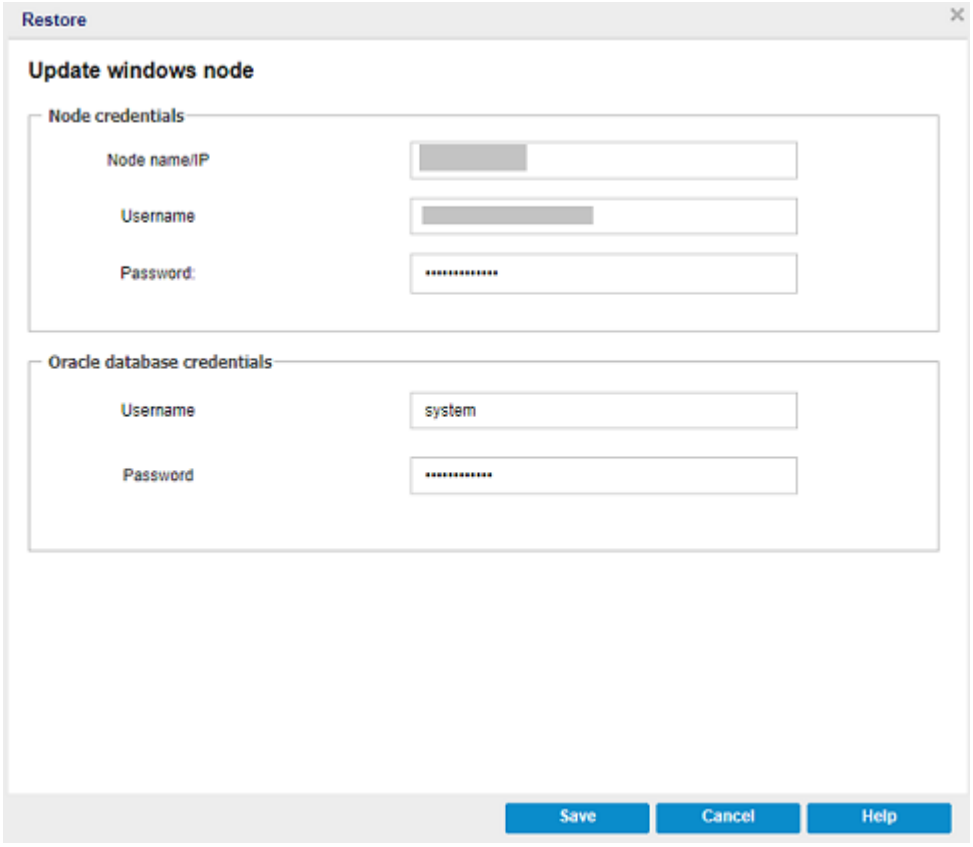
To restore to the original node, perform the following:

- The existing Node name/IP details appear automatically. In case there is a requirement to change the Node name/IP details, click **Modify**.

  Based on the node type, the Update windows node or Update linux node dialog appears.

  Update windows node

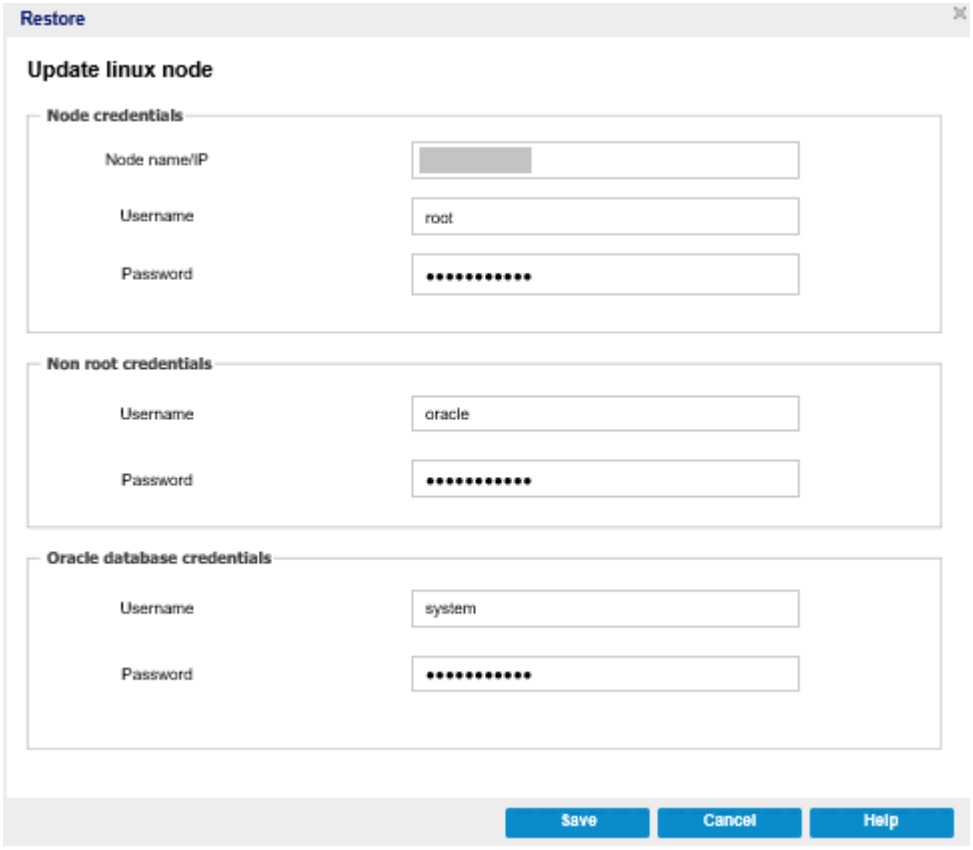
  On the Update windows node dialog, update the following details, and then click **Save**.



  - For Node credentials, enter the Node name/Ip, Username, and Password as needed.
  - For Oracle database credentials, enter the Username and Password as needed.

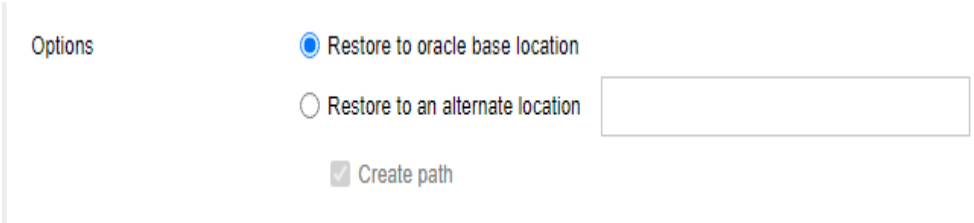
  Update linux node

  On the Update linux node dialog, update the following details, and then click **Save**.



- For Node credentials, enter the Node name/Ip, Username, and Password as needed.
- For Non root credentials, enter the Username and Password as needed.
- For Oracle database credentials, enter the Username and Password as needed.

- For options, select one of the following:



The available options are:

Restore to oracle base location

Restores the selected database items to the oracle base location from where you took a backup.

Restore to an alternate location

Restores to an alternate location. When you select this option, you need to add a path of the alternate location.

Notes:

- The **Create path** check box is selected automatically when you select the **Restore to an alternate location** option.
- If you restore the database items on a RAC setup, provide the shared storage path *(+DATA/new/destination)*.

Restore to the alternate node

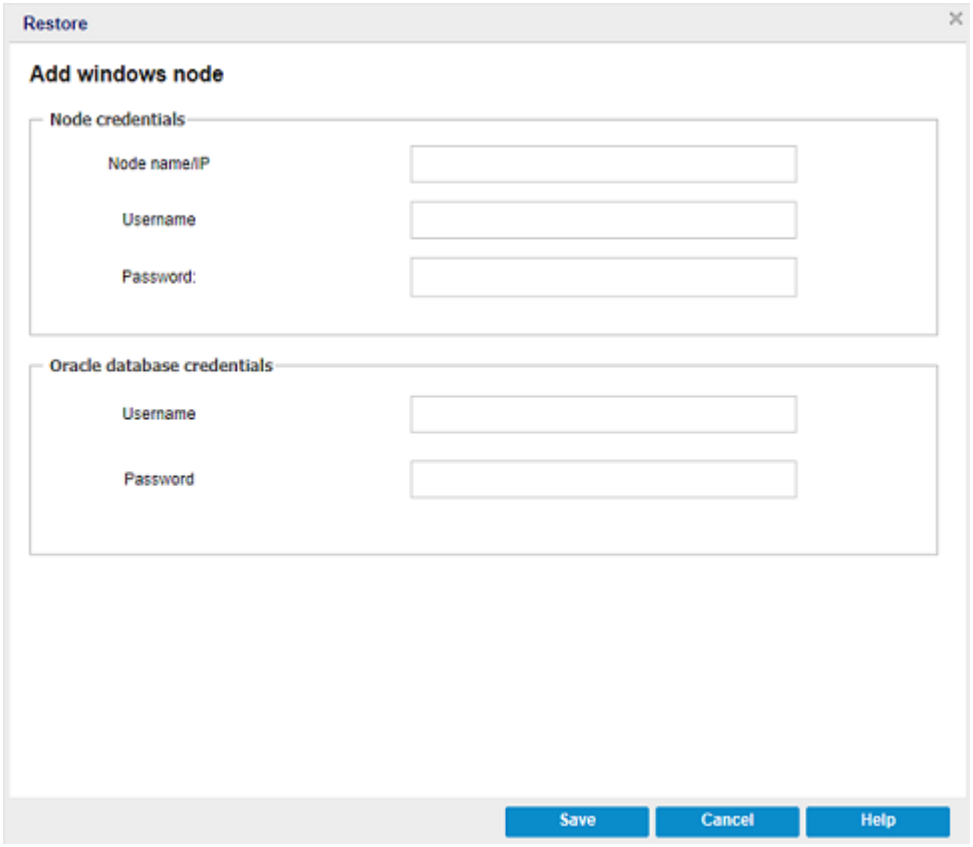
To restore to the alternate node, perform the following:

- To add a new Node/IP, click the **Add** button.

  Based on the node type, the Add windows node or Add linux node dialog appears.

  Add windows node

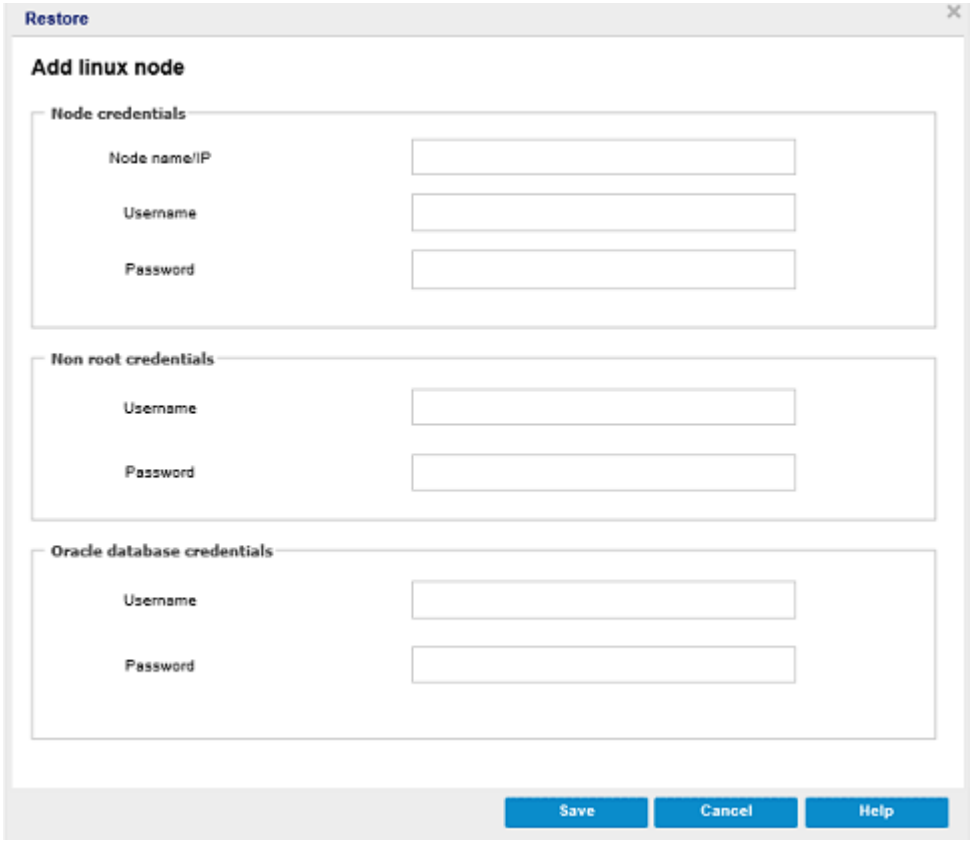
  On the Add windows node dialog, add the following details, and then click **Save**.



- For Node credentials, enter the new Node name/Ip, Username, and Password as needed.
- For Oracle database credentials, enter the new Username and Password as needed.

Add linux node

On the Add linux node dialog, add the following details, and then click **Save**.



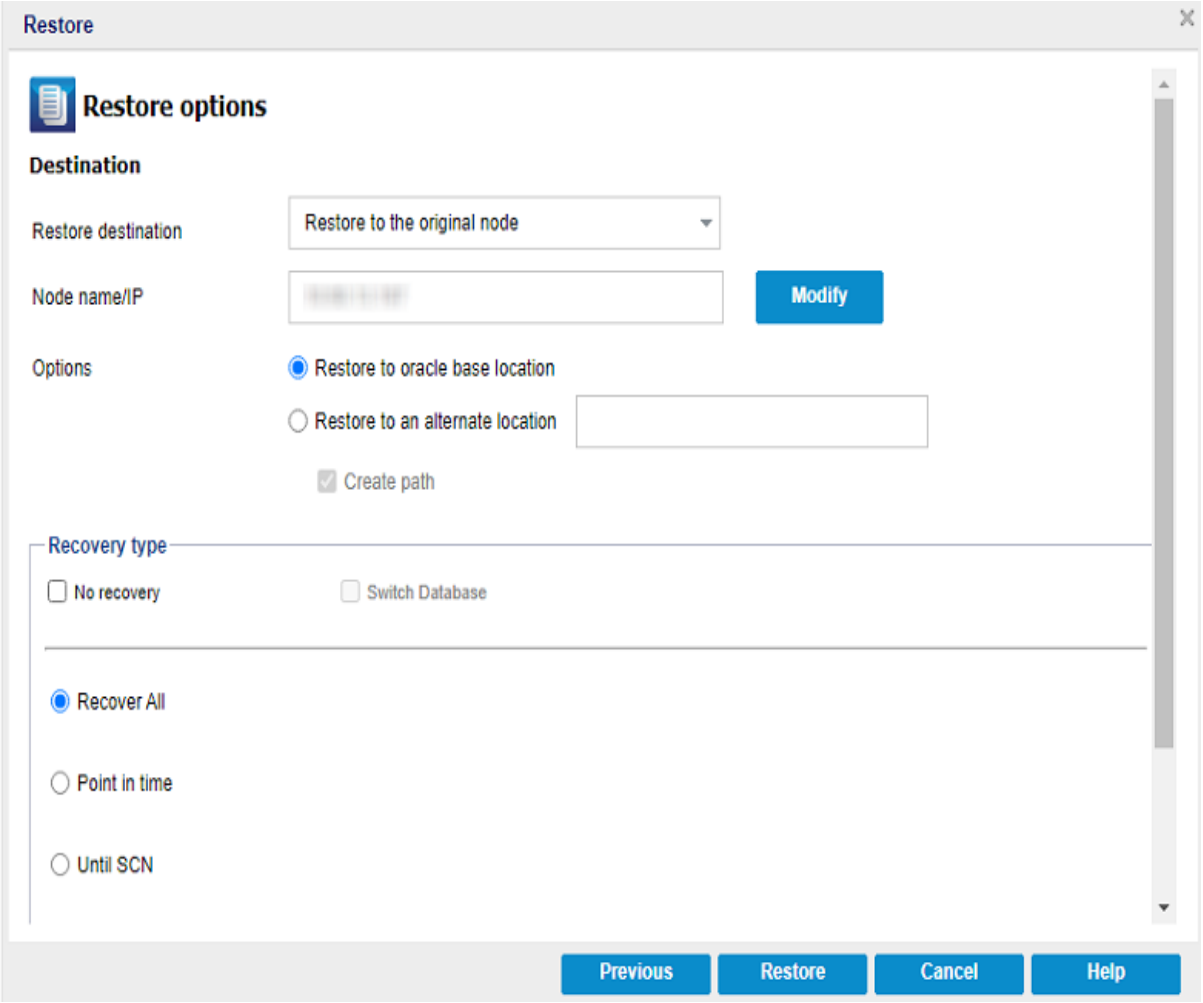
- For Node credentials, enter the new Node name/Ip, Username, and Password as needed.
- For Non root credentials, enter the new Username and Password as needed.
- For Oracle database credentials, enter the new Username and Password as needed.

- For Options, the **Restore to an alternate location** option is selected by default. Enter a path of the alternate location.



Note: If you restore the database items on a RAC setup, provide the shared storage path *(+DATA/new/destination)*.

2. In the Recovery type section, do the following:



- If you do not want to run the recovery, select the **No Recovery** check box. Based on selection, the following option is enabled.

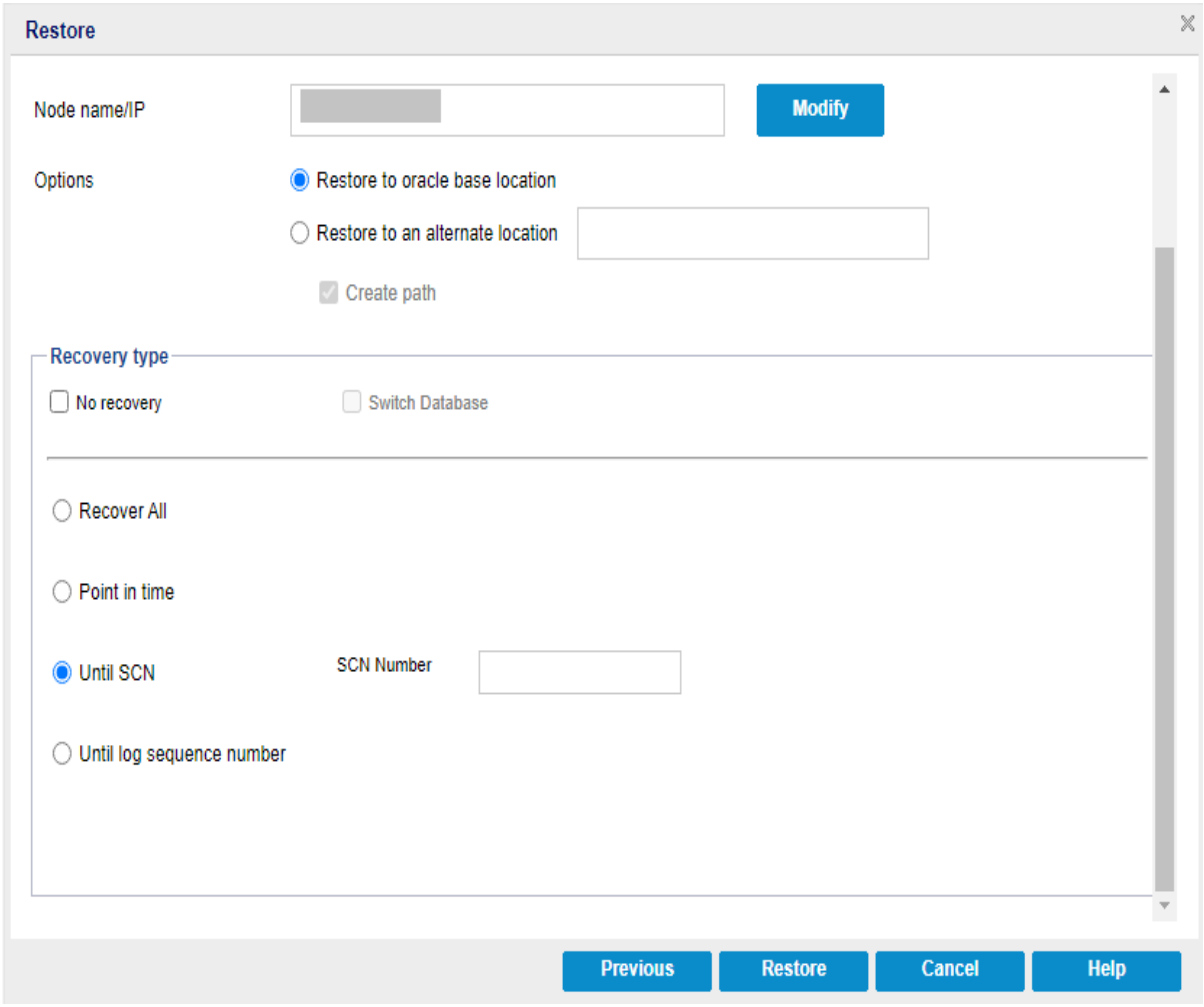
  Switch Database

  If you want to switch the database items to the new location, select the **Switch Database** check box.

- If you want to perform the point in time restore, select **Point in time**, and then provide the following:
  - Select a date to run the Point in time restore.
  - Select the start time as needed.

  Note: The date and start time for Point in time restore must be less than the selected backup recovery point.

- If you want to recover all the database items until the given System Change Number (SCN) timestamp, select **Until SCN**, and then type the value for SCN Number as needed.



- If you want to recover all the database items until the given log sequence number, select **Until log sequence number**, and then type the values for Log sequence and Thread Number as needed.

3. Click **Restore**.

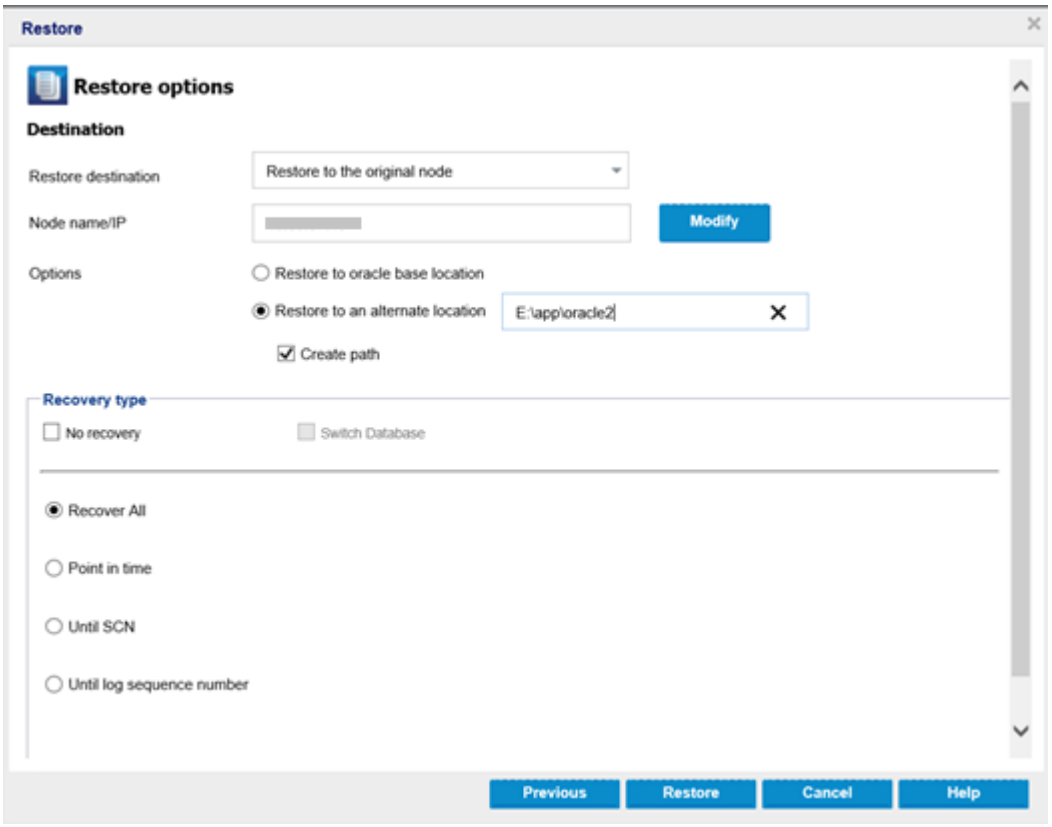
The Oracle database items get restored.

How to Restore Database to Alternate Location on Windows

This section provides information about how to restore database to alternate location on Windows.

Follow these steps:

1. From the UDP Console, select the recovery point from where you need to perform restore. For more information, see Define the Restore Options.
2. On the Restore Options dialog, from the Restore destination drop-down, select **Restore to the original node**.



3. For Node name/IP, enter the node name or IP address.
4. For Options, select the **Restore to an alternate location** option, and then add a path of the alternate location.

   Notes:

   - Make sure the path added does not have any special characters or space.

- The **Create path** check box is selected automatically when you select the **Restore to an alternate location** option.
- If the database is not already shut down, a notification appears and prompts you to shut down the database. To confirm, click **OK**.

5. Confirm that Restore Job is successful from the UDP Console.
6. Connect to **sqlplus** and confirm that the database is up and running.

   Restored database is operational with specified alternate location. However, after database restarts, the database might go back to previous copy or might not come back properly because of a known issue. To avoid any further problems after restoring to alternate location, it is recommended to follow the steps below. These steps help to create or place right copy of spfile/pfile.

   a. Go to **%ORACLE_HOME%\database and rename spfile<SID>.ora**.

   b. Connect to **sqlplus**.

   c. Go to SQL prompt and run the following command:

      sqlplus > create pfile from spfile.

   d. Go to **%ORACLE_HOME%\database**, open **init<SID>.ora** in any editor, and then do the following:

      1. Update “db_create_file_dest” value with the path specified for alternate restored location.
      2. Remove additional ‘\’ from end of path under the *diagnostic_dest* parameter.

         Change the following:

         From *diagnostic_dest='E:\app\oracle2\ORA19C280220211833\'*

         To *diagnostic_dest='E:\app\oracle2\ORA19C280220211833'*

      3. Save the file.

   e. Go to SQL prompt and run the following commands:

      Sqlplus > shutdown immediate

      sqlplus > startup

      sqlplus > create spfile from pfile

      Note: If DB is not started from the *spfile*, restores from the corresponding backups are limited to the original node only. To avoid it, restart the database again after creation of the *spfile*. To restart, run the follow commands:

Sqlplus > shutdown immediate

sqlplus > startup

How to Restore Database to Alternate Node on Windows

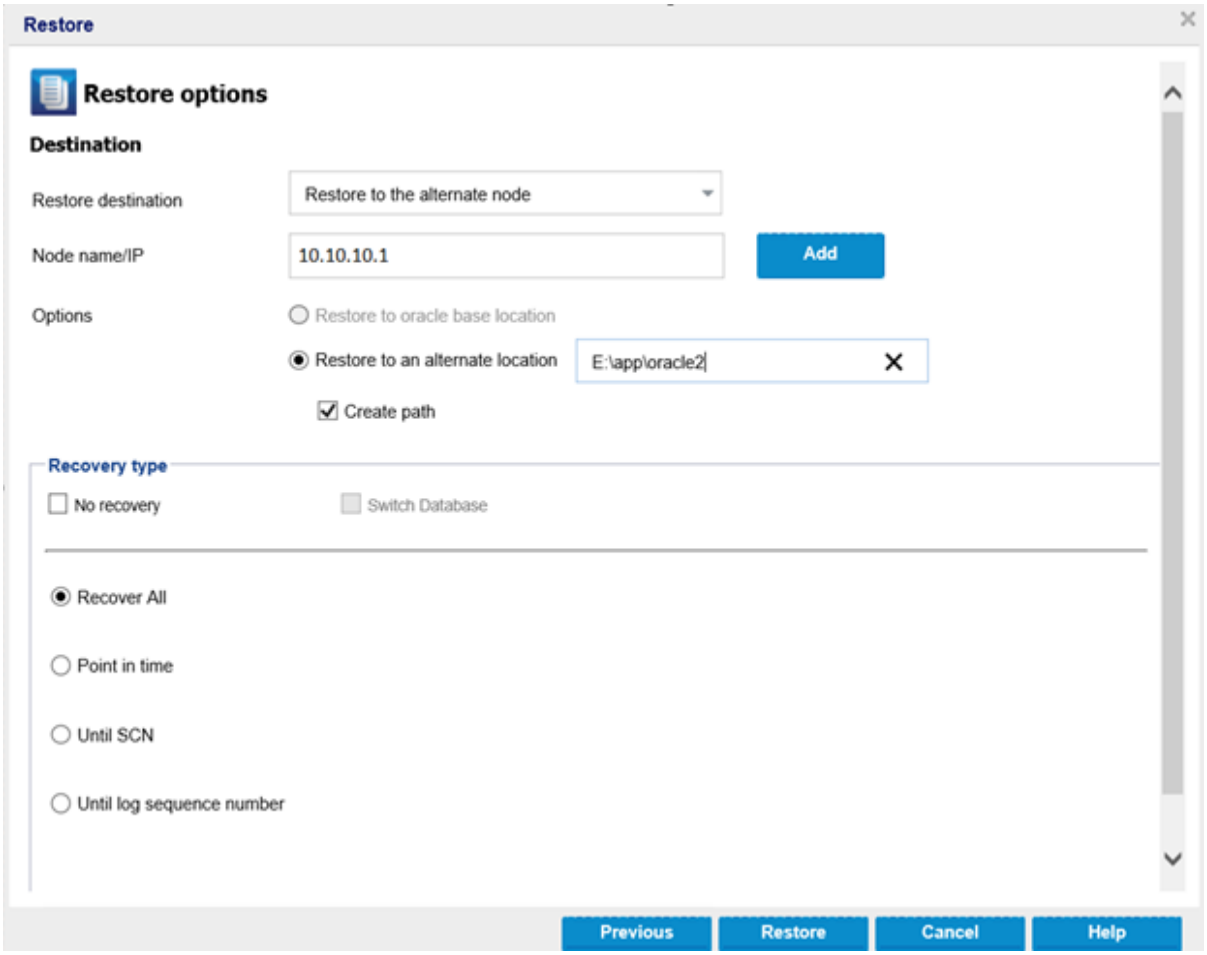
This section provides information about how to restore database to alternate node on Windows.

Notes:

- When **Change Block Tracker** is disabled and alternate node restore performed, the restore might fail. Hence, select the restore only option when you have the required logs available with you and you can recover database manually after restore. As a workaround, enable **Change Block Tracker** before performing alternate node restore on source node.
- Make sure to install same version of Oracle software and create database with same SID on alternate node where you planned to perform restore as Original node (source node). For example, when source node SID is orcl, create orcl on alternate node before starting restore. This is needed when Oracle software is installed with domain account. This is not required while using virtual account.

Follow these steps:

1. From the UDP Console, select the recovery point from where you need to perform restore. For more information, see Define the Restore Options.
2. On the Restore Options dialog, from the Restore destination drop-down, select **Restore to the alternate node**.



3. For Node name/IP, enter the node name or IP address.
4. For Options, select the **Restore to an alternate location** option, and then add a path of the alternate location.

   Notes:

   - Make sure the path added does not have any special characters or space.
   - The **Create path** check box is selected automatically when you select the **Restore to an alternate location** option.
   - If the database is not already shut down, a notification appears and prompts you to shut down the database. To confirm, click **OK**.
5. Confirm that Restore Job is successful from the UDP Console.
6. Connect to **sqlplus** on alternate node and confirm that the database is up and running.

   Restored database is operational with specified alternate location. However, after database restarts, the database might go back to previous copy or

might not come back properly because of a known issue. To avoid any further problems after restoring to alternate location, it is recommended to follow the steps below. These steps help to create or place right copy of spfile/pfile.

a. Go to **%ORACLE_HOME%\database and rename spfile<SID>.ora**.

b. Connect to **sqlplus**.

c. Go to SQL prompt and run the following command:

   sqlplus > create pfile from spfile.

d. Go to **%ORACLE_HOME%\database**, open **init<SID>.ora** in any editor, and then do the following:

   1. Update “db_create_file_dest” value with the path specified for alternate restored location.

      Example: **.db_create_file_dest-='E:\app\oracle2\ORA19C280220211833'*

   2. Remove additional ‘\’ from end of path under the *diagnostic_dest* parameter.

      Change the following:

      From *diagnostic_dest='E:\app\oracle2\ORA19C280220211833\'*

      To *diagnostic_dest='E:\app\oracle2\ORA19C280220211833'*

   3. Save the file.

e. Go to SQL prompt and run the following commands:

   Sqlplus > shutdown immediate

   sqlplus > startup

   sqlplus > create spfile from pfile

   Note: If DB is not started from the *spfile*, restores from the corresponding backups are limited to the original node only. To avoid it, restart the database again after creation of the *spfile*. To restart, run the follow commands:

   Sqlplus > shutdown immediate

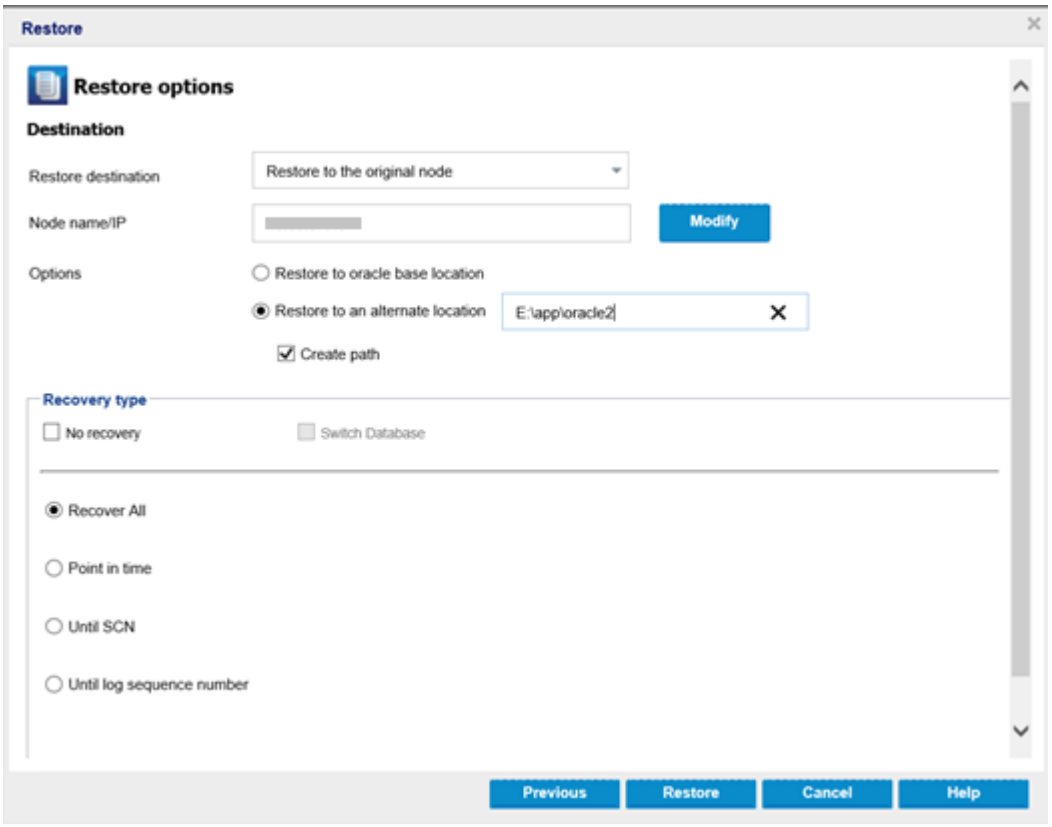
   sqlplus > startup

How to Restore Database to Alternate Location on Linux

This section provides information about how to restore database to alternate location on Linux.

Follow these steps:

1. From the UDP Console, select the recovery point from where you need to perform restore. For more information, see Define the Restore Options.
2. On the Restore Options dialog, from the Restore destination drop-down, select **Restore to the original node**.



3. For Node name/IP, enter the node name or IP address.
4. For Options, select the **Restore to an alternate location** option, and then add a path of the alternate location.

   Notes:

   - Make sure the path added does not have any special characters or space.

- The **Create path** check box is selected automatically when you select the **Restore to an alternate location** option.
- If you restore the database items on a RAC setup, provide the shared storage path (*+DATA/new/destination*).
- If the database is not already shut down, a notification appears and prompts you to shut down the database. To confirm, click **OK**.

5. Confirm that Restore Job is successful from the UDP Console.
6. Connect to **sqlplus** and confirm that the database is up and running.

   Restored database is operational with specified alternate location. However, after database restarts, the database might go back to previous copy or might not come back properly because of a known issue. To avoid any further problems after restoring to alternate location, it is recommended to follow the steps below. These steps help to create or place right copy of spfile/pfile.

   a. Go to **$ORACLE_HOME/dbs** and rename **spfile<SID>.ora**.

   b. Connect to **sqlplus**.

   c. Go to SQL prompt and run the following command:

      sqlplus > create pfile from spfile.

   d. Go to **$ORACLE_HOME/dbs**, open **init<SID>.ora** in any editor, and then do the following:

      1. Update **.db_create_file_dest* value with the path specified for alternate restored location.

         Example: **.db_create_file_dest-='/u02/oracle2/ORA19C280220210850'*

      2. Remove additional '/' from end of path under the *diagnostic_dest* parameter.

         Change the following:

         From *diagnostic_dest='/u02/oracle2/ORA19C280220210850//'*

         To *diagnostic_dest='/u02/oracle2/ORA19C280220210850'*

      3. Save the file.

   e. Go to SQL prompt and run the following commands:

      Sqlplus > shutdown immediate

      sqlplus > startup

      sqlplus > create spfile from pfile

Note: If DB is not started from the *spfile*, restores from the corresponding backups are limited to the original node only. To avoid it, restart the database again after creation of the *spfile*. To restart, run the follow commands:

Sqlplus > shutdown immediate

sqlplus > startup

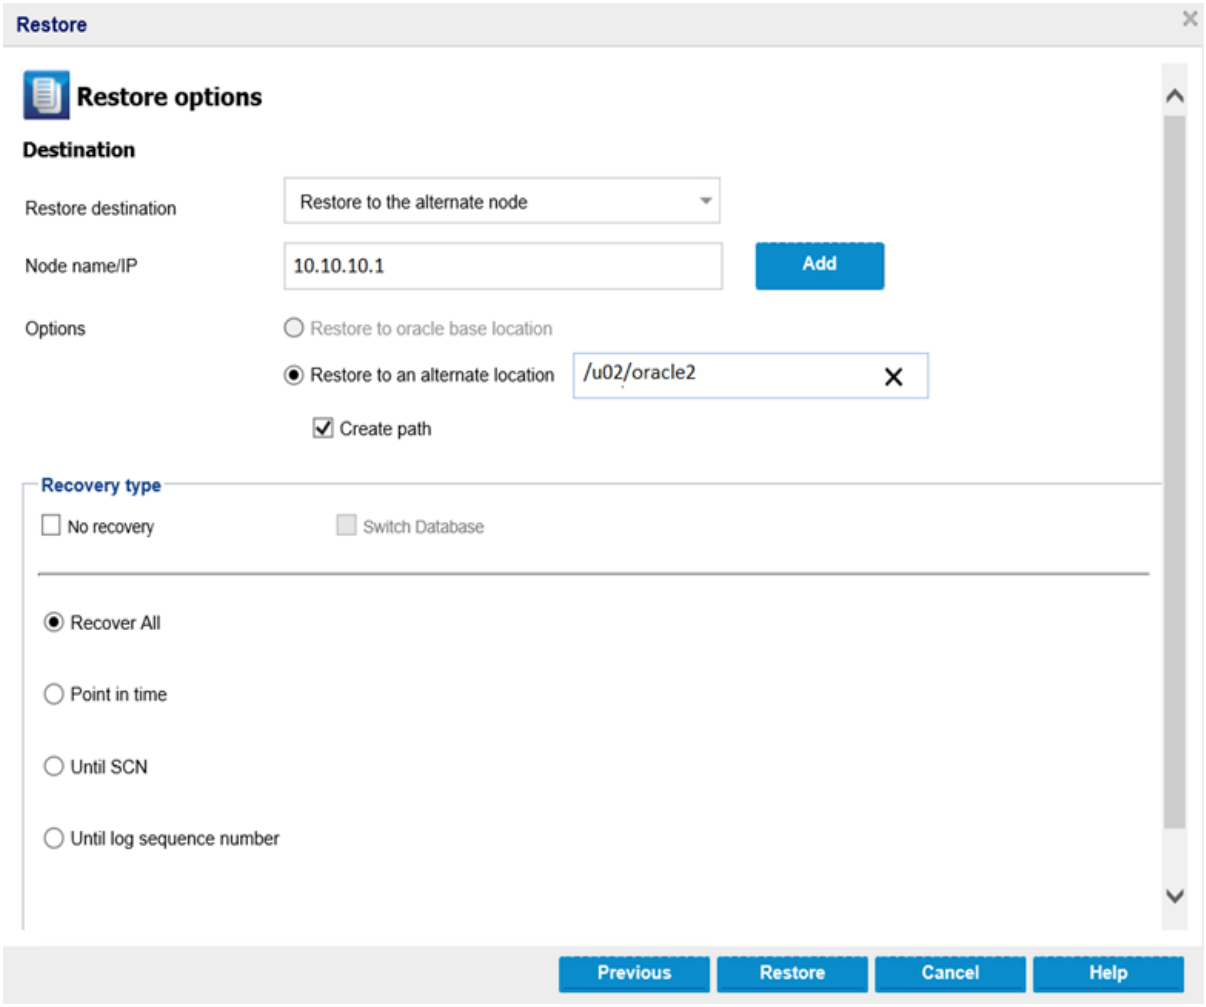
How to Restore Database to Alternate Node on Linux

This section provides information about how to restore database to alternate node on Linux.

Note: Make sure to install the same version of Oracle software on the alternate node that you planned to perform restore as the original node (source node).

Follow these steps:

1. From the UDP Console, select the recovery point from where you need to perform restore. For more information, see Define the Restore Options.
2. On the Restore Options dialog, from the Restore destination drop-down, select **Restore to the alternate node**.



3. For Node name/IP, enter the node name or IP address.

4. For Options, select the **Restore to an alternate location** option, and then add a path of the alternate location.

   Notes:

   - Make sure the path added does not have any special characters or space.
   - The **Create path** check box is selected automatically when you select the **Restore to an alternate location** option.
   - If you restore the database items on a RAC setup, provide the shared storage path (*+DATA/new/destination*).

5. Confirm that Restore Job is successful from the UDP Console.
6. Connect to **sqlplus** and confirm that the database is up and running.

   Restored database is operational with specified alternate location. However, after database restarts, the database might go back to previous copy or might not come back properly because of a known issue. To avoid any further problems after restoring to alternate location, it is recommended to follow below steps. These steps help to create or place right copy of spfile/pfile.

   a. Connect to ***sqlplus***.

   b. Go to SQL prompt and run the following command:

      sqlplus > create pfile from spfile.

   c. Go to **$ORACLE_HOME/dbs**, open **init<SID>.ora** in any editor, and then do the following:

      1. Update **.db_create_file_dest* value with path specified for alternate restored location.

         Example: **.db_create_file_dest='/u02/oracle2/ORA19C280220210850'*

      2. Update **.db_recovery_file_dest* value with right path.

         Example: **.db_recovery_file_dest='/u02/oracle2/ORA19C280220210850/ fast_recovery_area'*

      3. Remove additional '/' from end of path under the parameter *diagnostic_dest*.

         Change the following:

         From *diagnostic_dest='/u02/oracle2/ORA19C280220210850/'*

To *diagnostic_dest='/u02/oracle2/ORA19C280220210850'*

4. Save the file.

d. Go to SQL prompt and run the following commands:

Sqlplus > shutdown immediate

sqlplus > startup

sqlplus > create spfile from pfile

Note: If DB is not started from the *spfile*, restores from the corresponding backups are limited to the original node only. To avoid it, restart the database again after creation of the *spfile*. To restart, run the follow commands:

Sqlplus > shutdown immediate

sqlplus > startup

Limitations

- UDP 8.0 does not support backup of databases from different Oracle versions installed on the same server.
- Table names with special characters are not supported while restoring table with 'Assign a new table name'.
- Table recovery as per Oracle limitations:
  - Tables created from SYSTEM and SYSAUX tablespace cannot be recovered.
  - Tables created from system or system user cannot be recovered.
- RPS acts as a Backup, Restore and Assured Recovery Agent.
- If a table is dropped, then performing a granular restore of tablespaces does not recover the dropped table as Arcserve do not support the UNTIL clause in the RMAN command during recovery. It brings the tablespace to the latest recovery point.
- If multiple tables are selected while performing restore of tables, one or more tables cannot be restored successfully. It is recommended to select one table at a time.
- Point-in-time recovery is only supported from the latest full db incremental recovery points.
- Granular restore is only supported from the latest recovery points.

- After a Full DB restore, the next incremental backup will be auto converted into a full backup.
- If the WRP size changes, the next backup will be a full backup.
- The maximum number of channels supported are 10.

Chapter 17: Protecting Microsoft SharePoint Environment

This section contains the following topics:

Installation Consideration for Microsoft SharePoint Environment

This topic provides the information that you need to install and configure the Arcserve UDP Agent for Microsoft SharePoint Environment. The information in this topic assumes that you are familiar with the characteristics and requirements of your Microsoft SharePoint Server Farm.

Environmental Considerations

SharePoint environments can be complex and distributed across a number of machines. The server farm configuration must be supported by Microsoft. For example, a distributed SharePoint environment could contain the following components:

- One or more Web-Front-End Servers.
- One or more Database Servers used by SharePoint Server Farm.
- In most of the cases, the SharePoint is installed in a domain environment. There will be one or more Domain Controllers and DNS servers.

Installation Consideration

Consider the following when you install the agents:

- SharePoint is a distributed environment. We recommend to install UDP Agent on every server in the SharePoint Farm including Web-Front-End Servers and Database Servers.
- We also recommend to protect the entire Domain environment, including Domain Controller, DNS Server. Those are required to restore Farm, Farm Configuration and Web Application.

How to Back Up Microsoft SharePoint Servers

You can back up Microsoft SharePoint servers using Arcserve UDP. To ensure that restore works as expected, we recommend to protect the entire machine, and not selected volumes on all SharePoint servers.

For Domain Controller and DNS servers, as long as the Domain service and DNS service work properly after they are booted by Instance VM, you can back up the selected volumes.

- Review the Considerations
- Perform a Backup of the SharePoint Server

Review the Considerations

We recommend to add all Servers, including Web-Front-End Servers, Database Servers, Domain Controllers, and DNS Servers into one Plan with the same backup schedule. As a result, all servers are backed up at almost the same time. This is very important for a distributed environment. While restoring, you can get the same from the Recovery Points that are backed up at the same time for all servers.

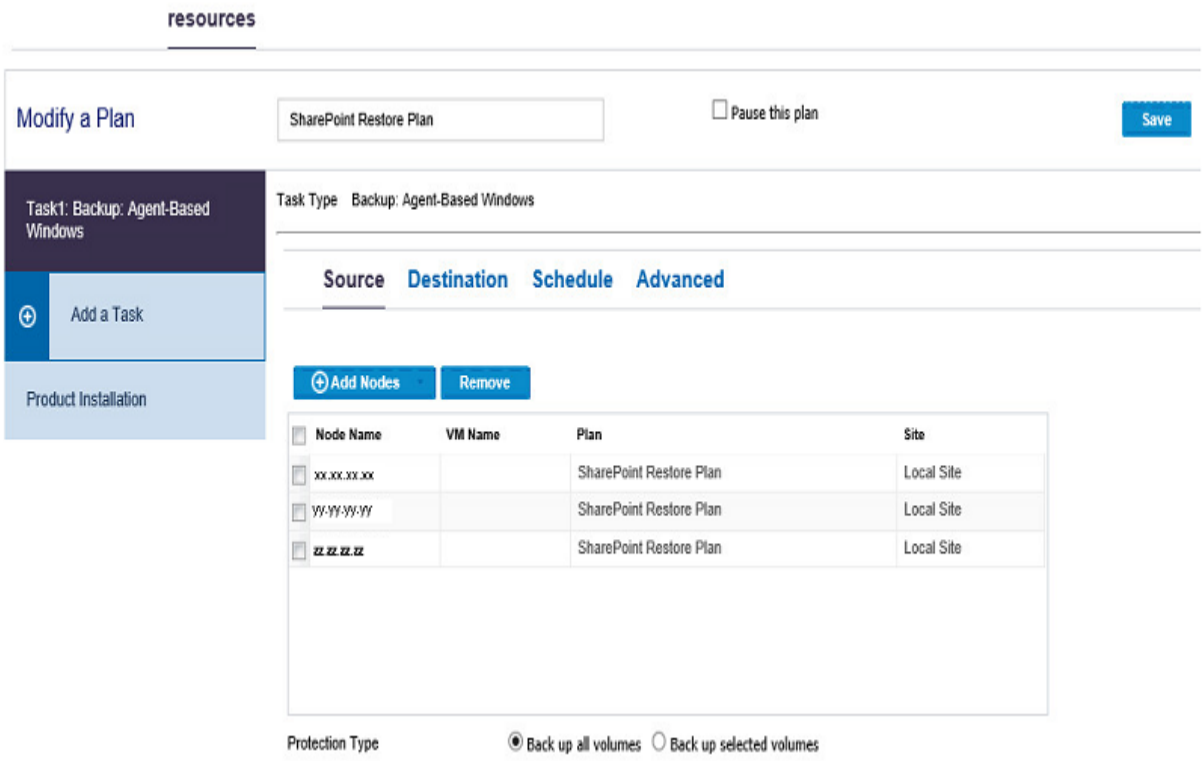
Note: If the SharePoint Farm topology is changed, new servers are added to perform load-balance improvement for SharePoint. You must add the new servers immediately to the same plan.

Perform a Backup of the SharePoint Server

Use the Arcserve UDP Console to perform the database-level backup on the SharePoint environment.

Follow these steps:

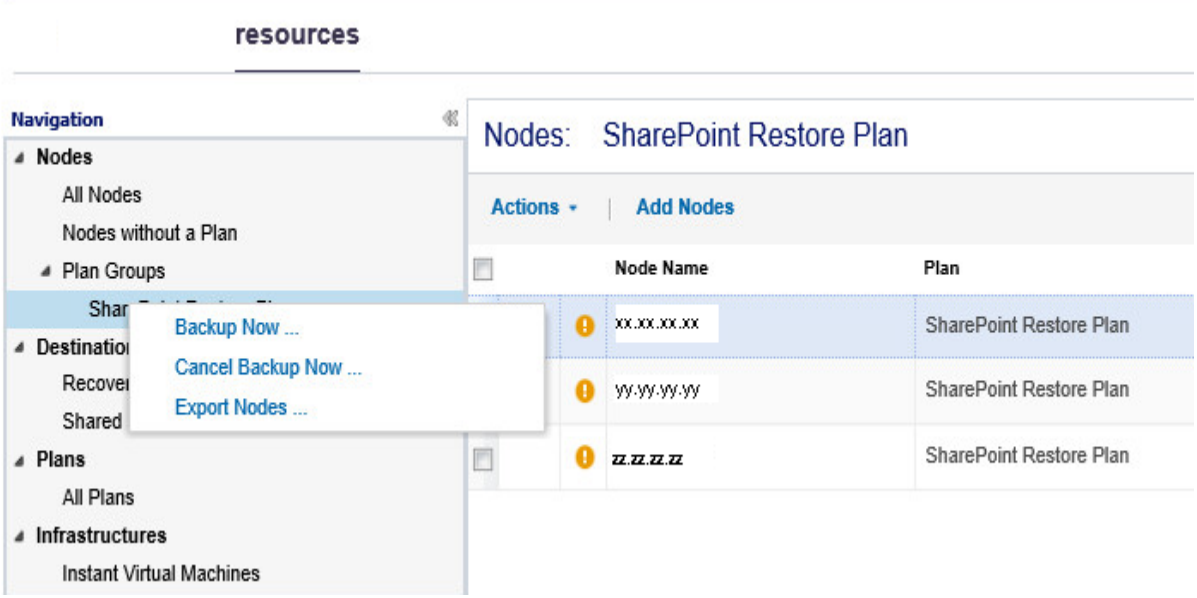
1. Log into the Console.
2. Click **resources** and navigate to **Nodes** in the left **Navigation** pane.
3. Click **All Nodes**.
4. Add all nodes in the SharePoint Farm environment.
5. (optional) Create a Data Store.
6. Expand **Plans** on the **Navigation** pane and click **All Plans**.
7. Click **Add** to create a new plan.
8. On the **Source** tab, add all the nodes that are in the SharePoint Farm environment.



9. Configure other settings and save the plan.

   The plan is created for the SharePoint Farm Environment. The settings of the Plan is deployed to all nodes. The backup job starts at the scheduled time.

You can manually start the backup job by clicking Backup Now. The backup job on all SharePoint Farm nodes start. The data is saved to the same data store.



How to Restore a SharePoint Environment

The following granularities restore are supported:

- Farm
- Farm with Configuration Only
- SharePoint Services
- Web Applications
- Content Databases
- Site Collections
- Sites
- Lists
- List Items (including Documents)

Different granularities are restored using different solutions:

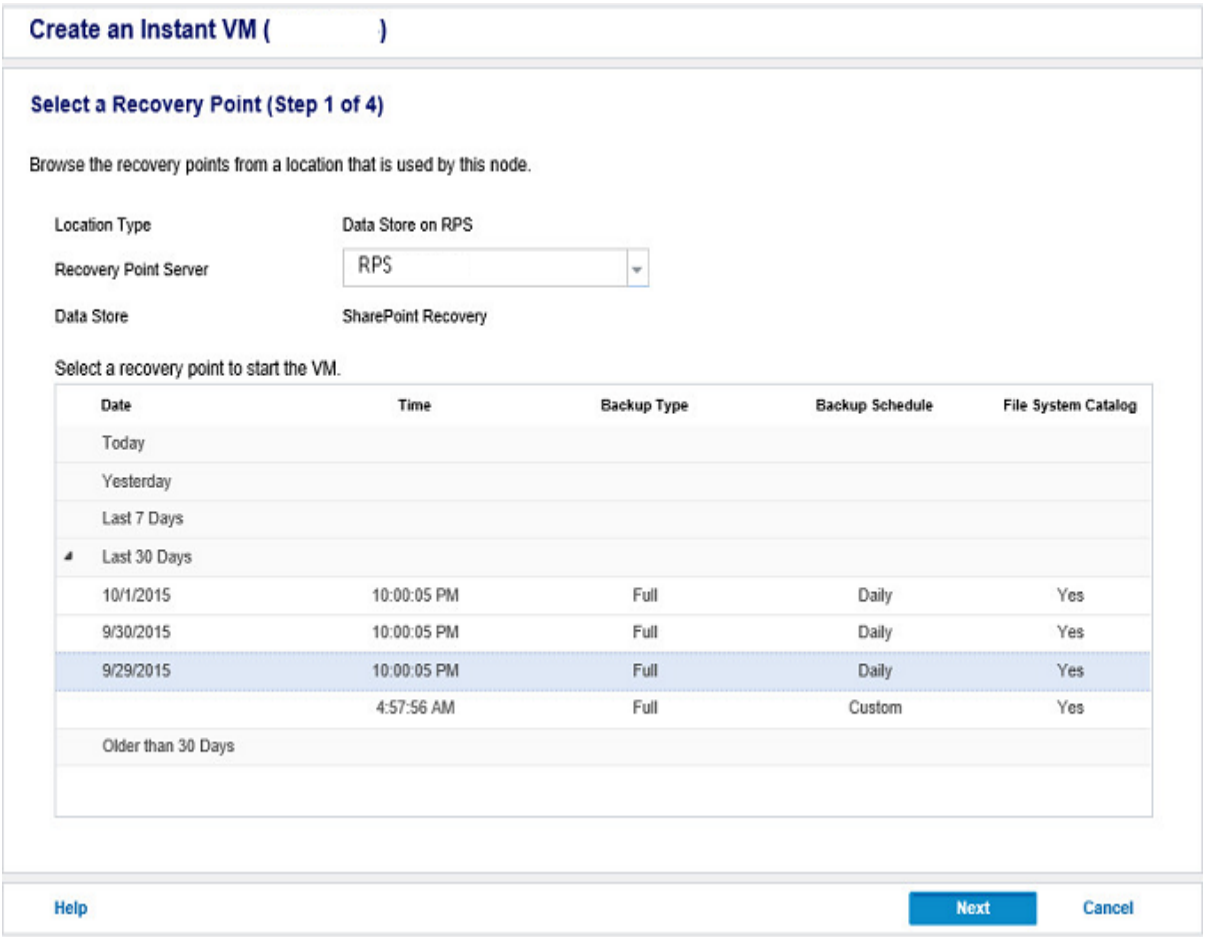
- Restore using Instant VM: Supports Farm, Farm with Configuration Only, SharePoint Services, and Web Applications.
- Restore using Arcserve UDP Agent UI: Supports Content Databases.
- Restore using Mount Database from Recovery Point: Supports Site Collections, Sites, Lists, and List Items.

Restore using Instant VM

Follow these steps:

1. Log into the Console.
2. Click **resources**, **Nodes**.
3. Right click the node, which is in the plan for SharePoint Environment.
4. Click **Create an Instant VM** to create instant virtual machines for SharePoint environment machine.

   The **Create an Instant VM** wizard opens.
5. Browse recovery points from a location and select a recovery point to start the instant VM.



6. Click **Next**.
7. Specify a location on VMware vSphere or Microsoft Hyper-V to host the instant VM. For example: Hypervisor Type is VMware vSphere.
8. Click **Next**.

9. Specify a machine to run instant VM.

   Example, use the current RPS.

10. Click **Next**.

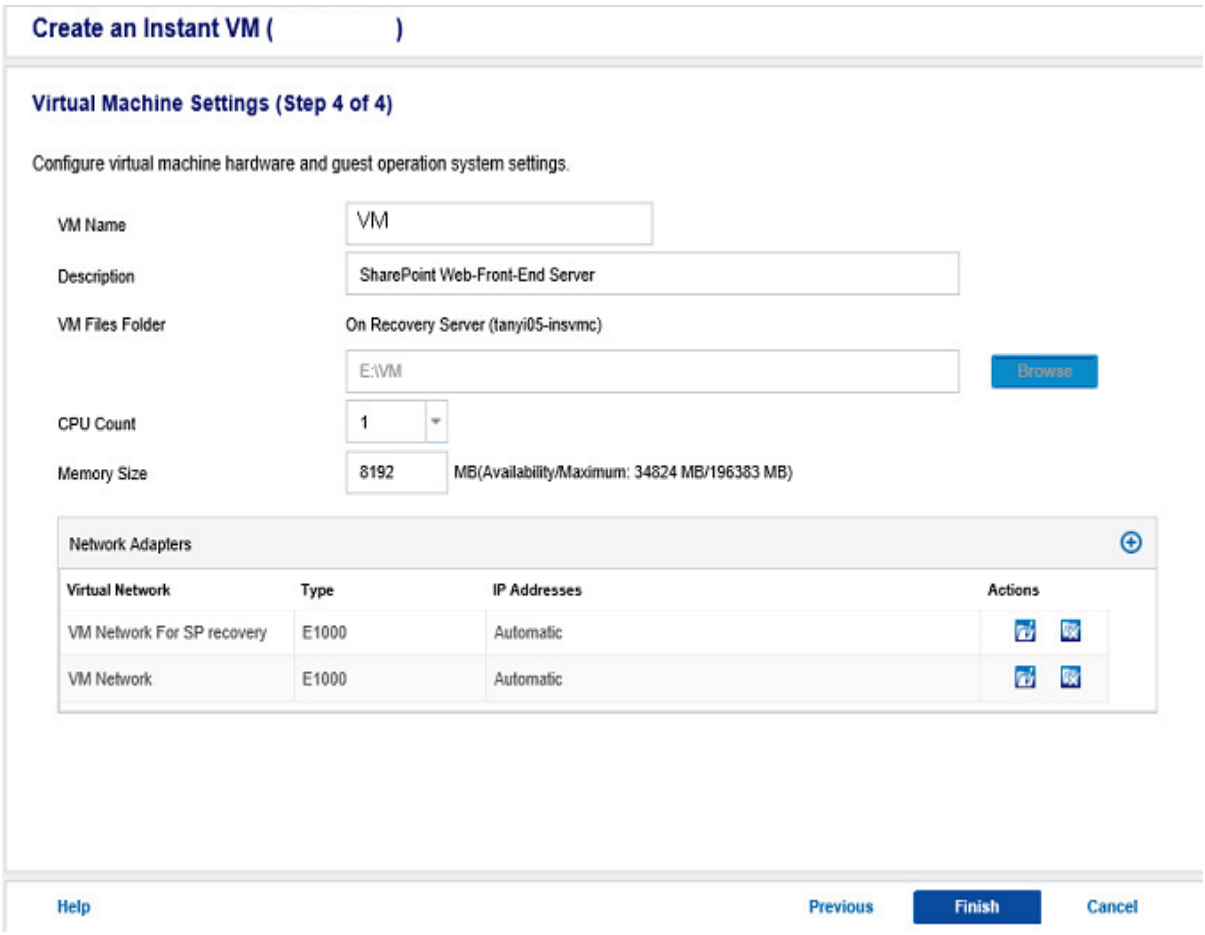
11. Configure instant VM hardware and system settings.

12. Create a new Virtual Network as an isolate network environment for SharePoint recovery. For more information on creating an isolate, see Create an isolated network for SharePoint recovery.

13. Click the (plus) sign to add network adapter.

14. Select the correct virtual network, which is an isolate network environment for SharePoint recovery, and use the default configuration of TCP/IP Settings "Source:XXX.XXX.XXX.XXX".

    Important! When creating an instant VM for a machine which is Web-Front-End Servers for SharePoint Environment, add one more Network Adapter for transferring the backup data file. The IP address of the new adapter should be in the same Virtual Network with the original Web-Front-End Servers for SharePoint, and in the same IP segment. Then, the original SharePoint Environment can use the network adapter to access any shared folder with the instant VM.



15. Click **Finish**.

    The Boot VM dialog opens.

16. Click **Boot Later**.

    A new Instant VM is created that appears in **Infrastructures**, **Instant Virtual Machines** on the Console.

17. Create Instant VM for all the nodes in the SharePoint Farm environment.

18. After the Instant VMs of all nodes are created, boot them up one by one.

    - Start the Domain Controller first, and then the DNS server.
    - After that, start the Database Server and at last, Web Front End Servers.

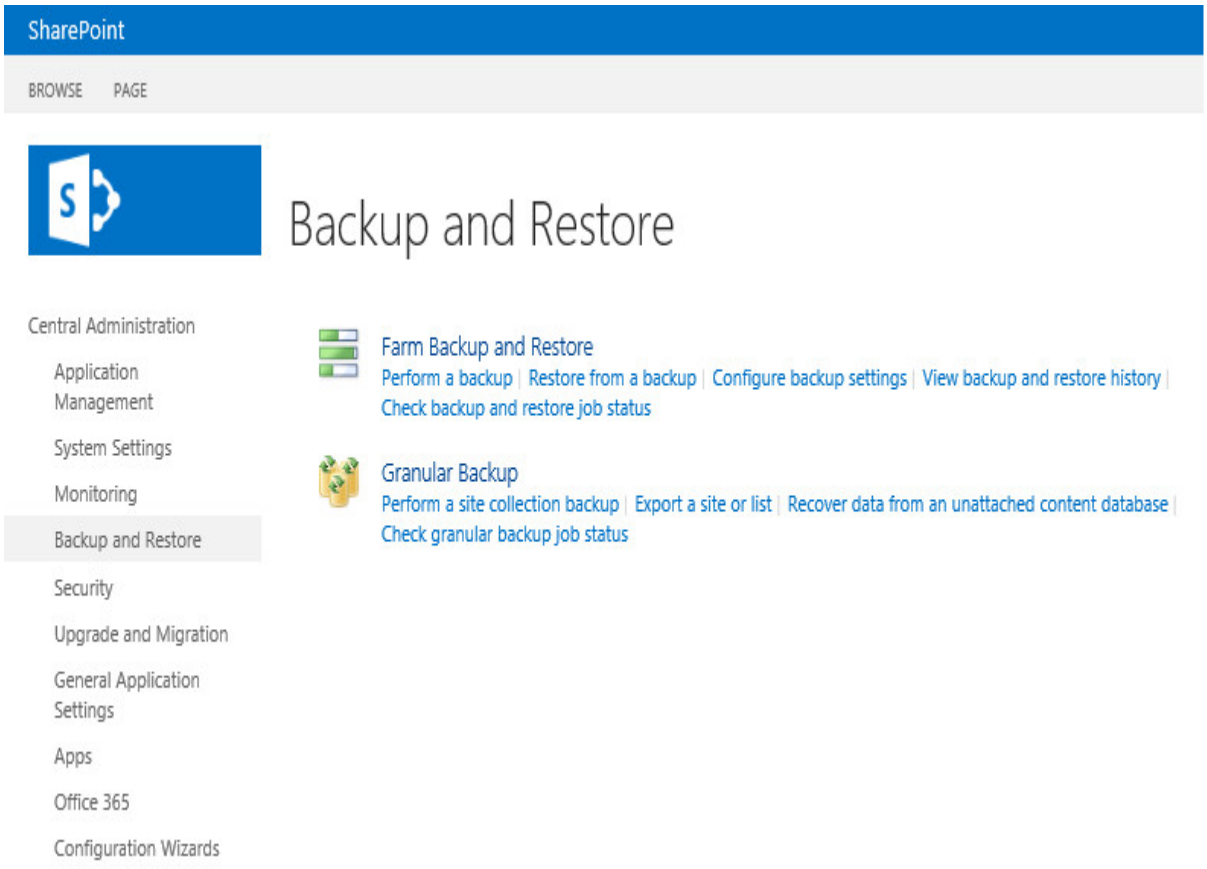
19. Right-click the Instant VM that you want to boot and select **Power on** to start the Instant VM.

20. Wait for these instant VMs for SharePoint Environment machines to power on.

    The temporary SharePoint Environment has been set up.

21. Log on to the instant VM where the machine is Web-Front-End Servers for SharePoint Environment. Open Central Administration, click Backup and Restore heading, and select the Perform a backup link.

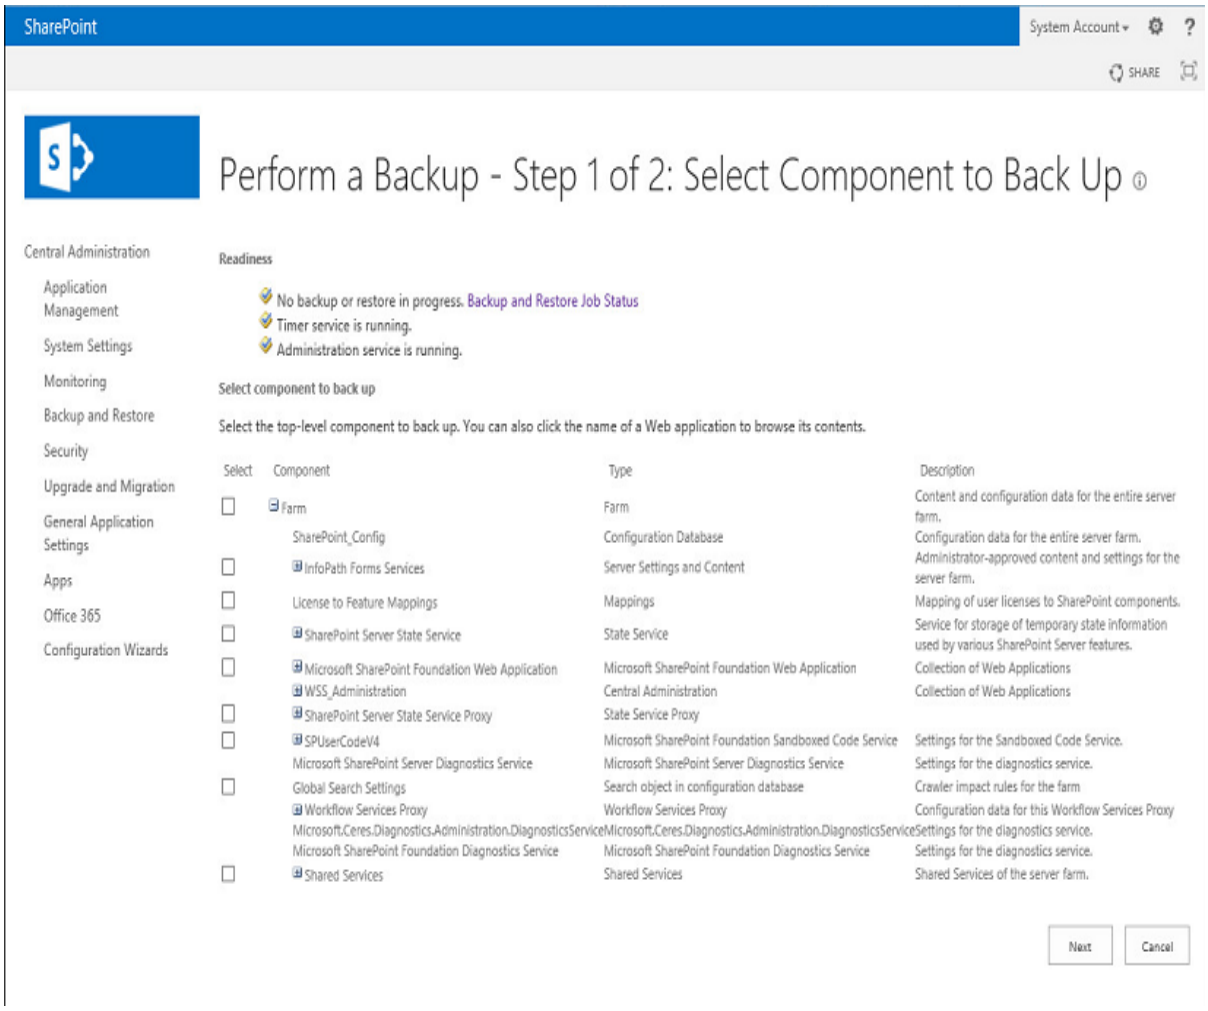
    Note: If you receive HTTP 404 error after opening Web Application such as Share-point - 80, disable IPv6 and clear the **Internet Protocol Version 6 (TCP/IP)** check box from the Local Area Network Properties window.



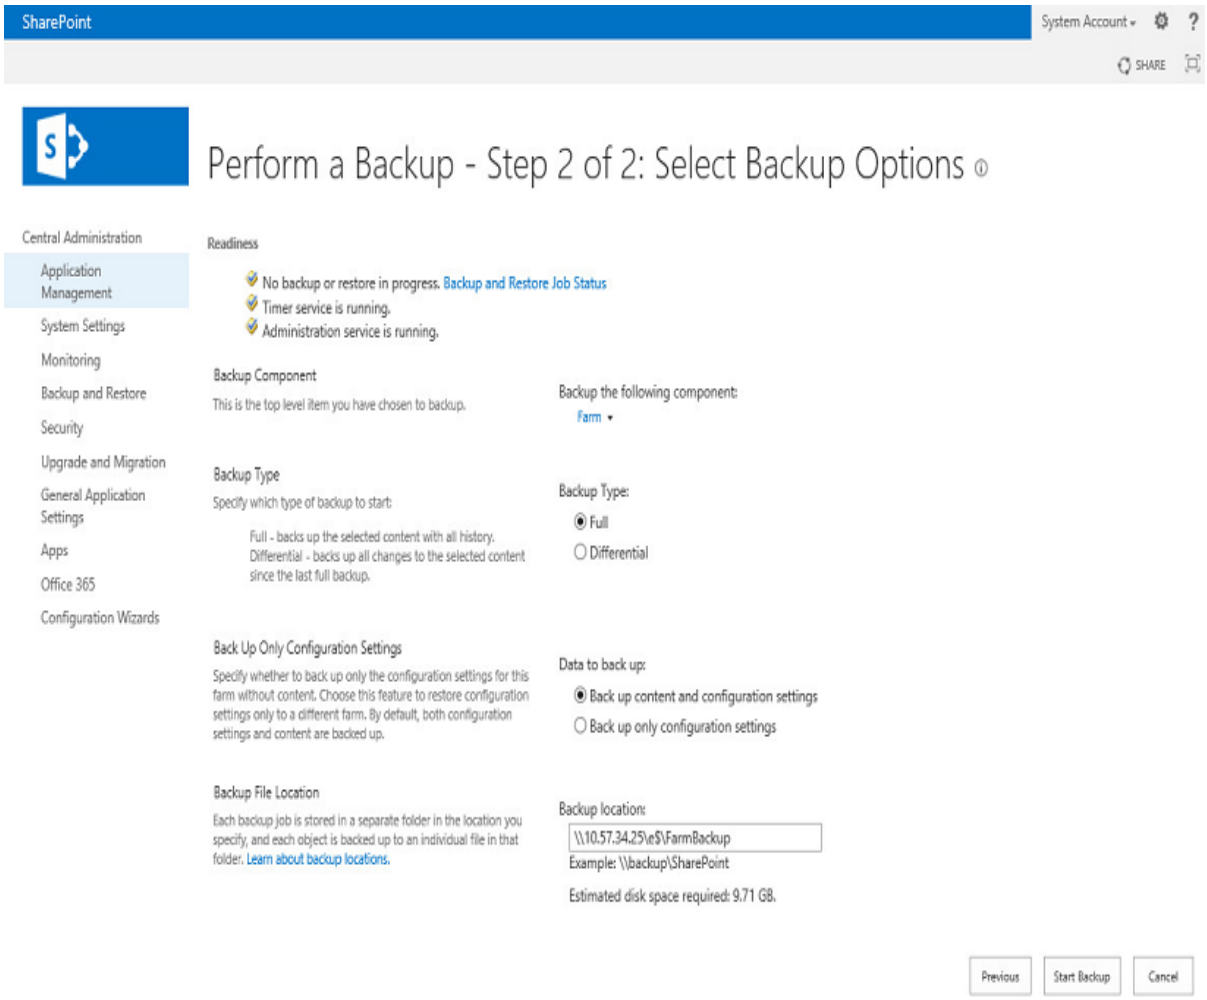
22. Select the component that you want to restore for Farm and then click **Next**.

    For example, Back up Web Applications SharePoint – 80. Expand node Microsoft SharePoint Foundation Web Application and select the component SharePoint – 80

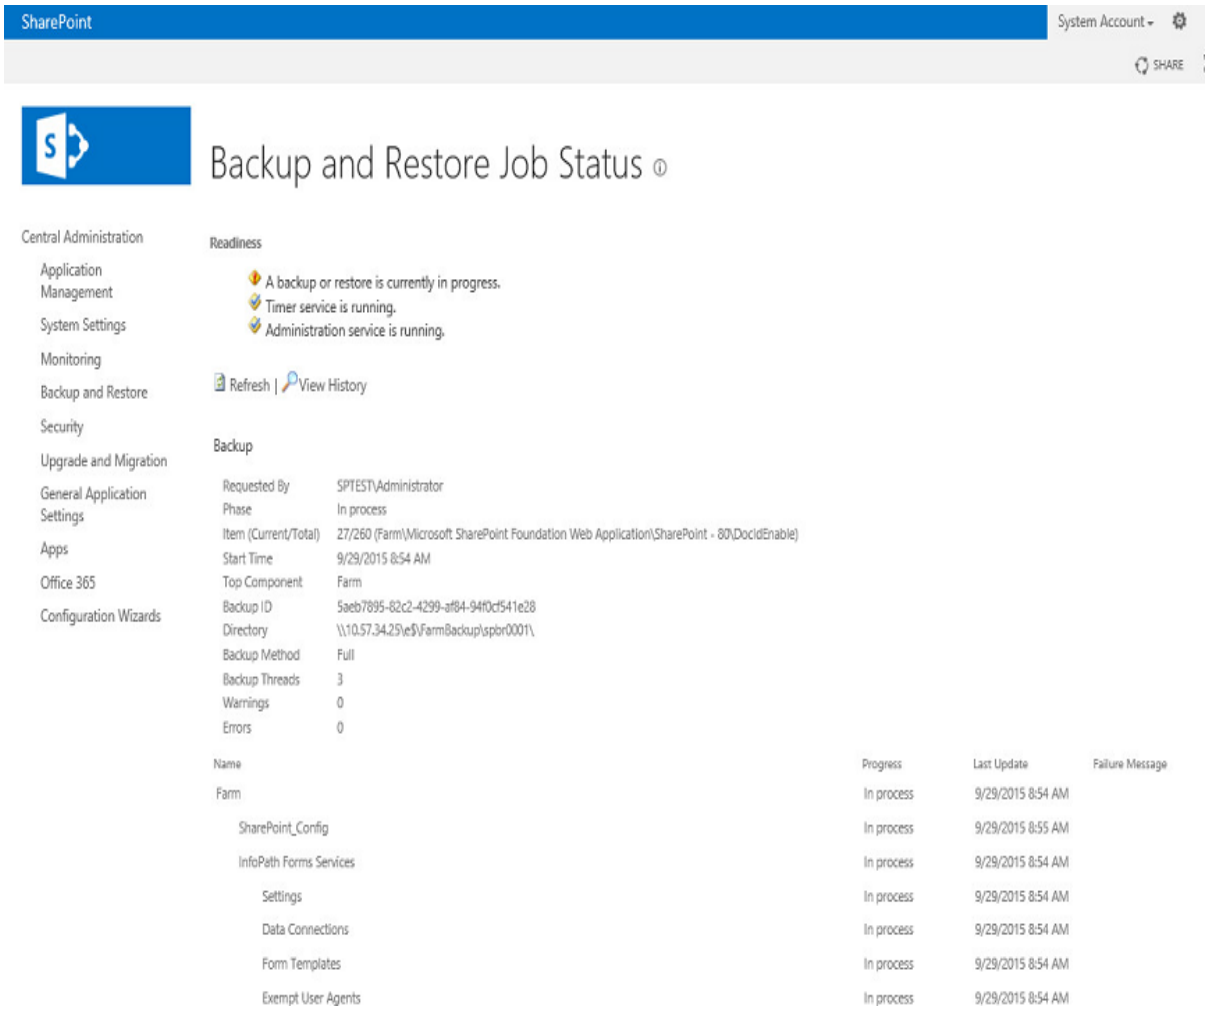
    Back up Access Services. Expand node Shared Services Applications under Shared Services and select the component Access Services 2010



23. Specify Backup Type to Full.
24. Create a Shared Folder on the instant VM, where the machine is Web-Front-End Servers for SharePoint Environment.
25. Provide Full Control access to everyone on the shared folder.
26. Provide the shared folder path in Backup File Location to store the backup file, and then click **Start Backup**.

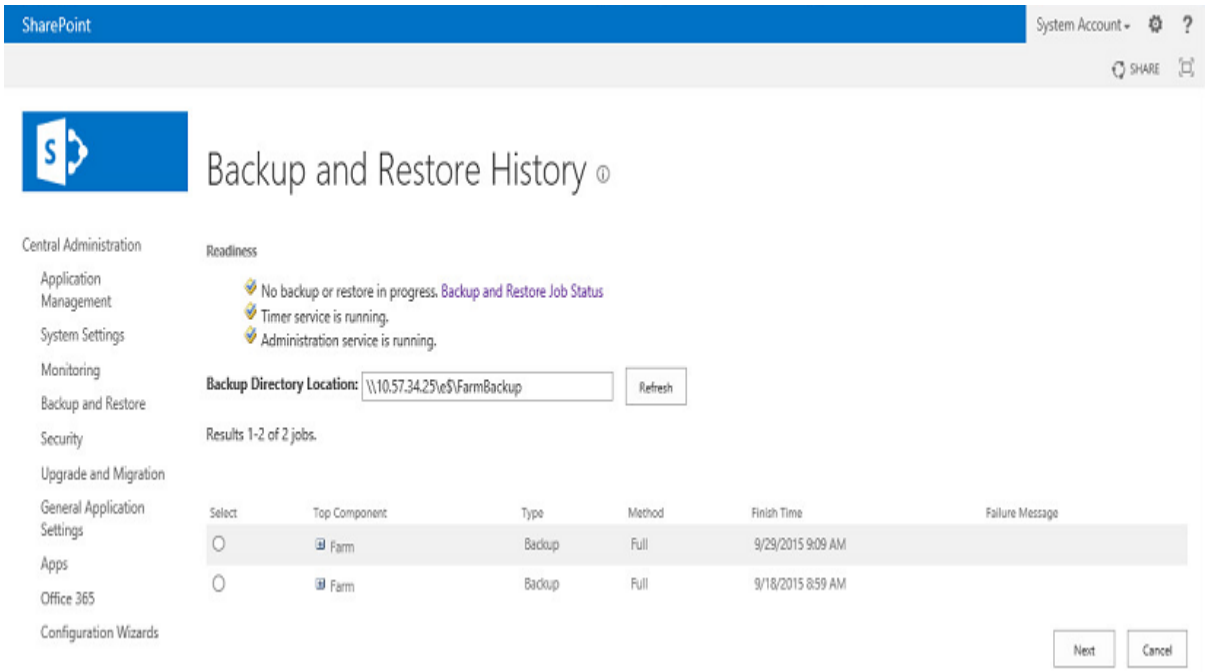


27. Wait for the backup job to complete.



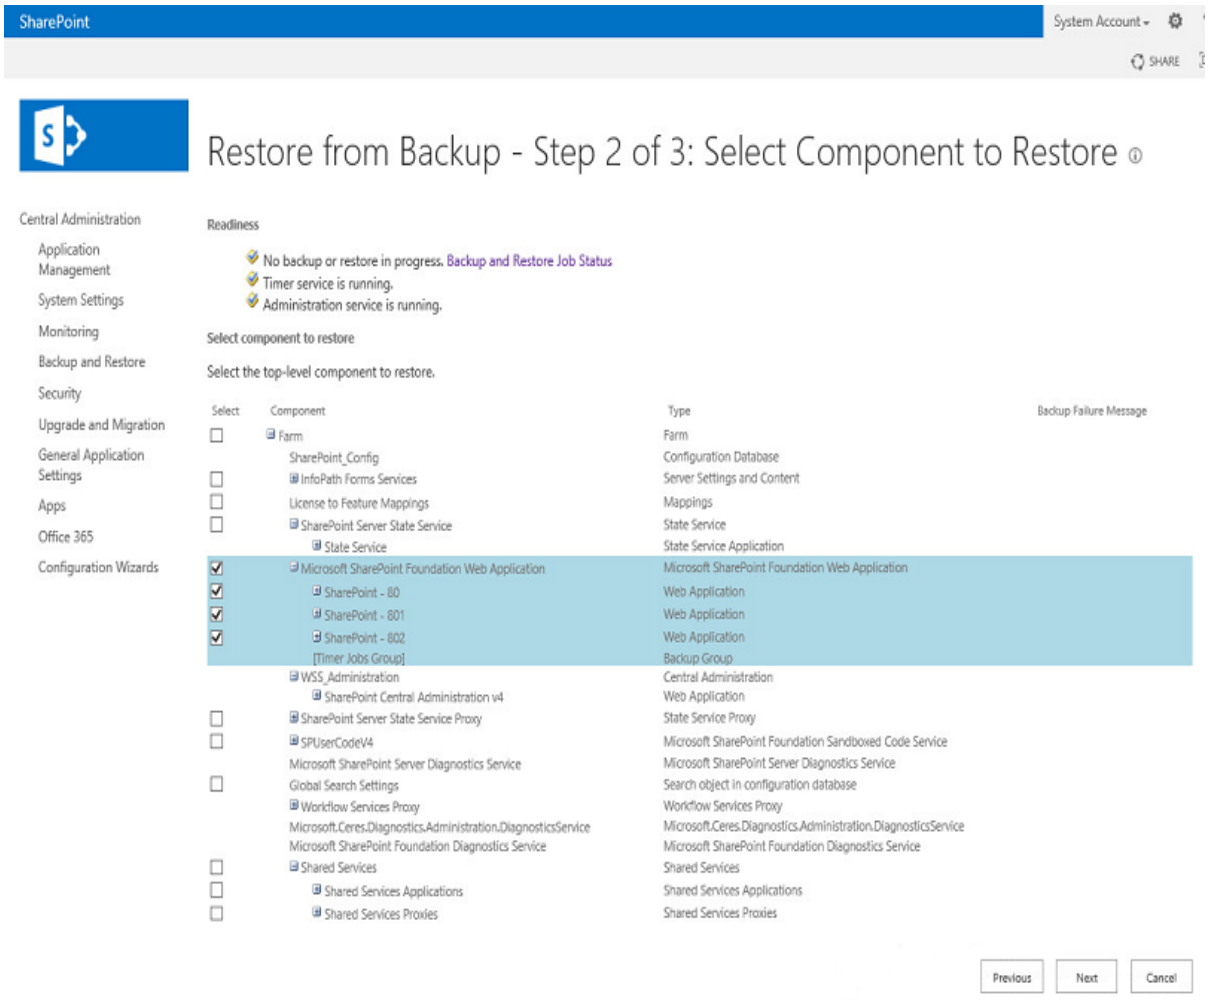
28. After the Farm backup job completes, log in to the machine that is Web-Front-End Servers for the original SharePoint Environment.
29. Open Central Administration, click Backup and Restore, and select the Restore from a backup link.
30. Provide the Shared Folder on Backup Directory Location field, and click **Refresh**.

    Note: Shared folder is already created in one of the previous steps.
31. Select the backup instance from the history list and click **Next**.



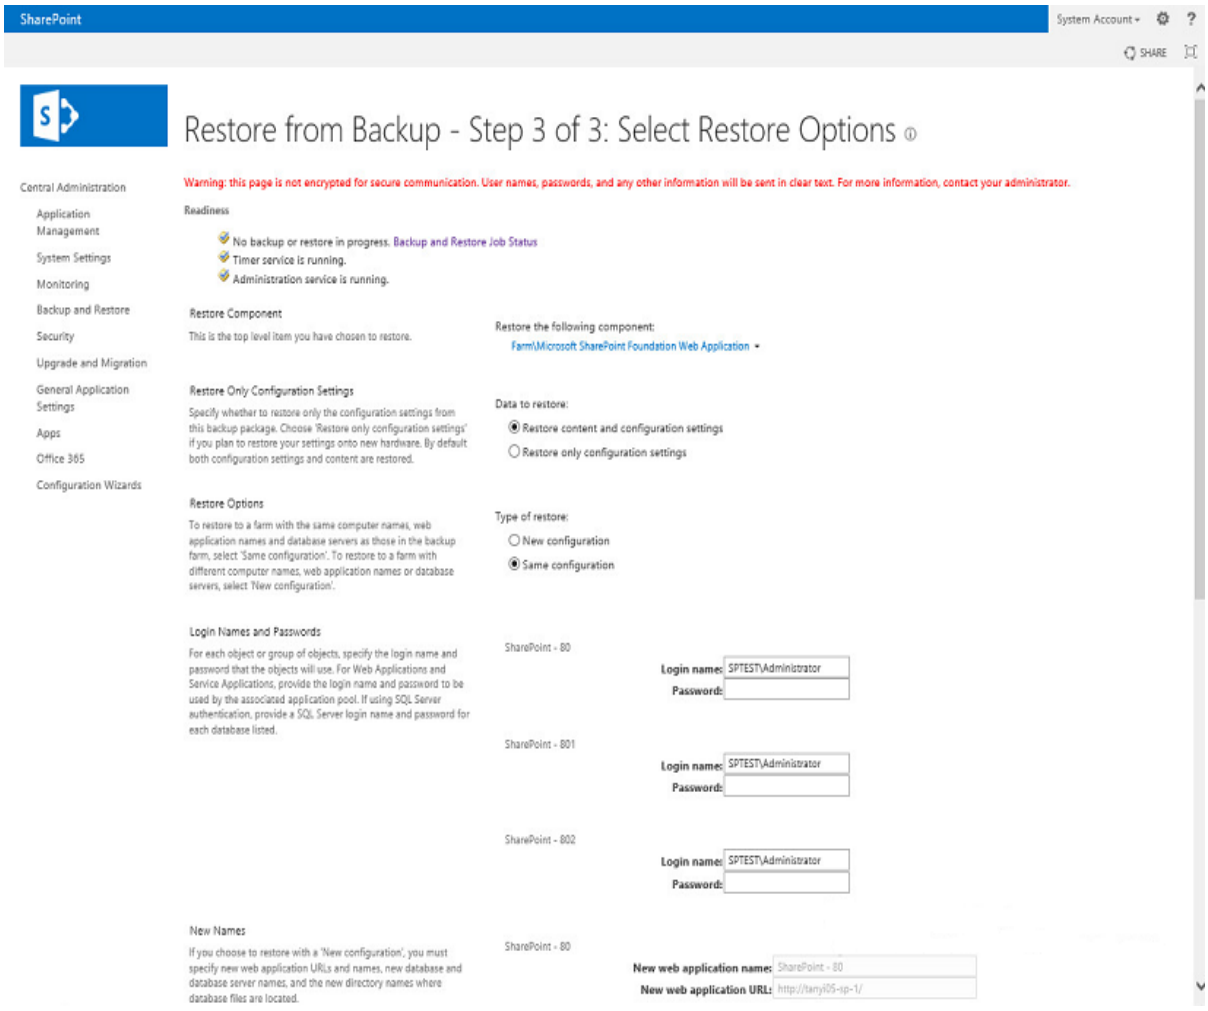
32. Select the configuration and content in the farm to restore, and then click **Next**.

    For example, select all Farm component or SharePoint Services component or Web Applications component.



The page displays various options for the selected service and content configuration.

33. Choose whether you want to:
    - Restore content and configuration settings or Restore only configuration settings
    - Overwrite configuration or create new in restore options.



34. Click the **Start Restore** button to begin the restore process.

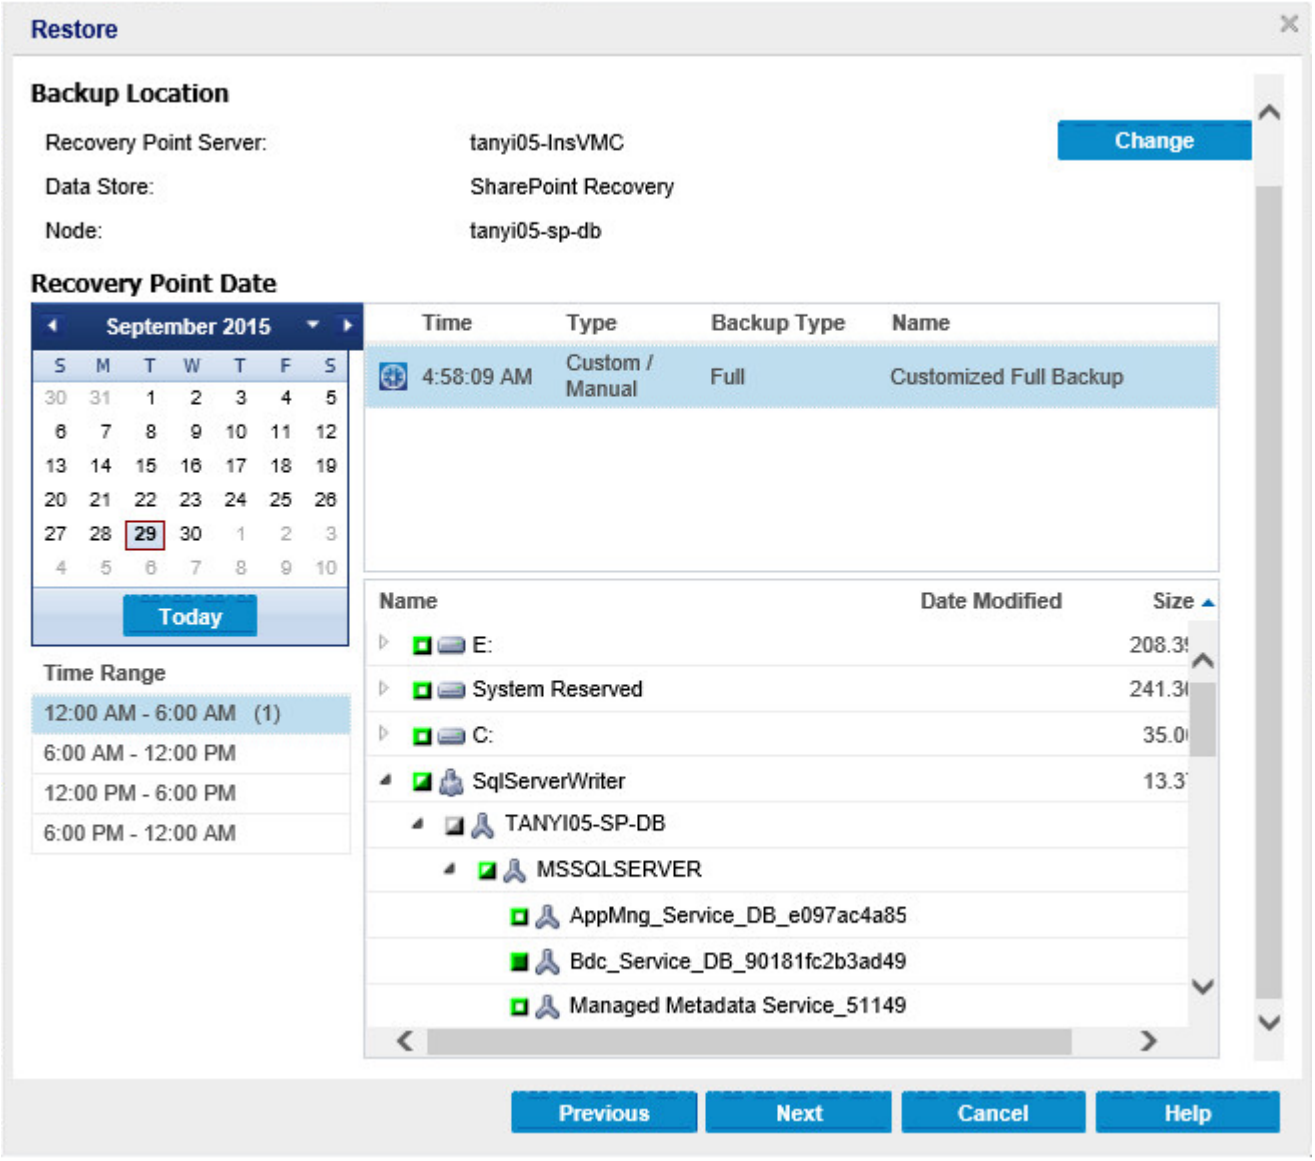
    After the restore job is completed, the selected components in Farm are restored.

Restore using Arcserve UDP Agent UI

Follow these steps:

1. Log in to the Console.
2. Right-click the node that is the Database Server used by the SharePoint.
3. Click Restore.

   The Arcserve UDP Agent UI opens that is hosted on the Database Server.
4. Select Browse Recovery Points.
5. Click the Recovery Point that includes the database, which will be restored.
6. Select database under SqlServerWriter/{SqlServerName}/{SqlServerInstantsName} to restore.
7. Click Next.



8. Select the restore destination and click Next.

   If you select "Restore to original location", the database is restored to the original location. If you select "Restore to alternative location", the database is restored to the specified location. Both the restored database will be attached in SQL Server automatically. If you select "Dump file only", the database data file and log file are saved to the specified location.

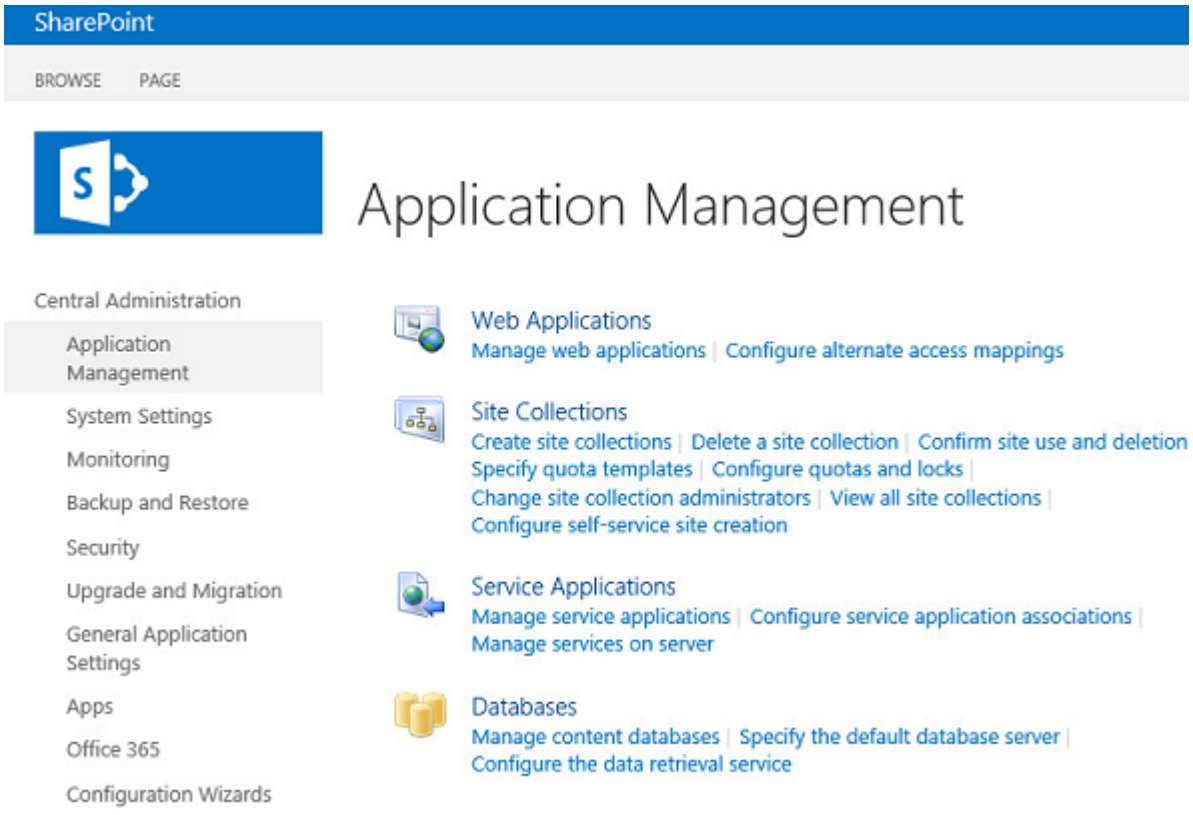
9. Click Finish and wait for the restore job to complete.

   Note: After select "Dump file only", attach the database first. Follow the steps in Restore using Mount Database from Recovery Point and verify if the new content database is associated with its original web application. If not, add the restored content database to original web application.

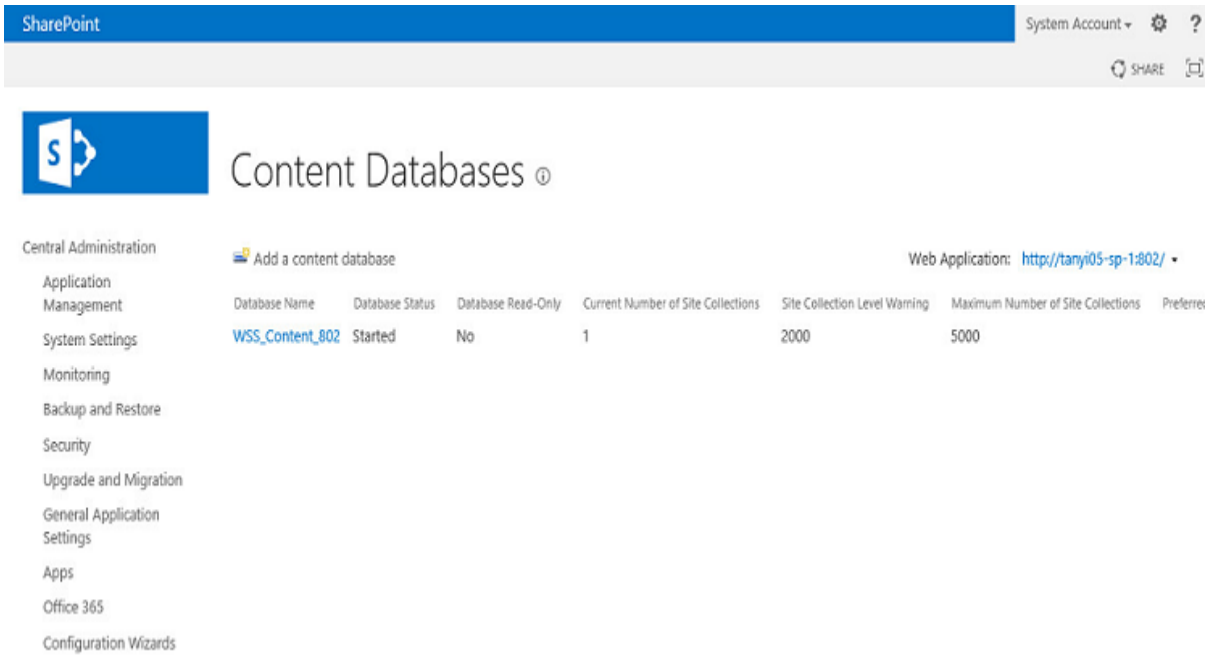
Add the Restored Content Database to Original Web Application

Follow these steps:

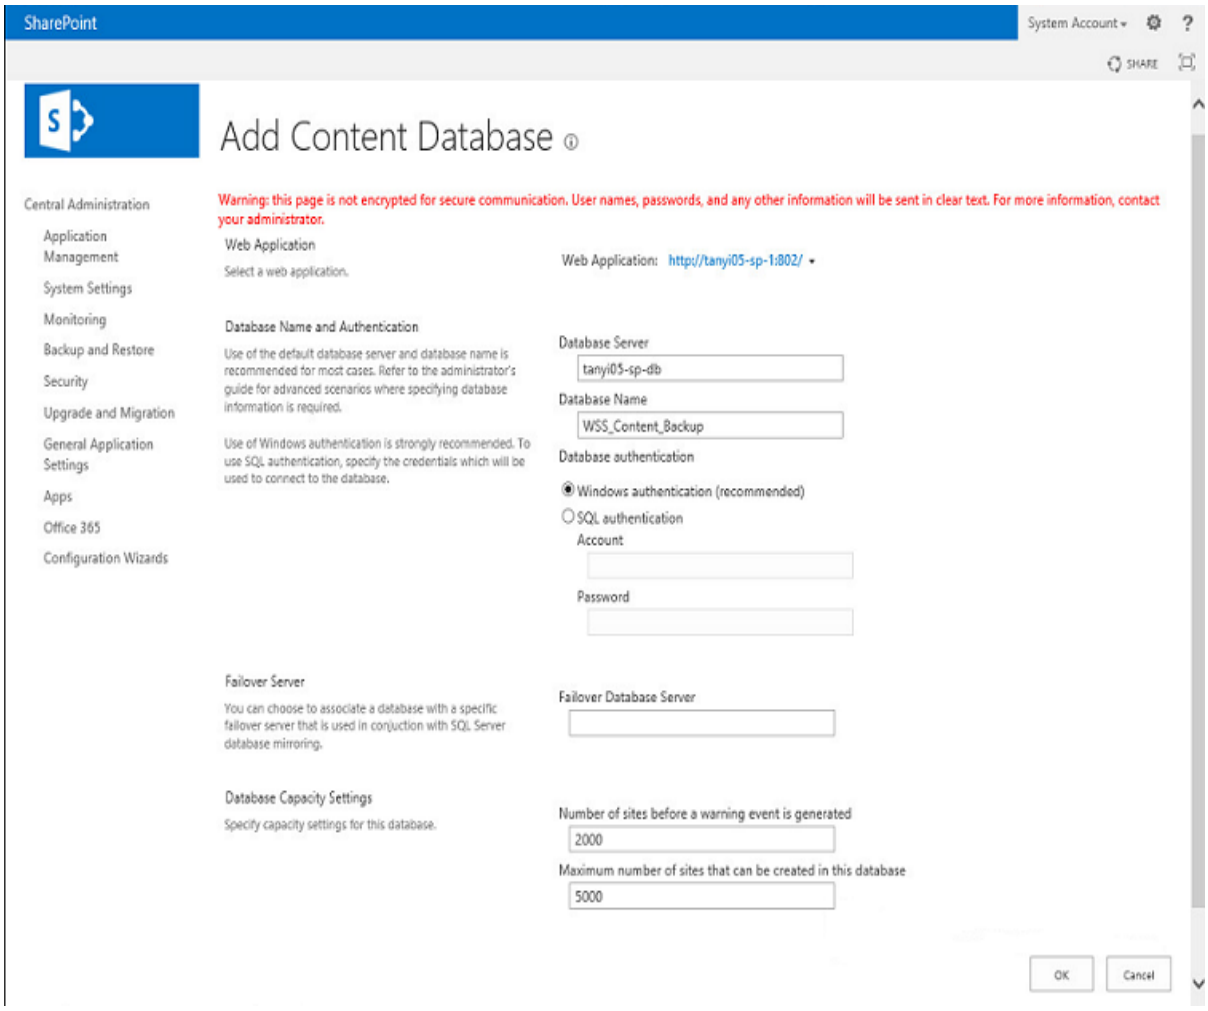
1. Open SharePoint Central Administration and select Application Management.



2. Select Management content databases.



3. Select the web application and click Add a content database.
4. Type Database Server and Database Name, (for example, WSS_Content_Backup) and then click OK.

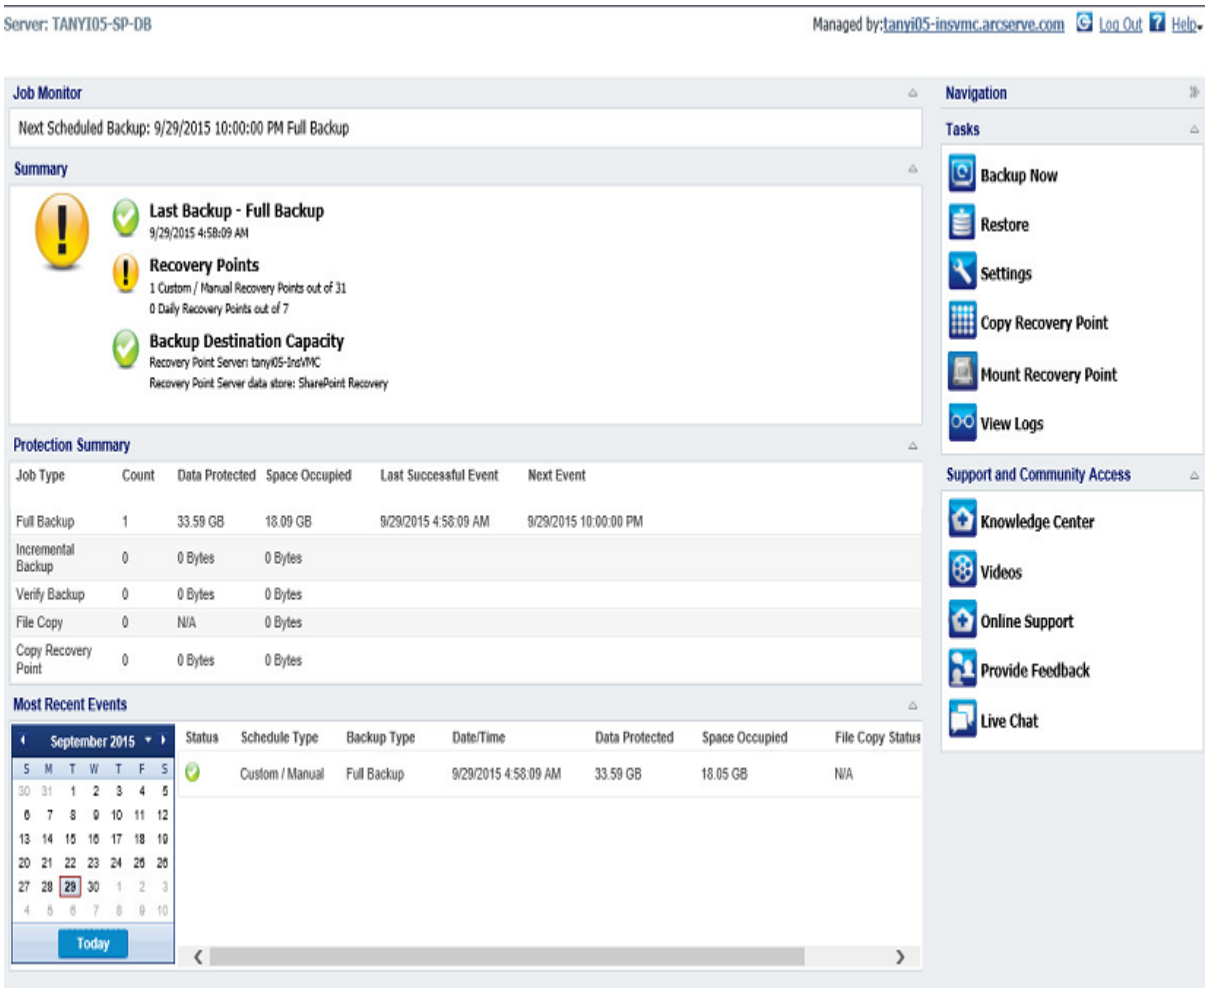


The content database is now associated with its original web application.

Restore using Mount Database from Recovery Point

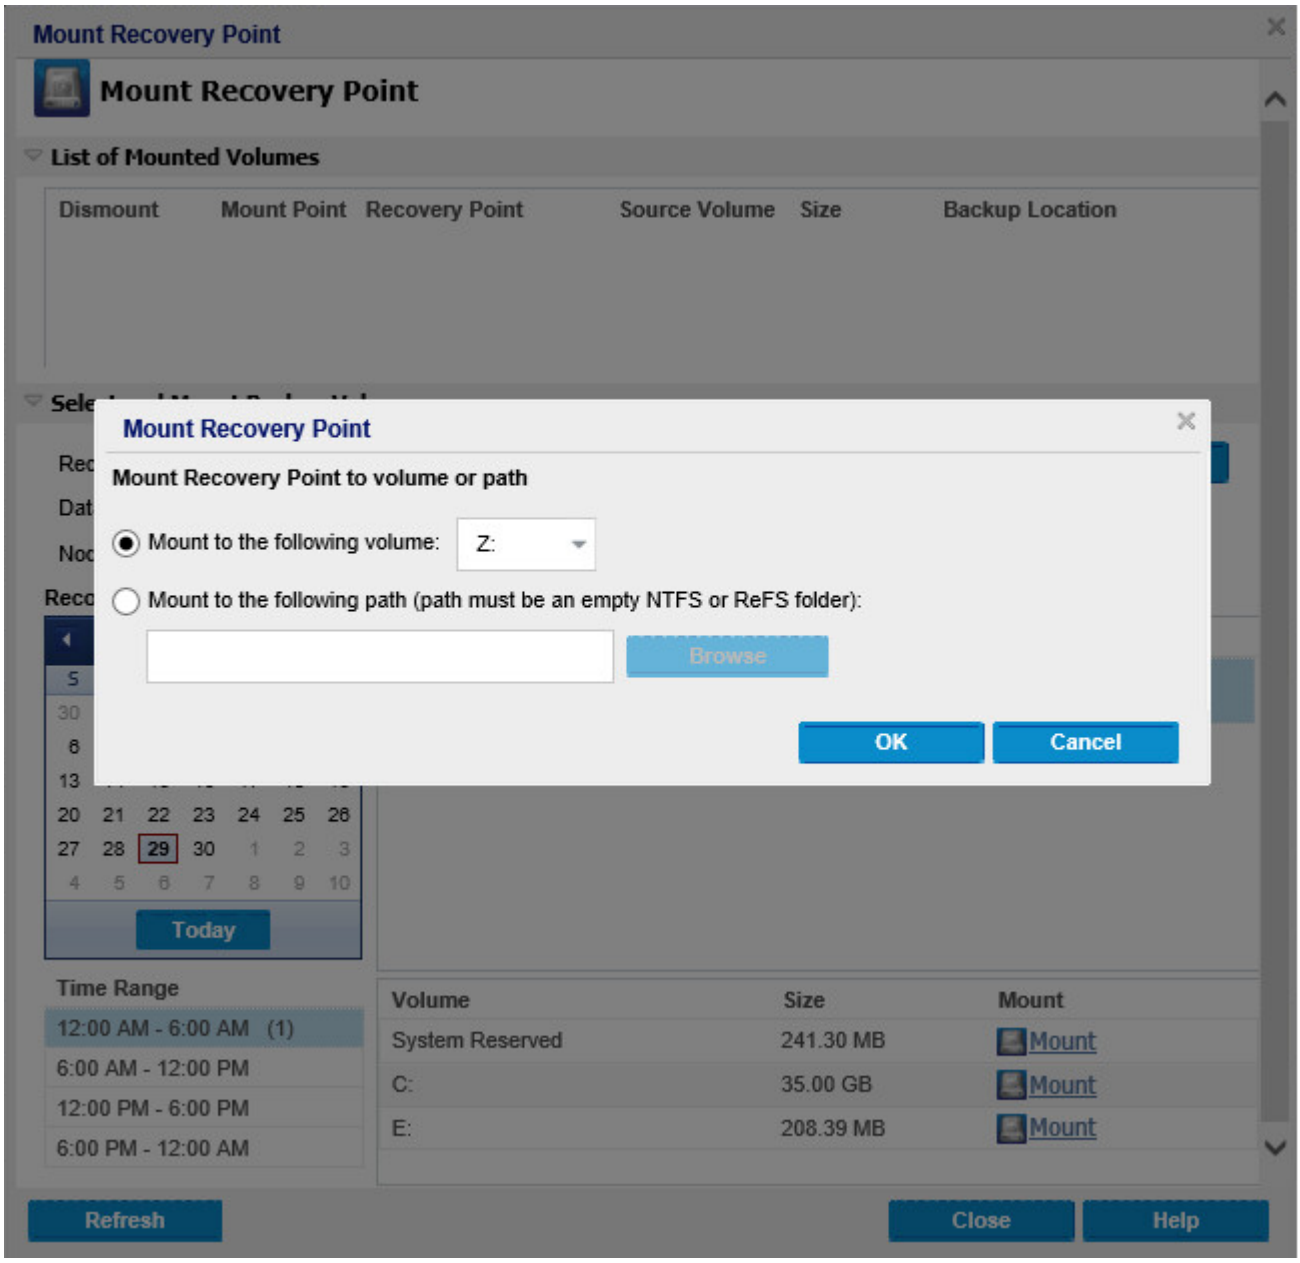
Follow these steps:

1. Open the Arcserve UDP Agent UI that is hosted on the Database Server.

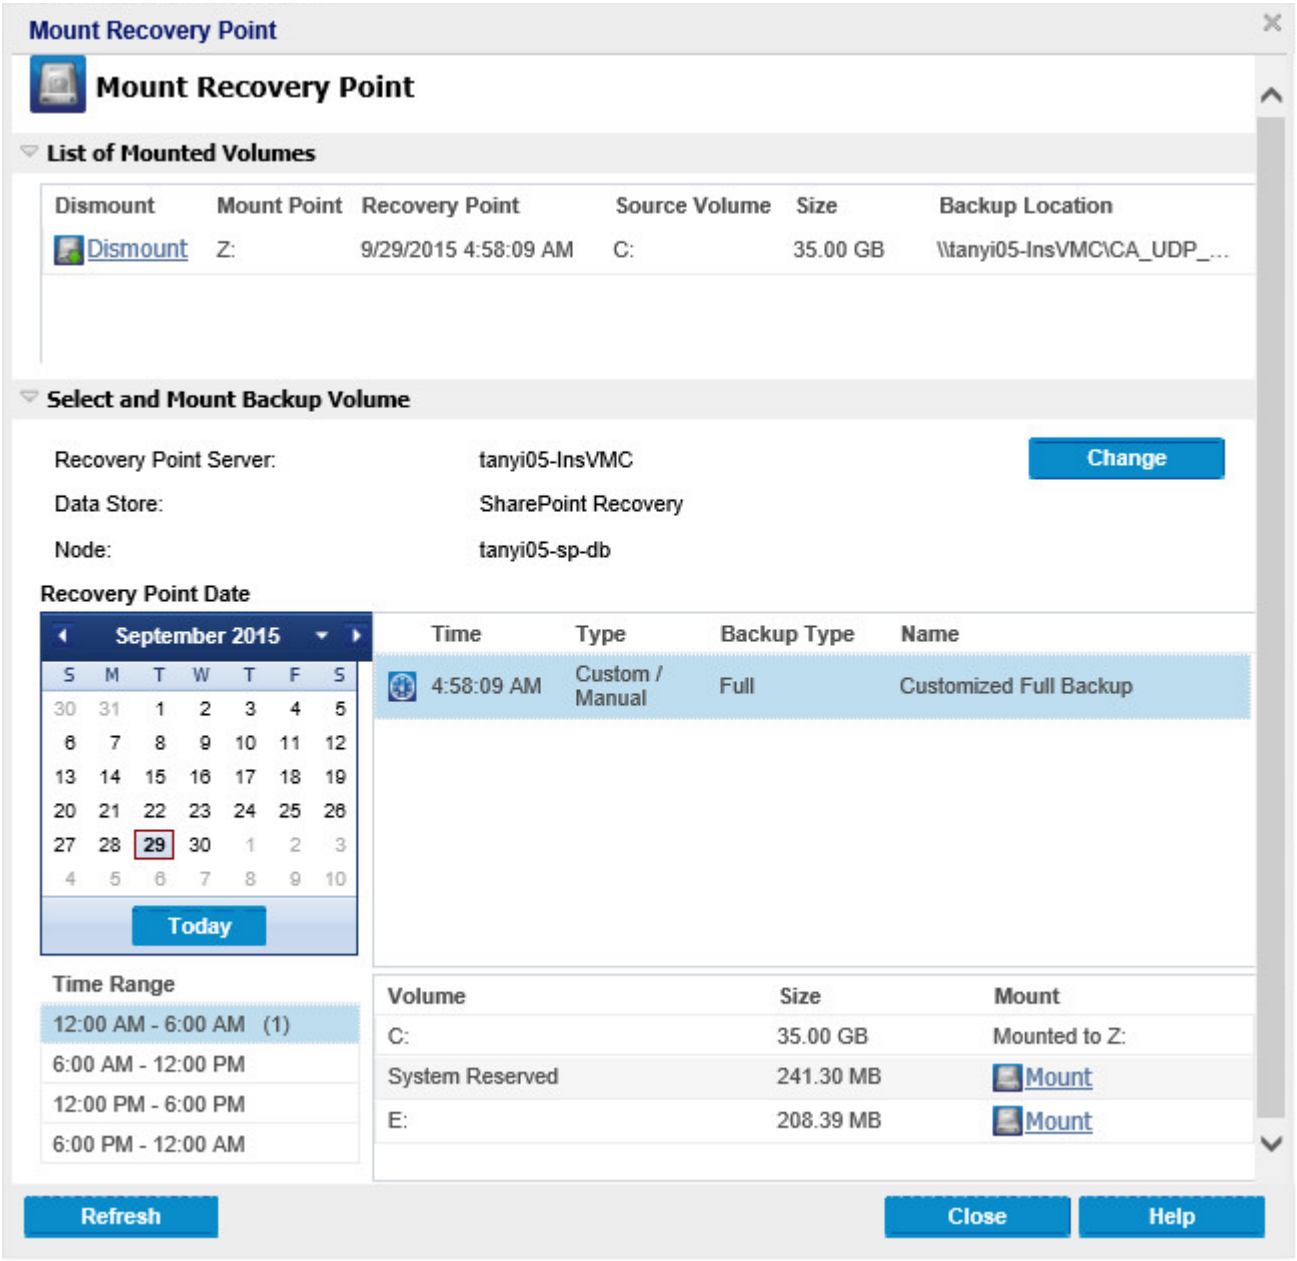


2. Click **Mount Recovery Point** from the Tasks pane.
3. Select the volume that includes SQL server Database for SharePoint.

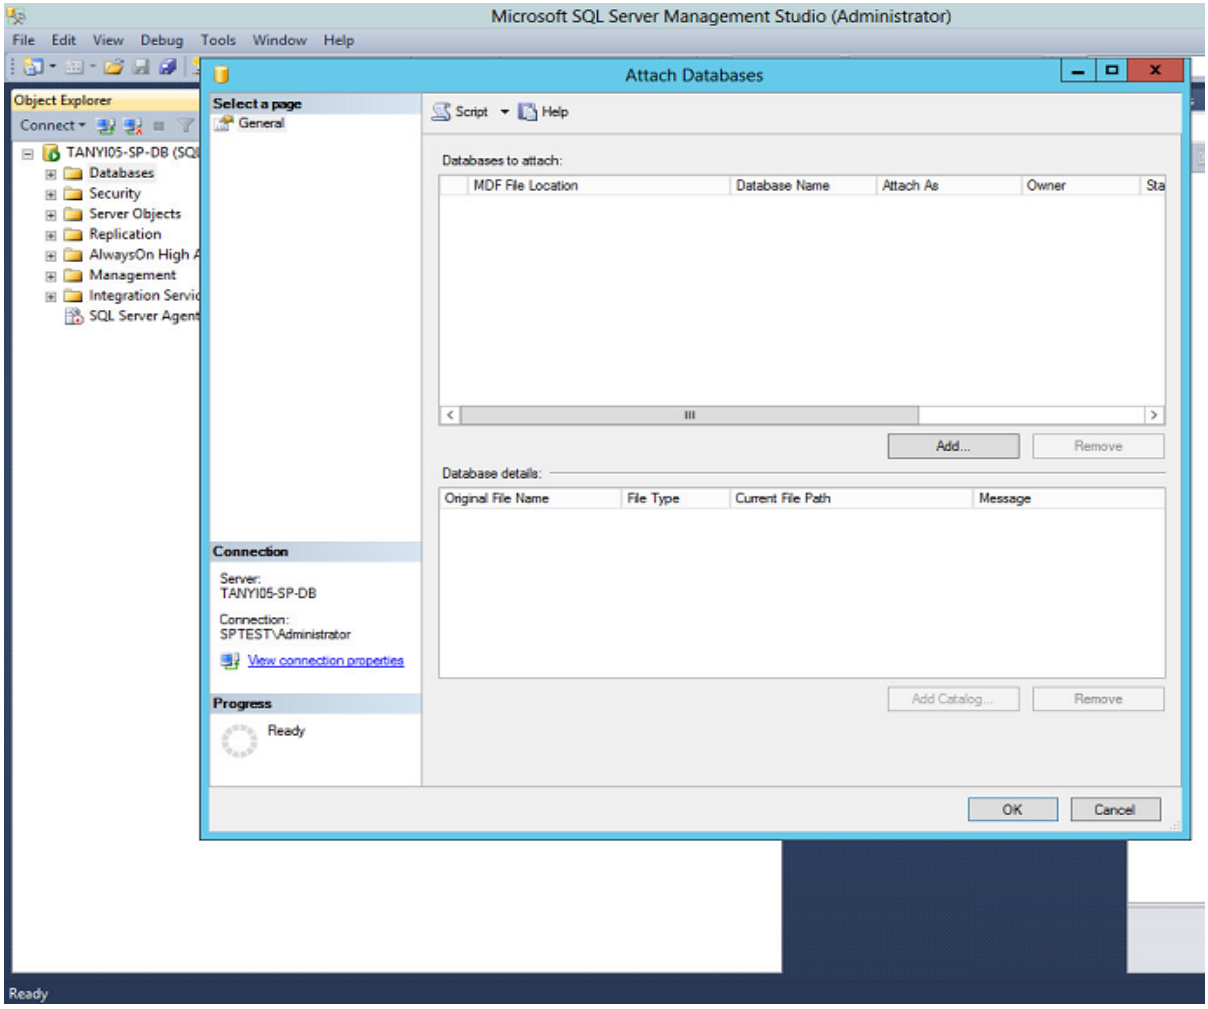
   For example, by default, the database file is saved on "C:\Program Files\Microsoft SQL Server\MSSQL12.MSSQLSERVER\MSSQL\DATA", and you select volume C.
4. Select a new volume name to mount in the following dialog and click **OK**.



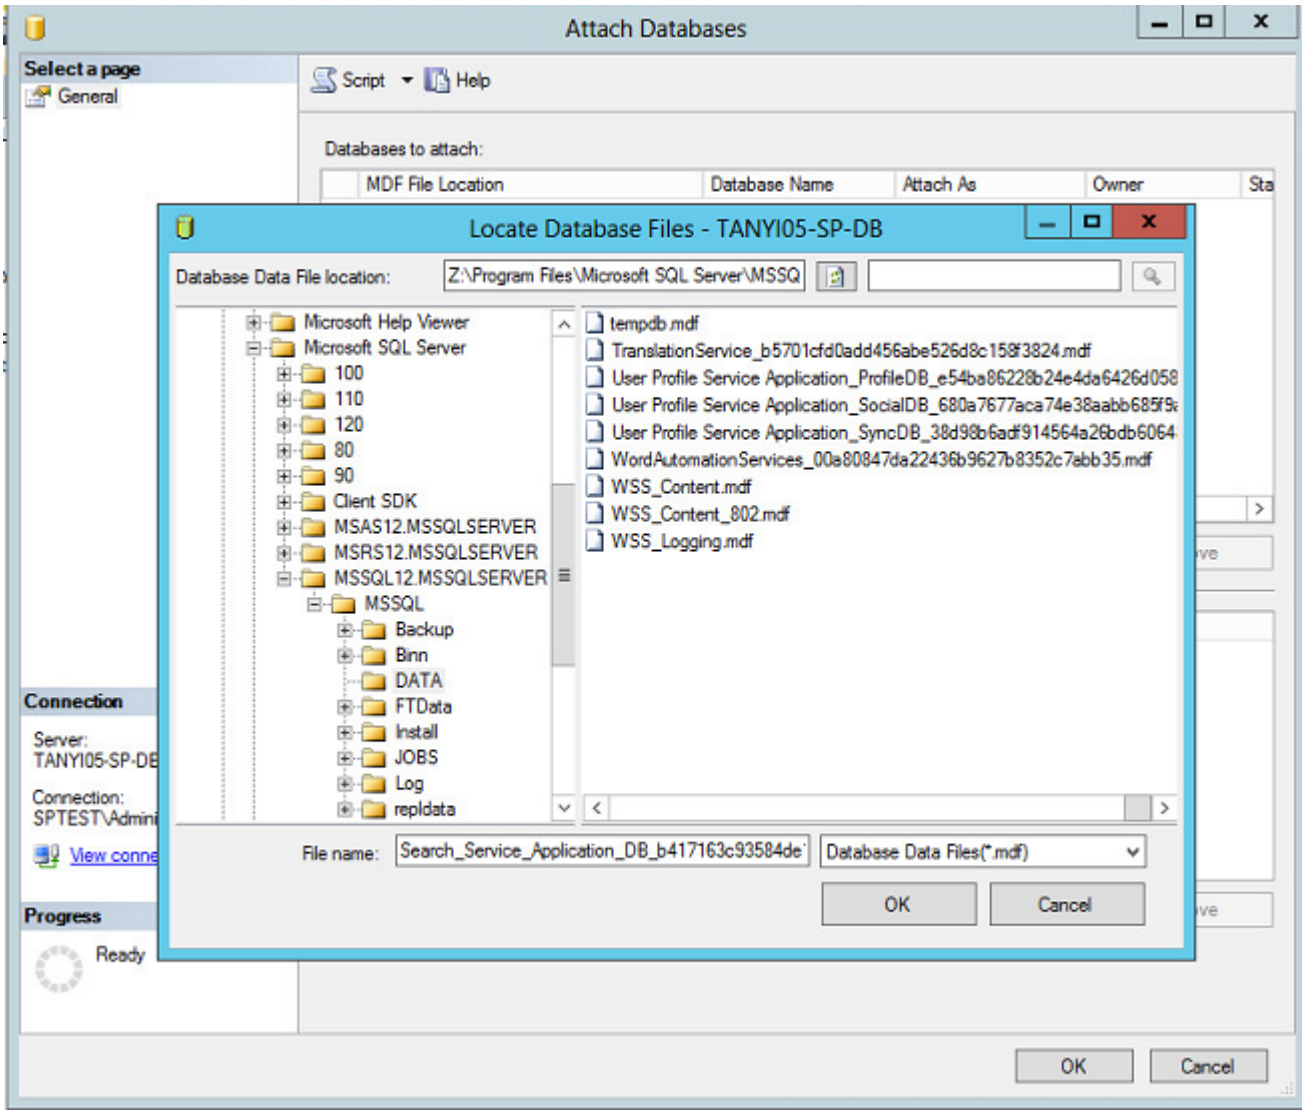
The new volume is listed in Mounted Volumes on the UI.



5. Open SQL Server Management Studio to attach the backup database.
6. Right-click the Databases folder and select Attach.



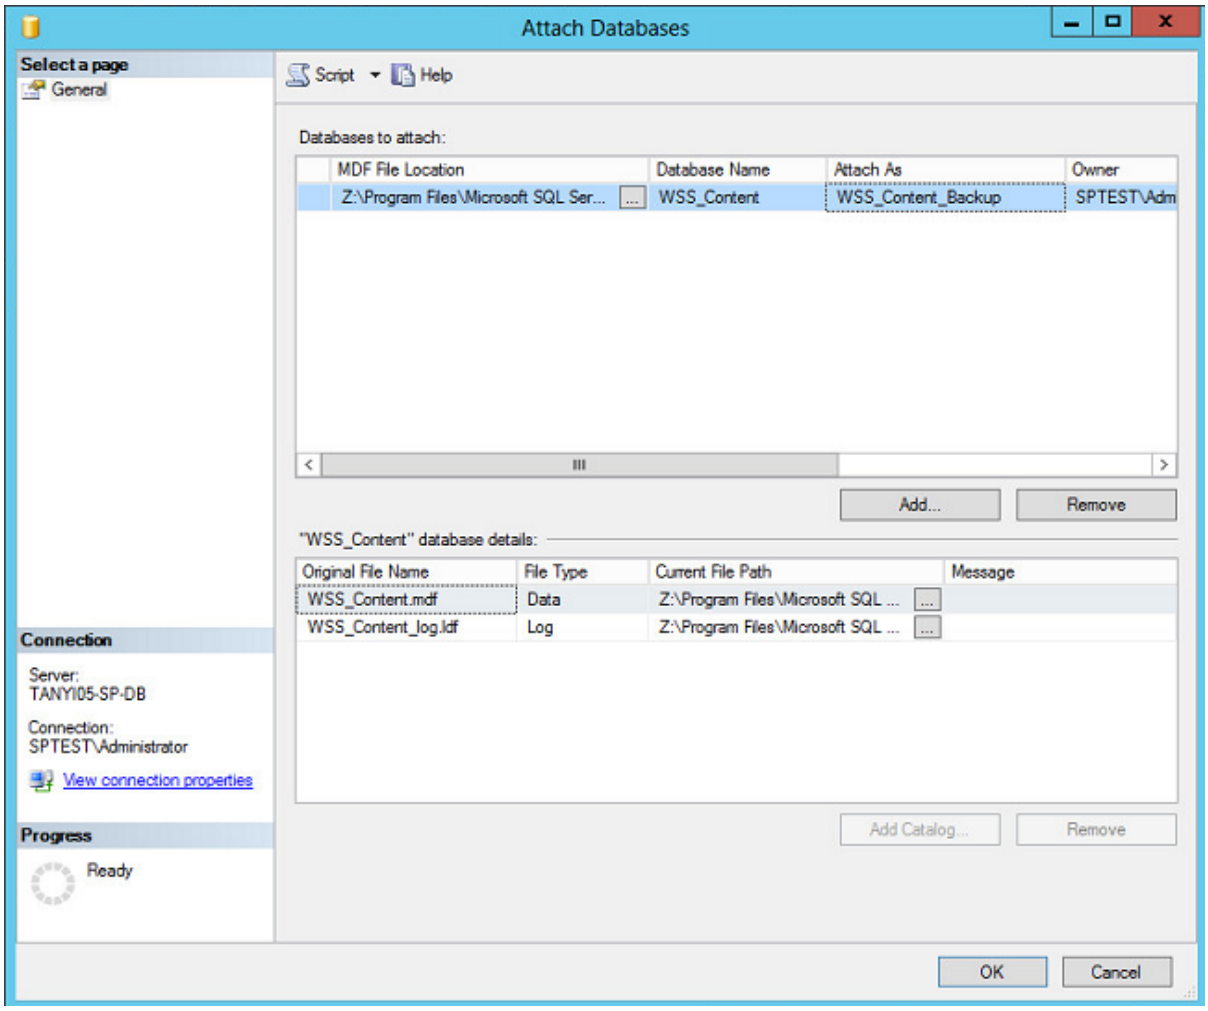
7. Click Add to select the database file that you want to attach.



8. If the mounted volume is Z, select the database data file location as "Z:\Program Files\Microsoft SQL Server\MSSQL12.MSSQLSERVER\MSSQL\DATA".
9. Select the file named "WSS_Content.mdf" and click OK.

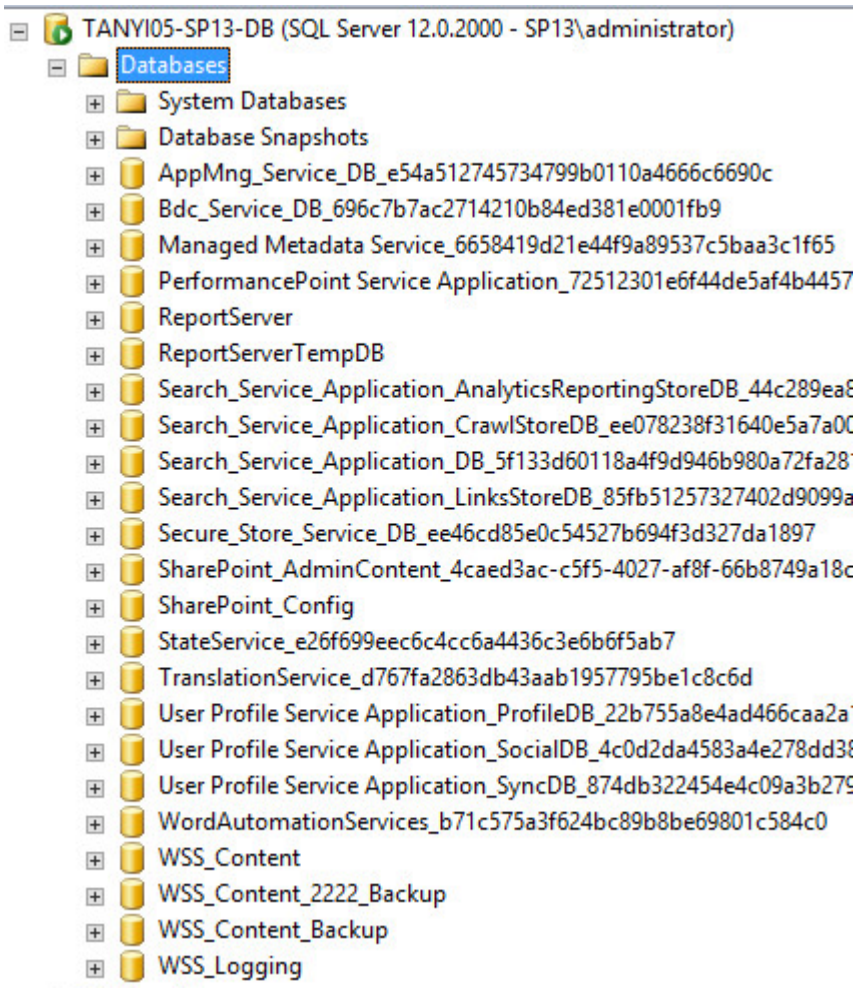
   Note: "WSS_Content.mdf" is the default database data file name for SharePoint Web Application. If you want to restore another database which is created by a new web application, use the related database data file name.
10. Click the Attach As column to type the database name (Example, "WSS_Content_Backup") and click OK.

    Note: Before you restore a content database, make sure the name of the database does not exist in any other web application.

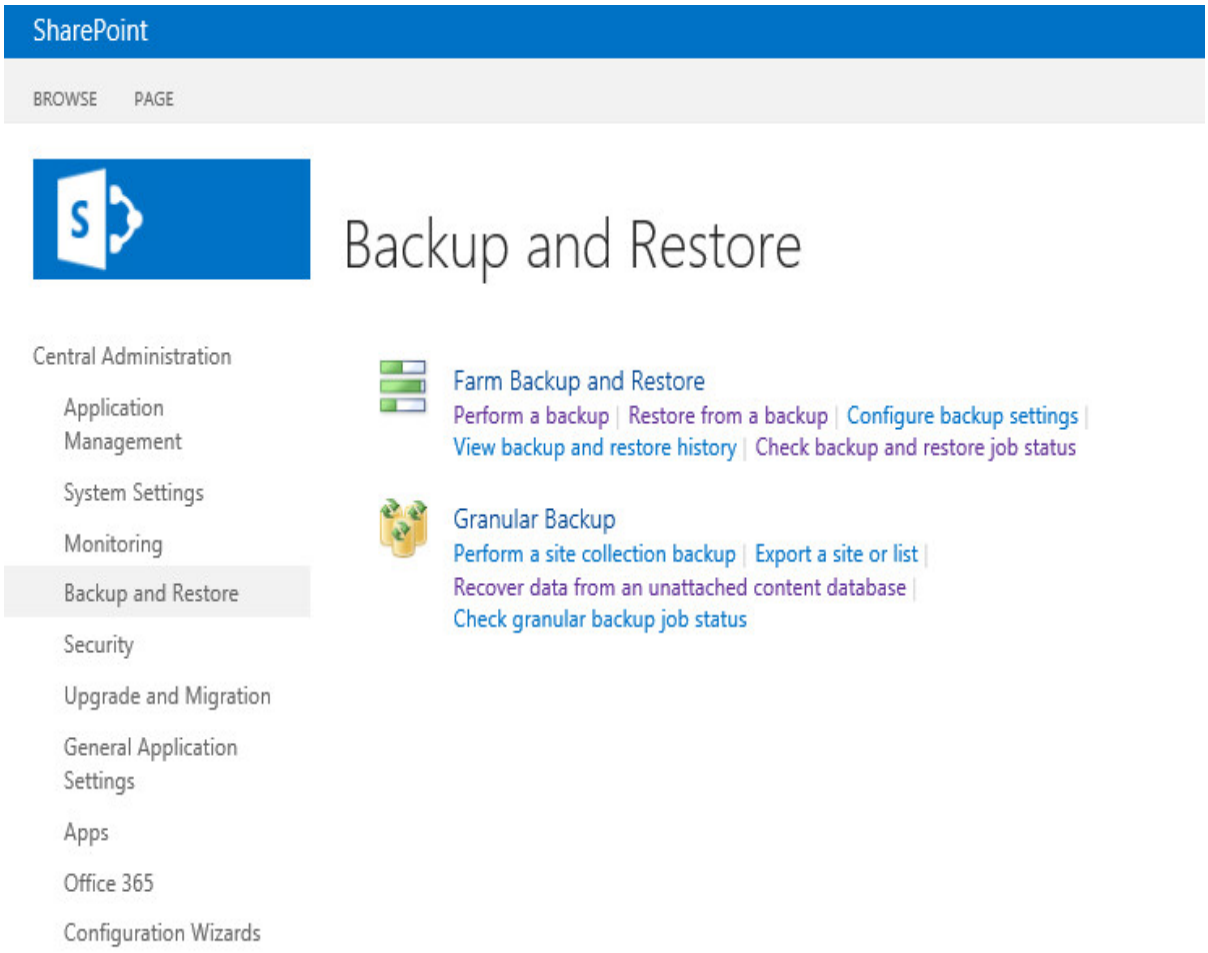


The new database is attached under the Databases folder.

Note: The new database is not associated with any web application.

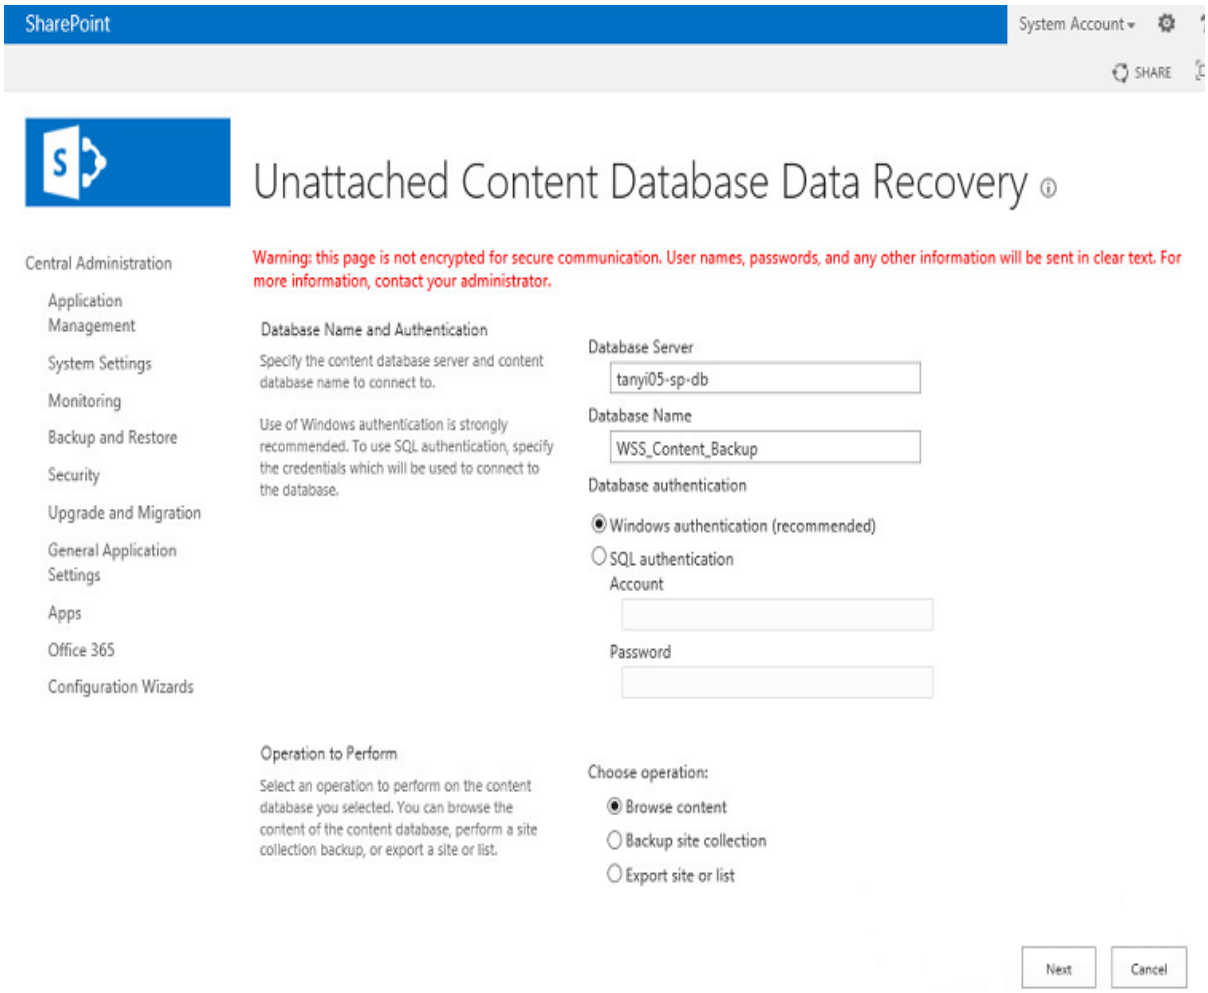


11. Log into the machine that is Web-Front-End Servers in SharePoint Server Farm.
12. Open Central Administration and click Backup and Restore heading.



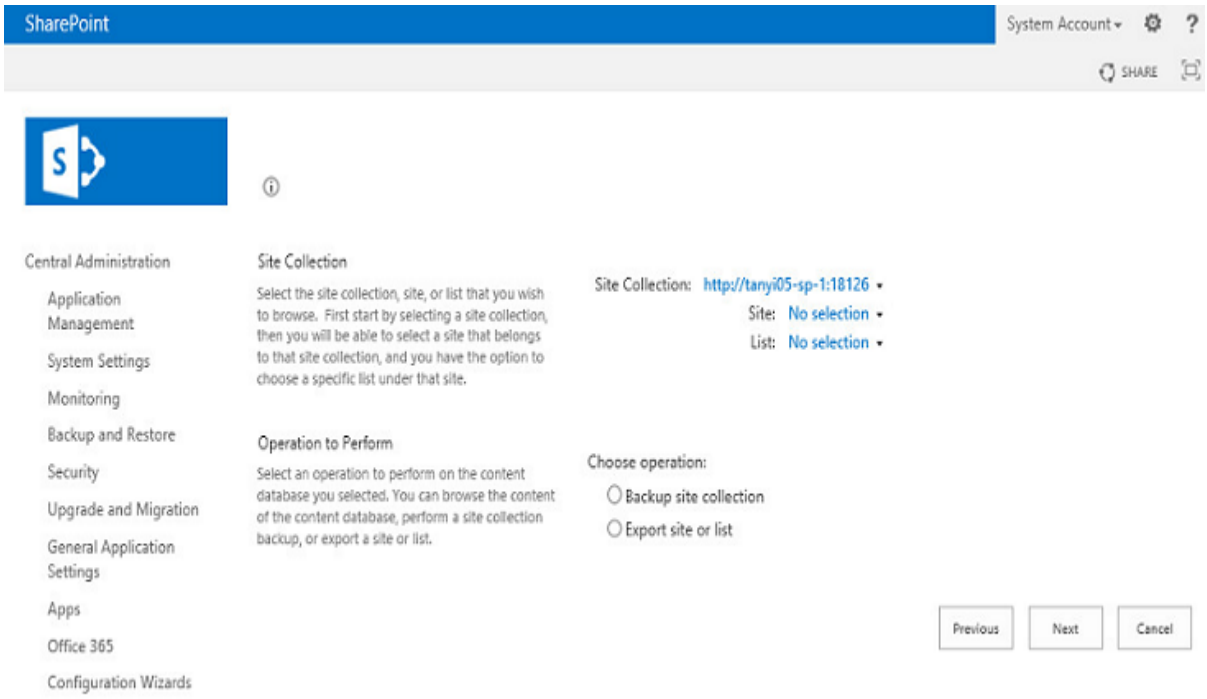
13. Click the Recover Data from an Unattached Content Database link, and provide the SQL Server name and database name for the unattached database, and use Windows authentication.
14. Select the Browse content option and click **Next**.

    Note: SQL Server name is the name of database servers used by SharePoint Server Farm, Database name is the name of the newly attached database.



15. Choose whether to back up the site collection or export the selected site and list.

    Note: By default, if the database is unattached content database, then the site collection URL still includes the central administration web application port number.



You have restored using mount database from a recovery point.

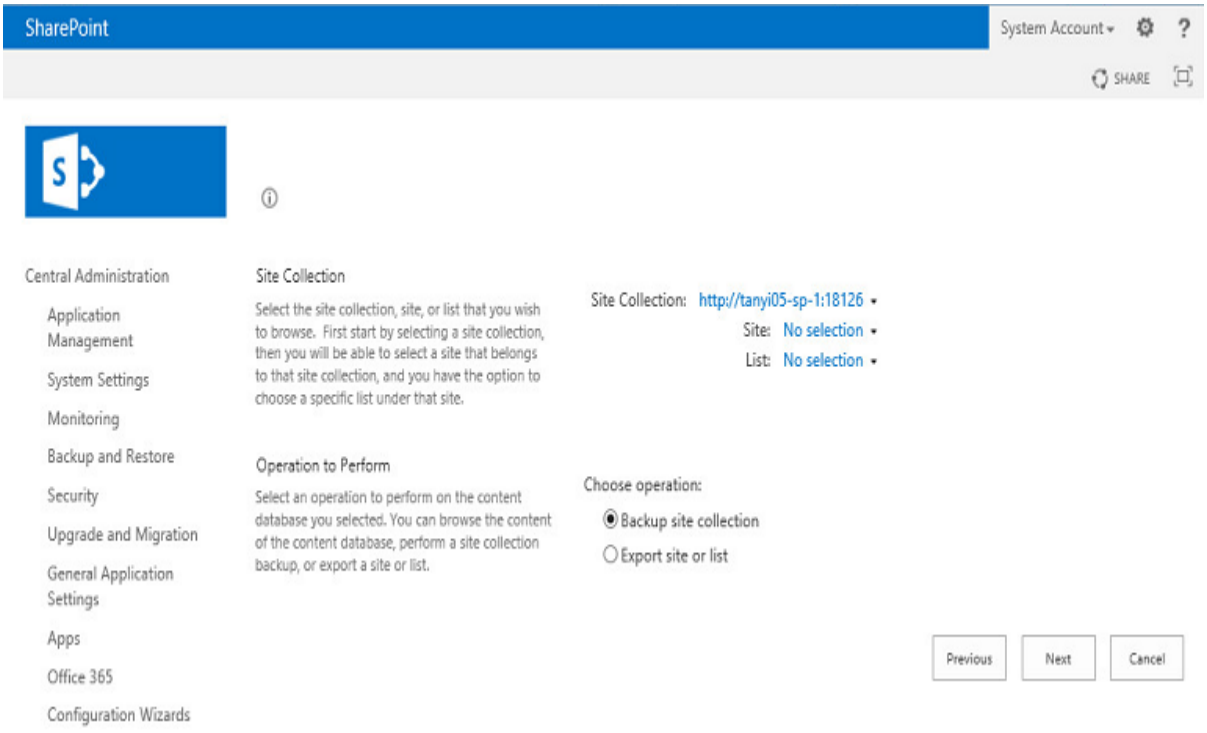
Restore a Site Collection

Follow these steps:

1. Recover the site collection content from an unattached content database.

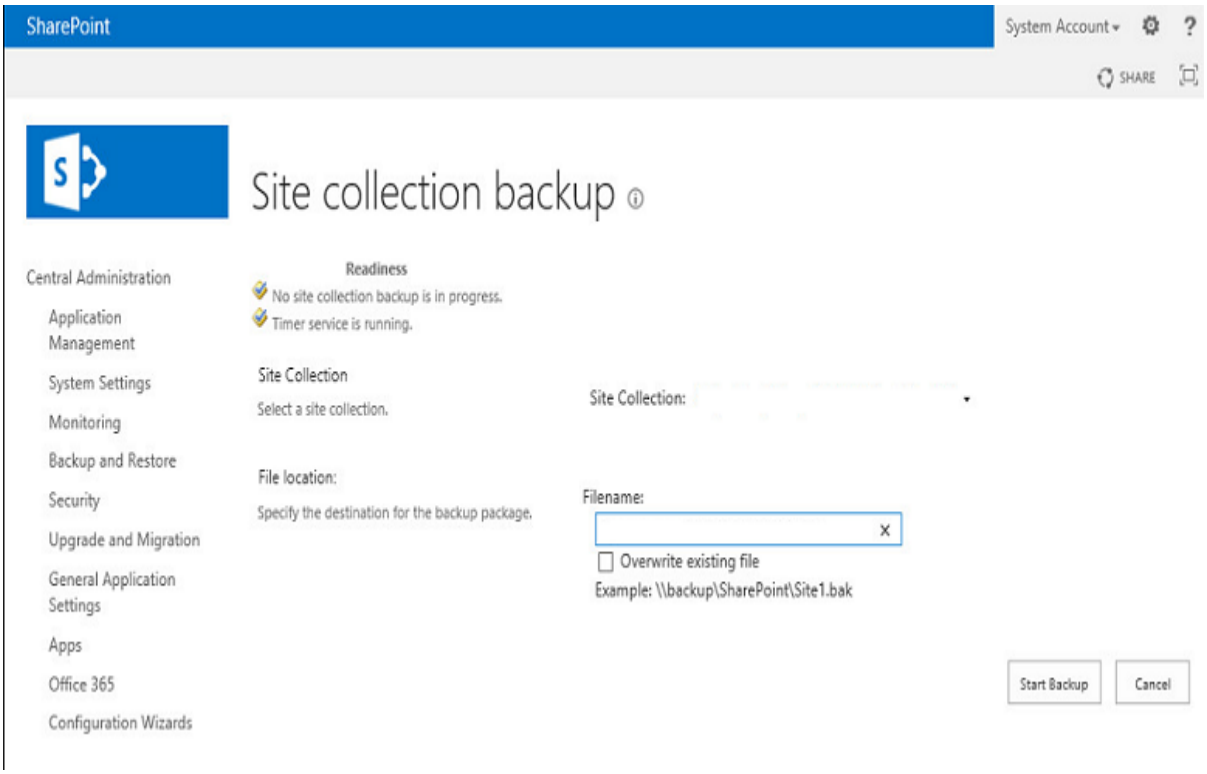
   From Central Administration

   a. Select the Backup site collection option, and click the Next button.



   b. Select the site collection and provide the file location for the backup package.

   c. Select Overwrite existing file.



d. Click the Start Backup button to begin the backup.

   The site collection is backed up to a file.

Use PowerShell commands

a. $database = Get-SPContentDatabase -ConnectAsUnattachedDatabase -DatabaseName xxxx -DatabaseServer xxxx

   ConnectAsUnattachedDatabase: Specifies that only unattached databases in the farm are returned.

   DatabaseName: Specifies the name of the content database.

   DatabaseServer: Specifies the name of the host server for the content database specified in the DatabaseName parameter.

   For more information, see the article from Microsoft.

b. Backup-SPSite -Identity xxxx -Path xxxx

   Identity: Specifies the URL or GUID of the site collection to be backed up.

   Path: Specifies the full path of the backup file (Example, C:\-backup\sitecollection.bak).

   For more information, see the article from Microsoft.

2. Click SharePoint Management Shell to launch the console.
3. Use PowerShell commands to restore the site collection.

```
Restore-SPSite -Identity xxxx -Path xxxx
```

Identity: Specifies the URL location to which the Site Collection is restored. (For example, http://www.contoso.com)

Path: Specifies a valid path of the backup location.(For example, C:\-backup\sitecollection.bak)

For more information, see the article from Microsoft.

Note : Restoring a site collection to its original location will fail. You can perform the following steps:

a. New-SPContentDatabase -Name xxxx -DatabaseServer xxxx -WebApplication xxxx

   Name: Specifies the new content database to create within the farm.

   DatabaseServer: Specifies the name of the host server for the content database specified in the Name parameter.

   WebApplication: Attaches the content database to the specified SharePoint Web application.

b. Restore-SPSite -Identity xxxx -Path xxxx -GradualDelete -DatabaseServer xxxx -DatabaseName xxxx

   Identity: Specifies the URL location to which the Site Collection is restored. (For example, http://www.contoso.com)

   Path: Specifies a valid path of the backup location.(Example, C:\-backup\sitecollection.bak)

   GradualDelete: Specifies that the site collection being overwritten with the Force parameter should be gradually deleted over time by a timer job, instead of deleting all at once. It reduces the impact on SharePoint 2010 Products and SQL Server performance.

   DatabaseName: Specifies the SQL Server content database where the site collection data will be stored.

   DatabaseServer: Specifies the name of the SQL Server containing the content database specified by the DatabaseName parameter.

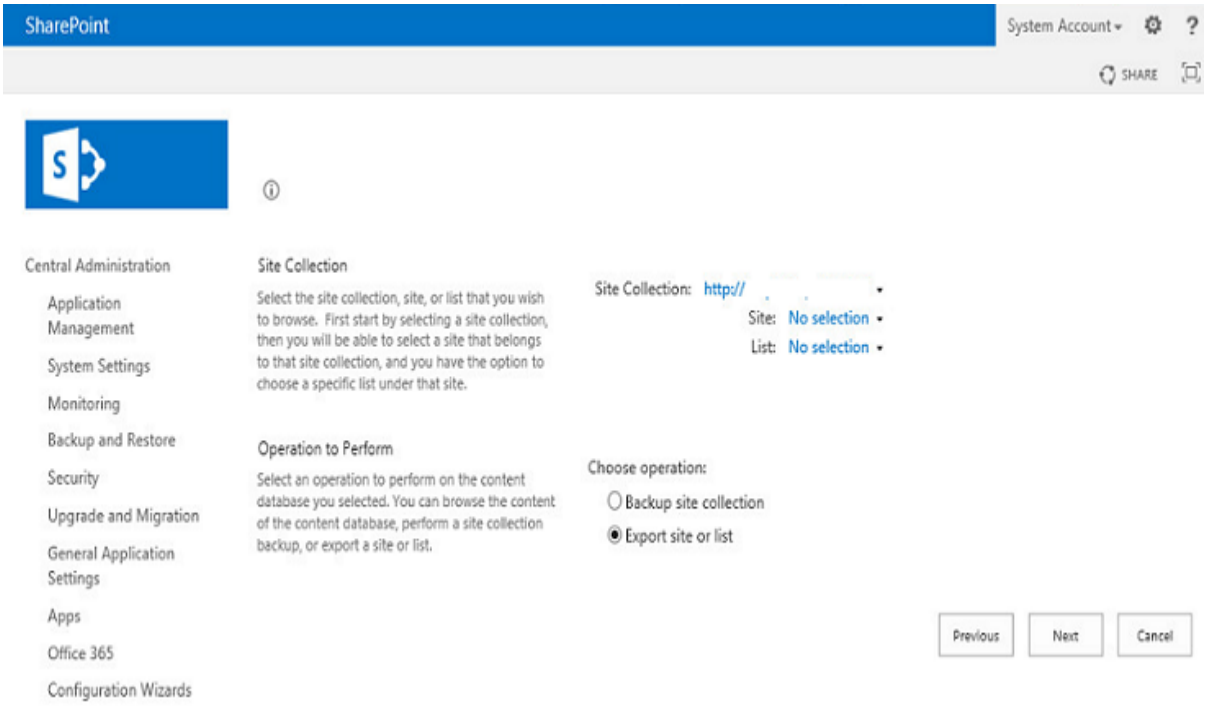
Restore a Site

Follow these steps:

1. Recover site content from an unattached content database.

   Use Central Administration

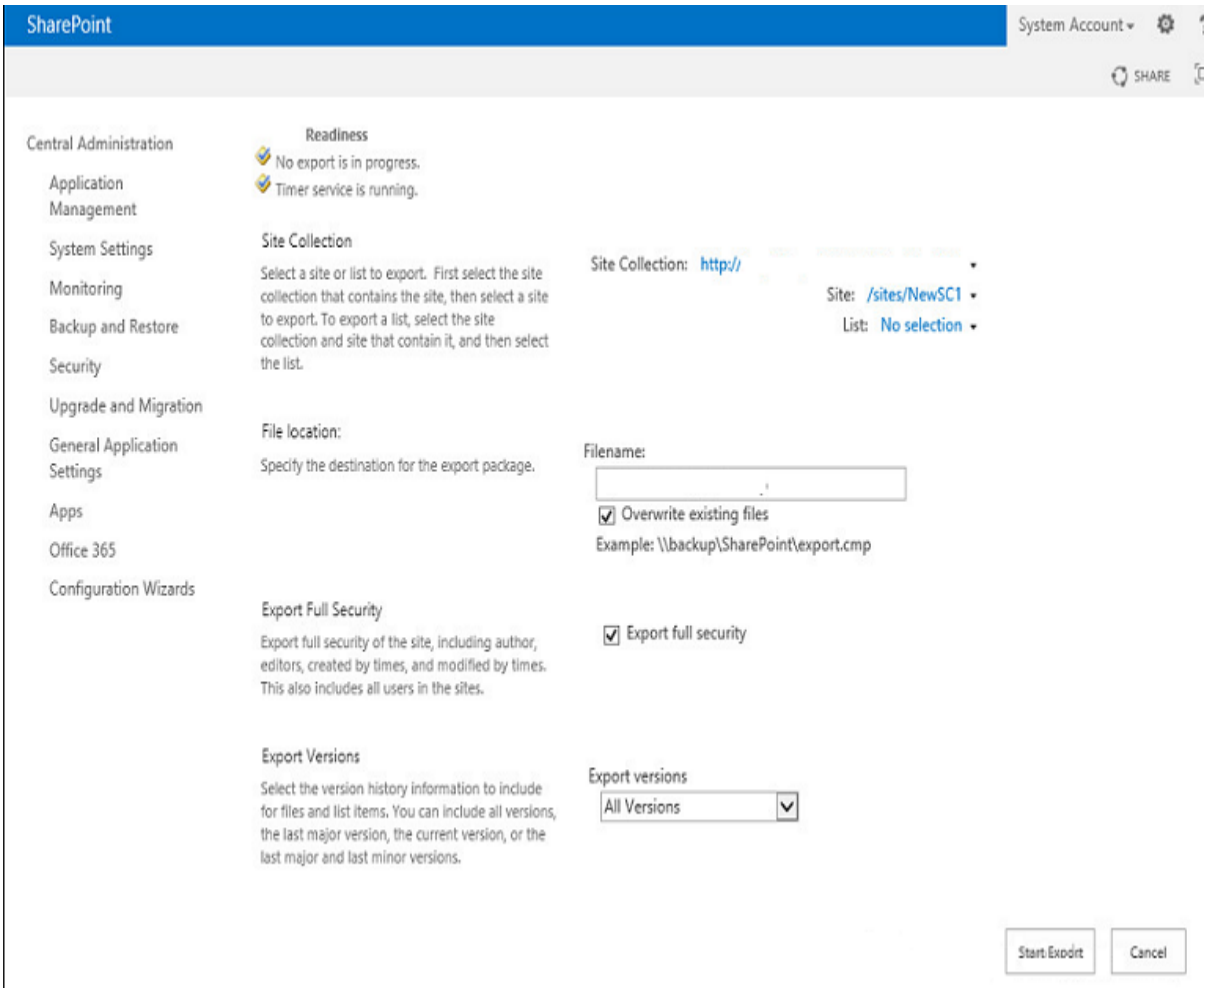
   a. Select the export site or the list option, and click the Next button.



   b. Select the site and provide the file location for the export package.

      Example: The name of site to be restored is TestSite1 and the URL is /TestSite1/.

   c. Select the options for security and versions (by default All Versions).



d. Click the Start Export button to begin the export, then the site is exported to a file.

Use PowerShell commands

a. $database = Get-SPContentDatabase -ConnectAsUnattachedDatabase -DatabaseName xxxx -DatabaseServer xxxx

ConnectAsUnattachedDatabase: Specifies that only unattached databases in the farm are returned.

DatabaseName: Specifies the name of the content database.

DatabaseServer: Specifies the name of the host server for the content database specified in the DatabaseName parameter.

For more details, see the article from Microsoft.

b. Setting object to export

$ExportObject = New-Object Microsoft.SharePoint.Deployment.SPExportObject

$ExportObject.Type = [Microsoft.SharePoint.Deployment.SPDeploymentObjectType]::Web

$ExportObject.Url = $SiteUrl

$SiteUrl: Specifies the URL location to which the site will be backed up.

c. Configuring Export Settings

$ExportSettings = New-Object Microsoft.SharePoint.Deployment.SPExportSettings

$ExportSettings.UnattachedContentDatabase = $database

$ExportSettings.SiteUrl = $CAUrl

$CAUrl: Specifies Central Administration Site Url.

$ExportSettings.FileLocation = $ExportPath

$ExportSettings.LogFilePath = $ExportPath

$ExportPath: Specifies the path to save the backup file (for example, C:\backup).

$ExportSettings.BaseFileName = $ExportFile

$ExportFile: Specifies the filename of the backup file (for example, site.cmp).

$ExportSettings.IncludeVersions = [Microsoft.SharePoint.Deployment.SPIncludeVersions]::All

$ExportSettings.ExportMethod = [Microsoft.SharePoint.Deployment.SPExportMethodType]::ExportAll

$ExportSettings.IncludeVersions = [Microsoft.SharePoint.Deployment.SPIncludeVersions]::All

$ExportSettings.ExportObjects.Add($ExportObject)

$ExportSettings.Validate()

$ExportJob = New-Object Microsoft.SharePoint.Deployment.SPExport($ExportSettings)

Back up the site to a file.

$ExportJob.Run()

For more details, see the article from Microsoft.

2. Click SharePoint Management Shell to launch the console.
3. Use PowerShell commands to restore the site to the origin location or the new location.

```
Import-SPWeb -Identity xxxx -Path xxxx -IncludeUser-
Security:$true -UpdateVersions:xxxx
```

Identity: Specifies the URL or GUID of the Web to import into. for example, http://www.contoso.com.

Path: Specifies the name of the import file. for example, C:\backup\site.cmp'

IncludeUserSecurity: Preserves the user security settings except for SPLists that have broken inheritance and item level permissions set.

UpdateVersions: Indicates how to resolve situations where a file version to be imported to a site already exists in that site. You can select one of the following options:

Add: Adds the file as a new version.

Overwrite: Overwrites the current file and all of its versions (delete then insert).

Ignore: Ignores the file if it exists on the destination. The new file is not added.

The default value is Add.

For more details, see the article from Microsoft.

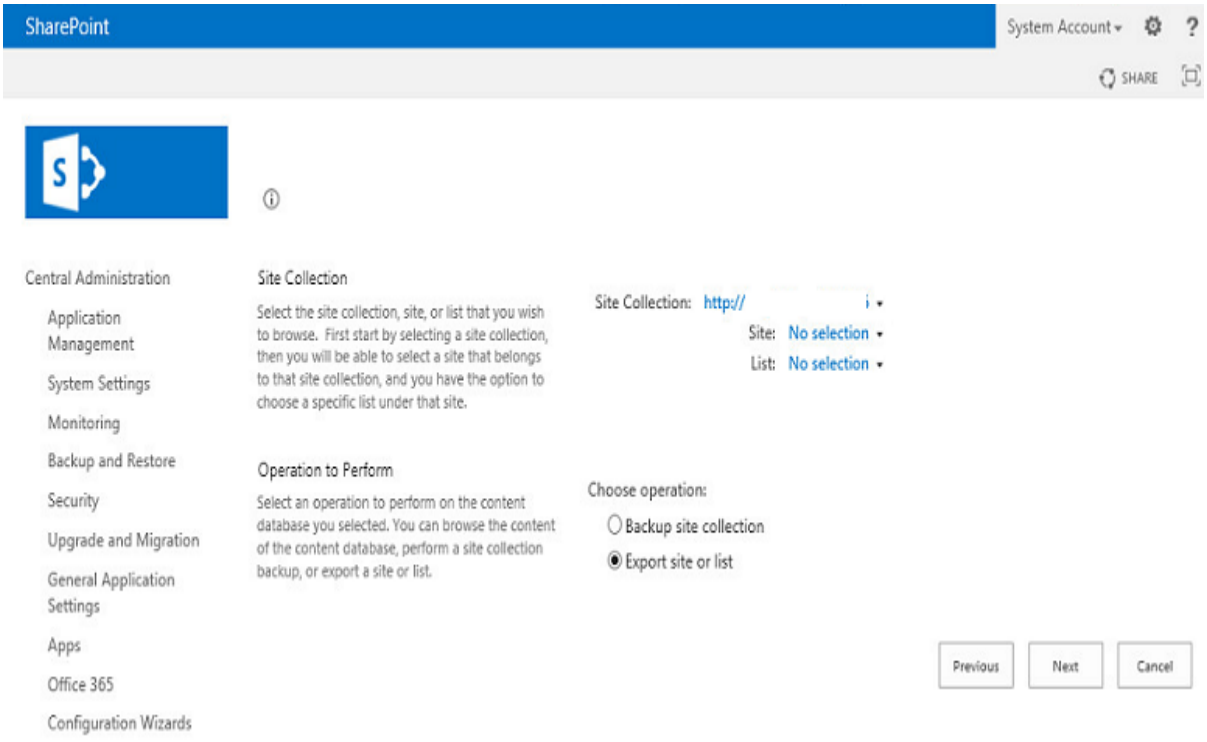
Restore a List or a Library

Follow these steps:

1. Recover the list or library content from an unattached content database.

   Use Central Administration

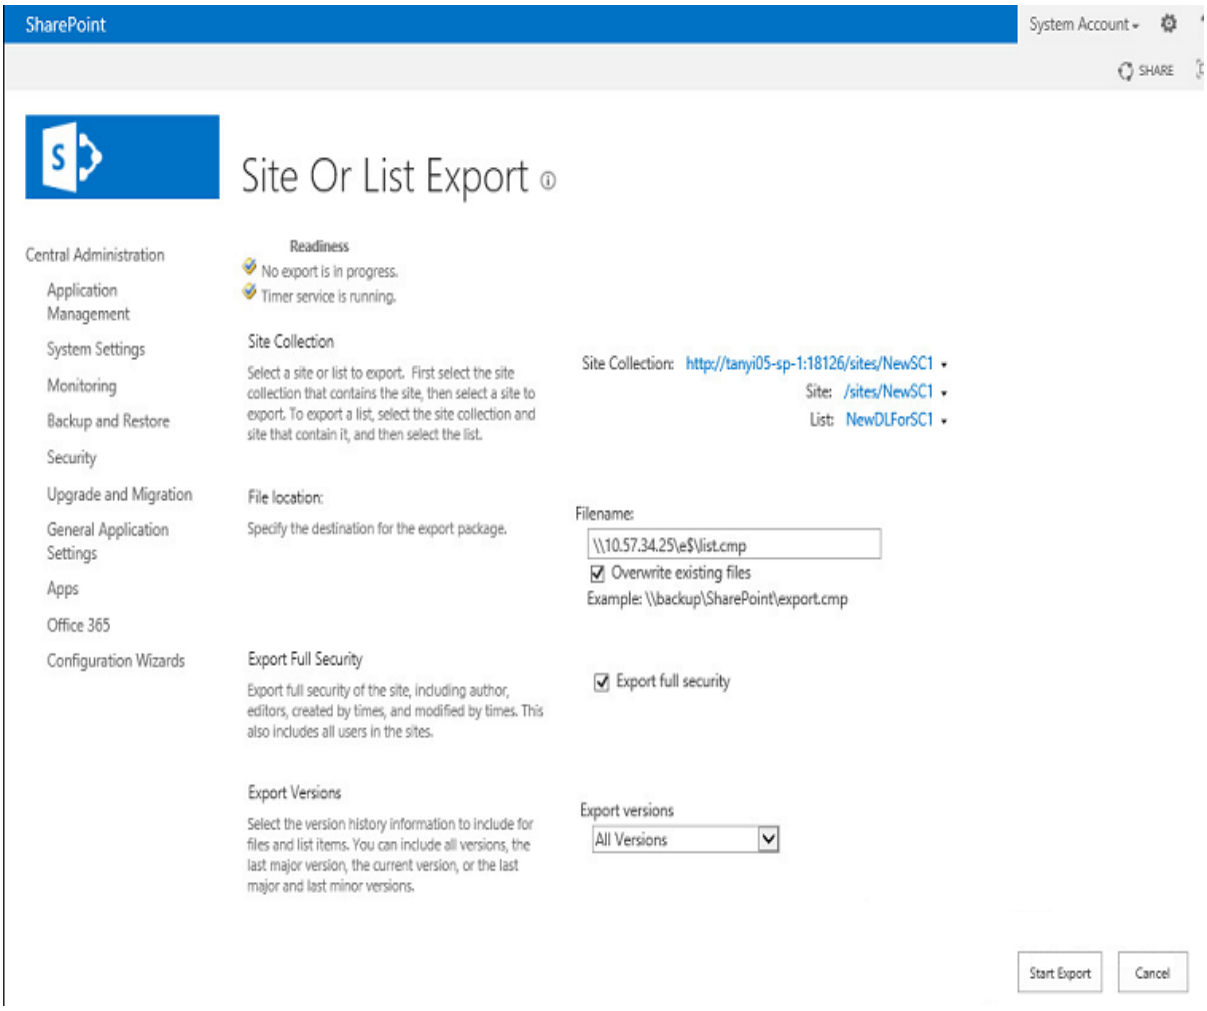
   a. Select the export site or list option, and click **Next**.



   b. Select the site and the list and provide the file location for the export package.

      Example: The name of list/library to be restored is NewList1 and the URL is /TestSite1/NewList1.

   c. Select options for security and versions (by default All Versions).



d. Click **Start Export** to begin the export.

   The list or library is exported to a file.

Use PowerShell commands

a. $database = Get-SPContentDatabase -ConnectAsUnattachedDatabase -DatabaseName xxxx -DatabaseServer xxxx

   ConnectAsUnattachedDatabase: Specifies that only unattached databases in the farm are returned.

   DatabaseName: Specifies the name of the content database.

   DatabaseServer: Specifies the name of the host server for the content database specified in the DatabaseName parameter.

   For more details, see the article from Microsoft.

b. Setting object to export.

   $ExportObject = New-Object Microsoft.SharePoint.Deployment.SPExportObject

$ExportObject.Type = [Microsoft.SharePoint.Deployment.SPDeploymentObjectType]::List

$ExportObject.Url = $ListUrl

$ListUrl: Specifies the URL location to which the list or library is backed up. If it is a list, use the parameter "/Lists/{ListName}". If it is a library, use the parameter "/{LibraryName}"

c. Configuring Export Settings

$ExportSettings = New-Object Microsoft.SharePoint.Deployment.SPExportSettings

$ExportSettings.UnattachedContentDatabase = $database

$ExportSettings.SiteUrl = $CAUrl

$CAUrl: Specifies Central Administration Site URL.

$ExportSettings.FileLocation = $ExportPath

$ExportSettings.LogFilePath = $ExportPath

$ExportPath: Specifies the path to save the backup file (for example, C:\backup).

$ExportSettings.BaseFileName = $ExportFile

$ExportFile: Specifies the filename of the backup file (for example, site.cmp).

$ExportSettings.IncludeVersions = [Microsoft.SharePoint.Deployment.SPIncludeVersions]::All

$ExportSettings.ExportMethod = [Microsoft.SharePoint.Deployment.SPExportMethodType]::ExportAll

$ExportSettings.IncludeVersions = [Microsoft.SharePoint.Deployment.SPIncludeVersions]::All

$ExportSettings.ExportObjects.Add($ExportObject)

$ExportSettings.Validate()

$ExportJob = New-Object Microsoft.SharePoint.Deployment.SPExport($ExportSettings)

d. Back up the list or library to a file.

$ExportJob.Run()

For more details, see the article from Microsoft.

2. Click SharePoint Management Shell to launch the console.

3. Use PowerShell commands to restore the list or library to the origin location or the new location.

```
Import-SPWeb -Identity xxxx -Path xxxx -IncludeUser-
Security:$true -UpdateVersions:xxxx
```

Identity: Specifies the URL or GUID of the Web to import into. For example, http://www.contoso.com

Path: Specifies the name of the import file. For example, C:\backup\list.cmp'

IncludeUserSecurity: Preserves the user security settings except for SPLists that have broken inheritance and item level permissions set.

UpdateVersions: Indicates how to resolve situations where a file version to be imported to a site already exists in that site. You can select one of the following options:

Add: Adds the file as a new version.

Overwrite: Overwrites the current file and all of its versions (delete then insert)

Ignore: Ignores the file if it exists on the destination. The new file will not be added.

The default value is Add.

For more details, see the article from Microsoft.

Restore a File

Follow these steps:

1. Restore the list or the library to a new location. For more information, see Restore a list or a library.

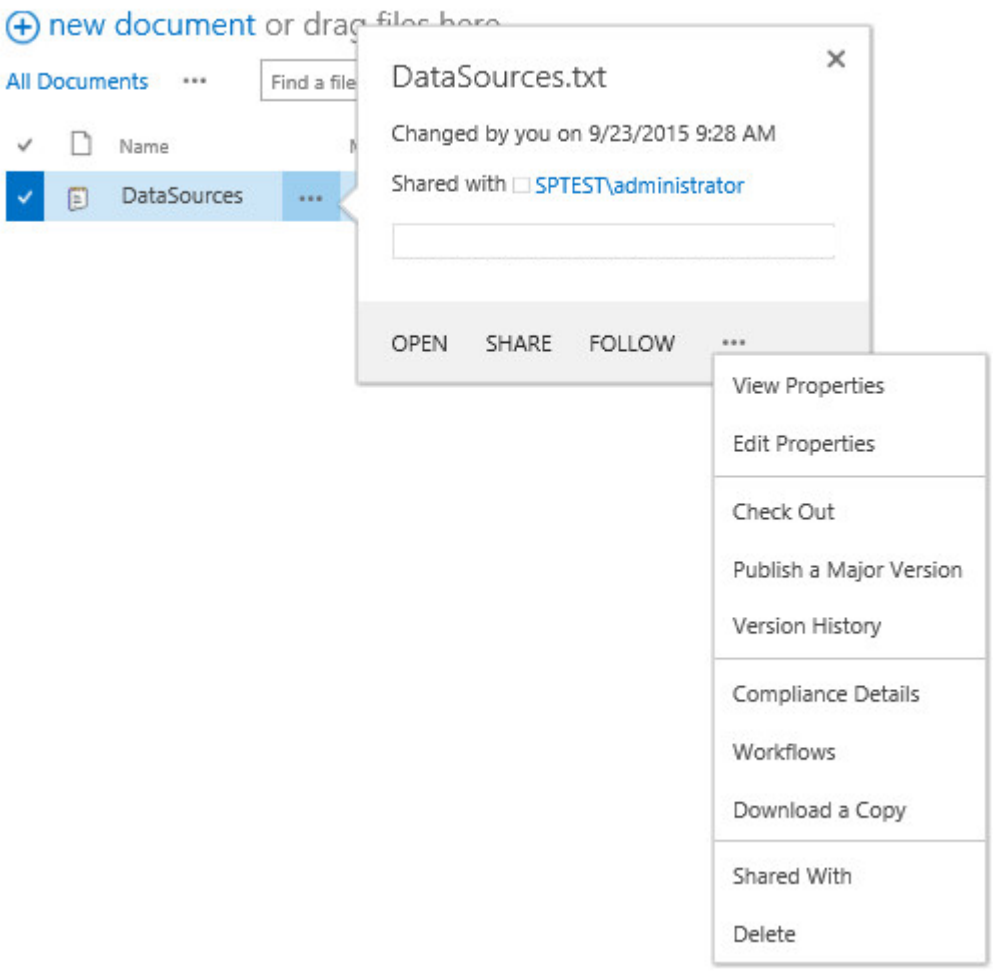
   For example: The origin list or library is named NewList1 and the URL is http://-contoso.com /TestSite1/NewList1

   - Using PowerShell commands to restore the list or library to new location. Example, http://contoso.com/TestSite2

     Import-SPWeb -Identity http://contoso.com/TestSite2 -Path C:\-backup\list.cmp -IncludeUserSecurity:$true -UpdateVersions:Overwrite

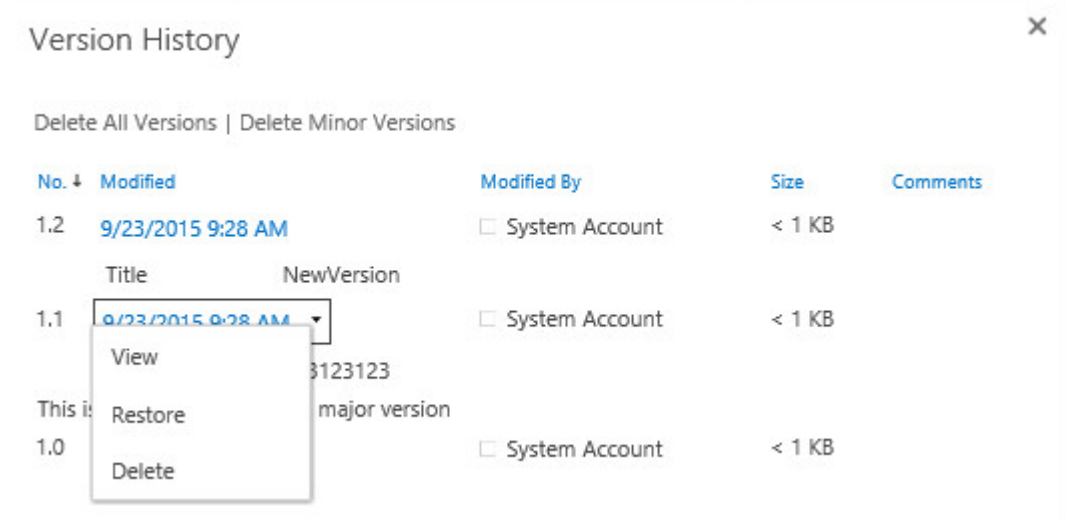
   - Navigate to the new URL of the list or library. All items are restored into http://contoso.com/TestSite2/NewList1.

2. Navigate to the new location URL of the list or library.
3. Check the file version history in the list or library.



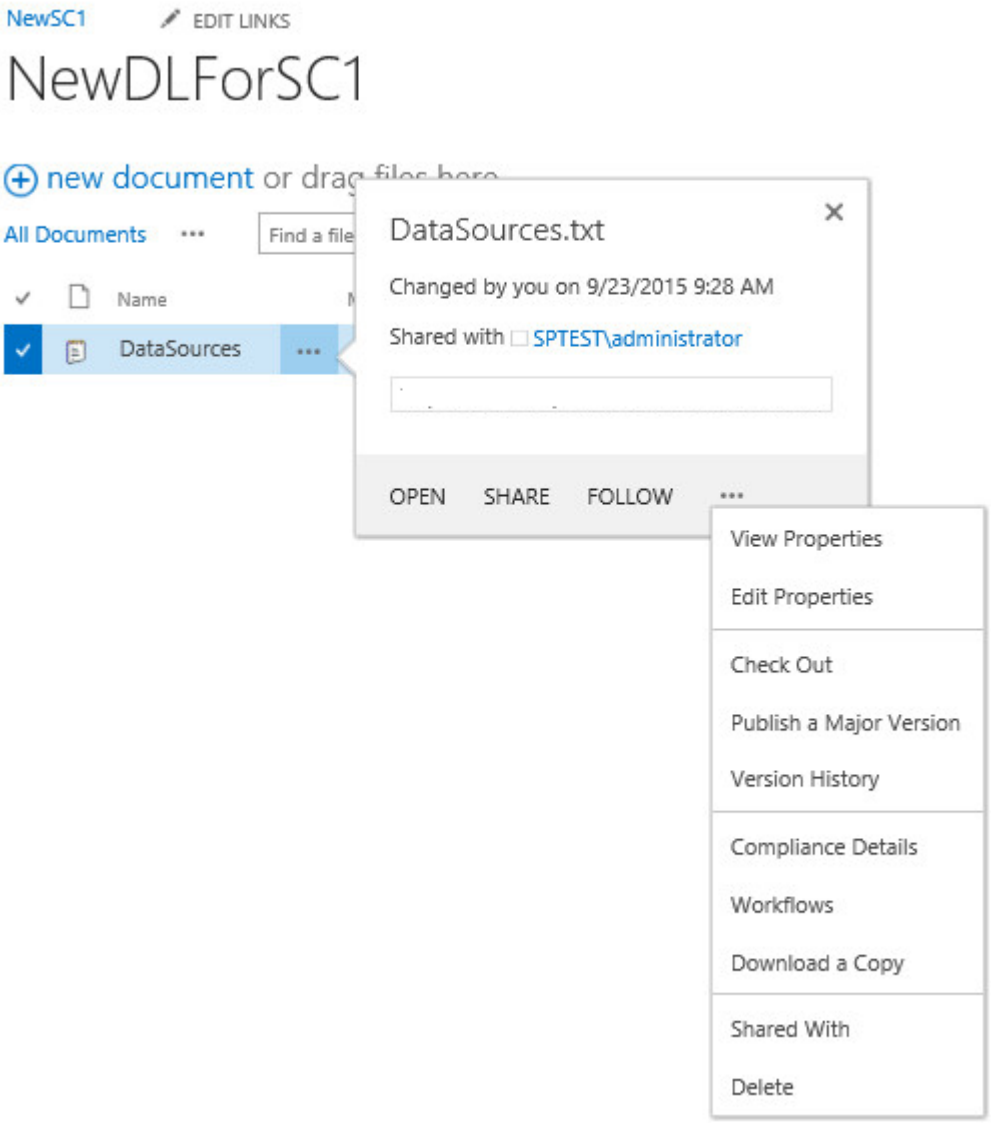
4. Select the specific version for the file and click Restore.

   For example, restore the specific version 1.1 for the file.



   The 1.1 version fie is restored.

5. Click "Download a Copy" to save the file with specific version 1.1 to a location.



The saved file is restore into origin list or library.

Create an Isolated Network for SharePoint Recovery

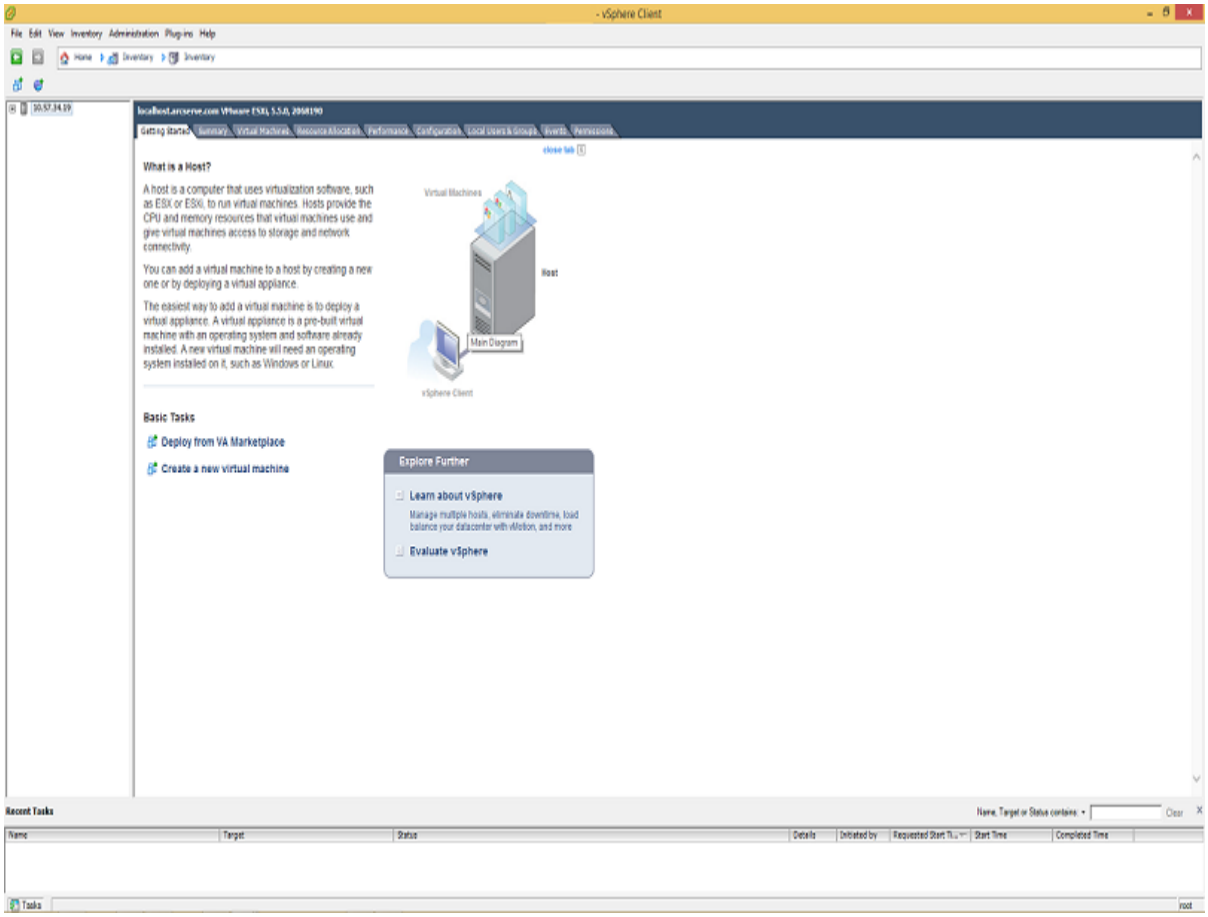
You can create an isolated network to recover SharePoint in Hyper-V and VMware machines.

- How to Create an Isolated Network for SharePoint Recovery for VMware VM
- How to Create an Isolated Network for SharePoint Recovery for Hyper-V VM

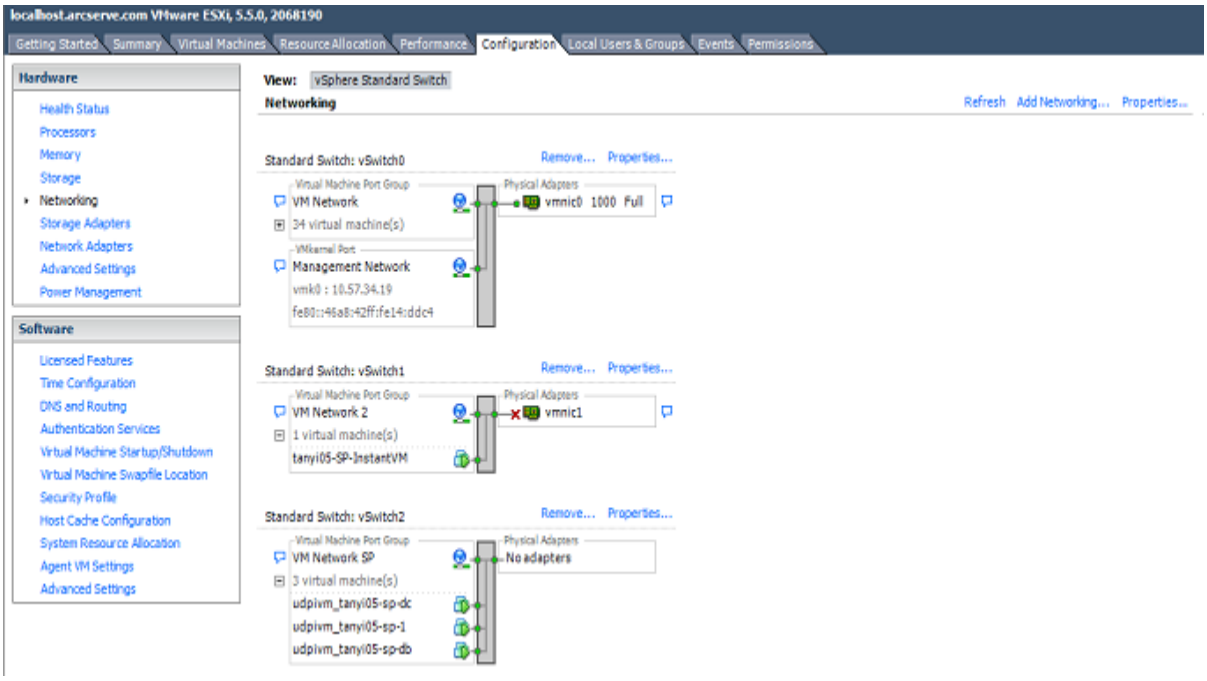
How to Create an Isolated Network for SharePoint Recovery for VMware VM

Follow these steps:

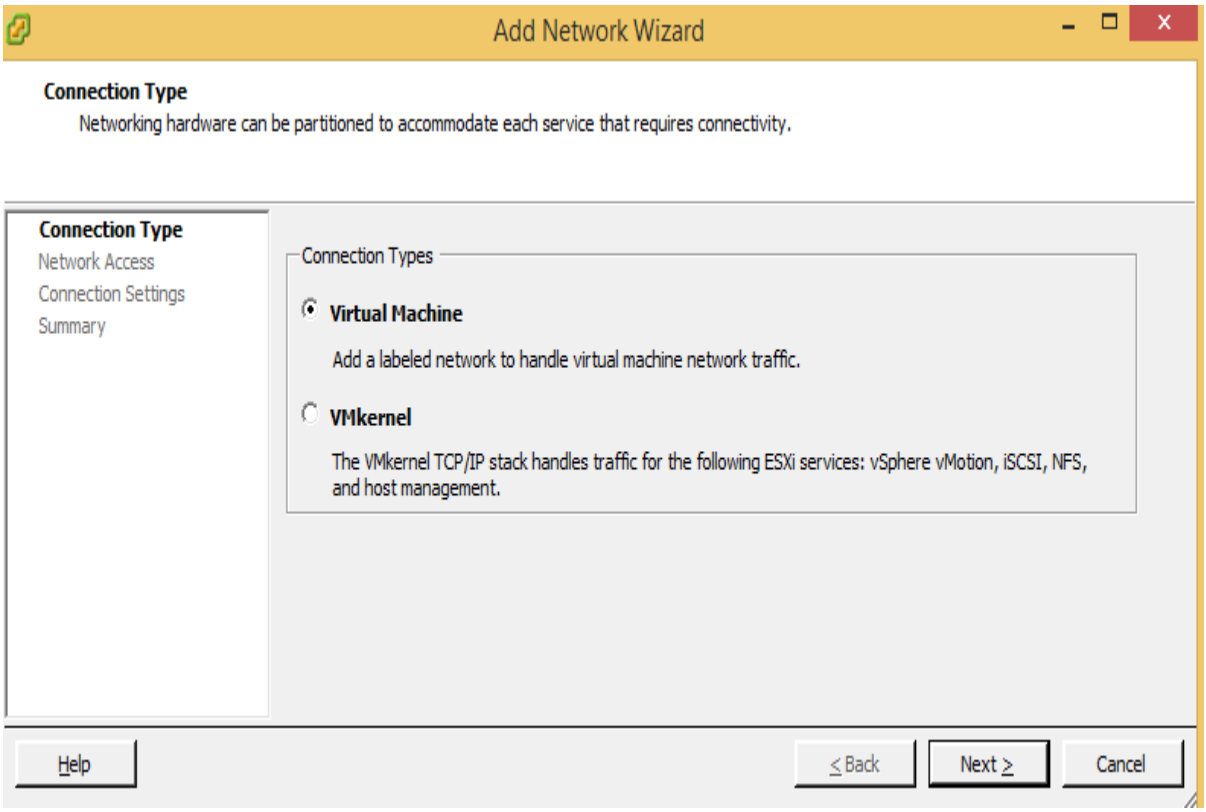
1. Log into the VMware ESXi server using the vSphere client.



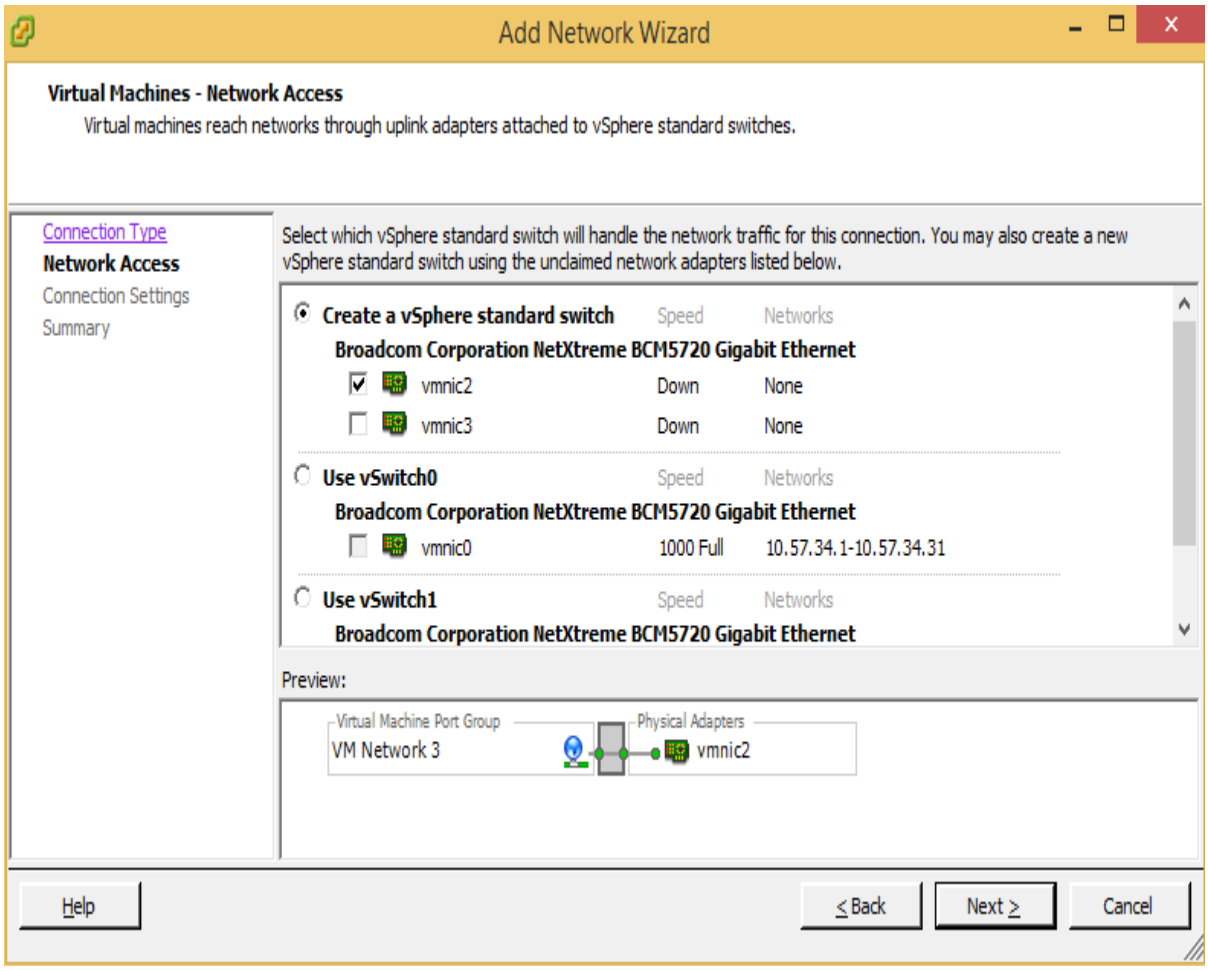
2. Click the Configuration tab.
3. Select Networking in the Hardware pane and click Add Networking.



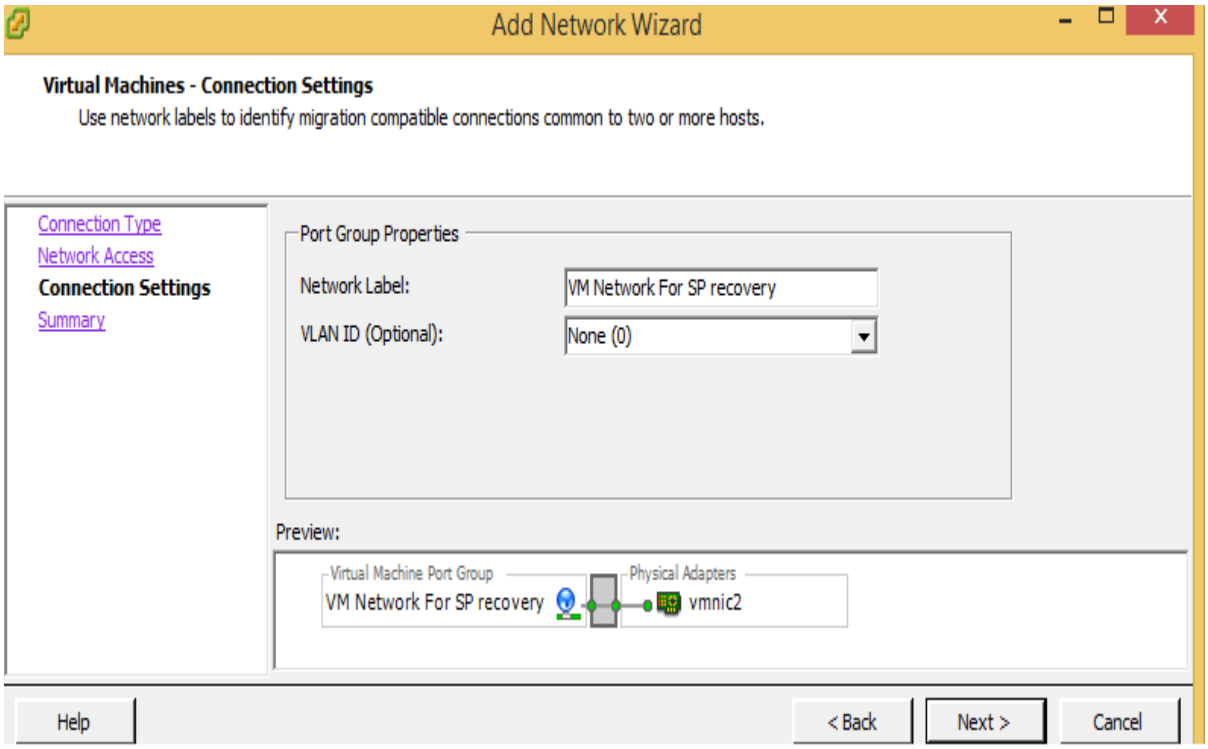
4. Ensure the Virtual Machine radio button is selected and click Next.



5. Select the physical NIC that you need to use to connect the virtual switch to other physical resources on the network and click Next.

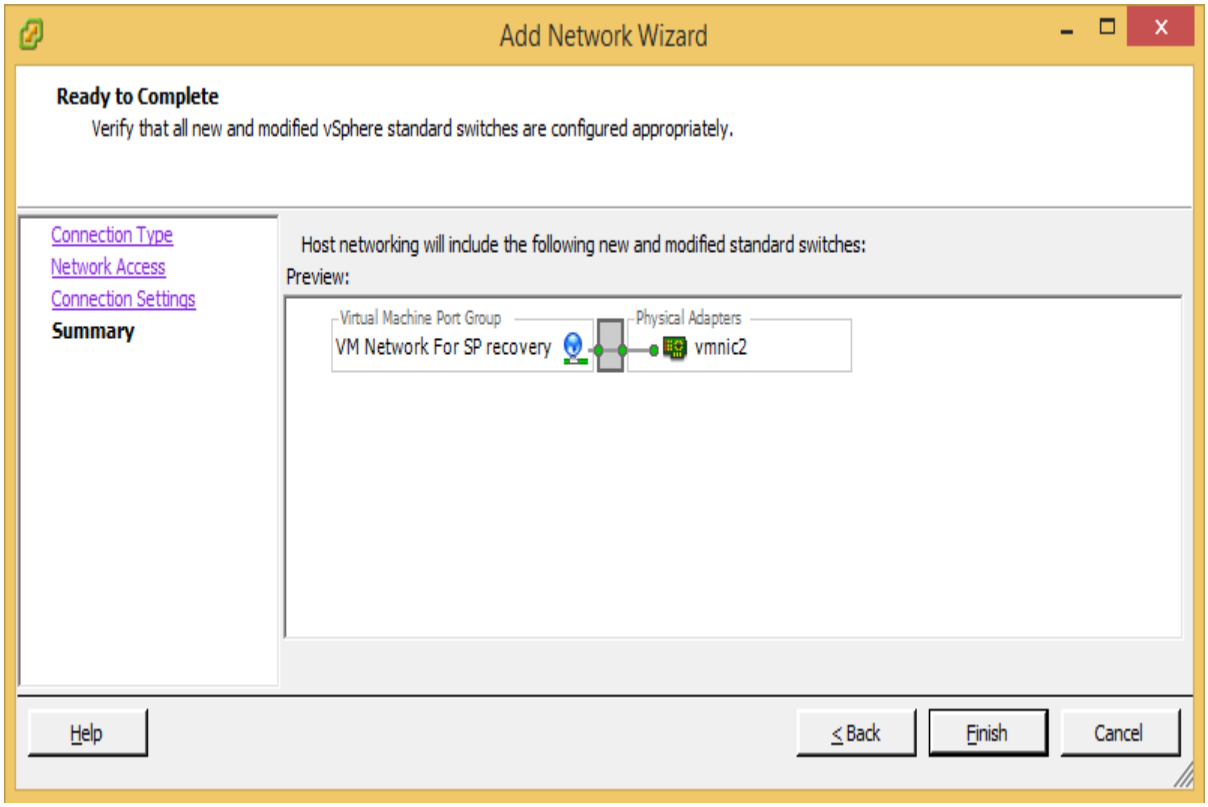


6. Assign a Network Label to the virtual switch, and a VLAN ID if necessary, and click **Next**.



7. Verify that your virtual switch settings are correct, and click **Finish**.

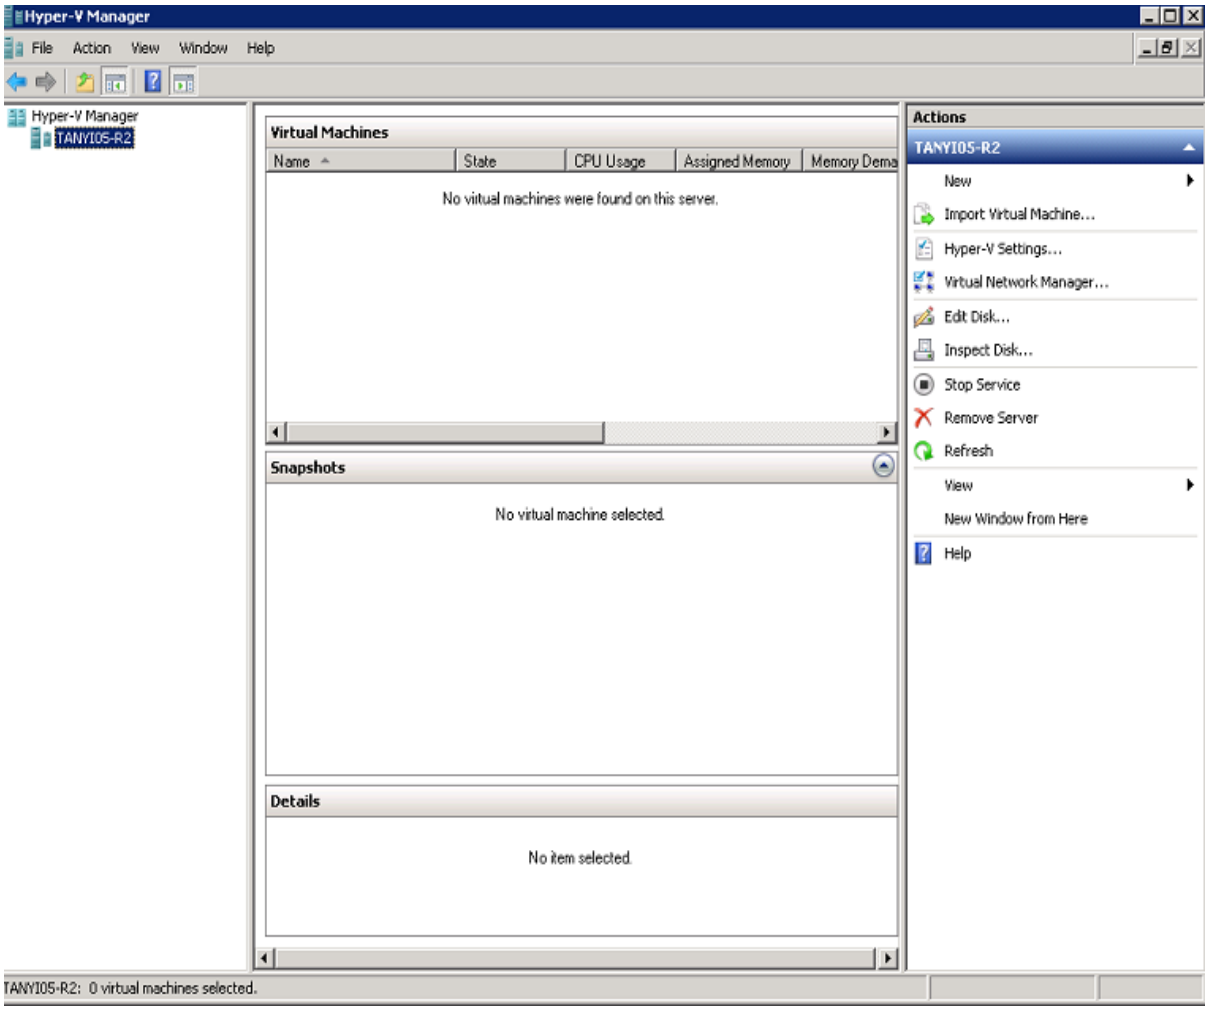
   When you return to the Networking Configuration tab, you can see that the new Virtual Switch is added.



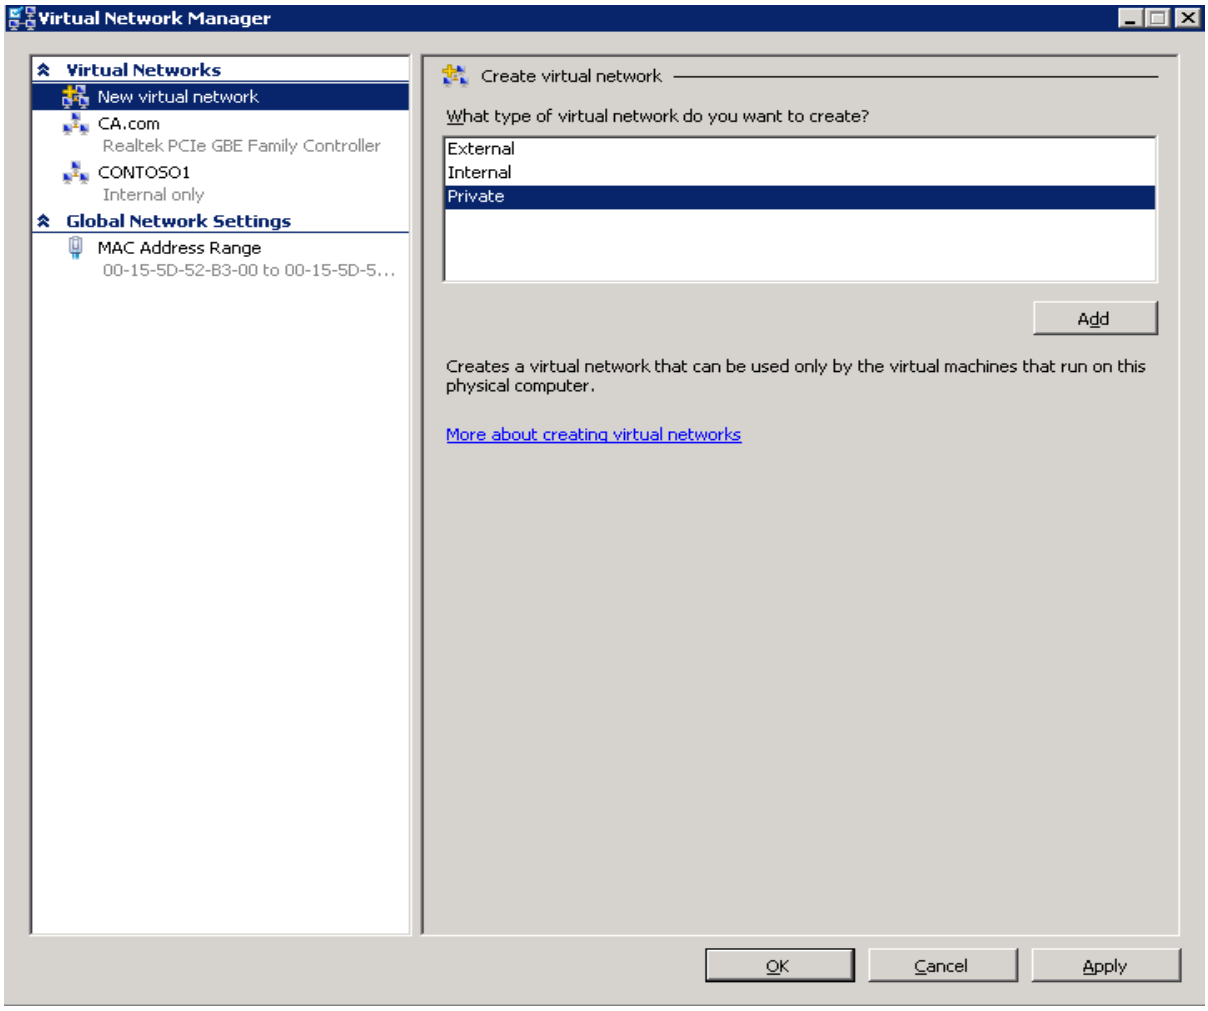
How to Create an Isolated Network for SharePoint Recovery for Hyper-V VM

Follow these steps:

1. Log into the Hyper-V Manager.
2. Click Virtual Network Manager.

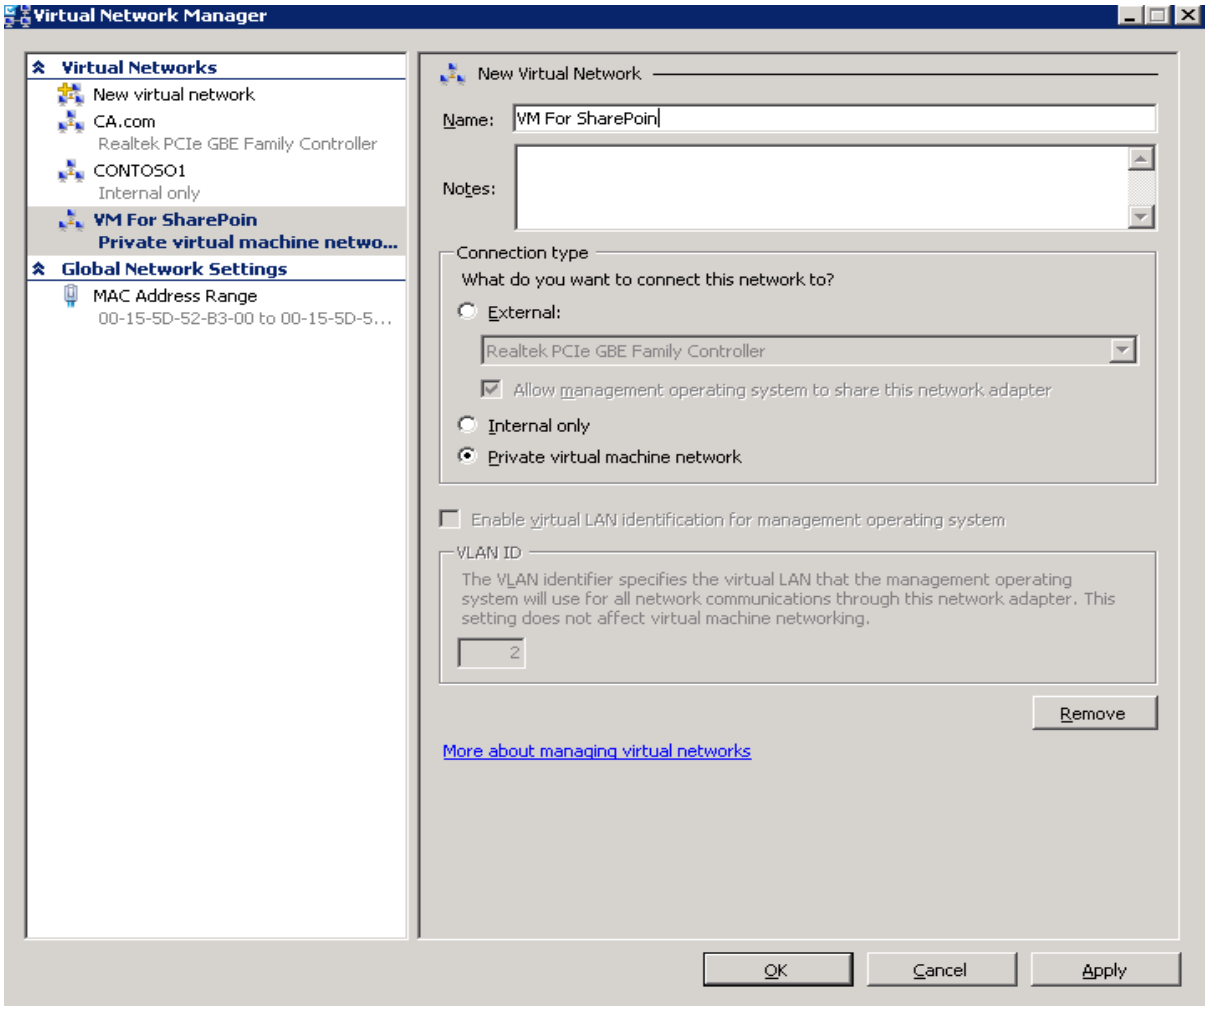


3. Click New virtual network, select the type "Private", and then click **Add**.



4. Type a name for the private virtual network, and click OK.

   The new virtual network is added.


Virtual Network Manager
Virtual Networks
New virtual network
CA.com
Realtek PCIe GBE Family Controller
CONTOSO1
Internal only
VM For SharePoin
Private virtual machine netwo...
Global Network Settings
MAC Address Range
00-15-5D-52-B3-00 to 00-15-5D-5...
New Virtual Network
Name:
VM For SharePoin
Notes:
Connection type
What do you want to connect this network to?
External:
Realtek PCIe GBE Family Controller
Allow management operating system to share this network adapter
Internal only
Private virtual machine network
Enable virtual LAN identification for management operating system
VLAN ID
The VLAN identifier specifies the virtual LAN that the management operating system will use for all network communications through this network adapter. This setting does not affect virtual machine networking.
2
Remove
More about managing virtual networks
OK
Cancel
Apply

Chapter 18: Generating Arcserve UDP Reports

This section contains the following topics:

Understanding Arcserve UDP Reports

The **reports** tab provides access to various types of reports such as Alerts, Data Trend, Backup Status, Data Distribution, and SLA. The left pane includes list of reports that you can generate. The center pane displays the details of the selected report and lets you configure various report settings. The report is generated for a group of nodes or servers. You can also filter the report to display detailed information for an individual node.

For more information about reports, see Arcserve UDP Reports.

The drill-down report includes the following items:

Job Nodes

Displays the node name where jobs of Arcserve UDP agents, Host-Based VM or Virtual Standby run.

Protected Nodes

Displays the name of agent node and node protected by the Arcserve Unified Data Protection agent or Host-based VM Backup or Virtual Standby or Arcserve Backup.

Product

Displays the product that is installed on the node. The product name could be Arcserve Unified Data Protection Agent, Arcserve Unified Data Protection Recovery Point Server, Host-based VM, or Arcserve Backup.

Filters/Actions

Displays the global and local options of filters and actions related to the reports. For more information, see Using Filters and Actions.

Arcserve UDP Reports

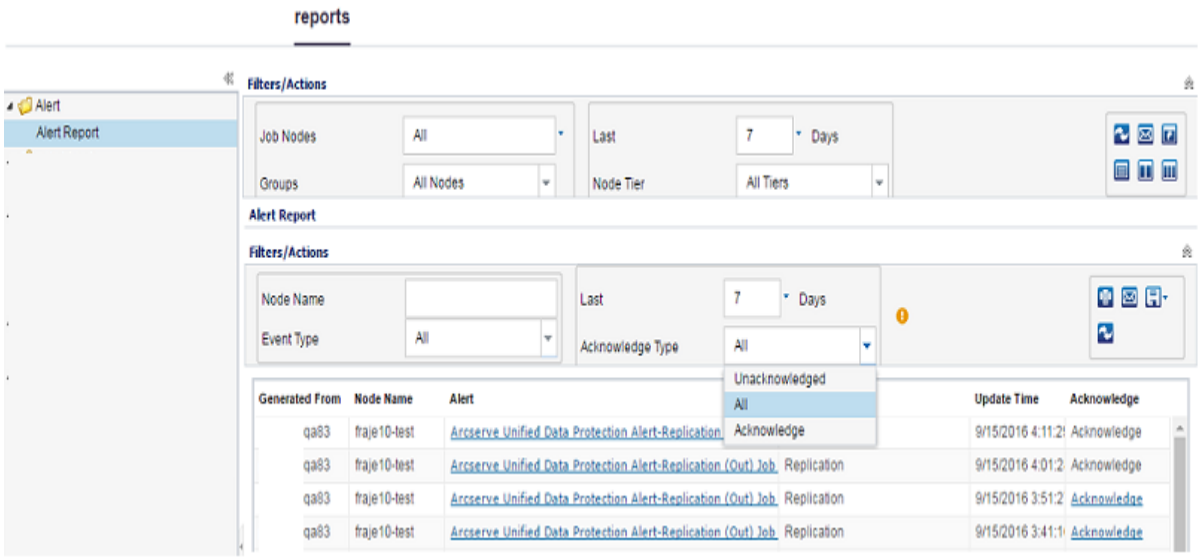
Arcserve UDP provides the following type of reports:

- Alert Report
- Backup Size Trend Report
- **Backup Related Report:** Arcserve UDP provides three types of Backup-related reports:
  - Node Backup Status Report
  - Virtualization Protection Status Report
  - Managed Capacity Report
- Data Distribution on Media Report
- **SLA Report:** Service Level Agreement (SLA) report displays compliance reports related to Recovery Point Objective (RPO) and Recovery Time objective (RTO)
  - RPO Report
  - RTO Report
- **Job Status report:** Arcserve UDP helps you generate job status report to retrieve details about all jobs performed in a defined period.

Alert Reports

Arcserve UDP displays the alert information for nodes. Besides the common filters and actions, alert report provides the capability to sort out your alert dashboard by acknowledging. By default, your Alert Report dashboard shows only unacknowledged alerts. You can click **Acknowledge** link placed with every alert to remove it from the dashboard view. If you want to view the acknowledged report, select **Acknowledge** in the **Acknowledge Type** option of local filter.

Alert Report

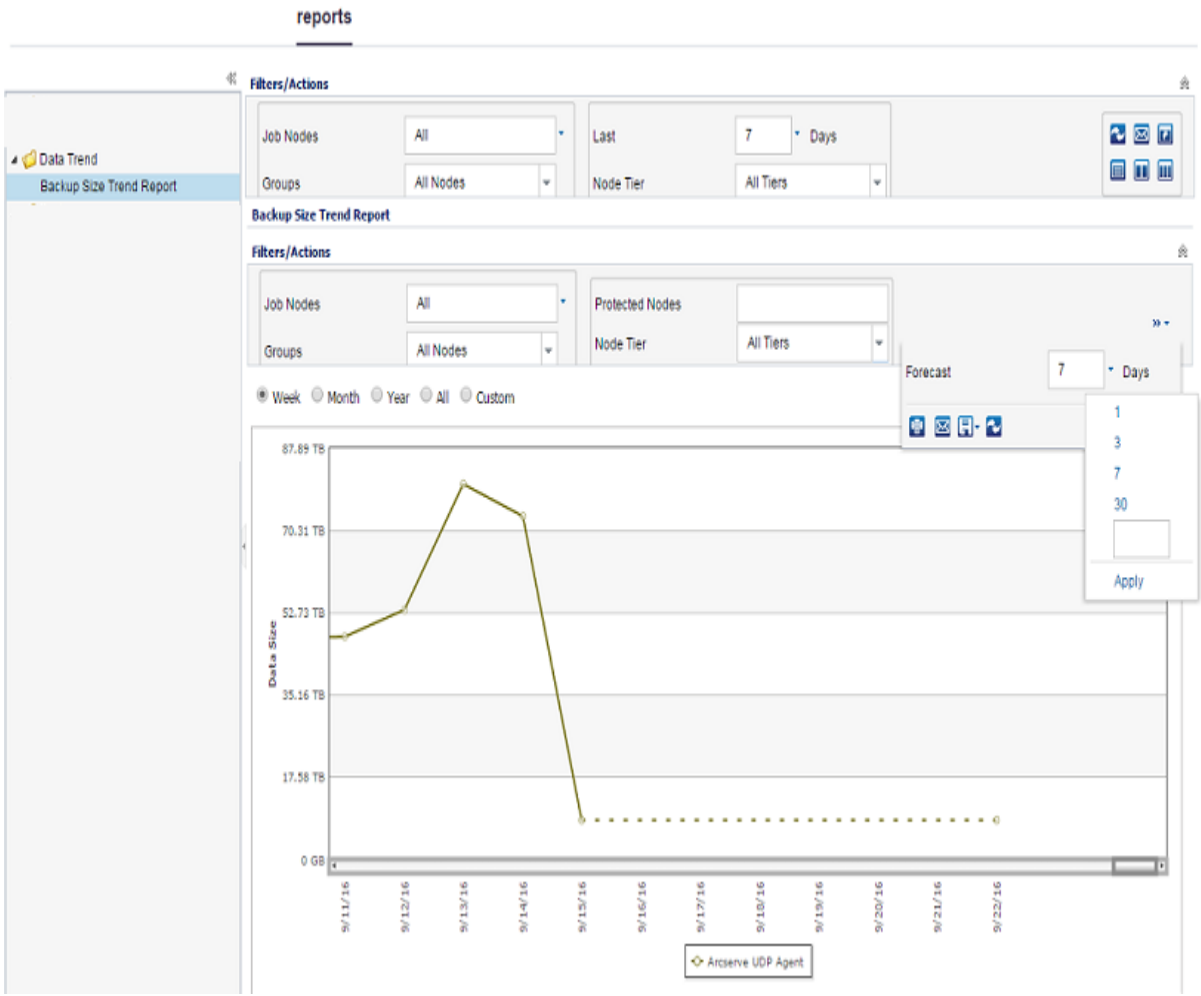


Backup Size Trend Reports

Arcserve UDP Backup Size Trend report displays the backup data size of Arcserve Backup and Arcserve Unified Data Protection agent in a historical view and then projects the growth trend that you can prepare for future storage space requirements. This report includes information for nodes which run on supported Windows and Linux operating systems and allows you to drill down to display more detailed information for an individual node.

Besides the common filters and actions, backup size trend report provides the capability to view results based on the number of **Days**. By Default, the filter for **Last** days is not applicable. Instead, you can use **View** mode filter(week/-month/year/custom).

Backup Size Trend Report



Node Backup Status Reports

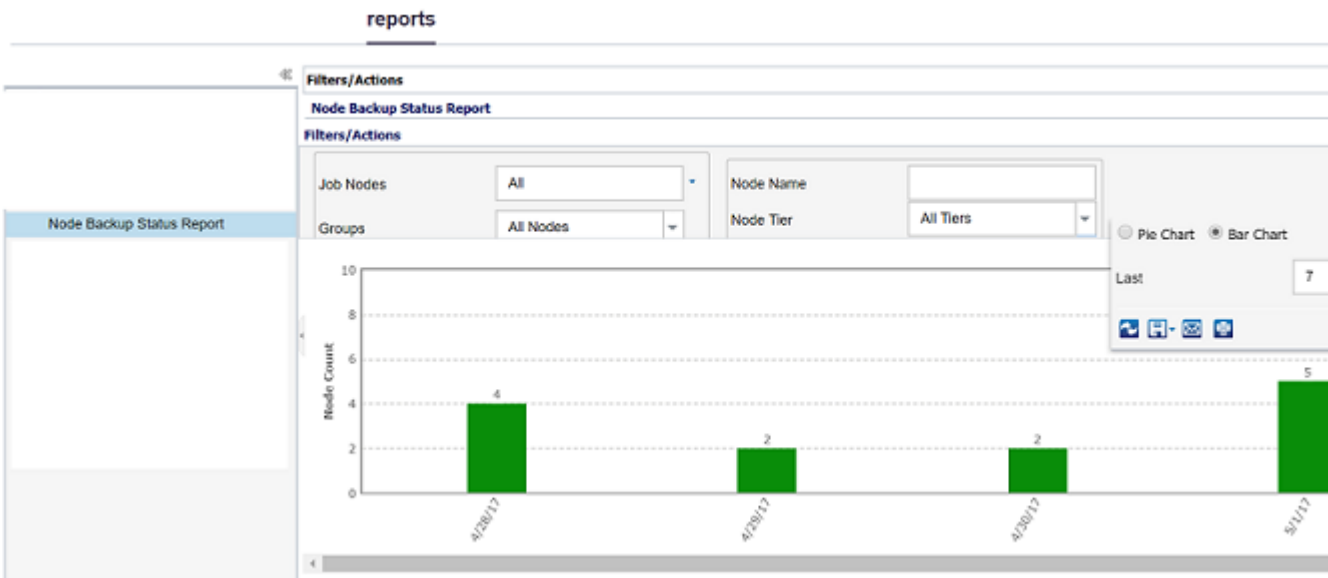
Arcserve UDP displays the latest backup status of all nodes during the specific time period. This report allows you to view detailed information about nodes based on categories such as selected type of Groups and Node Tier.

Besides the common filters and actions, node backup status report provides the capability to view results based on the number of **Days**. In the bar chart, number of node is visible on top of each bar.

The report displays the following job status:

- **Successful:** Provides a list of jobs that are completed successfully.
- **Failed:** Provides a list of jobs that failed.
- **Incomplete:** Provides a list of jobs that finished with incomplete status.
- **Canceled:** Provides a list of jobs that are canceled.
- **Missed:** Provides a list of jobs that are not attempted.
- **No Backups:** Provides a list of nodes without any assigned plan or nodes with plan assigned but still waiting for backups to run.

Node Backup Status Report:

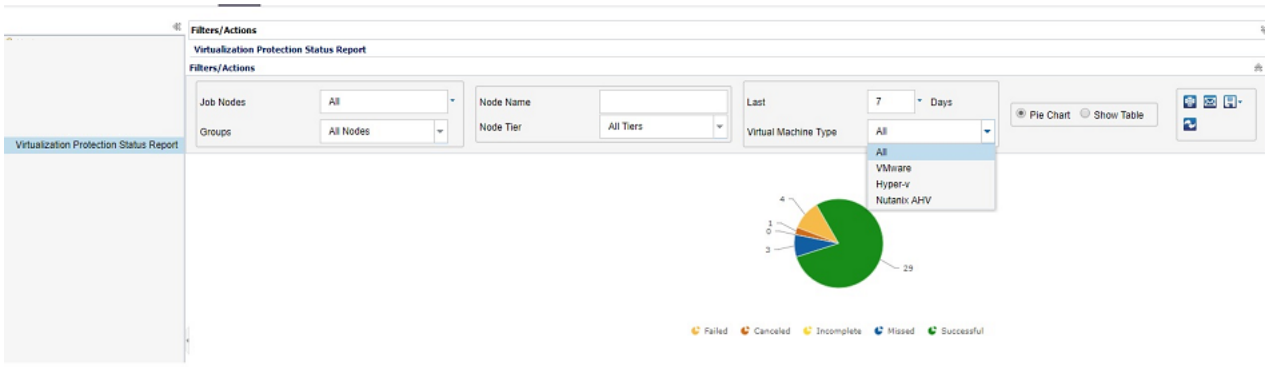


Virtualization Protection Status Reports

Displays the latest backup status of virtual machines that Host-Based VM Backup or Virtual Standby or Arcserve Backup protects. This report lets you view information for a specified time period and drill down to display more detailed information about each selected category.

Besides the common filters and actions, Virtualization Protection Status report provides the capability to view results based on the number of **Days** and view reports as pie chart or table.

Virtualization Protection Status Report



Managed Capacity Reports

Displays the raw data size of the last successful full backup for each node that Arcserve Backup, Arcserve Unified Data Protection Agent, and Host-based VM Backup protect.

For details about filters and actions, see the common filters and actions.

Notes:

- Entering a value in the **Last Days** filter does not change the result for global filter. Irrespective of the number of days you provide, applying filter displays the latest values.
- The Node backup related data in deleted data store may appear as Managed Capacity report is not updated unless the backup runs on the latest backup destination. For example, if you modified plan to run backups to data store DS2, and deleted data store DS1, then unless you run backup on DS2, the deleted data from DS1 may still appear in the report.
- In the grid area of the report, the following three columns are applicable only for VM nodes that are protected by host-based agentless backup:

  Used Space of volumes (VM)

  Indicates the sum of the used space of volumes inside the guest OS of the VM.

  For Windows: Refers to the sum of the used space of all NTFS volumes

  For Linux VM: Refers to the sum of the used space of all volumes.

  For Linux VM, only VMware VM is supported for this column. In this case, you need to update the VM node with root credentials. Whether it is a VMware Linux VM without credentials, or a Hyper-V Linux VM, this column is empty.

  For Linux VM support, VMware Red Hat or CentOS machine and Hyper-V Linux VMs face limitations. For more details, refer to Known Issue in Release notes of Arcserve UDP v6 Update 1.

  Synthetic Read Size

  Indicates the total size read during the backup.

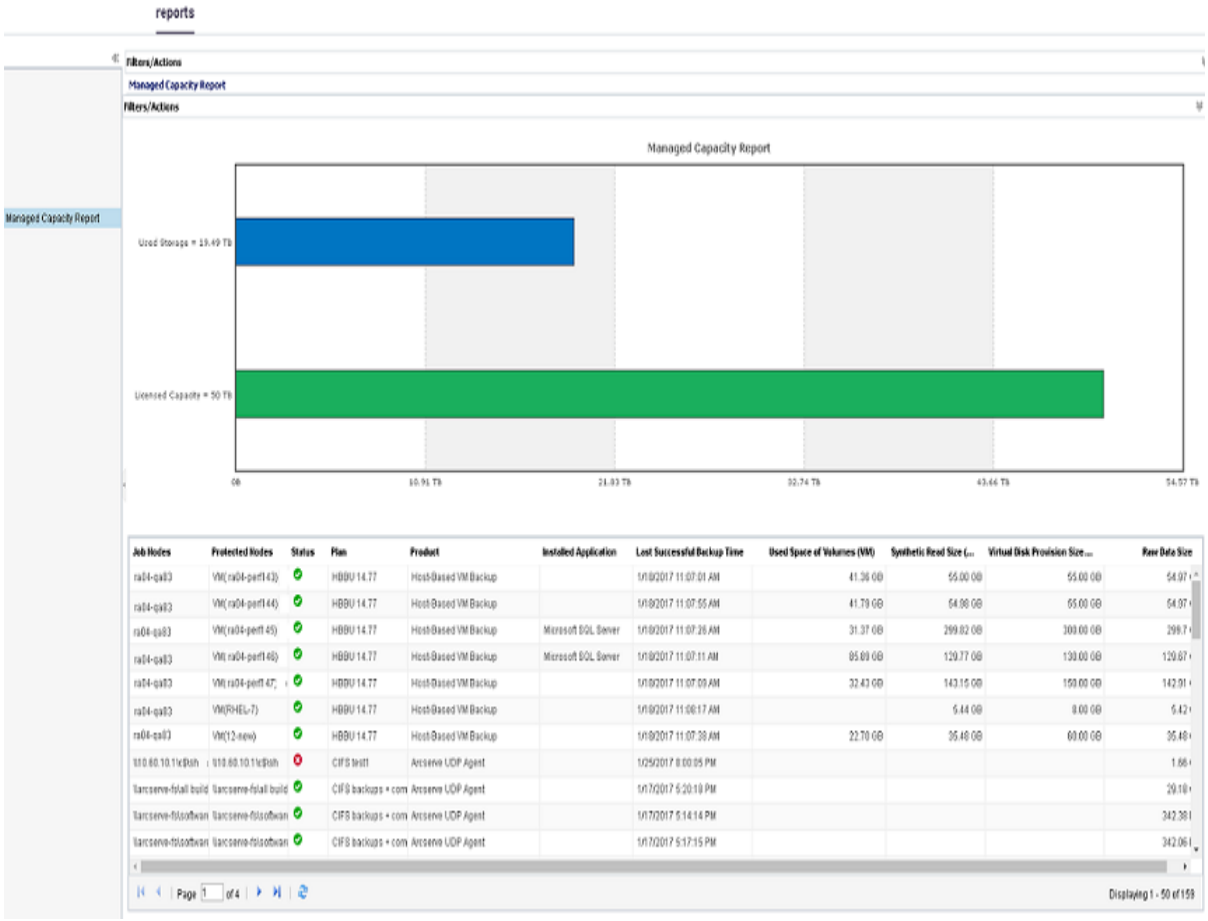
  Virtual Disk Provision Size

  Indicates the sum of the provision size of all virtual disks of the VM.

- Typically, Raw Data Size is the size of the data written into the backup destination. For host-based agentless backup, it may or may not be equal to

Synthetic Read Size because Arcserve UDP does not write all-zero data blocks to the backup destination. In other words, all-zero data blocks are skipped during the backup. In addition, for VM nodes that are protected by host-based agent-less backup, you can customize the data displayed for Raw Data Size by configuring few registry values. For more information on configuring the registry values and associated behaviors, see Understanding Raw Data Size in Managed Capacity Report for Host-based Agentless VM Backup.

Managed Capacity Report

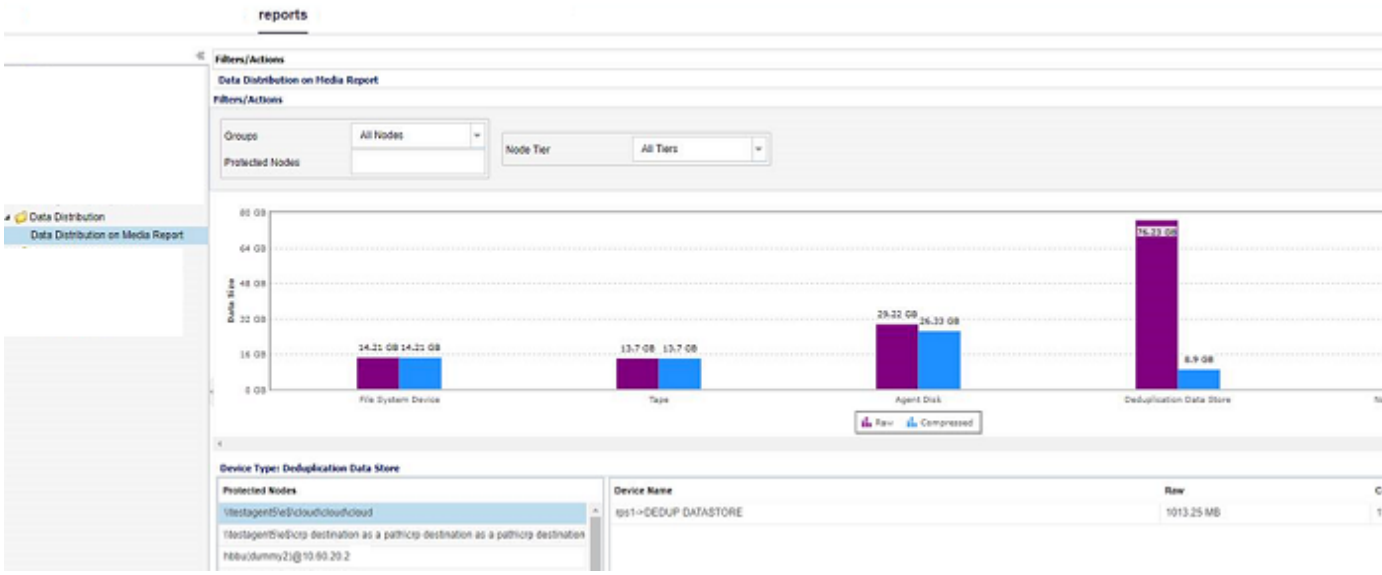


Data Distribution on Media Reports

Displays the compressed and actual (raw) backup data size for different destination types (deduplication or nondeduplication and local destination). This report lets you view all nodes, including nodes of Arcserve Backup and Arcserve UDP Agent. To view Arcserve Backup data information, enable **Arcserve Backup Data Synchronization Schedule** in the **settings** tab. To view latest information, click **Run Now** in the settings tab.

For details about filters and actions, see the common filters and actions.

Data Distribution on Media Report



Note: The Data Distribution on Media report depends on the recovery point data available on the backup destination. You can use the **Refresh** option to initiate a synch on demand and get the latest status in the reports.

RPO Reports

Recovery Point Objective (RPO) report is the compliance report that displays how the recovery points are distributed in the backup environment. The report helps assessing in case of a disaster what is the oldest and latest point in time that the node can return to. As the RPO report is for backup destination, the report gets populated with the data when you have some backup ready. The report gets populated either through on demand refresh or according to schedule.

Note: The report also gets populated directly from Arcserve UDP Dashboard. Clicking inside the RPO bar graph on a month at Dashboard displays RPO report screen for that specific month on the RPO page.

Following types of information is provided:

- Monthly distribution of recovery points in the backup destination.
- Age of the newest recovery point available for each node .
- Age of the oldest recovery point available for each node.
- Numbers provided on top of every bar.
- *Destination type* lets you select type of destination. For example, Cloud destination, Recovery Point Server, and Local Share.
- *Destination Name* lets you select a destination from the list of destinations that appears as a result of the option you select as Destination Type.
- Clicking on specific bar inside any of the following three graphs displays specific results: *Number of Recovery Points* for every month, *Age of Newest Recovery Point*, and *Age of Oldest Recovery Point*

Besides the common filters and actions, RPO reports provide unique option of **Refresh** in local filter.

RPO Report



Note: The report depends on the recovery point data available on the backup destination. You can use the **Refresh** option to initiate a sync on demand and get the latest status in the reports.

RTO Reports

Recovery Time Objective (RTO) report is the compliance report that displays if the defined recovery time objective is met for all the executed recovery type of jobs. The RTO report displays the following types of status:

- **RTO Met:** Recovery Job has met the defined objective.
- **RTO Not Met:** Recovery Job has not met the defined objective.
- **RTO Not defined**: Objective is not defined for the recovery job.
- **RTO Not Tested:** Recovery time Objective is defined but job has not been tested yet.

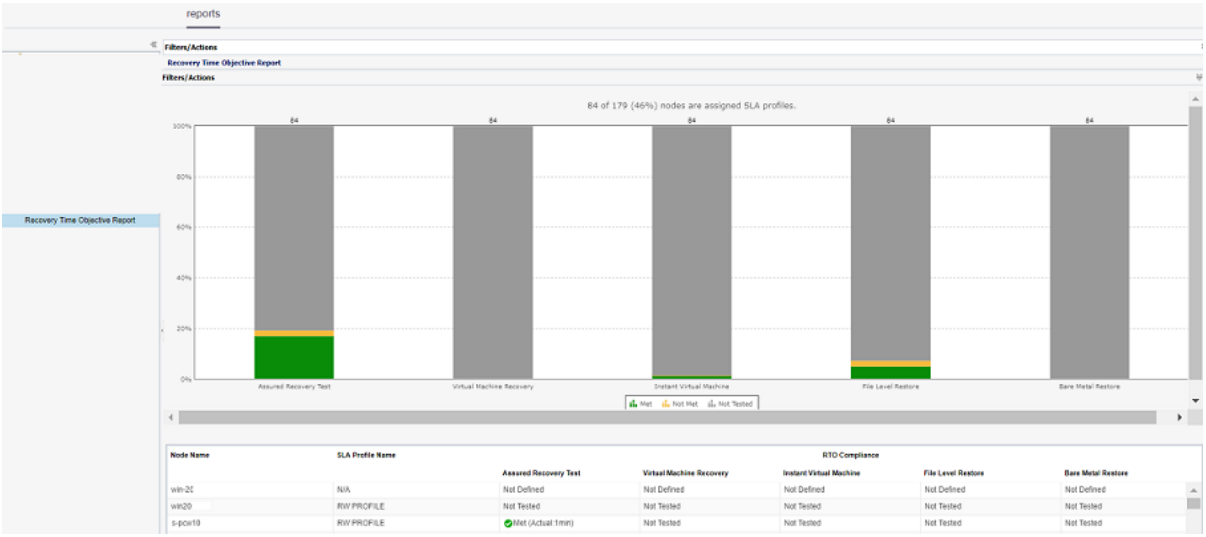
You can define recovery time objectives for all recovery jobs. For more information, refer to creating the SLA profile.

Besides the common filters and actions, RTO reports provide unique option of **SLA Profile** and RTO Status in the local filter.

Note: RTO report is not supported for Bare Metal Restore jobs that run for backup jobs configured on the local destination.

Dashboard of RTO Report

RTO Report Considerations	Description
Parameters considered for generating RTO chart	• Met • Not Met • Not Tested
Formula to define RTO status Percentage (%)	SLA Profile assigned nodes/Total nodes*100
Color coding used for chart	• Met = Green • Not Met = Red • Not Tested = Grey


reports
Filters/Actions
Recovery Time Objective Report
Filters/Actions
Recovery Time Objective Report
84 of 179 (46%) nodes are assigned SLA profiles.
84
84
84
84
84
100%
80%
60%
40%
20%
0%
Assured Recovery Test
Virtual Machine Recovery
Instant Virtual Machine
File Level Restore
Bare Metal Restore
Met
Not Met
Not Tested
Node Name
SLA Profile Name
RTO Compliance
Assured Recovery Test
Virtual Machine Recovery
Instant Virtual Machine
File Level Restore
Bare Metal Restore
win-2(
N/A
Not Defined
Not Defined
Not Defined
Not Defined
Not Defined
win20
RW PROFILE
Not Tested
Not Tested
Not Tested
Not Tested
Not Tested
s-pcw10
RW PROFILE
Met (Actual 1min)
Not Tested
Not Tested
Not Tested
Not Tested

Job Status Reports

Job status report helps you generate the overall status for all the job types that are executed in console. You can find this report useful to meet audit requirements. Besides the common filters and actions, job status report provides the capability to view results based on Job Type, Job ID, Job Status and Plan Name (Both active and deleted). Also, for duration display time format in Hours:Minutes:Seconds.

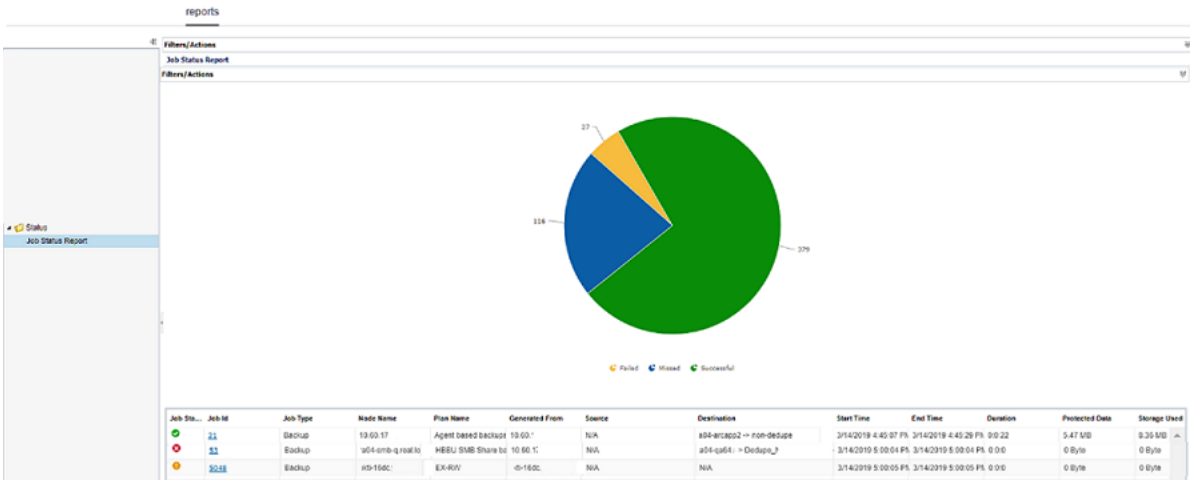
You can click the Job ID to view details from the log. You can also view the Job status of each job.

You can view the Actual Values at Source and Destination for Job Types. The report displays the following job status:

- **Successful:** Provides a list of jobs that are completed successfully.
- **Failed:** Provides a list of jobs that failed.
- **Incomplete**: Provides a list of jobs that finished with incomplete status.
- **Canceled:** Provides a list of jobs that are canceled.
- **Missed:** Provides a list of jobs that are not attempted.

Note: The default retention duration to retrieve job status is 180 days. You can modify the configuration to customize the Default retention Days for purge.

Job Status Report



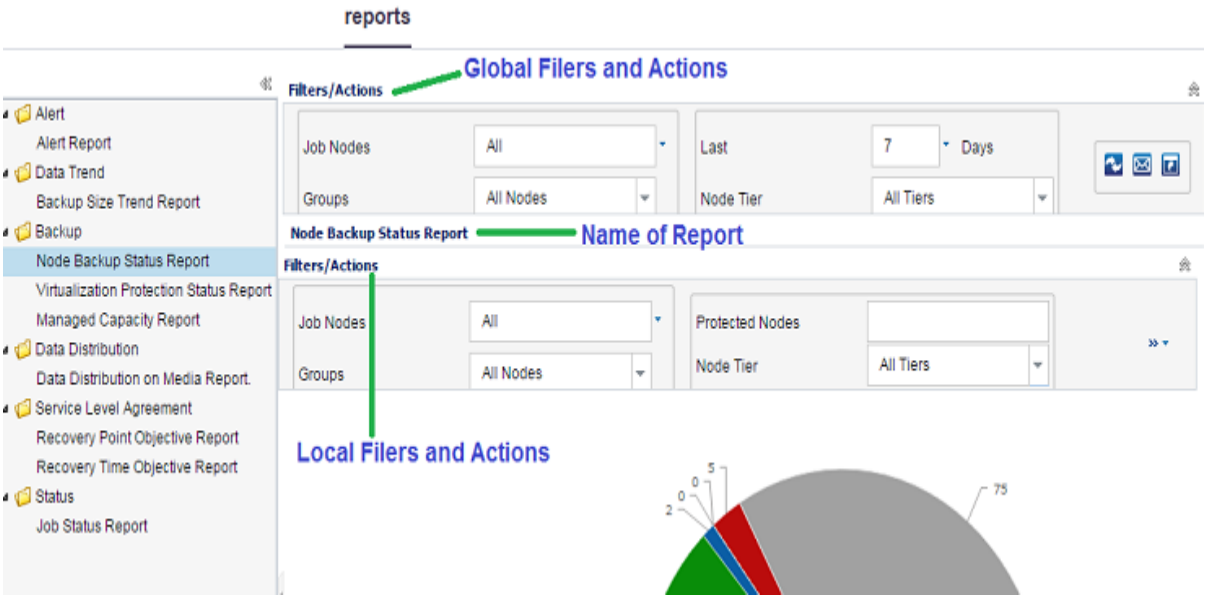
Note: When there are multiple jobs running, the Job status report generates. However, when you click on the pie chart, the displayed drill-down data may be inconsistent. To fix this inconsistency, refresh the report.

Using Filters and Actions

Every report page contains two options of Filters/Actions. The first option is the global option that appears on top of the report page. The other option is the local option that appears below the name of the report on the report page and provides solutions related to a particular report.

As a prerequisite, install Microsoft .NET Framework (version 3.5) on the machine where you have installed the Console for the Report Chart export feature to export images in a report successfully.

Note: Adobe Flash is no longer required and installing .NET Framework as a pre-requisite allows you to view the charts.



Filters

Global and local options contain filters where you can enter data to set report viewing options. The available options for global filters are similar for all the reports. The available options for local filters vary for different reports.

Actions

For Reports using Global Option:

- **Refresh:** Lets you update the information related to the page.
- **Schedule Reports to send by Email:** Lets you create a schedule for reports to send using Email. For more information, see Schedule Emails.

  Note: A schedule email exports maximum of 5000 records.
- **Reset:** Lets you change all filter parameters to the default values.

- **Report view show only one report:** Lets you view one report in a single pane.
- **Report view show multiple reports in two columns:** Lets you divide report viewing pane into two columns to view multiple reports.
- **Report view show multiple reports in three columns:** Lets you divide report viewing pane into three columns to view multiple reports.

  For Reports using Local Option:

- **Print:** Click the icon to print the report.

  Note: Using Print, you can only access the first 50 nodes or alerts from the complete list of data.

- **Refresh:** Click to update the report related information.
- **Email:** You can email the report. For more information, see Send Report by Email.

  Note: An email exports maximum of 10000 records.

- **Save:** You can use the option to export a report. Select one of the formats from **CSV**, **PDF**, and **HTML**, and then click **Open** or one of the options of **Save** from the dialog displayed at the bottom of the page to export the report.

  Note: Using Save, you can only access the first 10000 nodes or alerts from the complete list of data.

Working with Arcserve UDP Reports

This section contains the following topics:

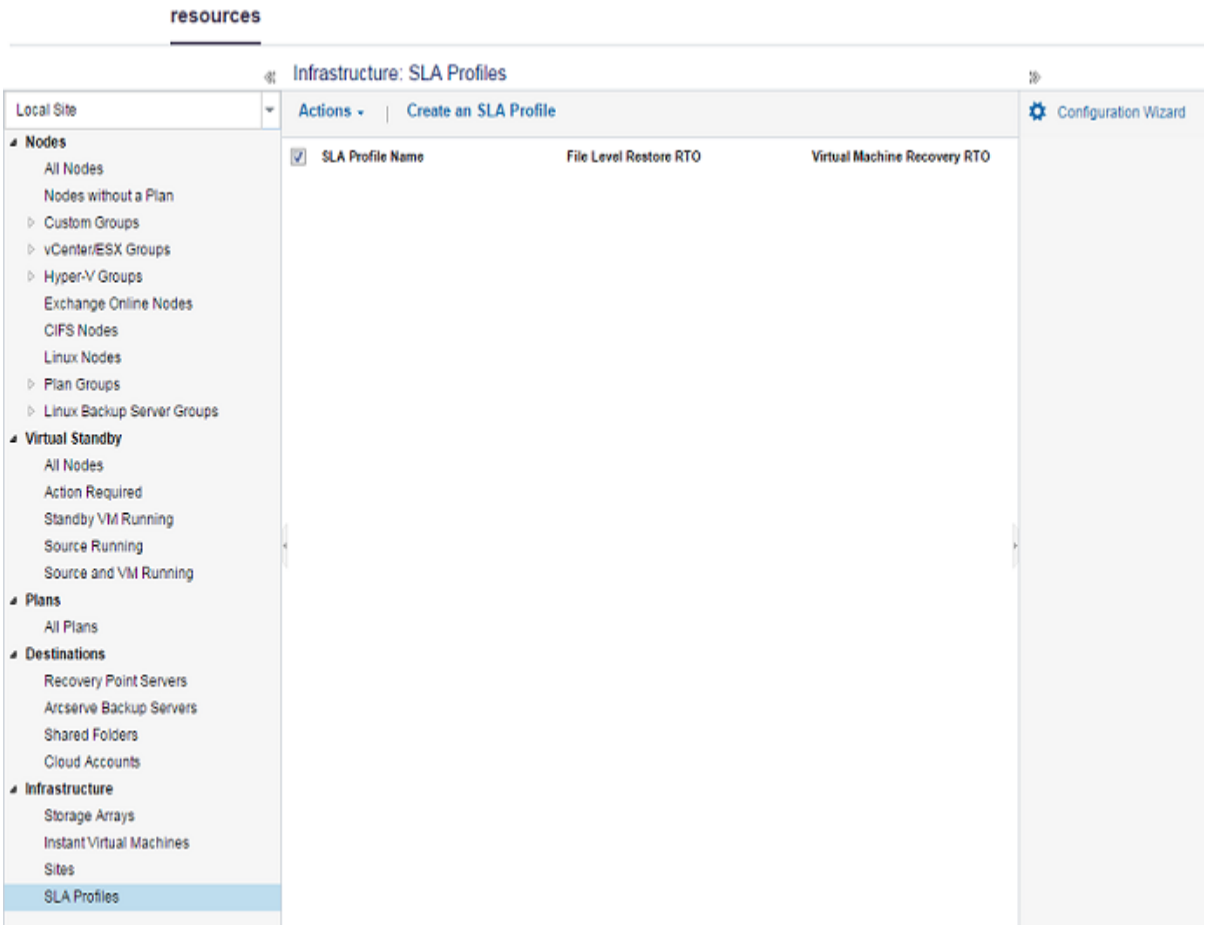
1. Create an SLA Profile
2. Schedule Emails
3. Send report by Email
4. Generate a Report
5. Customize Retention Days for a Job Status Report
6. Raw Data Size in Managed Capacity Report for Host-based Agentless VM Backup
7. View Actual Values at Source and Destination for Job Types
8. Enable Adobe Flash Player on Windows Server 2016

Create an SLA Profile

You need to create an SLA profile to generate RTO reports.

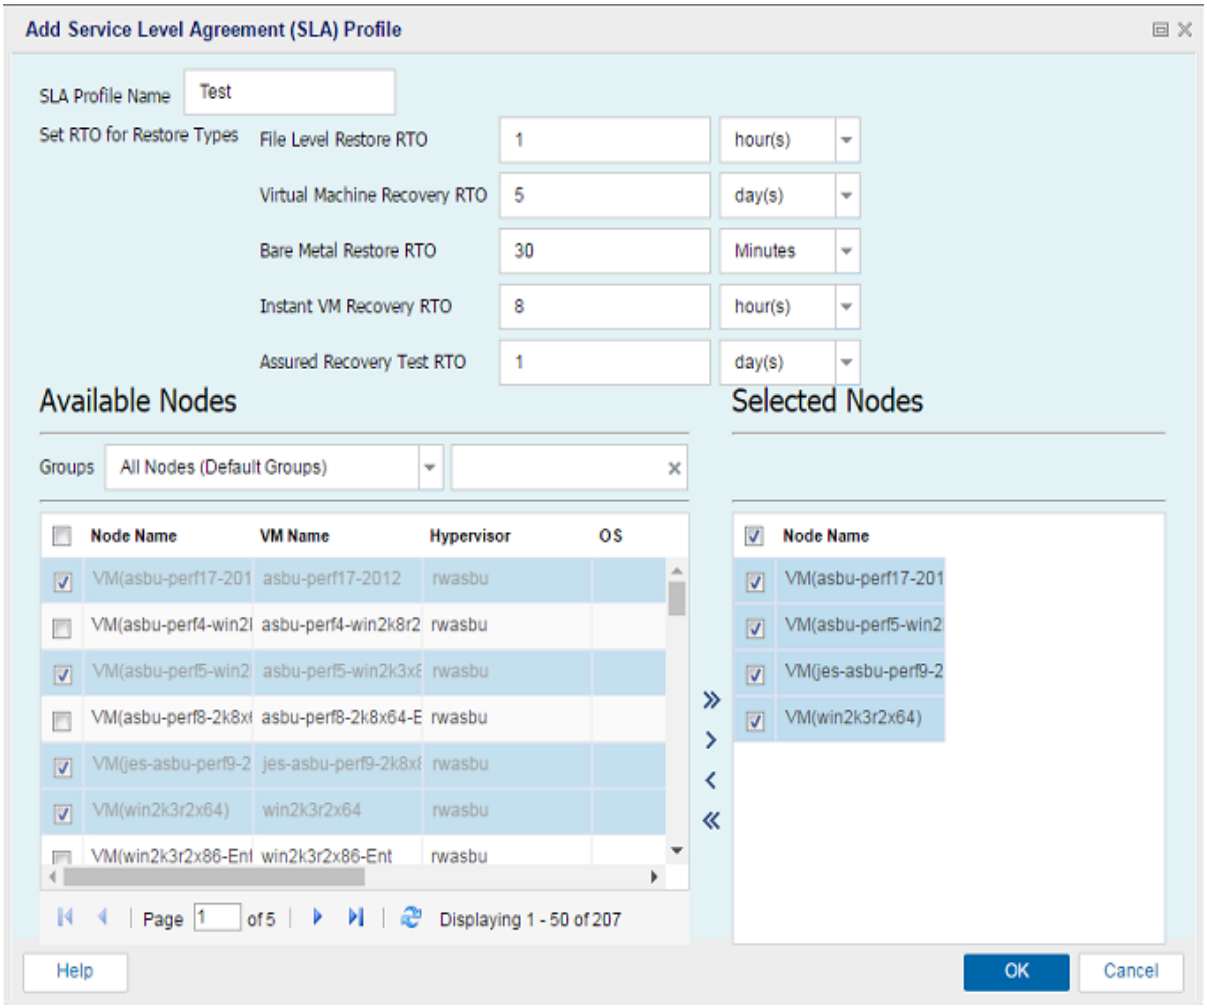
Follow these steps:

1. Navigate to the left pane of the **resources** tab, click **Infrastructure**>**SLAProfile**.



2. Click **Create an SLA Profile** from the center pane.

   Add Service level Agreement (SLA) Profile opens.



3. From the Add Service level Agreement (SLA) Profile pane, perform the following steps:

   a. Enter an **SLA Profile Name**

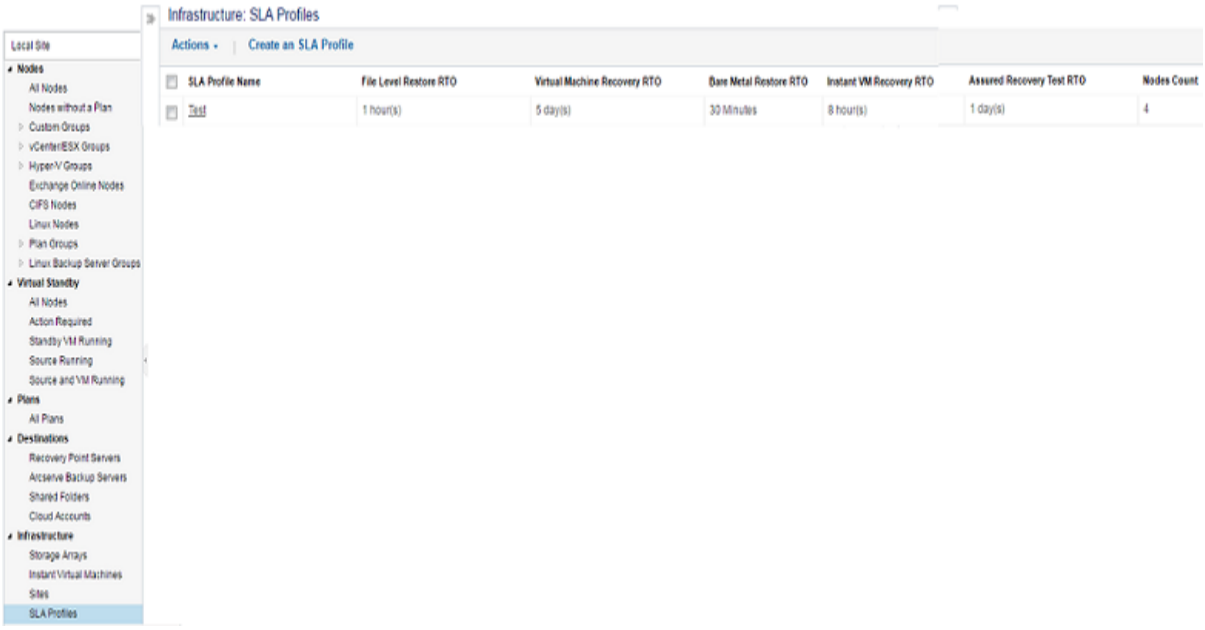
   b. Enter type of time and duration of time for desired options to **Set RTO for Restore Type**.

      Note: For all the options, you can select time in days, hours, and minutes.

   c. From the **Available Nodes** section, select the check boxes of nodes for which you want to generate reports, and move to **Selected Nodes** section.

   d. Select one or more nodes under **Selected Nodes**, and click **OK**.

   The SLA Profile is created and added under SLA Profile Name.



4. To modify or delete the existing SLA profile, select the check box of desire profile and click **Actions**.



You can view the **RTO** report using the **reports** tab to understand the job status for all the defined SLA profiles.

Schedule Emails

Using Arcserve Unified Data Protection, you can create a schedule to send reports by email to specified recipients.

Note: Before creating a schedule to send an email, configure the email settings. For more information about how to configure, see Configure Email and Alert.

You can create a schedule, and edit the schedule.

Create Schedule

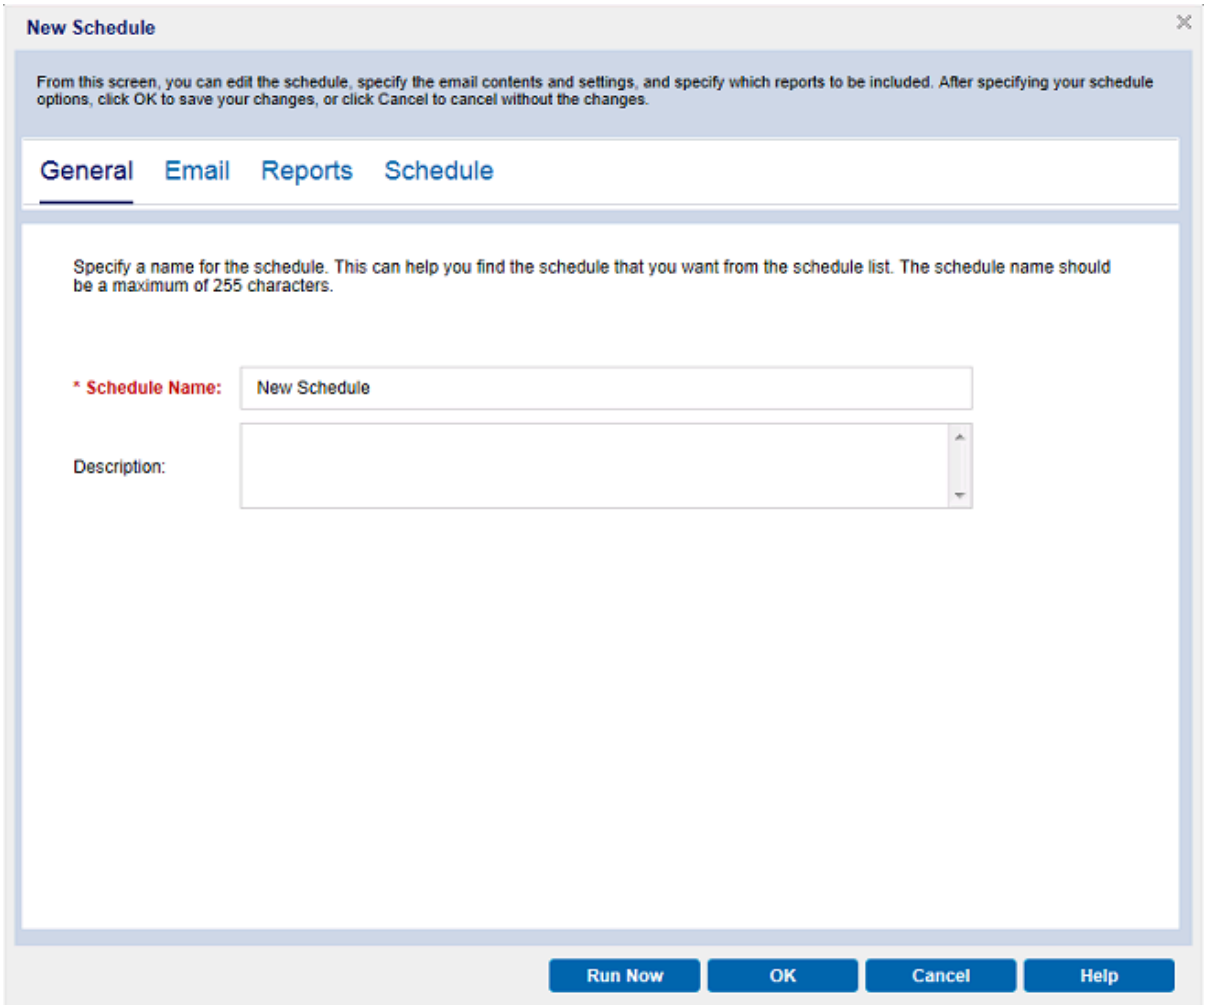
You can add new schedule for the emails report. These report emails are automatically updated, generated, and sent as scheduled. You can customize the schedule of the report email messages. The application lets you define the email content, the reports to attach, to whom to send the reports, and the date and time to send the report. The selected reports display detailed information in table format within the email.

Follow these steps:

1. Log into the Arcserve Unified Data Protection.
2. Click **reports** on the Navigation bar.
3. From the upper right corner of any report, click the global **Filters/Actions** section.
4. From the expanded list, select the email icon to open the **Schedule reports to send by Email** dialog.

   The **Schedule Emails** dialog is displayed.
5. Click **New** on the **Schedule Emails** dialog.

   The **New Schedule** dialog is displayed.



The following tabs are displayed:

- **General:** Specify a name and description (optional) for the new schedule.
- **Email:** Specify the mail settings, content, and attachment for the email schedule.
- **Reports:** Select the specific reports that you want to include in the email.
- **Schedule:** Specify a schedule for the email.

6. Complete the required fields in each tab.
7. Click **OK** to save the schedule.

   The new schedule is added to the **Schedule Emails** dialog.

   Note: Do not click **OK** if you want to view the report immediately.
8. (Optional) To view the report immediately, click **Run Now**.

   The report is sent to the recipients.

Edit Schedule

Using Arcserve UDP, you can update a schedule that you added using Create Schedule.

Follow these steps:

1. Log into the Arcserve Unified Data Protection.
2. Click the **reports** tab.
3. Click the global **Filters/Actions** section.
4. From the expanded list, select the email icon to open the **Schedule Emails** dialog.
5. Click **Edit** on the **Schedule Emails** dialog.

   The **Edit Schedule** dialog is displayed.
6. Update the schedule details, and click **OK**.

   The updated schedule is displayed at the **Schedule Emails** dialog.

   Note: Do not click **OK** if you want to view the report immediately.
7. (Optional) To view the updated report immediately, click **Run Now**.

   The report is sent to the recipients.

Send Report by Email

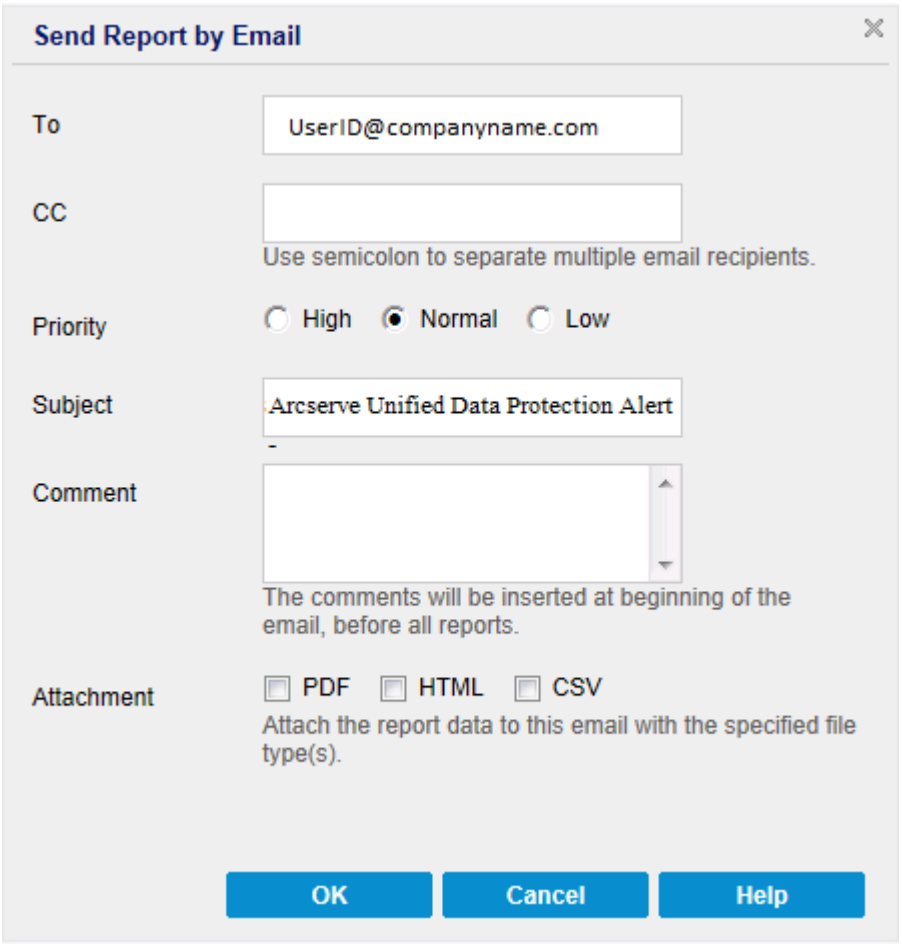
Using Arcserve Unified Data Protection, you can send individual reports to specific recipients. When you send a report by email, the content is the same as the printed content and all graphical charts are sent as embedded images.

Note: Before using the **Send Report by Email** option, configure the **Email settings**. For more information about how to configure, see Configure Email and Alert.

Follow these steps:

1. Log into the Arcserve Unified Data Protection.
2. Click **reports** on the Navigation bar, and select one of the reports.
3. Click the local **Filters/Actions** section, available below the name of the selected report.
4. From the expanded list, select the email icon to open the **Send Report by Email** dialog.

   Note: If the email configuration is not complete, a **Warning** dialog informs that the emails settings are not specified. For more information about how to configure, see Configure Email and Alert.



5. Complete the following fields:

   - **To**: Specify the recipient the email is sent to.

     Note: This field defaults to the email address specified in the Email Configuration module.

   - **CC**: Specify additional recipients, separated by semicolons, you would like to email the report to.
   - **Priority**: Specify the priority of the email. This field defaults to Normal.
   - **Subject**: Specify the subject of the email. This field defaults to the report you selected.
   - **Comment:** (optional) Enter any information that you want to share.
   - **Attachment**: Select the formats to attach the report data.

6. Click **OK**.

   The email is sent successfully.

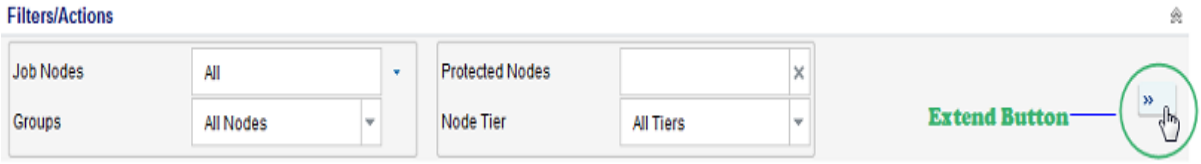
Generate a Report

You can generate predefined reports from the **reports** tab. You can generate the reports in the PDF, CSV, and HTML formats.

Follow these steps:

1. Navigate to the **reports** tab and select a report from the left pane.
2. Click the local **Filters/Actions** drop-down list.
3. Enter or select the details in the **Filters/Actions** drop-down options.
4. From the drop-down list of the **Save** button, click **CSV**, **PDF**, or **HTML**.

   Note: Large images or more data in a report page may hide some of the options, including the Save button. To view these options, click the Extend button.



The report is generated in the selected format.

Customize Retention Days for a Job Status Report

The default retention days for retrieving job status report is 180 days. You can modify the number of retention days that is available in the configuration mentioned in ConsoleConfiguration.xml from the following location:

Program Files\Arcserve\Unified Data Protection\Management\Configuration\ConsoleConfiguration.xml

According to your requirement, you can modify the number for retention days in the configuration.

```
<PurgeConf>
<!-- retentionDays, defalut value is 180 (180 days).
The unit is day.
Number of days to retent data for console database.-->
<retentionDays>180</retentionDays>
<!-- purgeHourOfDay, defalut value is 0 (0:00 a.m.).
Execute purge job on this time point every day.-->
<purgeHourOfDay>0</purgeHourOfDay>
</PurgeConf>
```

Raw Data Size in Managed Capacity Report for Host-based Agentless VM Backup

You can configure the following registries at the proxy level or VM level:

At proxy level for all VMs protected by the current proxy:

[HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll]

"CountNtfsVolumeSize"=dword:00000001

"ReportZeroIfHavingNonNtfsVolume"=dword:00000001

"BackupZeroBlock"=dword:00000001

At VM level for a specific VM:

[HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll\{VM UUID}]

"CountNtfsVolumeSize"=dword:00000001

"ReportZeroIfHavingNonNtfsVolume"=dword:00000001

"BackupZeroBlock"=dword:00000001

Note: The VM-level registry takes precedence over the proxy-level registry.

The following behaviors are observed:

Registry key	Default Setting	Possible Values	Purpose	Additional Info
GetVMGuestVolumeUsedSize	0	0 or 1	Specifies whether to count the used space of all the volumes of a VM. 0 - Counts the Virtual Machine VMDK file	

			size as the Raw data size of the VM. 1- Counts Only the used space of the volumes. **Note:** For Windows Guests, only Ntfs volumes are considered. For Linux guests all the volumes are considered.	
ReportZeroIfHavingNonNtfsVolume	0	0 or 1	Specifies the behavior for Raw Data size column of Managed Capacity Report when non-Ntfs volume (s) exist on the virtual machines. 0 - Size of Non Ntfs volume(s) is not con-	1. Ignored when GetVMGuestVolumeUsedSize=0 2. Applicable only to Windows VM

			sidered as part of RAW Data Size Column. 1 - Overall Raw Data Size is shown as 0 when non-Ntfs volume (s) exist on the Virtual Machine	
BackupZeroBlock	0	0 or 1	Specifies whether to write all zero data blocks to the backup des-tination. 1- Writes zero data blocks to the backup des-tination. 0- Ignores the zero data blocks as part of backup.	1. Ignored when GetVMGuestVolumeUsedS-ize=1 2. If changed, the value takes effect only after a full backup

Example

A VM has one thin-provision virtual disk with the provisioned size 1000 GB. The size of the virtual disk's VMDK file is 800 GB, among which 200 GB data blocks are all-zero data blocks. In the guest OS of this VM, there are 2 NTFS volumes whose used space are 100 GB and 200 GB respectively, and 1 FAT32 volume with the used space as 1 GB.

Key Names	Default Val-	Customized	Customized	Customized

	ues	Values	Settings	Values
GetVMGuestVolumeUsedSize	0	0	1	1
ReportZeroIfHavingNonNtfsVolume	N/A	N/A	1	1
BackupZero Block	0	1	N/A	N/A
Expected Raw DataSize	600 GB (excluding zero data blocks)	800 GB (including zero data blocks)	300 GB	0

View Actual Values at Source and Destination for Job Types

The following table describes the actual values at source and destination displayed for all the job types to help you understand the Job Status report:

Job Type	Actual Value at Source	Actual Value at Destination	Comments
FS Backup to Datastore (Dedupe / Non-Dedupe)- Windows, Linux nodes	Not Available (N\A)	Displayed as expected	
FS Backup to Network Share	N\A	Displayed as expected	
FS Restore from Network Share	N\A	N\A	
CIFS Backup jobs –Non-Dedupe Datastore	N\A	Displayed as expected	
CIFS Restore jobs – Non-Dedupe Datastore	Displayed as expected	N\A	
FS Restore from Datastore (Dedupe / Non-Dedupe)	Displayed as expected	N\A	
File copy from Datastore (Dedupe / Non-Dedupe), File copy from Replication Ds	Displayed as expected	N\A	**Note:** Here “Protected Data”, “Storage used” columns are also displayed as “N\A”
File Archive from Datastore (Dedupe / Non-Dedupe) , Flie Archive from Replication Ds	Displayed as expected	N\A	**Note:** Here “Protected Data”, “Storage used” columns are also displayed as “N\A”
File copy / File Archive to N\W Share	N\A	N\A	**Note:** Here “Protected Data”, “Storage used” columns are also displayed as “N\A”
File copy / File Archive Restore	N\A	N\A	**Note:** Here “Protected Data”, “Storage used” columns are also displayed as “N\A”
File Archive Delete	N\A		**Note:** Here “Protected Data”, “Storage used”

			columns are also displayed as “N\A”
FS Catalog job (Agent-based/A-gentless)	Displayed as expec-ted	N\A	
CRP from Datastore (Dedupe, Non-Dedupe)	Displayed as expec-ted	N\A	**Note:** Here “Protected Data”, “Storage used” columns are also displayed as “N\A”
CRP from N\W share, Local	N\A	N\A	**Note:** Here “Protected Data”, “Storage used” columns are also displayed as “N\A”
CRP Restore to N\w Share, Local	N\A	N\A	**Note:** Here “Protected Data”, “Storage used” columns are also displayed as “N\A”
Instant Virtual Machine jobs	Displayed as expec-ted	N\A	**Note:** Here “Protected Data”, “Storage used” columns are also displayed as “N\A”
Assured Recovery jobs	Displayed as expec-ted	N\A	
Bare Metal Recovery fromData-store (Dedupe/ Non-dedupe)	Displayed as expec-ted	N\A	**Note:** Here “Protected Data”, “Storage used” columns are also displayed as “N\A”
Virtual Standby jobs	Displayed as expec-ted	N\A	**Note:** Here “Protected Data”, “Storage used” columns are also displayed as “N\A”
Copy to Tape jobs	Displayed as expec-ted	N\A	**Note:** Here “Protected Data”, “Storage used” columns are also displayed as “N\A”
Replication jobs, Cross-site Rep-lication jobs	Displayed as expec-ted	N\A	**Note:** Here “Protected Data”, “Storage used” columns are also displayed as “N\A”
MSP Replication jobs	Displayed as expec-	Displayed as expected	**Note:** Here “Protected Data”, “Storage used”

	ted		columns are also displayed as “N\A”
RPS Jumpstart from N\w Share	Displayed as expec-ted	N\A	**Note:** Here “Protected Data”, “Storage used” columns are also displayed as “N\A”
RPS Jumpstart (Dedupe/ Non-dedupe Datastore)	Displayed as expec-ted	Displayed as expected	**Note:** Here “Protected Data”, “Storage used” columns are also displayed as “N\A”

Chapter 19: Managing High Availability

This section contains the following topics:

How High Availability Works

Using Arcserve Unified Data Protection, you can monitor and manage High Availability functions from the **high availability** tab. To manage these functions, you must first log in to Control Service. When you first click the **high availability** tab, the **Add Control Service** dialog opens. This dialog will not appear afterwards.

Manage High Availability Control Services

Follow these steps:

1. Click the **high availability** tab.

   The **Add Control Service** dialog opens.

2. Enter the Control Service details such as the IP address, account name, password, protocol, and port number.

3. Click **OK**.

   The specified control service is added below the Control Services and Scenarios heading in the left pane. To modify or delete a control service, select the Control Service and right-click to see the options. You can also select the Control Service in the center pane and click **Actions** to modify or delete a Control Service. Or, right-click the control service in the navigation pane

   Note: Expand the control service to see scenarios, groups, and other details.

Manage High Availability Licenses

Follow these steps:

1. Click the **high availability** tab.
2. On the left pane, click **Control Services and Scenarios**.

   The **Control Services and Scenarios** page is displayed.
3. Select the Control Service and click **Register**.

   The **Register** dialog opens.
4. Enter the registration key
5. Click **OK**.

   The license is registered.

Manage Scenarios

Arcserve Unified Data Protection lets you manage your existing HA scenarios and you can create Full System scenarios. You can also create scenario groups to organize your scenarios. The following sections describe how to manage HA scenarios:

- Manage Scenario Group
- Create Full System Scenarios
- Manage Scenarios
- Edit Scenarios
- Manage Scenario Hosts
- Operations on Scenarios
- BMR and Reverse Replication
- Monitor Scenarios

Manage Scenario Group

Arcserve Unified Data Protection lets you manage groups in a control service. You can add, rename, delete, flag, and post comments for a group.

Follow these steps:

1. Select a managed control service from the left pane.

   All the groups in the control service are listed in the center pane.

2. Click the **Actions** drop-down menu, and then click one of the following:

   Add Scenario Group

   Creates a group.

   Select a group to perform the following actions:

   Rename Scenario Group

   Renames the group.

   Remove Scenario Group

   Deletes the group. You cannot remove a group if you have scenarios within the group.

   Flag and Comment

   Flags the group in various colors and lets you add comments for the flag. Use flags to personalize and easily identify your group.

3. Optionally, right-click a group in the left pane to add, delete, or rename a group from the selected control service.

   The group is added or updated based upon your selection.

Create Full System Scenarios

Follow these steps:

1. From the left pane, click **Control Services and Scenarios** and then click a managed control service.

   All the scenario groups in the control service are listed.

2. Click a scenario group.

   The **Scenarios** page is displayed in the center pane.

3. On the center pane, click **Create Scenario**.

   Note: Optionally, on the left pane, you can right-click and then click **Create Scenario**.

   The **Create Full System** wizard opens and the **Select Server and Product Type** dialog opens.

4. Enter the scenario name, select the product type and specify whether you want AR testing.

5. Click **Next**.

   The **Master and Replica Host** dialog opens.

6. Enter the Master and Replica details.

7. Click **Next**.

   The engines are verified on the hosts if you selected the **Verify Engine on Hosts** option. You can also install engines to the hosts or uninstall the engines from the hosts.

8. Click **Next** after the engines are verified.

   The **Volume Setting** dialog opens.

9. Select the volumes you want to protect.

   Note: When you select the **Enable Exclude Directory and files** option, the pagefile.sys, hyberfil.sys, system volume information, Recycler, and Recycled files and folders are filtered by default.

10. Click **Next**.

    The **Resource Pool Selection** dialog opens. You can select the resource pool where the VM is located after switchover or during AR testing.

11. Click **Next**.

    The **Storage Selection** dialog opens.

12. Select the data store to store the virtual machine. Optionally, select **Allocated and commit space on demand (Using Dynamic Disk)**. The generated VM uses thin provision for its virtual disk if you select this option.
13. Click **Next**.

    The **Scenario Properties** dialog opens.
14. Expand the properties and modify as desired and then click **Next**. For more information, see the Arcserve RHA Administrator Guide.

    The **Master and Replica Properties** dialog opens.
15. Review the master and replica properties and then use **Click to edit physical network mappings**.

    The **High Availability Network Adapter Mapping** dialog opens.

    Note: When there is only one virtual network adapter in both the master and replica servers, they are mapped automatically.
16. Perform the following steps:
    a. From **Replica Network Adapter**, click to choose the adapter to map to the adapter listed in the Master Network Adapter column.
    b. From **Apply master adapter information**, (Default) select if the Master Adapter is in the DHCP mode.
    c. From **Customize adapter information**, select to enable the IP, Gateways, DNS Servers, and WINS Servers setting. Add or remove IP address, Gateways, DNS Servers, and WINS servers as required.
17. Click OK to close the **Network Adapter Mappings** dialog and then click **Next** to continue.

    The **Switchover Properties** dialog opens.
18. Expand the **Network Traffic Redirection** and other properties to verify the values and then click **Next**.

    The **Switchover and Reverse Replication Initiation** dialog opens.
19. Specify the switchover type. For full system scenarios, the reverse replication is manual.
20. Click **Next**.

    Wait while the Scenario Verification process completes and opens the **Scenario Verification** dialog.

If the Scenario Verification process lists any errors, you must resolve them to continue. If any warnings are listed, you should also resolve them to successfully continue. After making changes, click **Retry** to repeat verification.

21. Click **Next**.

    The **Scenario Run** dialog opens.

22. Click **Finish** to save the current settings and run the scenario later.

    Optionally, to run the scenario instantly, select **Run Now after clicking Finish button** and then click **Finish**.

    For full system scenarios, choose **Volume Synchronization**.

    The scenario is created.

Manage Scenarios

When you select a managed control service from the left pane, all scenarios in the control service are displayed in the center pane. The scenarios are listed with its type, state, product, mode. The statistics of RPO/RTO, master spool usage and synchronization progress are also listed here. You can perform various operations such as delete, rename, flag, or comment by selecting a scenario.

Follow these steps:

1. From the left pane, click **Control Services and Scenarios** and then click a managed control service.

   All the scenario groups in the control service are listed in the center pane.

2. From the left pane, click a scenario group.

   Scenarios in the scenario group are listed in the center pane.

3. Select a scenario.

4. Click the **Actions** drop-down menu, and then click one of the following:

   Rename Scenario

   Renames the scenario.

   Remove Scenario

   Deletes the scenario.

   You cannot remove a group if you have scenarios within the group.

5. Optionally, from the left pane, right-click a scenario to delete or rename the scenario.

   The scenario is updated.

Edit Scenarios

Arcserve Unified Data Protection lets you edit scenario properties when the scenario is in the stopped state. You can insert, rename, or delete hosts or modify the topology of a scenario.

Follow these steps:

1. From the left pane, click **Control Services and Scenarios** and then click a managed control service.

   All the scenario groups in the control service are listed in the center pane.

2. From the left pane, click a scenario group and then click a scenario.

   The <scenario group>:<scenario> page is displayed.

3. Select a host from the scenario.

4. Click the **Properties** tab and select one of the following options from the drop-down list.

   Scenario Properties

   Updates the scenarios properties.

   HA Properties

   Updates the High Availability properties.

   Host Properties

   Updates the host properties.

   Root Directories

   Updates the root directories.

   Note: This is applicable only to Full System scenarios.

5. Click **Save** from the **Action** drop-down menu.

   The scenario properties are updated.

   Edit the virtual platform settings for a stopped Full System scenario:

   Follow these steps:

1. Select a replica host from the scenario.

2. Click the **Properties** tab and **Host properties** from the drop-down list.

3. Expand **Virtual Machine** and click **Click here to edit virtual platform setting**.

   The **Virtual Platform Setting** wizard opens.

4. Select the **Virtual Platform Type** and the related IP address or host name.

5. Select the resource pool for ESX and vCenter, or Host server for Citrix Xen.
6. Select the storage. For Hyper-V, browse the directories and select the location of the VM on the Hyper-V server.
7. Click **Finish**.

Edit network adapter mapping for High Availability or Assured Recovery:

Follow these steps:

1. Select the replica host from the scenario.
2. Click the **Properties** tab and select **Host properties** from the drop-down list.
3. Expand **Virtual Machine** and **Virtual Machine Setting**.
4. Click **Click to edit physical network mappings** for either the **High Availability Network Adapter Mapping** or **Assured Recovery Network Adapter Mapping** property.

   The **High Availability Network Adapter Mapping** dialog opens.
5. Select the replica network adapters to map the master network adapter.

   You can customize the adapter information of the replica adapter by including the IP address, gateway, DNS servers and WINS servers.
6. Click **Okay**.

   The mappings are modified and saved.

Manage Scenario Hosts

You can insert, delete, and rename hosts in a scenario.

Follow these steps:

1. From the left pane, click **Control Services and Scenarios** and then click a managed control service.

   All the scenario groups in the control service are listed in the center pane.

2. From the left pane, click a scenario group and then click a scenario.

   The <scenario group>:<scenario> page is displayed.

3. Select a host from the scenario.

4. Click the **Edit** drop-down menu, and then click one of the following options:

   Insert Host

   Inserts a child host to the selected host in the scenario.

   Remove Host

   Deletes the selected host from the scenario.

   Rename Host

   Renames the selected host in the scenario.

   Save

   Saves all the modification to the scenario properties.

   Refresh

   Refreshes all modifications.

   The scenario properties are modified.

Operations on Scenarios

You can run various operations on the scenarios.

Follow these steps:

1. From the left pane, click **Control Services and Scenarios** and then click a managed control service.

   All the scenario groups in the control service are listed in the center pane.

2. From the left pane, click a scenario group and then click a scenario.

   The <scenario group>:<scenario> page is displayed.

3. Click the **Actions** drop-down menu, and then click one of the following options:

   Run

   After you create a scenario, you need to run it to start the replication process. Normally, before data changes on the Master can be replicated on the Replica, the Master and the Replica must be synchronized. Therefore, the first step in initiating replication is synchronizing the Master and Replica servers. After the servers have been synchronized, online replication starts automatically, continuously updating the Replica with all of the changes that occur on the Master.

   Run (Assessment mode)

   The assessment mode enables you to assess the accurate bandwidth usage and compression ratio benchmarking that is needed for replication, without actually replicating data. When you run this command, no replication occurs but statistics are gathered. A report is provided once the assessment process is stopped.

   Stop

   You stop a running scenario to set or change properties. You could stop the scenarios in running state or assessment mode.

   Synchronize

   Synchronization is a process to make data consistent in the Master and Replica. Activate the synchronization process (whether replication is running or not).

   Difference Report

   A Difference Report compares the differences between the Master and the Replica at a certain point in time. The comparison is performed using the same algorithms that are used in the synchronization process, but no data is transferred. A Difference Report is generated for each Replica and sent to the Manager at the end of the process. This report can be produced at any time.

Perform Switchover

Switchover (or failover) is the process of changing roles between the Master and Replica. This means, making the Master server the standby server, and the Replica server the active server.

Recover Active Server

When the switchover process did not complete properly, you could manually select the server that acts as the active server through a process called Recover Active Server.

Suspend Is Alive Check

Suspend the Is Alive check that verifies that the active server is operational. You can manually suspend/resume Is Alive checking for a running HA scenario.

Replica Integrating Testing

The Assured Recovery option enables you to perform a full transparent test of the recoverability of your data on the Replica server. The Replica server that is tested is the one that would take over the production server if it is down. The Assured Recovery option is a true test of the actual server, applications and actions that will be required in the event the Replica server will have to switch, become the Active server, and carry out its functions.

Start/Stop VM

Use this operation to start or stop a virtual machine from its latest system status or from a bookmark. You can start or stop a virtual machine after you create a scenario and synchronize the master and replica. Use this feature when the scenario is not running. This feature is available for Full System DR and HA scenarios. The Start/Stop is a toggle menu item.

Suspend Replication

Suspend replication updates on the Replica host in order to perform system maintenance or some other form of processing that does not modify the replicated data there. Changes continue to be recorded for update on the suspended Replica, but are not actually transferred until replication is resumed. You cannot suspend replication during synchronization.

Delete all VM resources

When you run a full system scenario, some temporary resources are created such as disk files, snapshots, and other files. This operation lets you delete these resources and is available when the scenario is not running.

Restore Data

Recover lost or corrupted Master data from any Replica by activating a synchronization process in the reverse direction.

Set Rewind Bookmark

A bookmark is a checkpoint that is manually set to mark a state back to which you can revert. This manual setting is called setting rewind bookmark. We recommend that you set a bookmark just before any activity that can cause data to become unstable. Bookmarks are set in real-time, and not for past events.

The selected operation is performed.

BMR and Reverse Replication

Arcserve Unified Data Protection lets you process BMR and reverse replication for your full system scenarios.

Follow these steps:

1. Prepare a bare metal machine by booting the computer from the BMR CD.
2. Select the full system scenario and click **Restore** from the **Actions** drop-down menu.

   The **Restore Data Wizard** opens.
3. Follow the instructions on wizard screens to create and run the recovery scenario.

   Note: On the **Volume Mapping** page, if the volumes are mapped automatically for the source and destination, the custom volume mapping is disabled. To enable the custom volume mapping, click **Clear** to remove the previous mapping. Right click on the selected volume and select **Custom volume mapping** to open the **Resize volume size** dialog and change the size as required.

To run Reverse Replication, follow these steps:

1. Prepare a bare metal machine by booting the computer from BMR CD.
2. Select the full system scenario that performed switchover or failover and click **Run** in the **Actions** drop-down menu.

   The **Restore Data Wizard** opens.
3. Follow the instructions on wizard screens to create and run the recovery scenario.

   Data is restored to the bare metal machine. If you selected automatic switchover, the switchover process is initiated and the bare metal machine is ready. If you selected manual switchover, you have to manually initiate the switchover process.

Monitor Scenarios

Arcserve Unified Data Protection lets you monitor DR or HA scenarios by providing various statistics and reports.

Follow these steps:

1. From the center pane, select a scenario.

   The status of the running scenario is displayed with details such as sent data, sent files, received data, received files, etc.

2. Click the **Statistics** tab to see more details. The tab has the following two categories:

   Running Statistics

   Displays the detailed statistic data when the scenario is running.

   History Record

   Displays reports for synchronization, difference reports, and AR testing reports.

3. Click the **Events** tab to see all events of a selected scenario. To copy or delete the events, select the events and right click, and then select **Show Events** to open the show events dialog to copy or delete the events. Use Shift+Ctrl keys to select multiple events.

   Note: The events are automatically refreshed. The five recent critical events are displayed in the pane when you select a scenario.

4. Select the scenario group from the left pane. All scenarios in the group are listed in the center pane. You can check RPO/RTO, master spool usage, and synchronization progress in this list.

5. The details in the right pane displays scenario information such as the scenario name, scenario state, and synchronization progress.

   Note: On the right pane, the Spool usage (% of spool) indicates the spool usage of the master in the scenario.

Remote Installation

Arcserve Unified Data Protection lets you deploy the engine from a managed control service to the remote hosts. You could also manage the installation and verification from the host list.

Follow these steps:

1. On the left pane, click **Remote Installation**.

   The **Remote Installation** page is displayed in the center pane.

2. From the **Control Service** drop-down list, select a control service which you want to use to deploy the engine.

   The existing hosts where the engine was installed or verified earlier are listed in the center pane.

3. From the **Action** drop-down menu, click **Add Hosts**.

   The **Hosts to Install Engine** dialog opens.

4. Enter the host name or IP address of the host and click **Add**.

   The host is added to the list.

5. Click **OK**.

   The **Add Hosts** dialog opens.

6. Select one of the following options:

   Edit Hosts

   Opens the **Hosts to Install Engine** dialog to let you add hosts or manage the existing hosts.

   Change Installation Settings

   Opens the **Edit Installation Settings** dialog. You can specify the following details:

   Installation Account

   Service Account

   Port

   Use previous settings when reinstall or upgrade

   Upgrades or reinstalls an existing RHA engine.

7. Click **OK**.

8. The host is displayed on the **Remote Installation** page.

The **Status** column displays the installation status.

Note: Move the mouse on the status to get the details if the installation fails.

Remote Installation Actions

You can perform various operations on the added hosts.

Follow these steps:

1. From the center pane, select a host.
2. Click the **Action** drop-down list and then select one of the following:

Add Hosts

Opens the **Hosts to Install Engine** dialog. See Remote Installation for more details.

Install/Upgrade

Installs or upgrades the HA engine on the selected host.

Uninstall

Uninstalls the HA engine from the selected host.

Edit Settings

Opens the **Edit Installation Settings** dialog.

Check Host Status

Verifies the existence of the host.

Remove Hosts

Removes the host from the list.

View Logs

Opens the **Remote Installation Logs** dialog and displays the logs of all remote hosts. Click to refresh to view the latest logs.

The operation is successfully completed.

High Availability Reports

Arcserve Unified Data Protection provides various reports to monitor the High Availability status. You can apply filters to generate various types of report as required.

Follow these steps:

1. Click the **Control Services and Scenarios** tab.
2. From the left pane, click **Reports**.
3. The **Reports** page is displayed in the center pane.
4. From the center pane, select a control service from the **Control Service** drop-down list.
5. Enter details and apply filters as required.

   The HA report is generated.

Chapter 20: Using the VSphereConfig Utility

This section contains the following topics:

Enable or Disable the VSPhereBackupNFC tag Using Command Prompt

Important! Enable the VSphereBackupNFC tag on specific kernel adapter for backup (instead of Specific Backup Network) when a Source VM hosts ESXi Hypervisor and the Backup uses VDDK7 or above versions.

Arcserve UDP allows you to enable or disable the *VSphereBackupNFC* tag using Command Prompt. The command-line tool (*AFVSphereConfigTool.exe*) lets you enable or disable the VSphereBackupNFC tag on a specific kernel adapter.

This section provides the information about how to enable or disable the *VSPhereBackupNFC* tag using command prompt.

Follow these steps:

1. Open command prompt and navigate to <Engine>\BIN>.

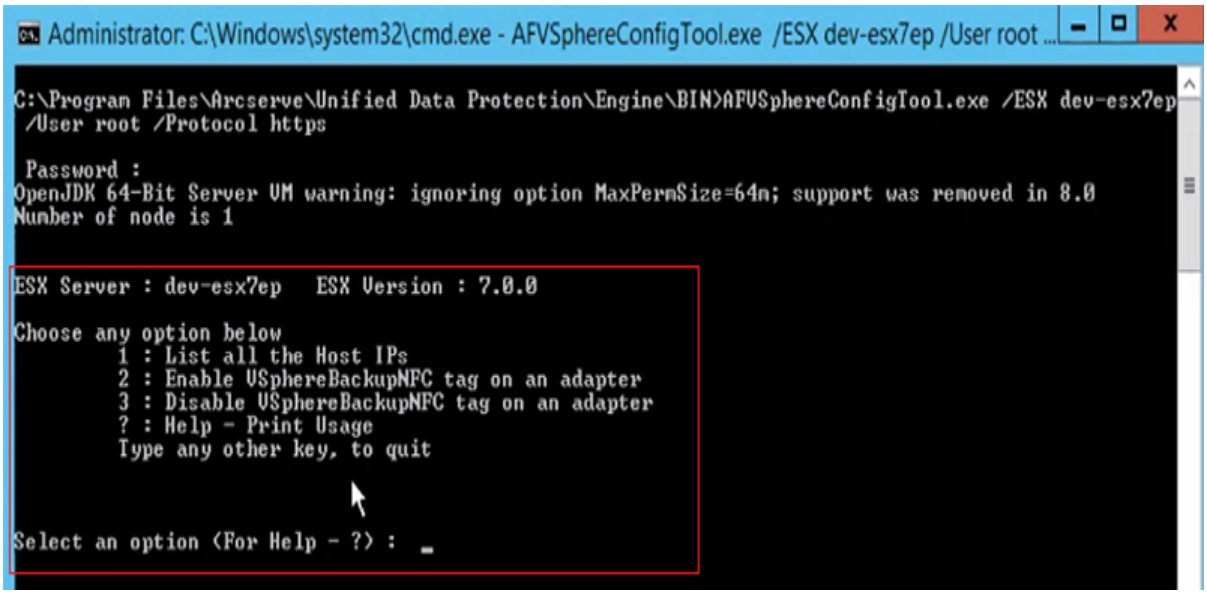
   Example: C:\Program Files\Arcserve\Unified Data Protection\Engine\BIN>

2. Connect to the ESX server using the command-line tool (*AFVSphereConfigTool.exe*), and then type a password when it prompts.

   AFVSphereConfigTool.exe /ESX dev-esx7ep /user root /protocol https

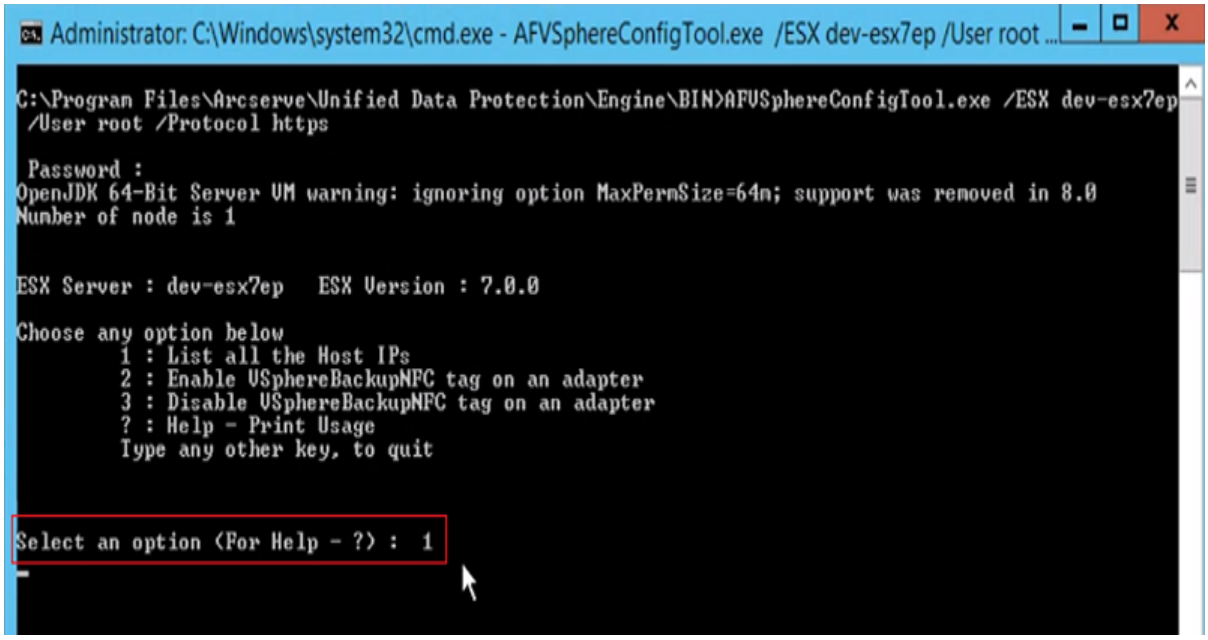


   The selected ESX server (dev-esx7ep) is now connected.

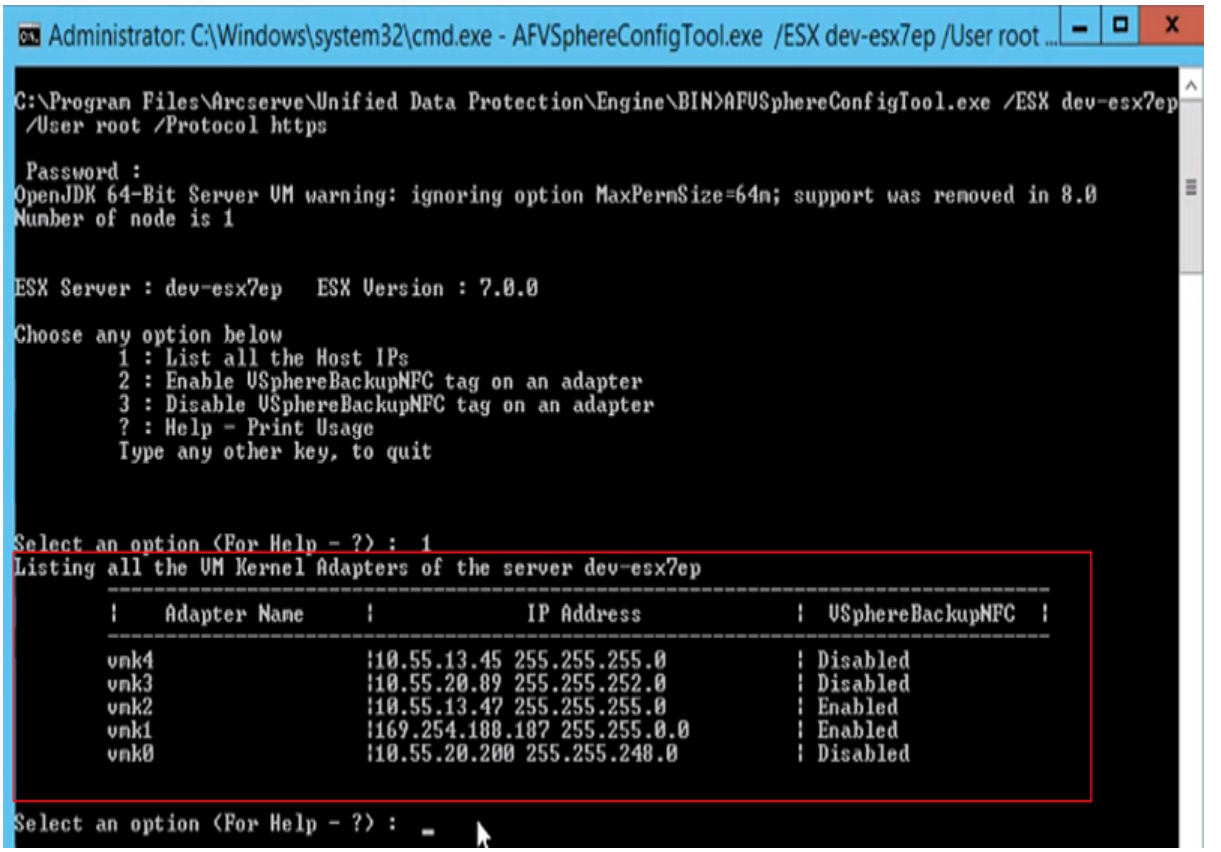


3. After connecting to the ESX server successfully, when you are prompted to select an option, do the following:

   - To list all the VM Kernel Adapters of the ESX server, type the option number. For example, type 1 for the option.

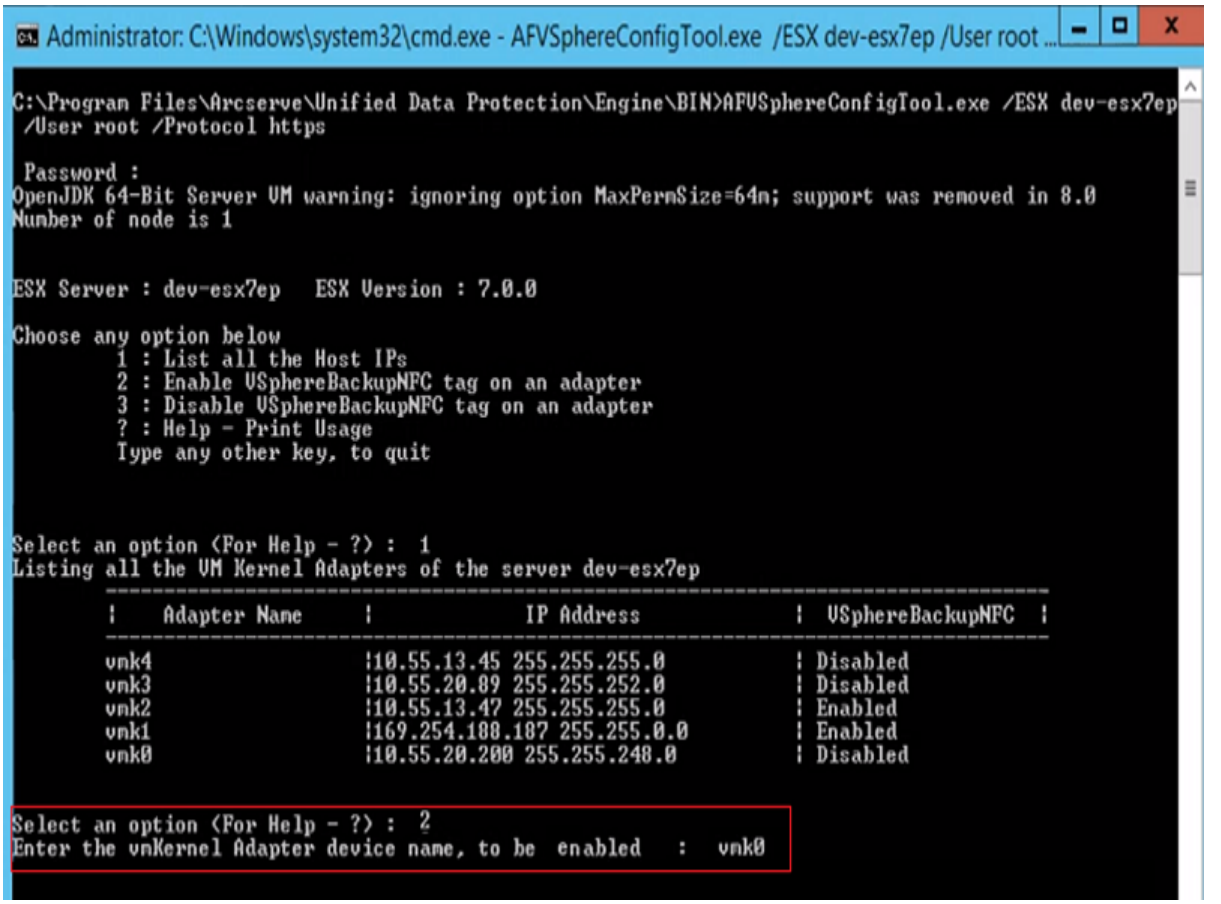


   A list of all VM Kernel Adapters is displayed.

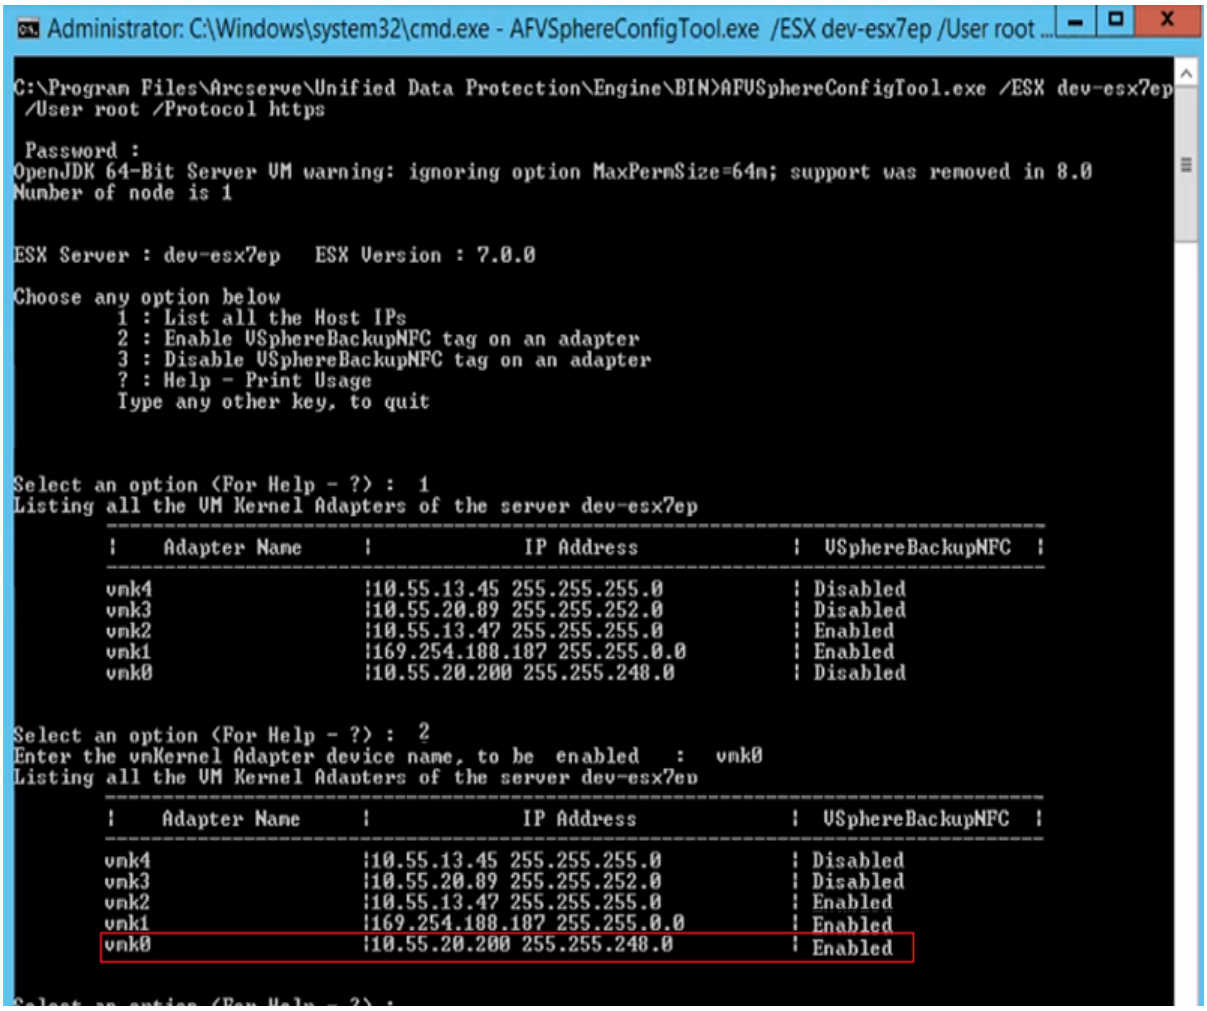


- To enable the VSphereBackupNFC tag on an adapter, type the option number. For example, type 2 for the option.

  When you are prompted to enter the name of the VM Kernel Adapter device, type the device name.

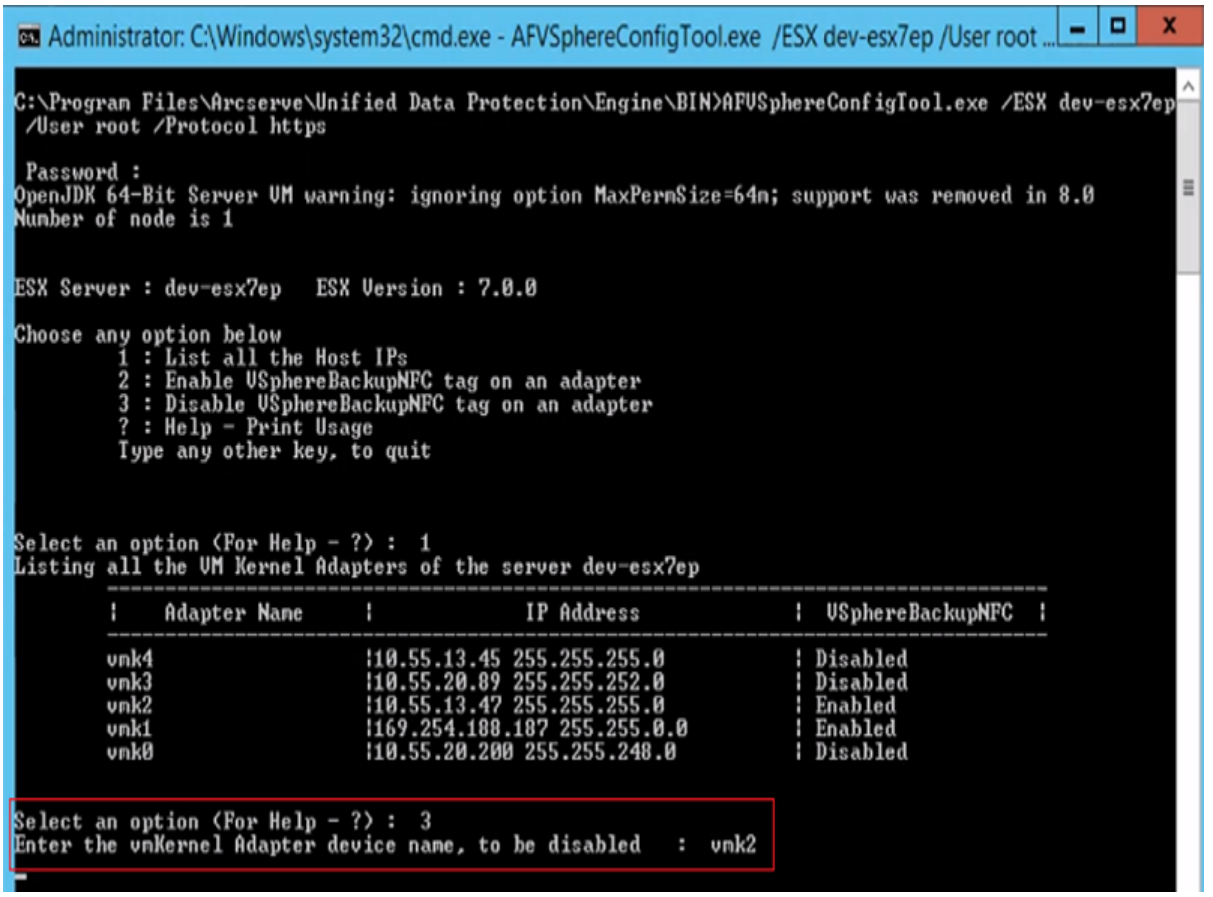


The VSphereBackupNFC tag gets enabled on the selected adapter.

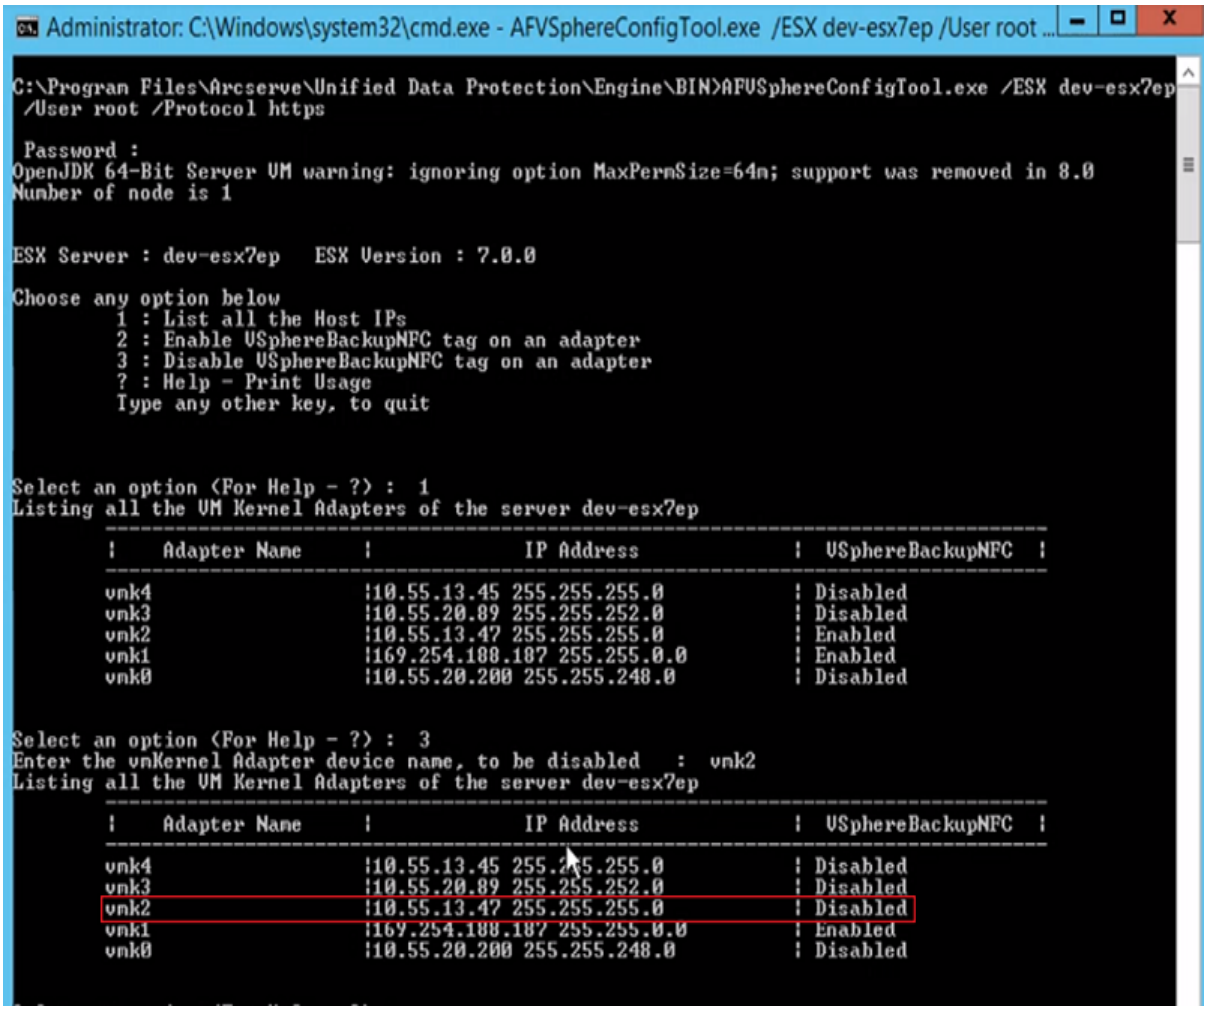


- To disable the VSphereBackupNFC tag on an adapter, type the option number. For example, type 3 for the option.

  When you are prompted to enter the name of the VM Kernel Adapter device, type the device name.



The VSphereBackupNFC tag gets disabled on the selected adapter.



4. Type any key (except 1, 2, 3, and ?) to close the command prompt.

Chapter 21: Using the Diagnostic Utility

The diagnostic utility lets you collect logs from your machines. When you contact the Arcserve Support team regarding any issue, the Support team uses the logs to investigate and fix the issue.

This section contains the following topics:

Collect the Diagnostic Information

Diagnostic information is a collection of logs, events, registry, and application information of the product and system that Arcserve Support team requires to investigate any error. Arcserve UDP lets you collect all such information at one location, typically a network share path. When you contact Arcserve Support, you have these information readily available. You can collect the diagnostic data for Windows, Linux, VMware, and Hyper-V machines.

Note: For Linux backup servers, the **Collect Diagnostic Data** option is available only from the **<Site_name> Nodes: Linux Backup Server Groups** view.

Follow these steps:

1. From the Console, click **resources**.
2. Follow one of these steps depending on the type of node:

   For Linux Backup Servers

   - From the left Navigation pane, navigate to **Nodes**, and click **Linux Backup Server Groups**.
   - From the center pane, select all the Linux nodes.

   For All Other Nodes and Servers

   - From the left Navigation pane, navigate to **Nodes** and click **All Nodes**.
   - From the center pane, select all the required nodes.
3. Click **Actions**, **Collect Diagnostic Data**.

   The **Collect Diagnostic Information** dialog opens.
4. (Optional) Select the check box.
5. Provide the network share path to store the data.

   Notes:

   - When collecting the diagnostic information for a remote site, you have to provide the destination as the gateway server or any other machine that can access the agent or RPS in that site.
   - If you want to specify the local path as the destination, then convert the local path to the UNC path and provide the UNC path. For example, **C:\test** can be specified as **\\<LocalmachineName>\C$\test**.
   - For host-based agentless backup (VMs), the Collect Diagnostic Data gathers data from the Arcserve UDP proxy server.

- Collect Diagnostic Data gathers data from the machines that have the Arcserve UDP agent installed.

6. Click **Submit**.

   A job is submitted to collect the data.

   On successful completion of the job, you will see the data in the shared folder. The name of the zip file is suffixed with the current timestamp.

Collect Diagnostic Information from a Standalone Agent

Diagnostic information is a collection of logs, events, registry, and application information of the product and system that Arcserve Support team requires to investigate any error. Arcserve UDP Agent lets you collect all such information at one location, typically a network share path. When you contact Arcserve Support, you have these information readily available.

Follow these steps:

1. Using the command prompt, navigate to the following path:

```
%ProgramFiles% \Arcserve\Unified Data Protection\Engine\BIN\DiagnosticUtility
```

2. Run the following command to learn how to run the batch file:

```
arcserveAgentSupport.bat -help
usage: arcserveAgentSupport.bat [OPTIONS]
-help print help
-pass <arg> usrPass (If export path is a remote share, user password to access it)
-path <arg> export path (Can be a remote share)
-user <arg> usrName (If export path is a remote share, user name to access it)
-xmlConfig <arg> xmlConfigurationFile (Optional)
```

3. Use the following command to collect the diagnostic information:

```
arcserveAgentSupport.bat -path <remote share path> -user <username> -pass <password>
```

Example: arcserveAgentSupport.bat -path \\remote_share\data -user abc -pass xyz

where, \\remote_share\data is the path, abc is the user name, and xyz is the password

You can find the diagnostic information zip file in the remote share.

Upload Diagnostic Information to Arcserve Website Using FTP

You can upload any logs or files to the Arcserve support FTP using a file transfer protocol (FTP). However, users (the ticket requester) can enable the FTP link during the initial ticket creation process or while updating the ticket online. Once enabled, FTP generates the FTP link with login credentials. This information is provided to the user through an automated email. You can use the user name and password sent through automated email to upload and download files from FTP.

When you log into FTP using ftp://supportftp.arcserve.com, the Home folder is displayed. If you create a ticket in Arcserve Support portal, a sub-folder is created with the ticket number within the Home folder. For example, if you create a ticket having the ticket number Ticket-00XXXX30, a sub-folder named Ticket-00XXXX30 is created in the Home folder.

ftp://supportftp.arcserve.com/Ticket-00XXXX30/

User Name and Password: Use auto generated username and password that is sent through automated email.

Folder: Ticket-00XXXX30

Note: For Japanese users, the password is the requester's email address without the domain name. For example, if registered email address is abc@yahoo.jp, the password is abc.

When the ticket is resolved, the FTP server receives a notification. The folder is then compressed and the original folder is deleted. The compressed folder is available for the next three months and later gets deleted permanently.

Important! Do not share your user name with others.

Follow these steps to upload any files to the Arcserve website using FTP:

1. Either the ticket requester (user) or Arcserve Support agent logs into Arcserve Support portal and creates a support ticket.

   The Arcserve Support agent or the ticket requester selects the check box in the support ticket and updates the ticket.

   An FTP link is automatically generated for the FTP Home folder. The permission is set exclusively for the requester who opened the ticket.

   For example, the following folder is your FTP folder:

   ftp://supportftp.arcserve.com/<Ticket_number>

2. Once the FTP link is created, Arcserve Support sends an email with the FTP link and the login credentials to the ticket requester (user) automatically.
3. The user logs into the FTP link and uploads the files.

   You have successfully completed the process of uploading files to Arcserve FTP server.

Unzip Agent Logs

The log files generated by the diagnostic utility are in a ZIP format. To view the log files, you must unzip the files.

Follow these steps:

1. Copy the .arcZIP file to the machine that has the UDP agent installed.

   The .arcZIP file is created using the diagnostic utility.

2. Using the command prompt, navigate to the following path:

   ```
   %ProgramFiles% \Arcserve\Unified Data Protection\Engine\BIN\DiagnosticUtility
   ```

3. Run the following command to learn how to run the batch file:

   ```
   arcserveAgentSupportInternal.bat -help
   BaseOperation loadDefaultValue
   INFO: Load Agent install path C:\Program Files\Arcserve\Unified Data Protection\Engine\
   usage: arcserveAgentSupportInternal.bat [OPTIONS] rawfile
   -help          print help
   -keepFile      Keep temp file
   -path <arg>    path where the content needs to be unzipped
   ```

   The help section is displayed.

4. Use the following command to unzip the file:

   ```
   arcserveAgentSupportInternal.bat -path <should_be_the_same_machine_where_you_want_to_unzip> <name_of_the_zip_file>
   ```

   The agent logs are unzipped.

Unzip the Console Logs

The log files generated by the diagnostic utility are in a ZIP format. To view the log files, you must unzip the files.

Follow these steps:

1. Copy the .arcZIP file to the machine that has the UDP Console installed.

   The .arcZIP file is created using the diagnostic utility.

2. Using the command prompt, navigate to the following path:

   ```
   %ProgramFiles% \Arcserve\Unified Data Protection\Management\BIN\DiagnosticUtility
   ```

3. Run the following command to learn how to run the batch file:

   ```
   arcserveConsoleSupportInternal.bat -help
   BaseOperation loadDefaultValue
   INFO: Load Agent install path C:\Program Files\Arcserve\Unified Data Protection\Management \
   usage: arcserveConsoleSupportInternal.bat [OPTIONS] rawfile
   -help              print help
   -ignoreFailed      Ignore failed import table
   -keepFile          Keep temp file
   -noClean           Do not clean DB
   -path <arg>        path where the content needs to be unzipped
   -u                 Only unzip file
   ```

   The help section is displayed.

4. Use the following command to unzip the file.

   ```
   arcserveConsoleSupportInternal.bat -path <should_be_the_same_machine_where_you_want_to_unzip> <name_of_the_zip_file>
   ```

   Note: The command does not overwrite the Console database. If you want to overwrite the Console database, edit the *arcserveConsoleSupportInternal.bat* command and remove "*-u*" from the file, and then save the file.

   The Console logs are unzipped.

Skip Log History Folder Contents Using Registry

You can use the UDP Agent to collect all logs, which includes the log history, from the Engine. Sometimes, while collecting the log history, the UDP Agent processes the files in a slow pace and causes the file size to appear bigger. To avoid these discrepancies, you can exclude the LogHistory folder.

Follow these steps:

1. Add the key in registry at the location:

   HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine

2. Enter following details:

   KeyName: *SkipLogHistory*

   type: *DWORD(32-bit)*

   Value : 1 (value 1 is used to skip the LogHistory and 0 is used to collect LogHistory)

Collect Log from Gateway Machine Using Command Line

You must have Arcserve UDP gateway product installed to collect logs from command line.

Follow these steps to collect logs from command line:

1. Open command prompt and navigate to *<Gateway installation directory>\BIN\Diagnostic Utility*\.

   Example: C:\Program Files\Arcserve\Unified Data Protection\Gateway\BIN\Diagnostic Utility\

2. To zip the logs, run the below command:

   arcserveGatewaySupport.bat -Path "<Destination path>"

3. To unzip the logs, run the below command:

   arcserveGatewayInternalSupport.bat -Path "<Destination path>" "<source path>"

Collect Logs from RPS using Command Line

You must have Arcserve UDP product installed with RPS component to collect RPS logs from command line.

Follow these steps to collect logs from command line:

1. Open command prompt and navigate to *<Agent installation directory>\BIN\Diagnostic Utility*\.

   Example: C:\Program Files\Arcserve\Unified Data ProtectionEngine\BIN\Diagnostic Utility\

2. To zip the logs, run the below command:

   arcserveAgentSupport.bat -path <remote share path> -user <username> -pass <password>

3. To unzip the logs, run the below command:

   arcserveAgentSupportInternal.bat –path <destination> <source(rawfile)>

Notes:

- RPS logs have the job logs which are performed on RPS.
- RPS logs can only be collected from command line and not from console UI.
- Some RPS related job logs cannot be collected from console UI, when collecting VM logs from Console.

Example: When an MSP Replication task is configured, you configure the tasks with two consoles, where Console-1 will have Backup as primary task and Replication as secondary task. In Console-2, Remotely managed Replication task is configured and shared with Console-1.

In this scenario, In Console-2 you cannot collect Replication job logs when collecting VM logs from console. You need to collect those Replication logs from the command line in the RPS machine that is the destination of Replication Job.

Collect Stub logs From Hyper-V Having Agent Installed

You can collect stub logs from Hyper-V having agent installed. Hyper-V can act as proxy also to help collect Hyper-V Stub logs.

When Hyper-V itself acts as a proxy (Hyper-V machine must have agent installed)

Open Command Prompt and navigate to *<Agent Installation Directory>\BIN\ Diagnostic Utility*\.

Example: *C:\Program Files\Arcserve\Unified Data Protection\Engine\BIN\Diagnostic Utility*

- To zip the logs: Run *arcserveAgentSupport.bat -Path "Destination"*
- To unzip the logs: Run *ArcserveAgentInternalSupport.bat -Path "Destination" "source"*

When Hyper-V does not act as a proxy (Collect Entire UDP Host based VM Backup folder)

Prerequisite: The machine must have Java installed.

Follow these steps:

1. Import registry from the below path on the machine that has agent installed:

   HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine.

2. Copy the entire DiagnosticUtility folder from a machine that has agent installed already and Paste the folder at the following path:

   C:\Program Files\Arcserve\Unified Data Protection\Engine\BIN

   Note: You need to manually create this directory structure.

3. Create an empty folder named "Logs" under the following path:

   C:\Program Files\Arcserve\Unified Data Protection\Engine

4. Copy Common folder from a machine that has agent installed already to the following path:

   C:\Program Files\Arcserve\Unified Data Protection

   - To zip the logs: Run *arcserveAgentSupport.bat -Path "Destination"*
   - To unzip the logs: Run *ArcserveAgentInternalSupport.bat -Path "Destination" "source"*

Collect Hyper-V Event Viewer Messages

You can collect Event Viewer Messages from Hyper-V having agent installed with a folder name *hyperVEventViewerFiles*.

Follow these steps to collect Event Viewer Messages from command line:

1. Open command prompt and navigate to <*Agent Installation Directory*>*BIN**Diagnostic Utility*\.

   Example: C:\Program Files\Arcserve\Unified Data Protection\Engine\BIN\Diagnostic Utility\

2. To zip the Event Viewer Messages, run *arcserveAgentSupport.bat -Path "Destination"*
3. To unzip the Event Viewer Messages, run *ArcserveAgentInternalSupport.bat -Path "Destination" "source"*

Collect Ca_lic Folder and OLF Files

CA_LIC and OLF files are collected in the logs, when you collect from the machine that has agent installed already, with a folder name CA_LIC.

- To collect using CLI: Open Command Prompt and navigate to *<Agent Installation Directory>\BIN\ Diagnostic Utility*

  Example: *C:\Program Files (x86)\Arcserve\SharedComponents\CA_LIC*
- To collect using Console: Navigate to Agent node, right click and click **Collect Diagnostic Information**.
- To zip the logs: Run *arcserveAgentSupport.bat -Path "Destination"*
- To unzip the logs: Run *ArcserveAgentInternalSupport.bat -Path "Destination" "source"*

Collect Directory Listing of the Backup Destination/Datastore Directories

Directory listing for Engine and Management Folders

Engine: This log is collected under *agentLogs* folder (after unzipping arczip) with the name *EngineDirectoryBrowseInfo.log*.

Management: This log is collected under *consoleLogs* folder (after unzipping arczip) with the name *ManagementDirectoryBrowseInfo.log*.

Directory Listing for DataStores :

- **Non-Dedupe:** Log is collected for non-dedupe datastore with name *CommonStorePathDirectoryBrowseInfo.log* (after unzipping arczip) in the folder *agentLogs*
- **Dedupe:** Four logs are collected for Dedupe datastores (after unzipping arczip) in the folder *agentLogs* with the following names:
    1. *CommonStorePathDirectoryBrowseInfo.log*
    2. *HashRolePathDirectoryBrowseInfo.log*

3. *IndexRolePathDirectoryBrowseInfo.log*
4. *DataRolePathDirectoryBrowseInfo.log*

Chapter 22: Manually Installing Sophos Intercept X Advanced for Server on Arcserve UDP

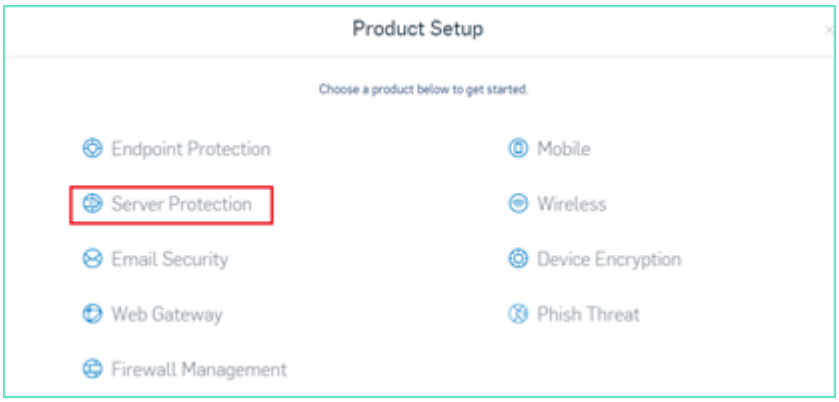
The installation of Sophos Intercept X Advanced for Server on Arcserve UDP enables the following:

- Protect data and system backups from ransomware and other attacks
- Endpoint protection that combines signature-based and signatureless malware detection.
- Deep learning neural network
- Anti-exploit technology
- CyptoGuard anti-ransomware and WipeGuard technologies, and more to stop the widest range of endpoint threats

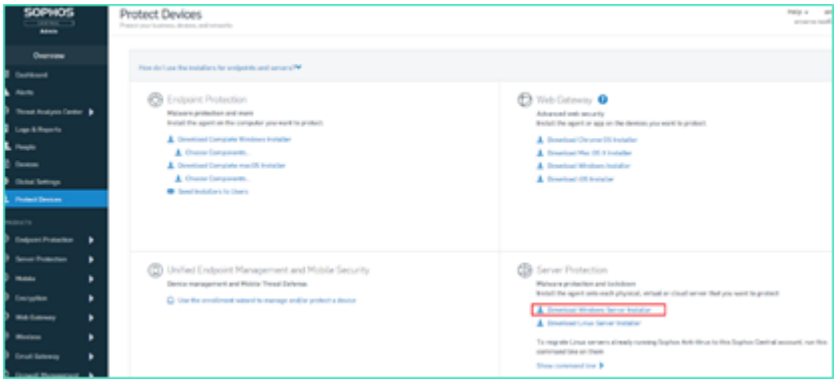
Follow these steps:

1. On the Arcserve Support Website, create an account.
2. To request for a free copy of Sophos Intercept X Advanced, provide all relevant details in the Sophos Request form and submit it to Arcserve Support. It is mandatory to share the details of Email ID and Order ID. You will receive an auto-generated email confirmation.

   After you confirm your email ID, Arcserve processes your request and creates an account on the Sophos Central and sends an email with instructions on how to create a password.
3. To create a password for your new account on Sophos Central, follow the instructions in the email.
4. Log into Sophos Central.
5. Open the Product Setup dialog, then select **Server Protection**.



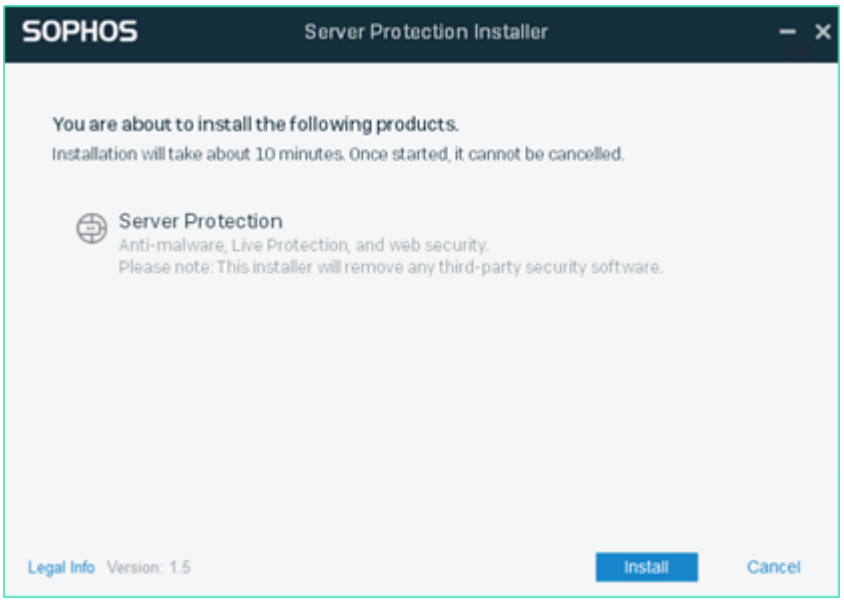
6. From the Server Protection section, click **Download Windows Server Installer**, and then save **SophosSetup.exe** installer to a folder on UDP.



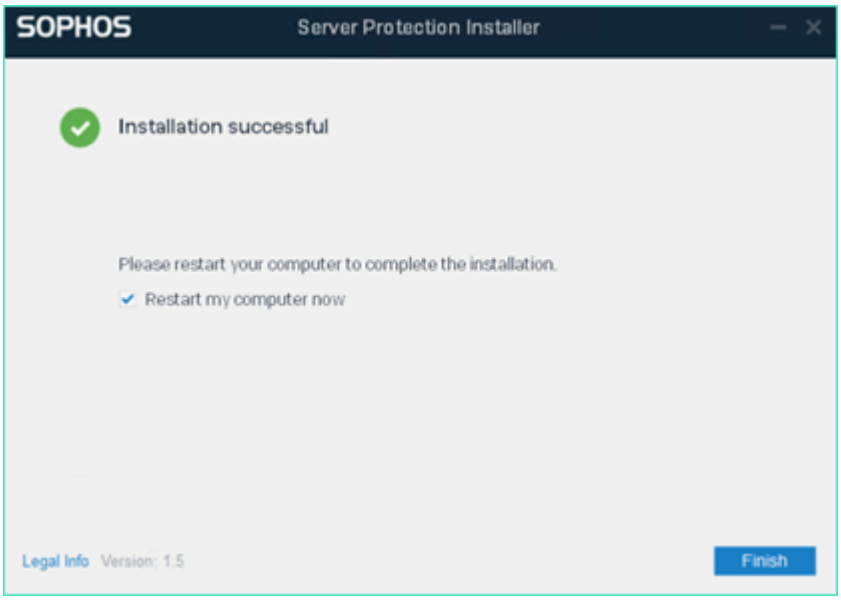
7. To start the installer, open the folder, and then double-click **SophosSetup.exe**.

   Note: To avoid unexpected behavior while installing the Sophos Intercept-X, disable the Windows Defender and\or other antivirus software from the server. Sophos will remove some non-Sophos security products from the Windows servers. For more information, see Sophos CRT: List of third-party security software removed by Sophos.

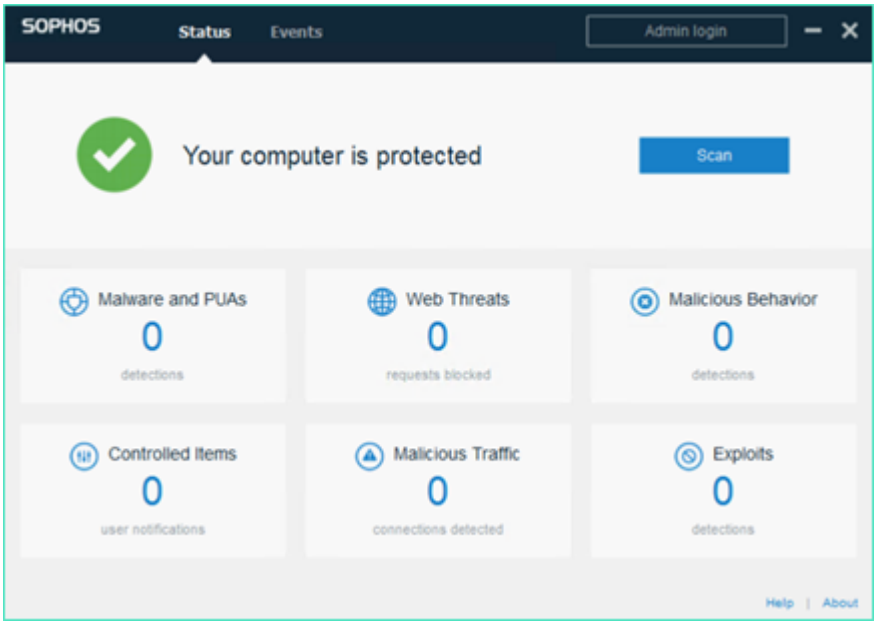
8. Click **Install**.



9. To restart your system immediately, click **Finish**. To restart later, uncheck the **Restart my computer now** option.



10. To view the protection status, open the **Sophos Intercept X** interface.



The status indicates that Arcserve UDP 8.0 is secured from ransomware attacks, malware, web threats, and zero-day exploits.

11. To access Sophos Central, click **Admin Login**. It allows you to manage Sophos Intercept X Advanced Server, set alerts and policies, and so on.

Notes:

- It is mandatory to have Internet access in Arcserve UDP to install *Sophos Intercept X Advanced* and any related updates. Sophos Intercept X Advanced is cloud based and there are no offline installers available.

- If you have already purchased another UDP previously and have a Sophos account through Arcserve, use the same account for all your Arcserve UDP 8.0.
- If you already have a Sophos account through any other purchase, such as directly from Sophos, provide a different email address for a separate account on Sophos Central.
- If the Sophos installation fails for any reason, follow the on-screen or email instructions that are provided along with the error message.
- To receive the Sophos Intercept X Advanced for Server updates such as malware definition updates and version upgrades, you must have a valid and active maintenance or subscription for your Arcserve UDP.

For further assistance, please contact Arcserve Technical Support on phone (+1.844.765.7043) or online, or contact your local Arcserve support office.

Chapter 23: Troubleshooting

This section contains the following topics:

Arcserve UDP Communication Failure Related

This section includes the following troubleshooting topics related to Arcserve UDP communication failure:

- Arcserve UDP Cannot Communicate with Windows Nodes
- Unable to Receive E-mail Alerts from Gmail Account
- Arcserve UDP Cannot Communicate with the Arcserve UDP Linux Backup Server on Remote Nodes
- Arcserve UDP Cannot Communicate with the Arcserve UDP Recovery Point Server on Remote Nodes
- Arcserve UDP Cannot Communicate with the Arcserve Backup Server on Remote Nodes
- Arcserve UDP Cannot Communicate with the Remote Site

Arcserve UDP Cannot Communicate with Windows Nodes

Valid on Windows operating systems.

Symptom

Arcserve UDP cannot communicate with Windows nodes.

Solution

The following table describe reasons why Arcserve UDP cannot communicate with Windows nodes and the corresponding corrective action:

Cause	Corrective Action
The network was unavailable or unstable when applying plans.	Verify that the network is available and stable, and then try again. Arcserve UDP can ping the remote node and remote node can ping back Arcserve UDP.
The network Admin$ share of the remote node was not available when Arcserve UDP tried to communicate with the node.	Verify the network Admin$ of the remote node is available and then try again.
The Arcserve Unified Data Protection Agent for Windows node could not handle the load when Arcserve UDP tried to communicate with the node.	Verify that the CPU on the remote Arcserve Unified Data Protection Agent for Windows node is in a normal state and then try again.
The Arcserve Unified Data Protection Agent for Windows service on the remote node was not running when Arcserve UDP tried to communicate with the node.	Verify that the Arcserve Unified Data Protection Agent for Windows service on the remote node is running and then try again.
Used a wrong protocol or port to communicate with the node.	Use the right protocol or port to add/update the remote node in the Arcserve UDP node view.
The Arcserve Unified Data Protection Agent for Windows service was not communicating properly.	Restart the Arcserve Unified Data Protection Agent for Windows service on the remote node and then try again.

Unable to Receive E-mail Alerts from Gmail Account

Gmail account is blocked by Google if you use Gmail account in your email settings.

Valid on Windows platforms.

Symptom

When you configure Gmail accounts, you do not receive email alerts. When you try to configure the Gmail account for email alerts and click “Send A Test Email”, you receive one of the following error messages:

The test mail failed due to incorrect user credentials.

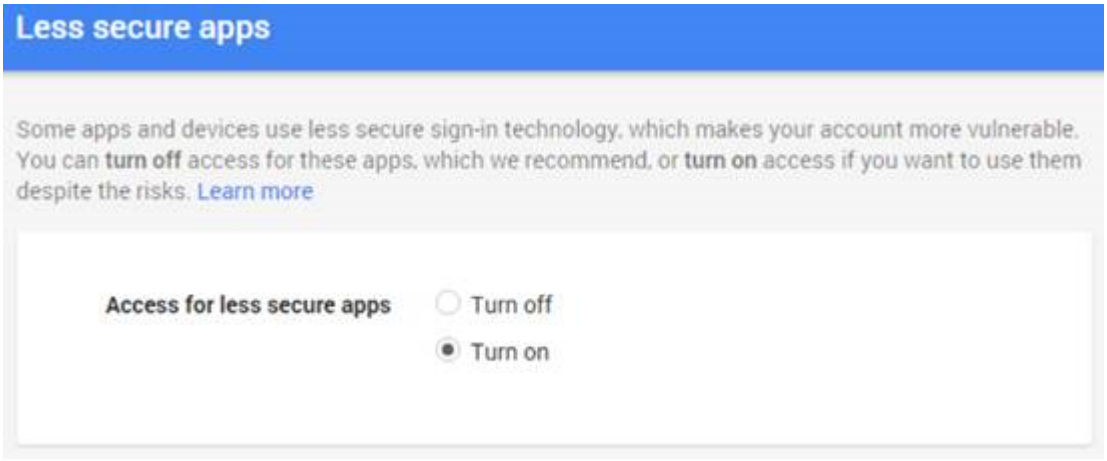
or

Failed to send test email: Invalid user credentials.

Solution

1. Provide the correct credentials and try again.
2. Google security blocks Emails from the Gmail account configured outside Google. To avoid, in the link below modify settings by selecting “Turn ON” to use Access for less secure apps:

   https://www.google.com/settings/security/lesssecureapps



Arcserve UDP Cannot Communicate with the Arcserve UDP Linux Backup Server on Remote Nodes

Valid on Linux operating systems.

Symptom

Arcserve UDP cannot communicate with the Arcserve UDP Linux Backup Server on remote nodes.

Solution

The following table describe reasons why Arcserve UDP cannot communicate with the Arcserve UDP Linux Backup Server on remote nodes and the corresponding corrective action:

Cause	Corrective Action
The network was unavailable or unstable when Arcserve UDP tried to communicate with the Linux Backup Server node.	Verify that the network is available and stable, and then try again. Arcserve UDP can ping the remote Linux Backup Server node and the remote Linux Backup Server node can ping back Arcserve UDP.
The Arcserve UDP Linux Backup Server node could not handle the load when Arcserve UDP tried to communicate with the node.	Verify that the CPU on the remote Arcserve UDP Linux Backup Server node is in a normal state and then try again.
The Arcserve UDP Linux Backup Server service on the remote node was not running when Arcserve UDP tried to communicate with the node.	Verify that the Arcserve UDP Linux Backup Server service on the remote node is running and then try again.
The Arcserve UDP Linux Backup Server service was not communicating properly.	Restart the Arcserve UDP Linux Backup Server service on the remote node and then try again.

Arcserve UDP Cannot Communicate with the Arcserve UDP Recovery Point Server on Remote Nodes

Valid on Windows operating systems.

Symptom

Arcserve UDP cannot communicate with the Arcserve UDP Recovery Point Server on remote nodes.

Solution

The following table describes reasons and the corresponding corrective action:

Cause	Corrective Action
The network was unavailable or unstable when Arcserve UDP tried to communicate with the Recovery Point Server node.	Verify that the network is available and stable, and then try again. Arcserve UDP can ping the remote Recovery Point Server node and the remote Recovery Point Server node can ping back Arcserve UDP.
The network Admin$ share of Arcserve UDP Recovery Point Server node was not available when Arcserve UDP tried to communicate with the node.	Verify the network Admin$ of Recovery Point Server node is available and then try again.
The Arcserve UDP Recovery Point Server node could not handle the load when Arcserve UDP tried to communicate with the server.	Verify that the CPU on the remote Recovery Point Server node is in a normal state and then try again.
The Arcserve UDP Agent Service, Arcserve UDP RPS Data Store Service, or Arcserve UDP RPS Port Sharing Service on the remote node were not running when Arcserve UDP tried to communicate with the node.	Verify that the Arcserve UDP Agent Service, Arcserve UDP RPS Data Store Service, and Arcserve UDP RPS Port Sharing Service on the remote node is running and then try again.
Used a wrong protocol or port to communicate with the Recovery Point Server node.	Use the right protocol or port to add/update the Recovery Point Server node in the Arcserve UDP destination node view.
The Arcserve UDP Agent Service, Arcserve UDP RPS Data Store Service, or Arcserve UDP RPS Port Sharing Service was not communicating properly.	Restart the Arcserve UDP Agent Service, Arcserve UDP RPS Data Store Service, and Arcserve UDP RPS Port Sharing Service on the remote node and then try again.

Arcserve UDP Cannot Communicate with the Arcserve Backup Server on Remote Nodes

Valid on Windows operating systems.

Symptom

Arcserve UDP cannot communicate with the Arcserve Backup Server on remote nodes.

Solution

The following table describe reasons why Arcserve UDP cannot communicate with the Arcserve Backup Server on remote nodes and the corresponding corrective action:

Cause	Corrective Action
The network was unavailable or unstable when Arcserve UDP tried to communicate with the Arcserve Backup Server node.	Verify that the network is available and stable, and then try again. Arcserve UDP can ping the remote Arcserve Backup Server node and the remote Arcserve Backup Server node can ping back Arcserve UDP.
The Arcserve Backup Server node could not handle the load when Arcserve UDP tried to communicate with the node.	Verify that the CPU on the remote Arcserve Backup Server node is in a normal state and then try again.
The Arcserve Backup Server related service on the remote node was not running when Arcserve UDP tried to communicate with the node.	Verify that the Arcserve Backup Server on the remote node is running and then try again.
Used a wrong protocol or port to communicate with the Arcserve Backup Server node.	Use the right protocol or port to add/update the Arcserve Backup Server in the Arcserve UDP destination node view.
The Arcserve Backup Server related service was not communicating properly.	Restart the Arcserve Backup Server related service on the remote node and then try again.

Arcserve UDP Cannot Communicate with the Remote Site

Valid on Windows operating systems.

Symptom

Arcserve UDP cannot communicate with the remote site.

Solution

The following table describe reasons why Arcserve UDP cannot communicate with the remote site and the corresponding corrective action:

Cause	Corrective Action
The network is unavailable or unstable.	Verify that the network is available and stable, and then try again.
Arcserve UDP is reinstalled and the remote site is not registered to Arcserve UDP.	Register the remote site to Arcserve UDP.
The host name or IP address of Arcserve UDP has changed and the remote site is not registered to Arcserve UDP.	Register the remote site to Arcserve UDP.
The host name or IP address of the remote site has changed and the remote site is not registered to Arcserve UDP.	Register the remote site to Arcserve UDP.

Plan, Job, and Settings Related

This section includes the following troubleshooting topics related to the backup job and settings:

- Backup Job Failure after Changing the Console Hostname/IP Address
- How to Add Encryption Password for an Existing Encrypted Destination
- Unable to Apply Backup Settings to Node
- Plan Deployment Fails after Changing Password of Agentless Backup Proxy Machine
- Settings Disabled when Opening Agent UI
- Pause and Resume Fails when Agent Is Not Connected to the Network
- The Arcserve UDP Agent Service Runs Slowly
- Configure the Registry to Rerun a Copy to Tape Job
- Configure the Registry to Copy Multiple Recovery Points of Same Type to Tape in the Same Job
- File/Folder Missed from Backup of NFS Shared Folder or File/Folder Name is Converted to Junk
- NFS Shared Folder Backup Fails
- Incorrect Server Name on License Request for UNC Paths

Backup Job Failure after Changing the Console Hostname/IP Address

Symptom

I installed the console and RPS server on the same machine. The backup was working fine but after I changed the hostname/IP address of the console, the backup job fails.

Solution

This problem occurs when you have plans assigned to nodes and then you modify the hostname/IP address of this machine.

To resolve this issue, manually update agent nodes and run the backup job again.

Follow these steps:

1. Navigate to Nodes: All Nodes page.
2. Select the node.
3. Right-click and select **Update**.
4. Click **OK**.

The nodes are updated.

How to Add Encryption Password for an Existing Encrypted Destination

Symptom

You forgot to add the encryption password for the File Copy destination.

Solution

You can add the encryption password now.

Follow these steps:

1. Open the Plan.
2. Open file copy destination in which you need to add the encryption password.
3. Change the destination type from Cloud Storage\Network Share to Network Share\Cloud Storage.
4. Provide any Network Share or Cloud Storage and save.
5. Again open the plan and go to the File Copy destination.
6. Change the destination to Cloud Vendor\Network Share.
7. Select the Cloud Vendor\Network Share and then select the bucket or Container\Provide path.
8. Provide the correct encryption password.
9. Save the plan.

Unable to Apply Backup Settings to Node

Symptom

I have two consoles, Console A and Console B. I add a recovery point server (RPS) to Console A and create a plan for the RPS. Then I add the RPS to Console B. Now this RPS is managed by Console B. But when i update the agent node from Console A that is backed up to the RPS, I get the following error:

Unable to apply 'backup settings' to node. (Failed to find the Arcserve UDP Recovery Point Server plan on this server.)

Solution

To avoid this error, follow these steps:

1. Select the plan from Console A,
2. From the center pane, click **Actions**, and then select **Deploy Now**.

   The plan gets redeployed and the backup settings are applied to the node.

Plan Deployment Fails after Changing Password of Agentless Backup Proxy Machine

Symptom

If the Console and the proxy server are different machines, then after changing the password of the proxy machine, the plan redeployment fails. The error message says credentials are not correct.

Solution

Follow these steps to resolve the issue:

1. From the node view in Console, update the proxy server with new credentials.
    a. From the left pane, navigate to **Nodes** and click **All Nodes**.
    b. Right-click the node and select **Update** to update the proxy server.
2. If the RPS and proxy server are the same machine, update the RPS with new credentials.
    a. From the left pane, navigate to **Destinations** and click **Recovery Point Servers**.
    b. Right-click the RPS from the center pane and select **Update**.
3. Restart Arcserve UDP Agent Service in the proxy machine.
4. Redeploy the plan.

Settings Disabled when Opening Agent UI

If Arcserve UDP Agent (Windows) nodes are not removed from the Arcserve UDP UI before uninstalling the Arcserve UDP console, the settings will be disabled when opening the agent UI on those Arcserve UDP Agent (Windows) nodes.

Symptom

The Arcserve UDP Agent (Windows) node is not notified that the Arcserve UDP Console is uninstalled. It assumes it is managed.

Solution

Remove the files "RegConfigPM.xml" and "BackupConfiguration.xml" under "<UDP_ENGINE_HOME>\Configuration" directory on the Arcserve UDP Agent (Windows) node, and then restart the Windows service "Arcserve UDP Agent Service".

If the problem persists, use Live Chat to contact Arcserve support. Live Chat lets you optimize intelligent conversation between you and the Technical Support team, allowing you to address your concerns and questions immediately, while still maintaining access to the product.

Pause and Resume Fails when Agent Is Not Connected to the Network

Symptom

If the agent is not connected to the network and I try to pause a plan, the plan is not paused. Similarly, if the agent is not connected to the network and I try to resume a plan, the plan is not resumed.

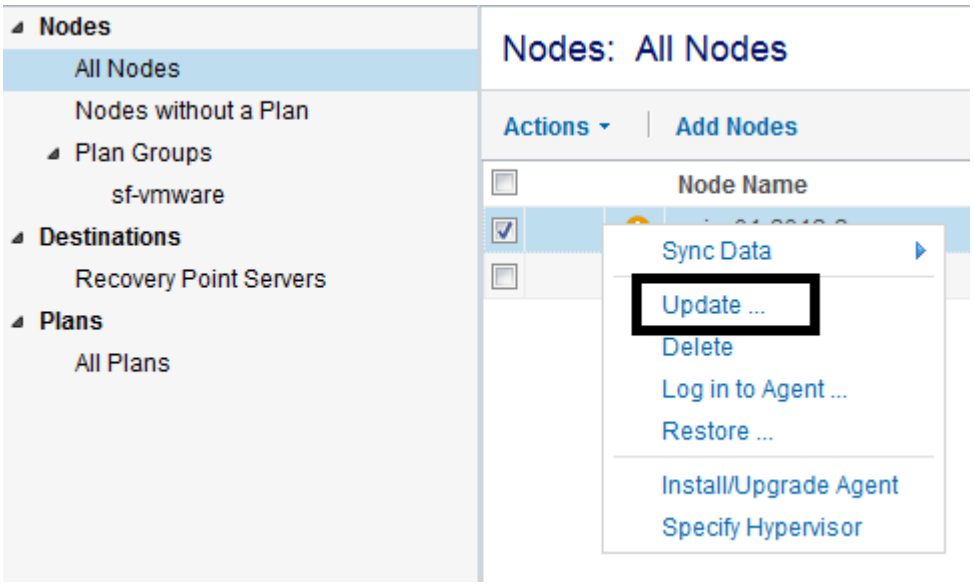
Solution

You can resolve this issue by manually updating the node from the Console.

Follow these steps:

1. Click the resources tab on the Console.
2. From the left pane, navigate to **Nodes**, and click All Nodes.

   All the added nodes are displayed on the center pane.
3. On the center pane, select the node.
4. Right-click and select **Update**.



The node is updated and the plan is refreshed.

The Arcserve UDP Agent Service Runs Slowly

Valid on Windows operating systems.

Symptom 1:

The Arcserve UDP Agent Service on Arcserve UDP Agent systems runs slowly. You can detect other symptoms such as:

- The Arcserve UDP Agent Service stops responding or occupies 100 percent of the CPU resources.
- Arcserve UDP Agent nodes perform poorly or cannot communicate with the web service.

Solution 1:

In various environmental configurations, you can discover that the Arcserve UDP Agent Service occupies too much CPU time, or the response is slow. By default, Tomcat is configured to allocate a limited amount of memory to the nodes, which may not be suitable for your environment. To verify this problem, review the following log files:

<D2D_home>\TOMCAT\logs\casad2dwebsvc-stdout.*.log
<D2D_home>\TOMCAT\logs\casad2dwebsvc-stder.*.log
<D2D_home>\TOMCAT\logs\catalina.*.log
<D2D_home>\TOMCAT\logs\localhost.*.log

Search for the following message:

java.lang.OutOfMemoryError

To correct this problem, increase the amount of allocated memory.

To increase the memory, perform the following steps:

1. Open Registry Editor and access the following key:

   x86 Operating Systems:

   - HKEY_LOCAL_MACHINE\SOFTWARE\Apache Software Foundation\Procrun 2.0\CASAD2DWebSvc\Parameters\Java

   x64 Operating Systems:

   - HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Apache Software Foundation\Procrun 2.0\CASAD2DWebSvc\Parameters\Java

2. Use one of the following options:

- If the message in the log file is the following:

  java.lang.OutOfMemoryError: PermGen space

  Append the following to the value of Options.

  -XX:PermSize=128M -XX:MaxPermSize=128M

  Note: You may need to increase the value of -XX:MaxPermSize to suit your environment.

- If the message in the log file is one of the following:

  java.lang.OutOfMemoryError: Java heap space
  java.lang.OutOfMemoryError: GC overhead limit exceeded

  Increase the value of the following DWORD:

  JvmMx

3. Restart the Arcserve UDP Agent Service.

Symptom 2

Scheduled backups are skipped and stop running.

Solution 2

When you configure the MAX value as 20 or less than 20 for concurrent backups, perform the following steps:

1. Increase the value of the following DWORD:

   JvmMx=256

   Note: This DWORD is referenced in Solution 1.

2. Append the following to the value of Options.

   -XX:MaxPermSize=128M

   Note: This DWORD is referenced in Solution 1.

When you configure the MAX value as more than 20 but less than 50 for concurrent backups, perform the following steps:

1. Increase the value of the following DWORD:

   JvmMx=512

   Note: This DWORD is referenced in Solution 1.

2. Append the following to the value of Options.

   -XX:MaxPermSize=256M

   Note: This DWORD is referenced in Solution 1.

Configure the Registry to Rerun a Copy to Tape Job

Symptom

The Copy to Tape job did not run for some media error and you want to rerun the job.

Solution

You can control the number of retry jobs and the time interval of retry jobs for the Copy to Tape task using the following two registry keys. Both the registry keys are in the machine where you have installed Arcserve Backup Server:

NumberOfRetryCopyToTapeJob

If one Copy to Tape job fails, Arcserve UDP retries the failed job. The retry number is configured using the **NumberOfRetryCopyToTapeJob** registry key. But if the to-copy node or recovery point information is changed in the job, Arcserve UDP resets the accumulated failed number as 0. It means once the to-copy node or recovery point information is changed in the job, the Copy to Tape job can be run without being limited by the retry number. Also, if Arcserve Backup web service is restarted, Arcserve UDP resets the accumulated failed number as 0.

The registry key is located in the Arcserve Backup Server at the following location:

HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\ComputerAssociates\CA ARCServe Backup\WebServiceInfo\NumberOfRetryCopyToTapeJob

Type is DWORD.

Default: 1

TimeIntervalOfRetryCopyToTapeJob

Controls the time interval of retry for the failed Copy to Tape job. This registry key is used in association with **NumberOfRetryCopyToTapeJob**.

The registry key is located in the Arcserve Backup Server at the following location:

HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\ComputerAssociates\CA ARCServe Backup\WebServiceInfo\TimeIntervalOfRetryCopyToTapeJob

Type is DWORD.

Default: 1 hour

Configure the Registry to Copy Multiple Recovery Points of Same Type to Tape in the Same Job

Symptom

By default, copy to tape job copies all of the qualified recovery points to tape though they are of the same type. As a result more tape space is consumed and takes long time to copy. You may want to copy only the latest one of the qualified recovery points.

Solution

You can configure registry on the Arcserve Backup Server node to enable copying the latest recovery point of same type to tape.

Follow these steps:

1. On Arcserve Backup Server node, add the below DWORD value to registry that is located under [HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\ComputerAssociates\CA ARCServe Backup\WebServiceInfo]:

   "CopyLatestRecoveryPointOfSameType"=dword:00000001

2. Restart Arcserve Backup Server Web Service for the changes to take effect.

   All the copy to tape jobs on the Arcserve Backup server node will copy the latest recovery point of same type to tape.

File/Folder Missed from Backup of NFS Shared Folder or File/Folder Name is Converted to Junk

Symptom

When backup of NFS Shared Folder is performed, some files/folders are missing from the backup data or the name of a file/folder is converted to junk.

Solution

This issue occurs, if a file/folder in the NFS Shared Folder has a name with an unsupported language encoding. As a solution, export an SMB Shared Folder instead of NFS Shared Folder. Also consider adding UNC or NFS path with SMB protocol.

For more information about how to add UNC path and supported language encoding, refer Add UNC Path.

NFS Shared Folder Backup Fails

Symptom

When the backup of NFS Shared Folder is performed, the job fails.

Solution

This issue occurs if the GID (Group Identifier) and UID (User Identifier) values are modified from the default values. Always maintain only the default GID and UID values for NFS Client and NFS Server.

Instant VM Related

This section includes the following troubleshooting topics related to Instant Virtual Machine:

- Failed to Create an Instant VM in VMware Due to Duplicate NFS Datastore Name
- Failed to Create an Instant VM with Windows 2008 as the Recovery Server for VMware or Windows 2008 R2 Hyper-V Server
- Failed to Boot an Instant VM if the source machine is a Windows 2008/2012/2016 AD Server
- Failed to Power on VM after Restoring Hyper-V
- Instant VM Job Fails Due to Windows NFS Service Error
- Instant VM File Folder cannot be Accessed or Deleted even with Administrator Privilege
- Instant Virtual Machine Fails to Boot in Hyper-V After Restarting the Recovery Server
- Failed to Create VMware NFS Datastore and Displayed Unable to Resolve Host-name Error
- Fail to Deploy Integration service to Guest VM in Hyper-V
- Linux Instant VM failed on Non-English Hyper-V Server
- Extra license consumption for UNC Paths residing on a VM under a licensed Hyper-V
- License failure error appears when you change the license edition/type and create an Instant VM

Failed to Create an Instant VM in VMware Due to Duplicate NFS Data store Name

Symptom

Instant VM creation fails with the following error:

Failed to create NFS-based data store [arcserve_UDP_<Hostnam/IP>] with the NFS Share Name [arcserve_UDP_IVM_{GUID}]. Error code: 12. Error message: The specified key, name, or identifier already exists (details:).

The reason for this error could be that an NFS data store with the same name already exist and it is not removed, or the vCenter/Host still has references of the NFS data store in its records. However, when you log into the host directly you see an NFS data store. This datastore is most likely marked as inactive and greyed out.

Solution

To resolve this issue:

1. Log in to the ESX host and remove the NFS data store.
2. Restart the management agents on that ESX host using the following command:

   /sbin/services.sh restart

For more information on working on the ESX host, see the VMware documentation.

Failed to Create an Instant VM with Windows 2008 R2 as the Recovery Server for VMware or Windows 2008 R2 Hyper-V Server

Symptom

The instant VM creation fails and the following error is seen in the job activity log:

Windows cannot verify the digital signature for this file. A recent hardware or software change might have installed a file that is signed incorrectly or damaged, or that might be malicious software from an unknown source.

Solution

This Instant VM fails to get created because the driver of the Instant VM is signed with a secured digital certificate that cannot be supported on Windows 2008 R2, until you apply a patch from Microsoft.

Apply the Microsoft patch 3033929 and try to create Instant VM again.

Failed to Boot an Instant VM if the source machine is a Windows 2008/2012/2016 AD Server

Symptom

Failed to boot instant VM if the VM is a Windows 2008 Active Directory Server.

When the source machine is a Windows Active Directory server performing Instant VM job, the Instant VM fail to boot and a blue screen displays with the following message:

STOP: c00002e2 Directory Services could not start because of the following error: a device attached to the system is not functioning. Error status: 0xc0000001.

Solution 1:

If the target hypervisor is a Hyper-V, add the following registry key on the HyperV hypervisor, and then trigger the IVM job.

If the target hypervisor is an ESX/Vcenter, add the following registry key on the Proxy machine, which is used to run Instant VM job, and then trigger the IVM job.

`HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\InstantVM`

Reg type: DWORD value

Name: RenameADLog

Value: 1

Limitations:

- If the NTDS database and the NTDS log file paths are on different volumes, then the above solution is not suitable.
- As of now, the above solution is valid only for Windows 2008 Active Directory Server.

Solution 2:

Manually solving the problem, by logging into the Instant VM.

Follow these steps:

1. Power on the Instant VM.
2. When the machine is powering on, press F8 before the OS begins to load and choose the **Directory Service Restore** mode.

3. Rename all *.log files in the folder, *C:\Windows\NTDS*. For example, rename the file edb.log to edb.log.old
4. Run the command:

   ```
   esentutl /p "C:\Windows\NTDS\ntds.dit"
   ```

5. Restart the system.

Note: The above solution is valid for all the versions of Windows Servers with Active Directory installed.

Failed to Power on VM after Restoring Hyper-V

Symptom

Unable to power on virtual machine after restoring Hyper-V virtual machine. Following error is displayed:

Synthetic Ethernet Port (Instance ID …): Failed to Power on with Error 'Attempt to access invalid address.'

Solution

This error occurs as the MAC address of the restored VM conflicts with the existing VMs. Remove the NIC adapter from the restored VM, then manually add a new one to resolve the issue and power on successfully.

Instant VM Job Fails Due to Windows NFS Service Error

Symptom

When the Instant VM Recovery Server has Arcserve UDP Agent and Arcserve Backup both installed, then when you create an Instant VM to VMware ESX(i) server, the Window NFS service does not start and the Instant VM job fails.

Solution

The reason the Windows NFS service fails to start because the default port number of the Windows NFS service is 111 and it is used by Arcserve Backup service **Remote Procedure Call Server**.

To resolve this issue, change the default port number of the Arcserve Backup service **Remote Procedure Call Server** to another port and try to create the Instant VM again. For more information on changing the default port, see Modify the Ports Configuration File and Primary Server and Member Server Communication Ports in Arcserve Backup documentation.

Instant VM File Folder cannot be Accessed or Deleted even with Administrator Privilege.

Symptom

The Instant VM file folder cannot be accessed or deleted due to the error "Require Administrator permission" even if the user has administrator privilege.

Solution

This is a problem of NFS. You need to remove the NFS share first, and then you can delete the folder. Use the following command line to delete the folder:

nfsshare /delete [nfs share name]

Instant Virtual Machine Fails to Boot in Hyper-V After Restarting the Recovery Server

Symptom

When I start the Instant Virtual Machine (IVM), and then restart the Hyper-V recovery server, the IVM fails to boot.

Solution

To resolve this boot failure, restart the IVM.

Failed to Create VMware NFS Data store and Displayed Unable to Resolve Hostname Error

Valid for VMware

Symptom

Instant VM creation fails and displays the following error message:

Failed to create VMWARE NFS DataStore by server <hostname of recovery server>.

Error Message: An error occurred during host configuration.

Details: Operation failed, diagnostics report: Unable to resolve hostname <one hostname but not recovery server>.

Example

Failed to create VMWARE NFS DataStore by server 'host1'.

Error Message: An error occurred during host configuration.

Details: Operation failed, diagnostics report: Unable to resolve hostname 'host2'.

Solution

The reason for this error is that the ESX server lists all the NFS data store when you create InstantVM NFS data store, even when some data store are no longer available. For example, the NFS dat astore created by **host2** still exists even if **host2** is not available because the machine was deleted. So the ESX server cannot resolve the host name.

To resolve this issue, delete the unavailable data store from the ESX server.

Follow these steps:

1. Log in to the ESX server using SSH.
2. Type the following command:

   ```
   esxcfg-nas –l
   ```

   The same error is displayed in the command line.

   ```
   Error performing operation: Unable to resolve hostname
   'host2'.
   ```

3. Add a mapping in the **/etc/hosts** file of the ESX server to resolve this issue.

   <IP address> <hostname>

   Note: The IP address should be reachable.

   Example: 10.57.X.X host2

4. List all the NFS data store using the following command:

```
esxcfg-nas -l
```

```
arcserve_UDP_<hostname> is /arcserve_UDP_IVM_{ESX_gen-
erated_number} from <hostname> unmounted unavailable
```

Example: esxcfg-nas –l

arcserve_UDP_host2 is /arcserve_UDP_IVM_{991555E6-09A4-4D80-A47E-522831A62Axx} from host2 unmounted unavailable

5. Use the following command to delete the unavailable data store:

```
esxcfg-nas -d arcserve_UDP_host2
```

6. Remove the <hostname> mapping in the **/etc/hosts** file from the ESX server.

   Now you can use the NFS function as usual.

   Note: For more information on this issue, see the VMware KB article.

Failed to Deploy Integration service to Guest VM in Hyper-V

Symptom

When backing up a virtual machine that has the application (SQL or Exchange) installed, the recovery point does not include the writer information and the activity log of the backup job displays the following warning message:

Failed to deploy integration service to the VM.

Solution

This issues happens often when the Windows Management Instrumentation (WMI) is disabled by firewall on the guest VM. To resolve this issue, use the following steps:

1. Log into the guest VM.
2. Open Control panel.
3. Open windows Firewall.
4. Click Allow an app or feature through Windows Firewall.
5. Enable Windows Management Instrumentation(WMI).
6. Click OK.

Linux Instant VM failed on Non-English Hyper-V Server

Linux Instant VM on Non-English Hyper-V server failed due to connection failure even when firewall is disabled.

Symptom

Instant VM creation fails displaying the following error:

Failed to connect to the Hyper-V host [Target Hyper-V ServerName]. Verify if the address of the host is correct or the credential is valid.

Solution

Incorrect configuration for Hyper-V server connection information may cause this issue.

For more details, refer to **Configure the Hyper-V server connection information for Instant VM** in the *Arcserve UDP Agent for Linux User Guide*.

Extra license consumption for UNC Paths residing on a VM under a licensed Hyper-V

Symptom

UNC paths residing on VM under a hypervisor consume extra socket license even when license is already applied before to the hypervisor host. This error results in consumption of extra licenses.

Solution

Hypervisor host backup and UNC paths residing on the same hypervisor consume same license. But, on backing up UNC paths residing in any VM that is under the same hypervisor, the UNC paths/shares consume an extra license. This issue occurs only if the VM is not added/imported as a node to the console and also when UNC paths are added with a name other than the name provided while adding the VM to the Console.

Follow these steps:

1. Either add the node (the VM having UNC paths) to the Console and Specify Hypervisor details to the node or import node from hypervisor.
2. Add UNC paths/shares with the same name as the one added in step 1.

License failure error appears when you change the license edition/type and create an Instant VM

Symptom

After you add or change a license type or edition in the console, and create Instant VM without running any backup job, a license failure error message appears. However, the next time you run an Instant VM, it successfully identifies the new license and creates an Instant VM.

Solution

Follow these steps:

1. After a new license is added to the UDP console, for the existing plan, run the backup job.
2. Now create an Instant VM.

Linux Agent Related

This section includes the following troubleshooting topics related to Arcserve UDP Linux Agent:

Backup Destination Settings Disable when Opening the Linux Agent UI

Job Status, Job History, and Activity Log are Not Visible

Backup destination settings disabled when opening the Linux agent UI

If the Linux Backup Server is not removed from the Arcserve UDP Console before uninstalling the Console, the backup destination settings will be disabled when opening the Backup Server UI.

Symptom

The Backup Server is not notified that the Arcserve UDP Console is uninstalled. The Backup Server assumes that it is still managed by the Console.

Solution

Log into the Backup Server and run the following command:

```
# /opt/Arcserve/d2dserver/bin/d2dreg --release
```

The Backup Server is released from the Console and now you can change the backup settings from the Backup Server UI.

If the problem persists, use Live Chat to contact Arcserve support. Live Chat lets you optimize intelligent conversation between you and the Technical Support team, allowing you to address your concerns and questions immediately, while still maintaining access to the product.

Job Status, Job History, and Activity Log are Not Visible

Symptom

I cannot see the job status, job history, and activity log for Linux nodes in Arcserve UDP Console.

Solution

Linux Backup Server is unable to connect to Arcserve UDP using the hostname.

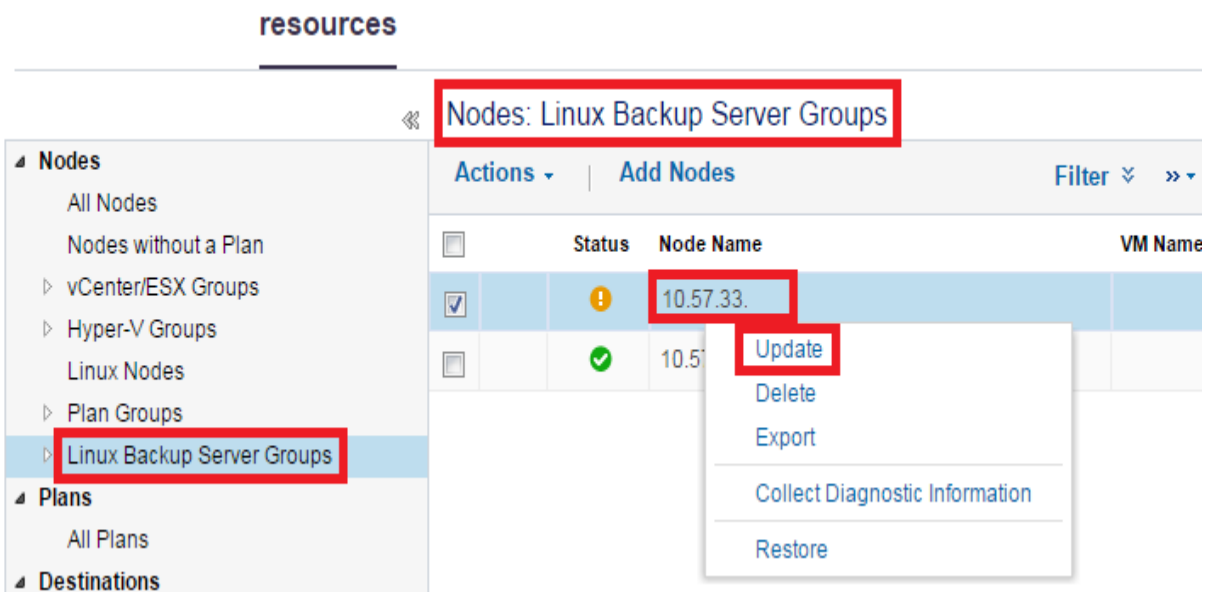
Follow these steps:

1. Create the server_ip.ini file at the following location of Arcserve UDP:

   "UDP installation path"\Management\Configuration\server_ip.ini

2. Enter the IP address of Arcserve UDP in this file.
3. Log into the Arcserve UDP Console and update Linux Backup Server and Linux nodes.

   Note: Linux Backup Server can be updated only from Linux Backup Server Groups, where all the Linux backup servers are listed.



The job status, job history, and activity log are visible.

System Volume of SUSE15 with XFS file system Failed to Restore

The system volume of SUSE15 with XFS filesystem failed to restore.

Symptom

Failed to perform restore job of SUSE15 node with XFS file system.

Solution

The issue occurs because system volume was not mounted. Create a CentOS 7.5 Live CD and use that Live CD to do BMR/instant BMR.

If the problem persists, use Live Chat to contact Arcserve support. Live Chat lets you optimize intelligent conversation between you and the Technical Support team, allowing you to address your concerns and questions immediately, while still maintaining access to the product.

Restore Related

This section includes the following troubleshooting topics related to restore:

- Unable to Restore Files
- Add the Restored Content Database to Original Web Application
- Database Unable to Mount while Restoring the Microsoft Exchange Database
- The From Field Does Not Display Correctly for Emails Sent by Users Having "On Behalf Of" Permissions for a Shared Mailbox
- Restore Jobs Fail After Light Integration Backups

Unable to Restore Files

Symptom

Due to a limitation from Microsoft, file data on the NTFS Deduplication volumes of a Windows 2012 R2 system cannot be read from a Windows 2012 system. As a result, if the UDP agent on a Windows 2012 system is being used restore a VM with guest Windows 2012 R2 OS and containing NTFS deduplication volumes, the following problem may occur. The problem occurs only during a file-level or mount recovery point restore operation.

The file or directory is corrupt and unreadable.

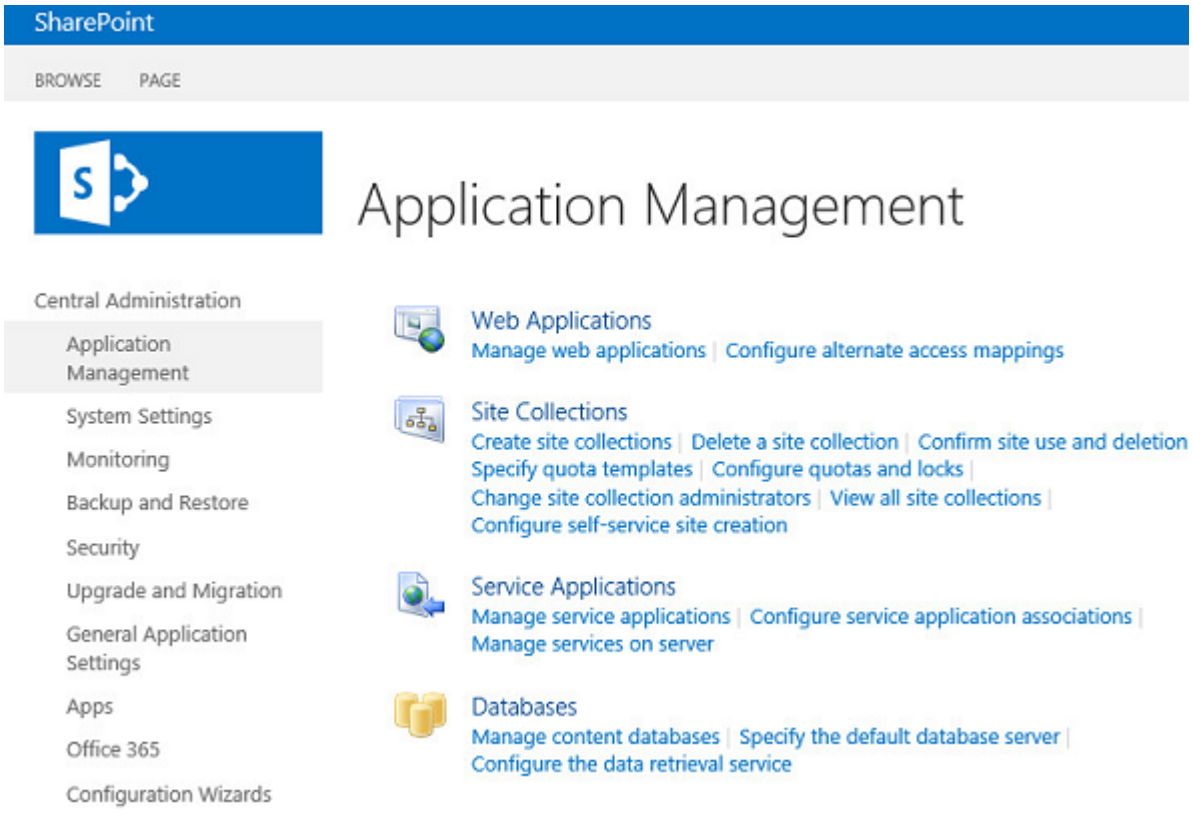
Solution

When this problem occurs, start the restore process from a UDP agent installed on a Windows 2012 R2 system.

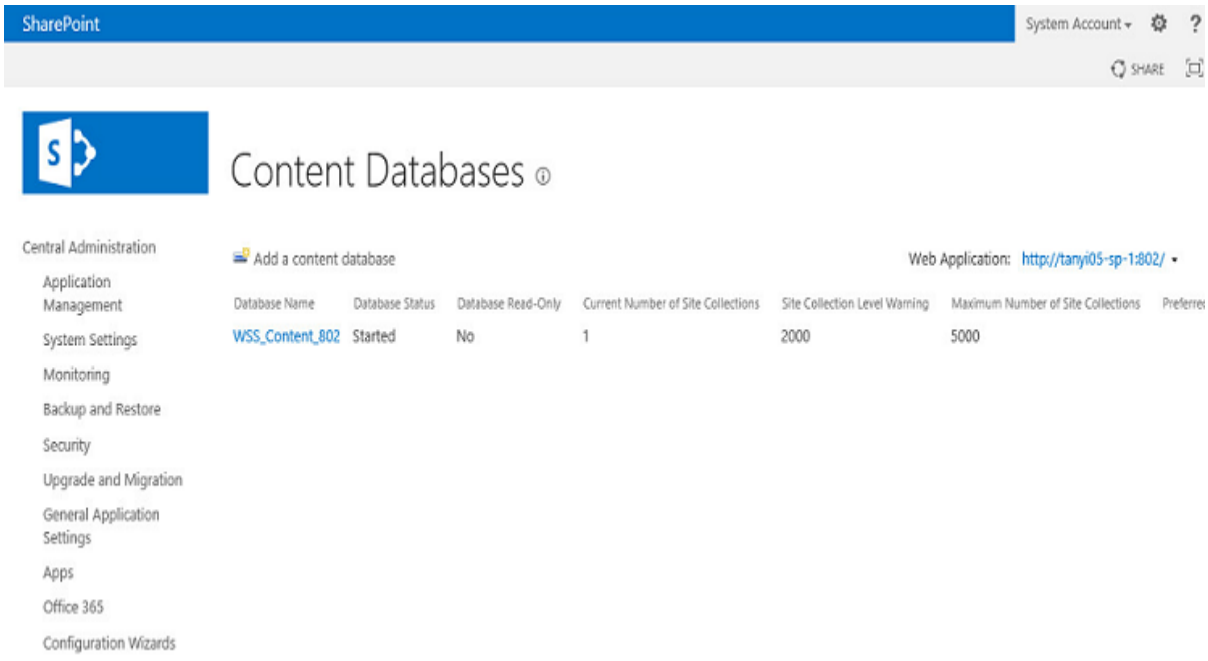
Add the Restored Content Database to Original Web Application

Follow these steps:

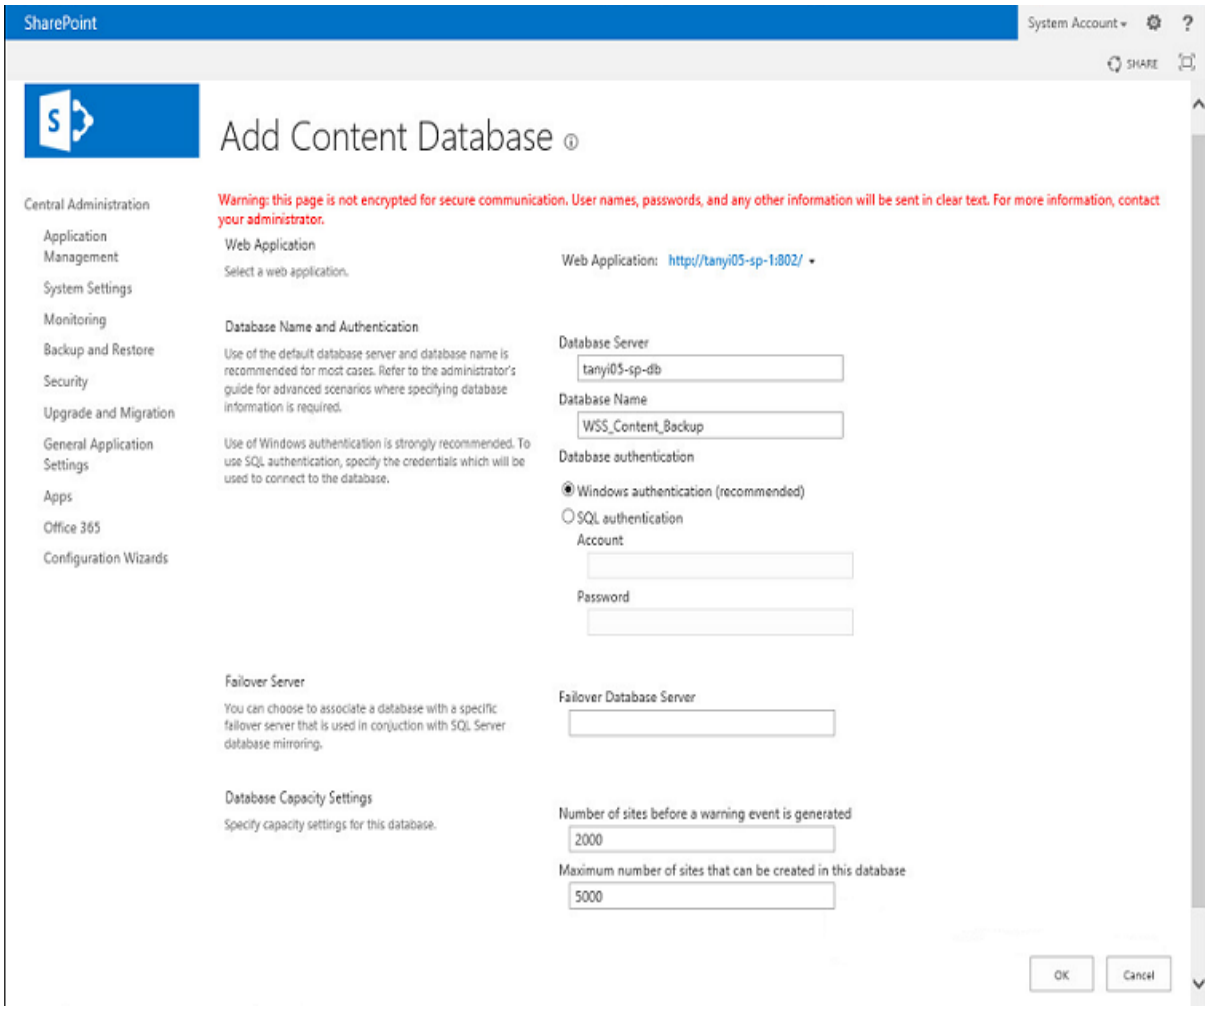
1. Open SharePoint Central Administration and select Application Management.



2. Select Management content databases.



3. Select the web application and click Add a content database.
4. Type Database Server and Database Name, (for example, WSS_Content_Backup) and then click OK.



The content database is now associated with its original web application.

Database Unable to Mount while Restoring the Microsoft Exchange Database

Symptom

When I restore a Microsoft Exchange database, the database is unable to mount. The required logs are missing or the transaction logs are not contiguous. There are event errors such as 454, 455, and 2006 in the event log.

The following two reasons could cause the database mount failure:

Reason 1: The Purge Exchange log option is enabled in the UDP settings and this setting deletes the transaction logs after every backup. Then user try to restore previous session after several backup which purge log operation occur.

Reason 2: Users have manually deleted the transaction logs or the logs are deleted by other programs such as an antivirus software.

Solution

Solution 1: If you have enabled the Purge Exchange log option and the transaction logs are not contiguous, then restore all the sessions one by one, starting from the latest session, until the session fails to restore. If the latest session fails to restore, then try Solution 2.

Solution 2: If Solution 1 does not work, use Solution 2. This solution resolves both the issues.

For example, you want to restore the Test database from Session 1. The following steps use the database name as Test.

1. Log in to the Exchange server on which the database is located.
2. Delete all the files (such as *.edb, *.log, *.jrs, *.chk) from the database folder.
3. Mount the database to create an empty database.

   Mount-Database -Identity Test
4. Restore the same session again to the original location.

   If the restore is successful, you do not have to perform the following steps. If the restore fails, continue with the following steps.
5. Mount the database again.

   Mount-Database -Identity Test
6. Create a temporary database.

   new-mailboxdatabase -name OtherDatabase
7. Move the mailbox to any other database.

get-mailbox –datatbase Test –resultsize unlimited | new-moverequest –targetdatabase OtherDatabse

8. Remove the mailbox database from the target machine.

   remove-mailboxdatabase –identity Test

9. Create a mailbox database with same name.

   new-mailboxdatabase –name Test

10. Restore the same session again to the original location.

    The database successfully mounts.

The From Field Does Not Display Correctly for Emails Sent by Users Having "On Behalf Of" Permissions for a Shared Mailbox

Symptom

When I restore Exchange mails, if the email is sent by a user having the "on behalf of" privilege for a shared mailbox, after restore, the "From" information is not displayed correctly. The "From" field displays only the <host sender> name.

Solution

Follow these steps to resolve the issue:

1. Perform one of the following actions:

   For agentless backup

   - On the HBBU proxy server, create a grtcfg.ini file in the Configuration folder:

     [product_installed_path]\ Engine\Configuration

   For agent-based backup

   - On the agent machine, create a grtcfg.ini file in the Configuration folder:

     [product_installed_path]\ Engine\Configuration

2. Add the following content in the grtcfg.ini file:

   [common]

   0xFF07_enable=1

3. Submit the restore job again.

Restore Jobs Fail After Light Integration Backups

Symptom

When you submit light integration backups from the Arcserve Backup Manager to backup exchange online node from RPS server, the node is already backed up with copy to tape session. The jobs fail when the source data includes Arcserve D2D sessions that were backed up previously by Arcserve Backup. Error message AW0813 appears in the Activity Log.

Solution

Arcserve Backup design behavior is responsible for this behavour. You need to modify this behavior to allow Arcserve Backup to back up Arcserve D2D sessions that were backed up before.

Follow these steps:

1. From the Arcserve D2D server (node) that you are backing up, open Windows Registry Editor.
2. Open the following key:

   HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA ARCserve Backup\ClientAgent\Parameters\AllowRedundantD2DBackups
3. Change the DWORD value of AllowRedundantD2DBackups to 1.

   Note: If the above-described Registry key is not present on the node that you are backing up, you must create the key.
4. Close Windows Registry Editor.

   For more details, refer to the KB article.

Restore to NFS Shared Folder Changes the File Name to Junk

Symptom:

When I perform restore to an NFS shared folder, the files/folders are restored with a junk name.

Solution:

UDP does not support restoring the files and folders into NFS shared folder. We recommend to export a Server Message Block (SMB) shared folder and then restore to the same folder.

Note: If the file/folder name is converted to junk during the NFS backup session, even after restore, the file/folder name remains junk.

Gateway, RPS, Data Store, Console, and Database Related

This section includes the following troubleshooting topics related to the recovery point server (RPS), data store, and database:

- Data Store Name is Already in Use
- Unable to Connect to data store due to DNS Issue
- Data Store Switches to the Restore Only Mode
- Error when RPS Version is Lower than the Console Version
- Adding the Same Resource in Different Sites is not Supported
- How to enable Log Truncations when SQL Database is in Full Recovery Mode
- Browse Recovery Points Does Not Show Available Recovery Points When RPS Configured with FQDN
- Access Denied when Adding RPS
- Unable to Change to UDP View for Recovery Point
- Access Denied when Adding or Updating Nodes
- UDP Console Does not Open when SQL Administrator Password is Changed
- Failed to Mount the Recovery Points Due to Timeout
- How to Update the Gateway Server Credentials
- How to Update the Gateway when the Gateway Proxy Credentials are Changed
- Console Displays Identity Service is Starting Message

Data Store Name is Already in Use

Symptom:

When I create a data store, sometimes the following message appears even though I specify a new data store name:

This name is already being used by another data store on the server. Please specify a different data store name.

Solution:

This happens when you have an existing data store but for some reason, the data store UUID at the registry is corrupt. You can delete the data store from GUI, but the name remains in the recovery point server registry.

To resolve, specify a new name.

Unable to Connect to data store due to DNS Issue

Symptom:

During BMR, I fail to connect to the shared folder exposed to RPS. Despite correct user name / password, I am not able to browse node on RPS.

Solution:

When Windows UAC is enabled on the RPS server, even an account belonging to local administrator group may not be able to access a share folder of data store, if this account is not granted access to this share explicitly.

Using built-in administrator of RPS, grant read / write privilege to the account used under BMR to the share folder exposed by the specified data store.

Data Store Switches to the Restore Only Mode

Symptom

I notice that a data store switched to the Restore Only mode, and does not allow me to back up any data.

Solution

When a disk that is used by a data store runs out of disk space, the data store switches to the Restore Only mode. In this mode, you can perform restore, but you cannot back up data to the data store. Also, when the specified memory allocation is utilized completely, you either increase the memory allocation or you change the data store from memory mode to SSD mode. Even in such cases, the data store switches to the Restore Only mode.

To resolve such issues, move the data store to a larger disk by importing the data store.

First copy the folders where the disk is full to a larger disk with more free space and then import the data store from the console.

The **Import Data Store** feature lets you add a data store to the recovery point server. You can import any existing data store to a recovery point server. The data stores that you have deleted earlier from a recovery point server are available to import.

Follow these steps:

1. From the Console, click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Recovery Point Servers**.

   The **Destinations: Recovery Point Servers** page is displayed.
3. Perform one of the following actions:
   - Right click a recovery point server.
   - Select a recovery point server, and from the center menu click the **Actions** drop-down list.
4. Click **Import Data Store**.

   The **Import a Data Store** page is displayed.
5. Perform the following actions, and click **Next**:
   - **Browse** to select the **Backup Destination Folder** from where you want to import the data store.
   - Enter **Encryption Password**.

Note: Leave it empty if the data store is not encrypted.

After authenticating the **Backup Destination folder**, the **Import a Data Store** page displays the details of the data store.

6. Modify the details, if necessary, and click **Save**.

   If you have copied folder of Data Destination, Index Destination, and Hash Destination for Deduplication data store, change the folder path.

   Note: You cannot enable or disable the encryption option for an existing data store.

   The data store is added to the recovery point server and displayed at the **Destinations: Recovery Point Servers** dialog.

   Now the data store is available for backups.

Error when RPS Version is Lower than Console Version

Symptom

There is a plan with a backup task or replication task and the destination is Recovery Point Server (RPS). The RPS is of older version and the Console is the latest version. When I create, modify, re-deploy, pause, or resume a plan and there are one or more older versions of RPS, the following error is displayed:

The version of destination Recovery Point Server 'rps1'is lower than the version of current console. To continue, you need to upgrade and update the Recovery Point Server.

Solution

This error occurs when your plan uses an older version of RPS. To resolve this error, upgrade the RPS that is used in the plan. If you upgrade the RPS manually (outside the Console), do not forget to update the RPS on the Console.

Upgrade the RPS in the following order of preference:

Replicate to a remotely managed RPS (RPS3) > Replicate task (RPS2) > Backup task (RPS1)

At first, upgrade RPS3, then upgrade RPS2. At last, upgrade RPS1.

Follow these steps to upgrade:

1. From the **resources** tab, navigate to the Destinations: Recovery Point Servers page.
2. Select the Recovery Point Server.
3. Right-click and select **Install/Upgrade Recovery Point Server**.
4. Click **OK**.

Follow these steps to update:

1. From the **resources** tab, navigate to the Destinations on the left pane, and click Recovery Point Servers.
2. Select the Recovery Point Server.
3. Right-click and select **Update**.
4. Click **OK**.

Adding the Same Resource in Different Sites is not Supported

Symptom

When I add a data store in remote site, I get the following error message:

Cannot submit job for a Hyper-V VM or ESX VM

Solution

The error occurs because the same resource (node, RPS server, ASBU server, Hyper-V server, ESX server, proxy server) is already present in another site. To resolve this error, delete the resource from all sites, then add the resource only in one site.

How to enable Log Truncations when SQL Database is in Full Recovery Mode

Symptom

When the database is in the Full mode and a full database backup is performed, the SQL truncation log cannot be truncated.

Solution

To resolve this problem, add two registry values to enable Arcserve UDP run the BACKUP LOG command to back up the transaction log. This command marks the space, which is already written to database file, as reusable.

Follow these steps to add the registry value:

1. Open the registry table editor on the agent machine using the following command:

   ```
   regedit
   ```

2. Navigate to the following keys depending on the agent-based or agentless backup:

   For agent-based backup for both 32-bit and 64-bit OS, navigate to the following key on the agent machine:

   HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll

   When using a version lower than Arcserve UDP v6.5 Update 2, then for agentless backup navigate to the below key. Create the registry table value inside the VM that you want to back up on the proxy server. If the registry table key is not available, then create the complete key path.

   - **32-bit OS:**

     HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll

   - **64-bit OS:**

     HKEY_LOCAL_MACHINE\SOFTWARE\WoW6432Node\Arcserve\Unified Data Protection\Engine\AFBackupDll

3. Create the following two registry values and for both set the value to 1:

   - dword value named BackupSQLLog4Purge
   - dword value named ForceShrinkSQLLog

   The registry value is added.

The solution is in effect when the next purge job occurs.

Browse Recovery Points Does Not Show Available Recovery Points When RPS Configured with FQDN

Valid on Windows Operating System

Symptom

When the RPS is not in a domain you configure the FQDN (by adding the DNS Suffix) in the UDP Console, the Browse recovery Points do not show accurate result. Even though you back up some sessions to RPS, the recovery point number is shown as zero.

The reason is, when the RPS is not in a domain, the RPS cannot identify itself using the FQDN.

Solution

To resolve this issue, add the DNS suffix to the RPS host.

Follow these steps:

1. Open the Control Panel and navigate to **System and Security**, **Systems**.
2. Click **Change settings** for Computer name, domain and workgroup settings.

   The System Properties dialog open.
3. In the **Computer Name** tab, click **Changes**.

   The **Computer Name/Domain Changes** dialog opens.
4. Click **More**.

   The **DNS Suffix and NetBIOS Computer Name** dialog opens.
5. In the Primary DNS suffix of this computer field, add the network DNS suffix and click **OK**.

   For example, add ABC.com.
6. Restart the system.

Access Denied when Adding RPS

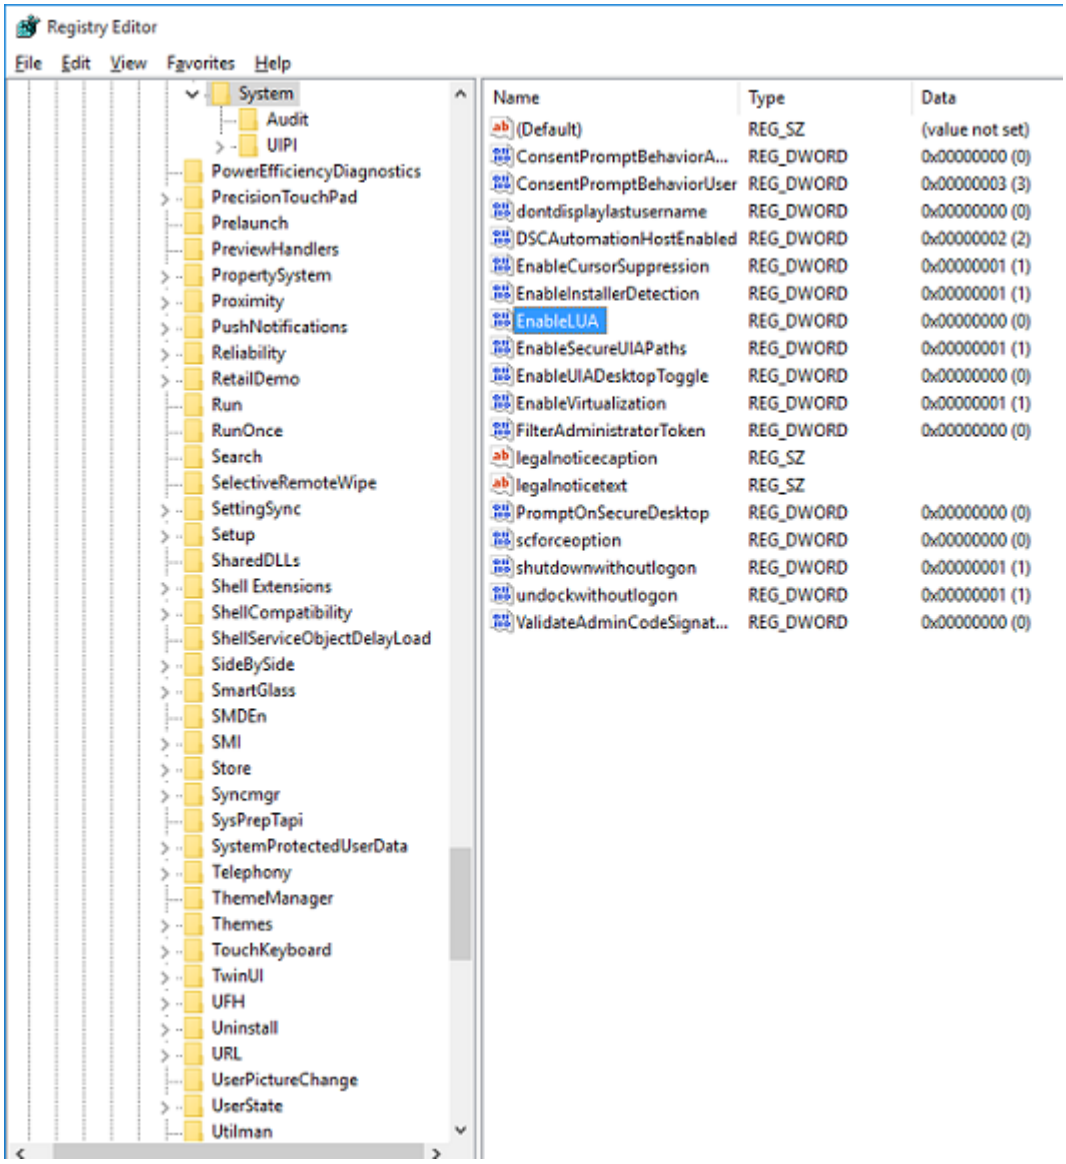
Symptom:

If you add Windows 10 as RPS and even when UAC is not running, an error message appears:

Access Denied. The account may not have Administrator privileges or the account is non-built-in administrator account and UAC is turned on.

Solution:

For Windows 10, to disable the UAC, change the value of the registry key, not just by switching the level to "Never notify" in the Control Panel, but also according to the registry diagram given below.



Set Registry "EnableLUA" as 0 to disable UAC completely and restart machine.

Then, add Windows 10 as RPS in console.

Unable to Change to UDP View for Recovery Point

Symptom:

When changing the view to UDP View for RPS recovery point, at times you receive error messages, such as access denied.

Solution:

Disable the UAC to fix the problem. For more information, view How to Disable a Remote UAC for a non-built-in Administrator.

Access Denied when Adding or Updating Nodes

Valid on Windows Operating System

Symptom

Sometimes when you add or update nodes, you get the following error:

Access denied. The account may not have Administrator privileges or the account is a non-built-in administrator account and UAC is turned on.

Solution

The error occurs in the following scenarios:

- You have logged in as a local user or domain user, who is not in the local administrators group of the node to add or update node.
- You have logged in as a user who is in the local administrator groups of the node but has a non-built-in administrator account of the node to add or update node.

To resolve this issue, follow these steps:

1. Add the local or domain user to the local administrator groups of that node.
2. Disable the UAC of that node.

   Follow these steps to disable the UAC:

   a. Click Start, type regedit in the Search programs and files field, and then press Enter.

   b. The Windows Registry Editor opens.

      Note: You may need to provide administrative credentials to open Windows Registry Editor.

   c. Locate and click the following registry key:

      HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System

   d. From the Edit menu, click New, and then click DWORD (32-bit) Value.

   e. Specify LocalAccountTokenFilterPolicy as the name for the new entry and then press Enter.

   f. Right-click LocalAccountTokenFilterPolicy and then click Modify.

   g. Specify 1 in the Value data field and then click OK.

   h. Exit the Registry Editor.

Notes:

- This procedure is not similar to disabling the UAC. Using this procedure, you can disable some of the functionalities of UAC.
- Considering that remote Windows Management Instrumentation (WMI) technology is used for import, ensure that WMI is not blocked by the firewall.

For more information about Windows behavior, see the Microsoft documentation.

UDP Console Does not Open when SQL Administrator Password is Changed

Symptom:

Arcserve UDP Console use SQL server as Database, and use SQL administrator "sa" to connect to the Database. If the "sa" password is changed, Console home page fails to open and the following message displayed:

SQLServer is not available now. Please check service status and then restart Arcserve UDP Management service.

Solution:

1. Run <homedir>\Management\BIN\DBAccountUpdate.bat
2. Type: updatePassword
3. Type the new password and press Enter

Failed to Mount the Recovery Points Due to Timeout

Symptom

When the RPS has a heavy load, the OS takes a longer time to attach the mounted volume and the mount recovery points fail. You get the following message in the Activity logs:

Mounting the volume takes longer time than expected (2 minutes), this may happen when your server is under heavy load. Please try again when the server load is less heavy or check the troubleshooting in online documentation to increase the time out value.

Solution

To resolve this issue, increase the timeout value.

Follow these steps:

1. Log in to the RPS and navigate to the following location:

   ```
   HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFStorHBAMgmt
   ```

   ```
   "WaitDeviceReadyTimeoutS"=dword:00000078
   ```

   Default: 120 seconds

2. Change the timeout value to a higher value.

   For example, change the timeout value to 600 seconds (10 minutes)

How to Update the Gateway Server Credentials

Symptom

If the gateway installation user name is changed or the password is expired, you get following error during the plan deployment:

Agent deployment failed.

Failed to impersonate the user who installed Arcserve Remote Management Gateway using the stored credentials. Please check if the credentials are still valid and redeploy.

Solution

To resolve this issue, update the gateway account user name or password.

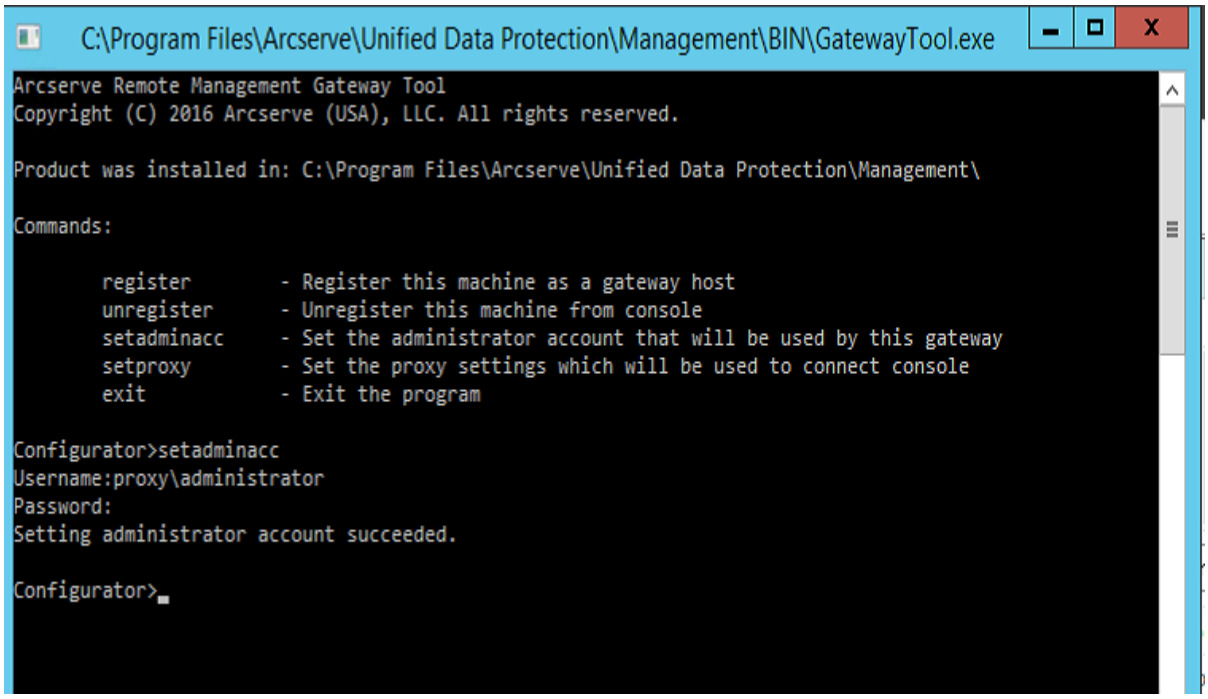
Follow these steps:

1. Log into the gateway server.
2. Navigate to the BIN folder in the Arcserve UDP installation folder.

   For example, C:\Program Files\Arcserve\Unified Data Protection\Management\BIN\GatewayTool.exe
3. Right-click GatewayTool.exe and click Run as Administrator.
4. In the command prompt window, type **setadminacc**.
5. Specify the new user name.
6. Specify the new password.

   If all the details are correct, you will see the following message:

   ```
   Setting administrator account succeeded
   ```



7. Redeploy the plan.

How to Update the Gateway when the Gateway Proxy Credentials are Changed

Symptom

When the proxy server credentials are changed, the connection to the gateway is broken. You get the following error message while trying to connect to the gateway:

Cannot connect to the gateway site proxy. Please check the status of the site and make sure that gateway is running.

Solution

To resolve this issue, update the parameters in GatewayTool.exe setproxy.

Follow these steps to update the gateway server:

1. Log in to the gateway server.
2. Navigate to the BIN folder in the Arcserve UDP installation folder.

   For example, C:\Program Files\Arcserve\Unified Data Protection\Management\BIN\GatewayTool.exe
3. Right-click GatewayTool.exe and click Run as Administrator.
4. In the command prompt window, type **setproxy**.
5. Type **2** to select Custom proxy because the IE proxy does not support credentials for now.
6. Specify the proxy server IP address.
7. Specify the proxy server port.
8. Type **Y** for authentication.

   Note: If you type N, it means disable the credentials.
9. Specify the new user name.
10. Specify the new new password.

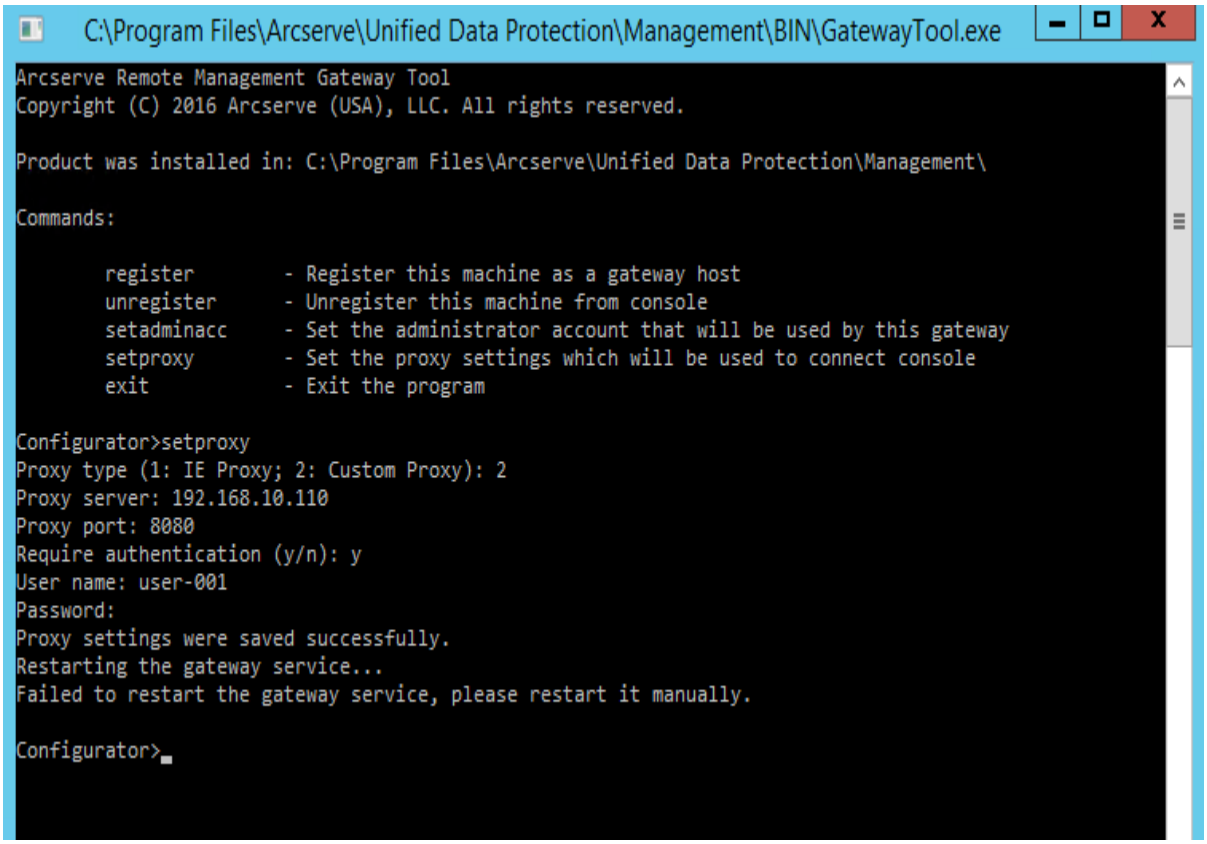
    If all the details are correct, you will see the following message:

    ```
    Proxy Settings were saved successfully. Restarting the gateway service.
    ```

Note: If the gateway proxy credential is enabled in a plan, update the proxy user name and password in that plan.

Similarly, you can change other settings such as Proxy Type, IP address, and Port

number in **setproxy**.



If you see the **Failed to restart the gateway service, please restart it manually** message in the command prompt window, follow these steps:

a. Run **services.msc** to look for **Arcserve remote management gateway service** and then restart the remote management gateway service.

b. If the remote service stops after you restart, use the Task Manager to end the **tomcat8.exe** task manually.

c. Refresh **services.msc** and restart **Arcserve remote management gateway service**.

Console Displays Identity Service is Starting Message

Symptom

Unable to log in to the Arcserve UDP Console. The Console displays the following messages even after five minutes of logging in:

Identity Service is starting

Solution

To resolve this issue, open the Windows Service Console and restart the Arcserve UDP Console service, **Arcserve UDP Management Service**.

Virtual Machine Backup and Restore Related

This section includes the following troubleshooting topics related to the recovery point server (RPS), data store, and database:

- Add Permission for VDDK at vCenter Server Level
- Backup job for VM template is always converted to full backup
- Independent disks are skipped by backup job of VM template
- Backup job for VM on SMB 3.0 share fails with error message
- VM recovery job fails when restoring VM to a Windows default file share
- Volume information unavailable for recovery point
- Permissions for Host-based Agentless Backup and Virtual Standby at vCenter Server Level
- Convert Incremental Backup to a Verify Backup Because the Virtual Machine Snapshot Either Changed from the Last Backup Job or Needs Consolidation
- Agentless backup for VMware VM fails for CD/DVD device of VM
- Agentless Host-based Backup for Hyper-V Fails after Upgrading Arcserve UDP
- Host-based Agentless Backup Fails in VMware ESXi 6.0
- Failed to Create a Snapshot for Hyper-V Virtual Machines when Multiple Jobs are Running
- Failed to Perform Backup of a Virtual Disk. System error=[The device is not ready(21)]
- Backup Job Fails
- Failing to import VMware VMs from vCenter
- Unable to Apply Backup Settings to Node
- Backups Fail Due to ESXi License
- Host-based Agentless Backup does not Use HotAdd Transport Mode
- HotAdd Transport Mode does not Work when Attempting to Back Up a VMware VM
- Host-based Agentless Backup or Restore Job Uses NBD or NBDSSL Transport Mode even when SAN Mode is Possible
- More Granular Control for VSS over Windows Guest OS Quiescings
- Recovery Operations Fail when Recovering Data Using the HotAdd or SAN Transport Mode

- Recover VM Operation fails when a non-default port is specified
- Scheduled Incremental or Full Backup Job Fails for Hyper-V VM
- Hyper-V VSS NTDS Writer Fails while Taking the VSS Snapshot in the VM
- MAC Address Changes are not Retained After VM Recovery
- Hyper-V Failed to Create VSS Snapshot
- Unable to Open VMDK Files
- Problems Caused by Duplicate VM UUID
- File System Catalog Job Fails or Recovery Point Check Fails for Host-Based Agentless Backup
- Incremental Backup Converts to Verify Backup or the Backup Size Increases in Hyper-V
- Host-Based Agentless Backup Fails Hyper-V VM That has a Special Differencing Disk Configuration
- Backup Job Fails for a VMware Virtual Machine
- Disable Rescan of Host Bus Adapters when the Source and Proxy are in Different VMware ESX Servers
- Disable the Consecutive Snapshot Creation in VMware VM for a Backup
- Agentless Host-based Backup crashes when using Windows 2003 R2 64-bit as Backup Proxy
- When Restored from a Higher Version of ESXi host to a Lower version ESXi host, VM gets stuck at booting stage
- RAM Utilization reaches 99% When Backup Jobs are Submitted to VM
- Hyper-V Restore Job Fails, Cannot Connect to Utility on Host
- Automatic Protection Fails to Detect and Protect VM
- Set Read block size when backing up VMDK file

Add Permissions for VDDK at vCenter Server Level

If you do not have the appropriate permissions, the backup job for a host-based virtual machines and virtual standby job fails.

To avoid this issue, verify that you have the appropriate permissions. If you are a vCenter user, you do not need an Administrator permission at the vCenter Server level but you must have an Administrator permission at the Datacenter level. In addition, you must have the following permissions at the vCenter Server level:

- Global, DisableMethods and EnableMethods
- Global, License

For more information, see the VMware KB article.

For more information on the permission, see Permissions for Host-based Agentless Backup and Virtual Standby at vCenter Server Level.

Permissions for Host-based Agentless Backup and Virtual Standby at vCenter Server Level

When you configure vCenter to manage virtual machines, generally you set up users or groups with vCenter administrator privileges. This approach helps to ensure that the vCenter accounts have unrestricted access to vCenter functionality and tasks. Optionally, you can create vCenter users and groups that can be used to facilitate only backup operations or only backup and restore operations.

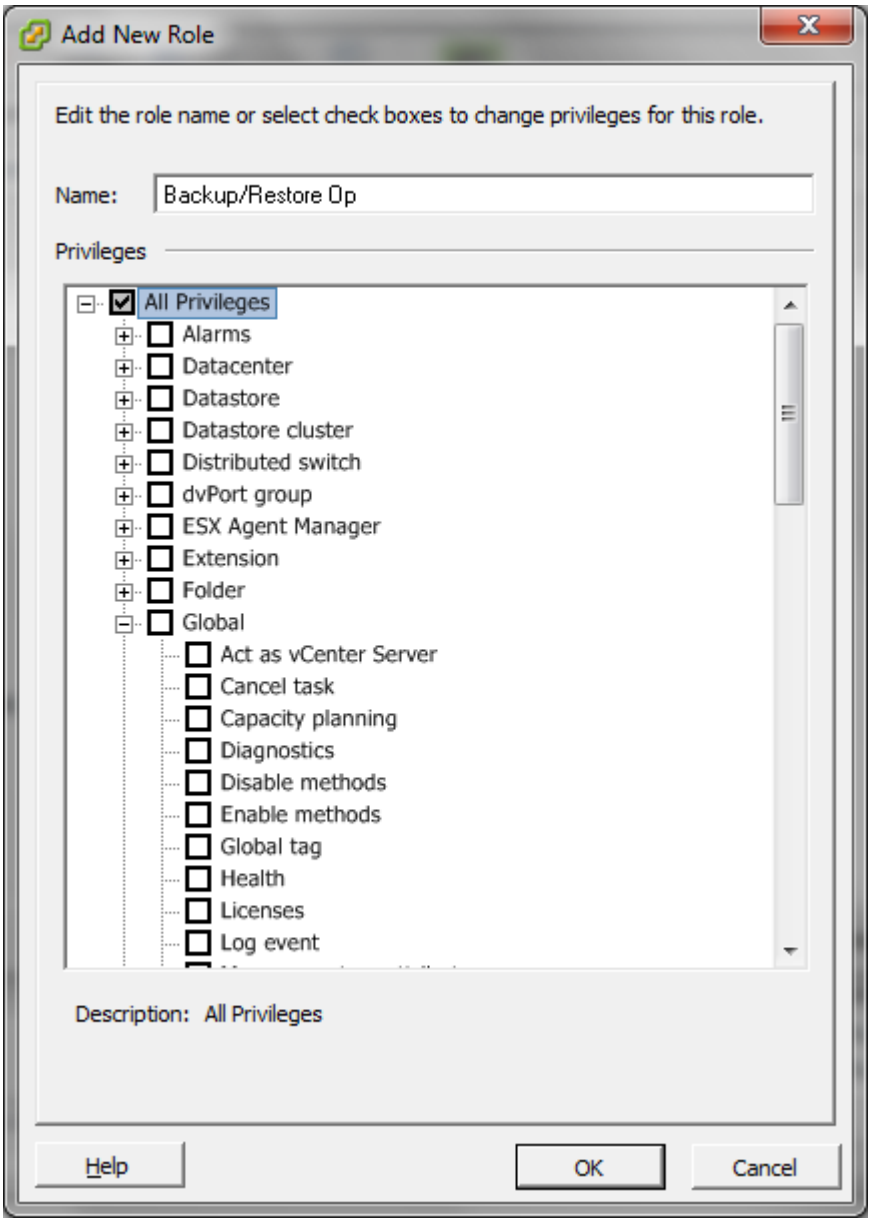
When using vCenter non-administrative accounts to facilitate backup and restore operations, you create vCenter roles, assign privileges to the roles, and then apply the role to individual users or groups.

Note: As a best practice, VMware recommends that you allow non-administrative vCenter user accounts to be members of the Windows local administrator group.

Important! The following steps assume that you are familiar with how to configure vCenter users, groups, roles, and permissions. Consult the vCenter documentation as needed.

Follow these steps:

1. Log in to vCenter using the VI Client.
2. Open the Add New Roles dialog and specify a name for the role.



3. Expand All privileges.
4. **(Optional)** To allow the role to **facilitate only backup operations**, specify the following privileges:

   Important! To allow the role to facilitate backup and restore operations, continue to the next step.

   - Expand Virtual machine and Configuration, and specify the following privileges:
     - Disk change tracking
     - Disk Lease
     - Add existing disk

- Add new disk
- Add or remove device
- Change resource
- Remove Disk
- Settings

- Expand Virtual machine and Provisioning, and specify the following privileges:
  - Allow read-only disk access
  - Allow virtual machine download
- Expand Virtual machine and specify the following privileges:

  vSphere 4: Expand State and specify Create Snapshot and Remove snapshot.

  vSphere 5: Expand Snapshot management, expand State and then specify Create Snapshot and Remove snapshot.
- Expand Global and specify the following privileges:
  - Disable methods
  - Enable methods
  - Licenses

  Go to Step 6.

5. To allow the role to **facilitate backup and restore operations**, specify the following privileges:
   - Expand Datastore and specify the following privileges:
     - Allocate space
     - Browse datastore
     - Low level file operations
   - Expand Global and specify the following privileges:
     - Disable methods
     - Enable methods
     - Licenses
   - Expand Host, expand Local Operations, and then specify Reconfigure virtual machine.

     Note: This privilege is only required when you need to perform backup and restore operations using the HotAdd transport mode.

- Expand Network and specify Assign Network.
- Expand Resource and click Assign Virtual Machine to resource pool.
- Expand Virtual machine and Configuration, and specify the following privileges:
  - Add existing disk
  - Add new disk
  - Add or Remove device
  - Advanced
  - Change CPU count
  - Change resource
  - Disk change tracking
  - Disk Lease
  - Host USB device
  - Memory
  - Modify device setting
  - Raw device
  - Reload from path
  - Remove disk
  - Rename
  - Reset guest information
  - Settings
  - Swapfile placement
  - Upgrade virtual hardware
- Expand Virtual machine and Guest Operations, and specify the following privileges:
  - Guest Operation Modifications
  - Guest Operation Program Execution
  - Guest Operation Queries (vSphere 5)
- Expand Virtual Machine and Interaction, and specify the following privileges:
  - Power off
  - Power on
- Expand Virtual machine and Inventory, and specify the following privileges:

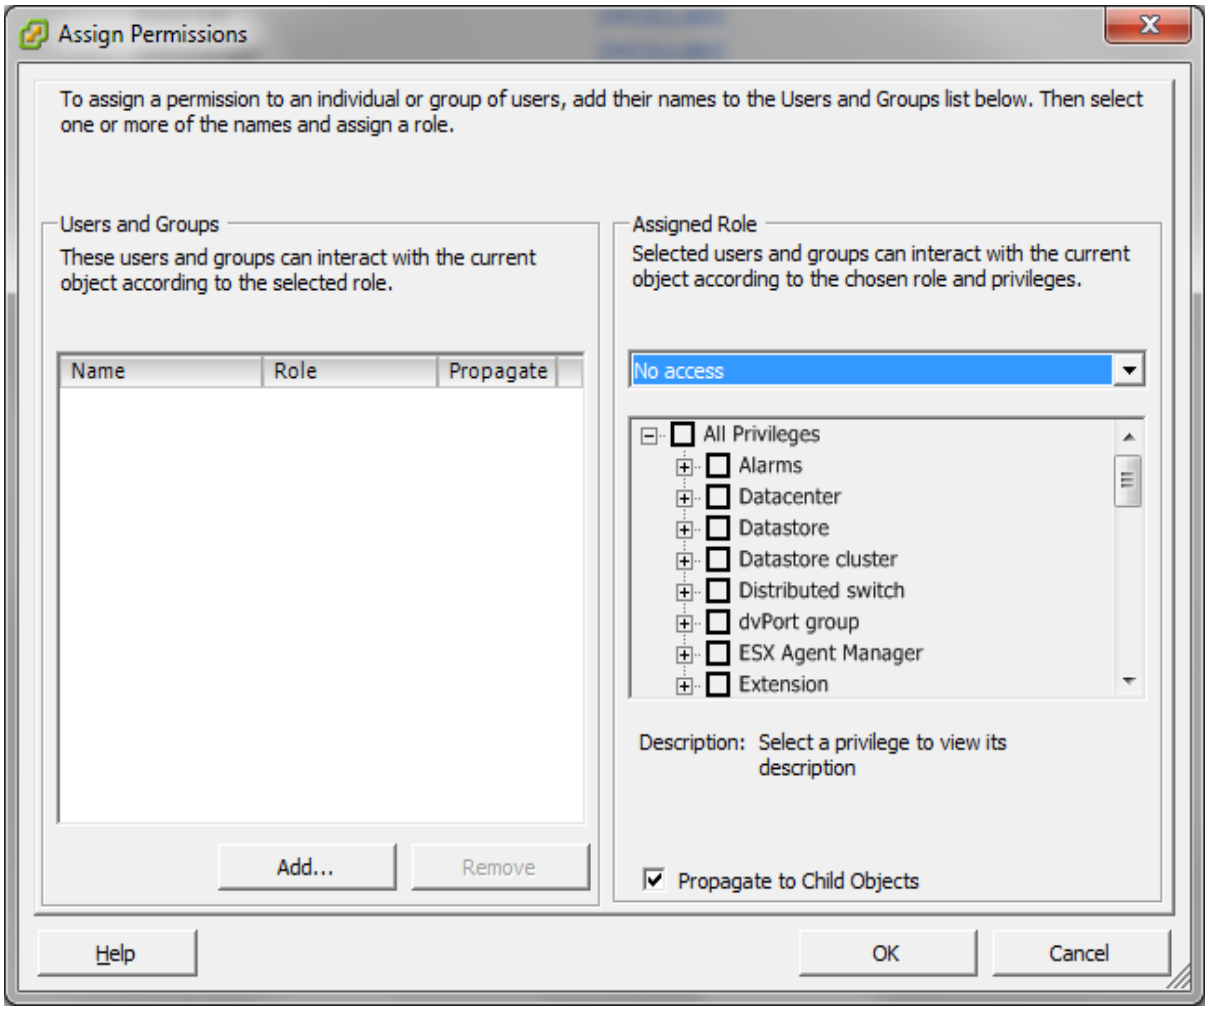
- Create new
- Register
- Remove
- Unregister

- Expand Virtual machine and Provisioning, and specify the following privileges:
  - Allow disk access
  - Allow read-only disk access
  - Allow virtual machine download
- Expand Virtual Machine and specify the following privileges:

  vSphere 4: Expand State and specify Create snapshot, Remove snapshot, and Revert to snapshot.

  vSphere 5: Expand Snapshot management, expand State, and then specify Create snapshot, Remove snapshot, and Revert to snapshot.

6. Click OK to create the role.
7. Open the Assign Permissions dialog, to assign the newly created role to users, groups, or both.



8. From the Users and Groups list, select the custom user that you want to use for backups and restores.

   From the Assigned Role drop-down list, specify the role that you want to apply to the users or groups.

9. Click OK to apply the role to the users or groups.

   The permissions are now defined for vCenter roles.

Backup job for VM template is always converted to full backup and backup data size is the provision size of virtual disk

Symptom

When backing up a VM template, backup job is converted to full backup and the processed data size is equal to provision size of the virtual disk. The following warning message appears in activity log.

As the virtual machine is configured as template, the job will be a full backup and the virtual disks will be backed up as entire disks.

Solution

This behavior is expected for the VM template backup. One workaround is to convert the template to VM at the beginning of backup, back it up, and convert it back to template at the end of backup job. If you prefer this workaround, follow the steps below to set a registry value in proxy machine.

1. Log onto proxy machine.
2. Create a registry value at proxy server level or VM level.

   Note: If you add the registry value at both the VM level and proxy level, the setting at the VM level registry holds priority over the setting at the Proxy level registry.

   At Proxy Server level (applicable to all backup jobs running in this proxy server)

   a. Open the registry key from the following location:

      [HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll]

   b. Add a DWORD value with name *TemplateDirectBackup* and specify its value as 0.

   At VM Level

   a. Open the registry key from the following location:

      [HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll\<vm instance uuid>]

   b. Add a DWORD value with name *TemplateDirectBackup* and specify its value as 0.

Note: After enabling this option, VM cannot convert again to template if backup aborts abnormally (for example, crashes during job or when proxy machine is rebooted).

Independent disks are skipped by backup job of VM template

Symptom

When backing up a VM template, backup job skips independent disks and the following warning message appears in the activity log:

Unable to backup virtual disk [datastore_720_4] shuli02-t235/shuli02-t235_1.vmdk because it is an independent disk.

Solution

This behavior is expected for the VM template backup. The root cause is a VMware limitation that prevents backup application from opening the VMDK of independent disks. One workaround is to set independent disks to dependent disks at the beginning of backup, back them up, and set them back at the end of the backup job. If you prefer this workaround, follow the steps below to set a registry value in proxy machine.

1. Log onto proxy machine.
2. Create a registry value at proxy server level or VM level.

   Note: If you add the registry value at both the VM level and proxy level, the setting at the VM level registry holds priority over the setting at the Proxy level registry.

   At Proxy Server level (applicable to all backup jobs running in this proxy server)

   a. Open the registry key from the following location:

      [HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll]

   b. Add a DWORD value with name *ConvertIndependentVMDK* and specify its value as 1.

   At VM Level

   a. Open the registry key from the following location:

      [HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll\<vm instance uuid>]

   b. Add a DWORD value with name *ConvertIndependentVMDK* and specify its value as 1.

Notes:

- After enabling this option, you cannot set back the independent disks in case backup aborts abnormally (For example, when job crashes or proxy machine is rebooted).
- This option does not work when the option "converting template to VM during backup" is enabled (the registry value TemplateDirectBackup is created with value 0).

Backup job for VM on SMB 3.0 share fails with error message

Note: Valid for Hyper-V.

Symptom

The VM resides on Hyper-V 2012 or 2012 R2 and has files on SMB 3.0 share. The backup job keeps failing with either of the following error messages:

Failed to take VSS snapshot. System error=[VSS_E_VOLUME_NOT_SUPPORTED_BY_PROVIDER]

Failed to take VSS snapshot. System error=[VSS_E_BAD_STATE]

Solution

1. If the SMB share is hosted by a Windows file server, add the role File Server VSS Agent Service on the file server and make sure that Microsoft File Server Shadow Copy Agent Service is installed.
2. If the SMB share is hosted by third-party NAS appliances or other similar solutions, verify if such appliances or solutions support SMB 3.0 and File Server Remote VSS Protocols. For details, contact third-party vendor.

Note:

- For VM residing in Hyper-V 2016, File Server Shadow Copy Agent Service is not needed.
- You must configure SMB 3.0 share properly before Arcserve UDP can back up the VM in it. For detailed requirements of SMB 3.0 share, refer the Requirements and supported configurations section of Microsoft documentation.

VM recovery job fails when restoring VM to a Windows default file share

Note: Valid for Hyper-V.

Symptom

When restoring VM by specifying Windows default file share (For example, \\hostname\C$\abc) as the destination path, restore job fails and the following error message appears:

VM recovery job was unable to create the new virtual machine.

Solution

The job failure is expected as VM files cannot be stored in Windows system default file share. Only Microsoft SMB 3.0 file share is supported. For more details, refer to *Microsoft documentation*.

Volume Information Unavailable for Recovery Point

Symptom

In the restoring VM/files wizard, while mounting recovery point or copying recovery point, no volume and file are displayed on the recovery point screen. Instead, the following message appears:

Volume information is not available for this recovery point.

Solution

This behaviour is expected if the source VM does not have Windows OS. Thus, Arcservev UDP is unable to parse virtual disks of VM to get volume information. You can restore the entire VM or copy the recovery point. If the source VM has Linux VM OS, you can also restore files from the recovery point by using a Linux backup server.

Convert Incremental Backup to a Verify Backup Because the Virtual Machine Snapshots Either Changed from the Last Backup Job or Needs Consolidation

Note: Valid on Windows platforms.

Symptom

The Incremental Backups for VMware virtual machines are converted to Verify Backups. The activity log states the following message:

"Convert Incremental Backup to a Verify Backup because the virtual machine snapshots either changed from the last backup job or needs consolidation."

Solution

Use the VMware vSphere Client to consolidate the virtual machine snapshots. For more details about consolidating snapshots, refer to VMware Knowledge Base article.

Note: Consolidating the snapshots for a virtual machine can fail due to locked files. If the backup job uses the HOTADD transport mode, then verify that the backup proxy virtual machine settings on the ESXi server does not contain the hot added hard disks. Then consolidate the virtual machine snapshots.

Agentless backup for VMware VM fails if CD/DVD device of VM is connected to an ISO image which resides on a disconnected NFS datastore

Symptom

Perform the following steps to observe the issue:

1. Prepare an ISO image on a NFS datatsore connected to the ESX host.
2. Attach the ISO image to the CD/DVD device of a VM.
3. Disconnect the NFS datastore from network.
4. Perform agentless backup for the VM.

In this case, backup job fails with error message such as “Could not take snapshot of the virtual machine. ESX Server/vCenter Server reported the following error: A general system Error occurred."

Solution

Due to the limitation of ESX snapshot cannot be taken when VM disconnects the attached ISO image attached. As a workaround, detach the ISO image from the CD/DVD device of the VM before backup.

Agentless Host-based Backup for Hyper-V VM Fails after Upgrading Arcserve UDP

Valid for Hyper-V

Symptom

After upgrading Arcserve UDP from Version 5.0 Update 2 or before to the latest version, the agentless host-based backup has started failing with the following error message:

The backup job is canceled. For a VSS snapshot, the Hyper-V VSS writer needs to save the virtual machine and this is not applied in the current plan. To restart the backup job, change the Hyper-V Snapshot Method setting in the plan. For details on how to set Hyper-V Snapshot Method in a plan, see the product documentation.

The agent-less host-based backup was working before the upgrade.

Solution

In Arcserve UDP Version 5.0 Update 2 or before, when the virtual does not support the online backup method, the default behavior is to adopt the offline backup method. The offline backup method saves the virtual machine while taking a snapshot. In the Saved state, the virtual machine is inaccessible. However, critical virtual machines need to be accessible all the time.

In Version 5.0 Update 3 and later versions, if the virtual machine needs to be placed into the Saved state, the default behavior is to cancel the backup job to avoid any downtime of the virtual machine. If you do not want the backup job to get canceled, change the **Hyper-V Snapshot Method** option in the plan. For more details about the Hyper-V Snapshot Method option in the plan, see How to Create a Host-Based Virtual Machine Backup Plan.

You can also refer to the Arcserve KB article for more details on this issue.

Host-based Agentless Backup Fails in VMware ESXi 6.0

Symptom

Arcserve UDP agentless backups may fail when you attempt to back up a virtual machine in VMware ESXi 6.0 and if the Change Block Tracking (CBT) function is enabled.

This is a known issue of VMware. When the backup fails, the following two behaviors can occur:

- Arcserve UDP may not connect to the CBT function of the ESXi host. As a result, Arcserve UDP cannot receive the used or changed data block information from the virtual machine.
- Arcserve UDP may fail to capture the quiesced snapshots of the virtual machine. (This can occur every time Arcserve UDP captures a snapshot or when you manually capture a snapshot in the vSphere client.)

Solution

VMware has resolved this issue in their latest build, ESXi 6.0 Build 2715440. You can install the ESXi600-201505001 patch to resolve this issue. For more information about downloading and installing the patch, see the VMware KB article.

If you cannot apply the patch, you can resolve the issue by making the following changes in the registry key:

Solution for CBT connection failure.

If Arcserve UDP cannot connect to CBT, then instead of failing the backup job, Arcserve UDP can continue the backup job. However, instead of performing an incremental backup, by default Arcserve UDP will perform a complete disk backup of the VM. If you do not want to perform a full backup automatically, you can add a registry key to change this default behavior. If you add the key and set the value to 1, then Arcserve UDP will fail the backup job when a CBT error occurs.

You can add this registry key in the Proxy server as follows:

At Proxy Server level (applicable for all backup jobs running in this proxy server)

1. Open the registry key from the following location:

   [HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll]

2. Enter the following dword:

   "BackupEntireDiskOnCBTBitmapFailure"=dword:00000001

3. Save the registry key.

At VM level

1. Open the registry key from the following location:

   [HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll\<vm instance uuid>]

2. Enter the following dword:

   "BackupEntireDiskOnCBTBitmapFailure"=dword:00000001

3. Save the registry key.

Note: If you add the registry key in both the VM and proxy level registry, then the setting in the VM level registry will have the priority over the setting in the Proxy level registry.

Solution for quiesced snapshot failure

Ensure that **Take snapshot without guest quiescence if quiescence snapshot fails** option is selected on the Source tab of the Host-Based Agentless backup plan.

Agentless Host-based Backup crashes when using Windows 2003 R2 64-bit as Backup Proxy

Agentless Host-based Backup for VMware VM crashes when using Windows 2003 R2 64-bit as backup proxy.

Valid for VMware

Symptom

When Windows 2003 R2 64-bit machine is used as the backup proxy server to protect VMware VM, sometimes backup job may crash. You can see error messages as following in the backup job debug log file:

[2016/01/21 10:18:11:316 00 03820 03336] [VDDKLOG] VixDiskLib: VixDiskLib_OpenEx: Open a disk. {AFBackend.exe::AFBackupVirtual.dll(1746.0)}
[2016/01/21 10:18:11:316 00 03820 03336] [VDDKLOG] VixDiskLibVim: VixDiskLibVim_GetNfcTicket: Get NFC ticket for [datastore1 (3)] VMname/VMware_1.vmdk. {AFBackend.exe::AFBackupVirtual.dll(1746.0)}
[2016/01/21 10:19:11:691 00 03820 03336] [VDDKLOG] VixDiskLibVim: Error 18000 (listener error GVmomiFaultInvalidResponse). {AFBackend.exe::AFBackupVirtual.dll(1746.0)}
[2016/01/21 10:19:11:691 00 03820 03336] [VDDKLOG] VixDiskLibVim: Login failure. Callback error 18000 at 2439. {AFBackend.exe::AFBackupVirtual.dll(1746.0)}
[2016/01/21 10:19:11:691 00 03820 03336] [VDDKLOG] VixDiskLibVim: Failed to find the VM. Error 18000 at 2511. {AFBackend.exe::AFBackupVirtual.dll(1746.0)}

Solution

In Arcserve UDP Version 7.0, VMware VDDK 6.x is built in. But VDDK 6.x does not officially support Windows 2003 R2. As a workaround, you can use one of the following options:

- Switch to use a proxy that is officially supported by VDDK 6.x. For example, proxy with Windows 2008 R2, Windows 2012 or Windows 2012 R2.
- Replace built-in VDDK 6.x by VDDK 5.5, which is also supported by UDP 7.0. For details on how to replace VDDK, refer to *How to apply different version of VDDK other than the built-in Version in Arcserve UDP.*

Host-based Agentless Backup Does Not Use HotAdd Transport Mode

Symptom

To back up data, host-based backup job does not use the HotAdd Transport Mode even when it is available. This happens when the source virtual machine is imported to Arcserve UDP Console from an ESX host (instead from the vCenter server) and the ESX host is managed by a vCenter server.

Solution

To resolve this error, perform one of the following tasks:

- Delete that virtual machine node from Arcserve UDP Console. Import the node again from the vCenter server that manages the ESX host.
- Disconnect the ESX from the vCenter server.

Host-based Agentless Backup or Restore Job Uses NBD or NBDSSL Transport Mode even when SAN Mode is Possible

Valid on Windows platforms. Valid for VMware VM only.

Symptom

Although SAN transport mode is possible, agentless backup or restore job still uses NBD or NBDSSL transport mode.

Solution

The following prerequisites must be completed to utilize the SAN transport mode by agentless backup and restore jobs.

- The proxy machine must be a physical machine (cannot be a virtual machine).
- The proxy machine must be connected to the SAN LUN in which the VM resides.
- On the proxy machine, SAN Policy of the SAN disk must be configured to OnlineAll.

  To configure the disk, do the following:

  1. Log in to the backup proxy system using an account with administrative privileges.
  2. Open Windows Command Line.
  3. From the command line, type the following commands
     a. Type "diskpart" and then press Enter.
     b. Type "SAN" and then press Enter.

        The current SAN policy displays.

     c. Type "SAN POLICY=OnlineAll" and then press Enter.
- If you want to perform a VM recovery by SAN transport mode, SAN disk must be configured to be writable.

  To clear the read only flag, performing the following steps:

  1. Log in to the backup proxy system using an account with administrative privileges.
  2. Open Windows Command Line.
  3. From the command line, type the following commands

a. Type "diskpart" and then press Enter.

b. Type "list disk" and then press Enter.

   The disks list displays.

c. Type "select disk xxx" and then press Enter to select the SAN disk that gets be configured to be writable.

d. Type "attribute disk clear readonly" and then press Enter.

- For VM recovery, SAN transport mode offers the best performance on thick disks but the worst performance on thin disks. So, the VM recovery uses NBD or NBDSSL transport mode for thin disks by default. If you want to use SAN transport mode even for thin disks, you can add a string value EnforceTransportForRecovery with the value SAN under HKLM\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFRestoreDll (create the key AFRestoreDll if it does not exist).
- When taking snapshot during backup, additional files are generated. So that certain free space is needed in storage device where VM's VMDK files reside. SAN transport mode needs more free space comparing with NBD/NBDSSL transport mode. So that make sure that there is enough free space in SAN LUN if you want to use SAN transport mode.

Failed to Create a Snapshot for Hyper-V Virtual Machines When Multiple Jobs are Running

Symptom

When running multiple jobs, the snapshot creation for Hyper-V CSV virtual machine takes too much time and then fails. It fails even after multiple tries. The following message is displayed in the activity log of the respective virtual machine.

The creation of snapshot is in progress, and only one snapshot creation operation can be in progress at one time.

Retry after 600 seconds.

Solution

The problem occurs because you can only run one snapshot creation at a time.

To resolve the issue, you can either increase the number of tries or increase the retry time interval. You can also increase the number of concurrent jobs that you can run.

Note: The default retry value is 3 and the default time interval value is 10 minutes.

To increase the number of retry, perform the following steps on the proxy server:

1. Open the Windows registry.
2. Navigate to HKLM\SOFTWARE\Arcserve\Unified Data Protection\Engine.
3. Create a key named as **VSSWrap**.
4. Right-click **VSSWrap**, select **New**, and then select **DWORD (32-bit)** value and specify the name as **VssAsynchMaxRetryTimes**.
5. Specify the value as required.

To increase the time interval between each retry, perform the following steps on the proxy server:

1. Open the Windows registry.
2. Navigate to HKLM\SOFTWARE\Arcserve\Unified Data Protection\Engine.
3. Create a key named as **VSSWrap**.
4. Right-click **VSSWrap**, select **New**, and then select **DWORD (32-bit)** value and specify the name as **VssAsynchRetryInterval**.
5. Specify the value as required.

To increase the number of concurrent jobs, perform the following steps on the proxy server:

1. Open the Windows registry.
2. Navigate to HKLM\SOFTWARE\Arcserve\Unified Data Protection\Engine.
3. Right-click **HyperVMaxJobNum,** select **Modify**, and specify the value as required.

Hyper-V failed to Create VSS Snapshot

The Hyper-V Host-Based VM backup fails on a Hyper-V host and displays the following message

Failed to take VSS snapshot.

Symptom

The following reasons are responsible for the backup failure:

- One or more volumes on the Hyper-V host is not formatted with NTFS/Refs.
- One or more volumes on the Hyper-V host has less than 100 MB of free space.
- Excessive disk activity during the time of the backup.

Solution

Resolve the environmental issues and perform the backup again.

Failed to Perform Backup of a Virtual Disk due to System error=[The device is not ready(21)]

Valid on Windows platforms.

Symptom

When a network error occurs or a Hyper-V server is rebooted while the backup is in progress, the activity log specifies that the error can be a network error or a file system error.

Solution

Restart the backup job again after the Hyper-V server restarts.

Backup Job Fails

Symptom

A backup job failed with the following error message in the activity logs:

Reconfiguration for backup cannot be performed in the current state. Shut down the virtual machine and attempt to run backup job again. (The virtual machine can be powered on during or after the taking snapshot phase).

Solution

Set the registry values to not reconfigure disk.enableUUID.

Follow these steps:

Applies at the proxy level and all the VMware VMs are impacted.

1. Log in to the Backup proxy server.
2. Open the registry editor and locate the following key:

   HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll
3. Add a DWORD value with name DoNotReconfigDiskUUID and specify 1 as its value.

Applies at the specific VM level and only the specified VM is impacted.

1. Log in to the Backup proxy server.
2. Open the registry editor and locate the following key:

   HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll\<VM-InstanceUUID>

   Note: Replace <VM-InstanceUUID> with the UUID value of the virtual machine which this setting applies for. You can find the value in the URL of the virtual machine that is used when connected to the Arcserve UDP Agent.
3. Add a DWORD value with name DoNotReconfigDiskUUID and specify 1 as its value.

Be aware of the following points:

- The VM level takes precedence if both VM and proxy level registries are configured.
- If the registry does not exist, the registry value implies 0, that is, you have to reconfigure disk.enableUUID.

- If you specified not to reconfigure the disk.EnableUUID parameter, the backed up data may not be in a consistent state.

For more information about this issue, see the VMware Knowledge Base article.

Failing to import VMware VMs from vCenter

Symptom

Arcserve UDP cannot import VMware VMs from the vCenter, although the vCenter server is functional and is able to connect with both the browser and vSphere client. In ARCAPP-Gateway.log of Arcserve UDP Console server, the error message appears is as follows:

com.sun.xml.ws.client.ClientTransportException: HTTP transport error: javax.net.ssl.SSLHandshakeException: java.security.cert.CertificateException: Certificates does not conform to algorithm constraints

But when you change the following two lines in the file "C:\Program Files\Arcserve\Unified Data Protection\Common\JRE\lib\security\java.security", and restart the Arcserve UDP Management service, it can connect to the same vCenter server by Arcserve UDP:

Existing lines:

jdk.certpath.disabledAlgorithms=MD2, MD5, RSA keySize < 1024

jdk.tls.disabledAlgorithms=SSLv3, RC4, MD5withRSA, DH keySize < 768

Modified Lines:

jdk.certpath.disabledAlgorithms=MD2, RSA keySize < 512

jdk.tls.disabledAlgorithms=SSLv3, RC4, DH keySize < 512

Cause:

vCenter Server has a certificate with a short public key length or its algorithm is disabled by the JRE in Arcserve UDP. A certificate with public key length less than 1024 bits is considered unsafe (same applies to the MD5 algorithm). They are disabled by the JRE used by Arcserve UDP.

Solution

Generate a new certificate for the vCenter Server. Ensure that the new certificate has a public key that is greater than 1024 bits in size and use a stronger algorithm.

Backups Fail Due to ESXi License

Valid on Windows platforms.

Symptom

The full, incremental, and verify backup jobs fail. The following message appears in the Arcserve UDP Activity Log:

VM server <server_name> does not have a paid ESX license

Solution

Due to a VMware limitation, virtual machines running on ESXi servers with a free license cannot be backed up. To protect these VMs, apply a purchased license.

HotAdd Transport Mode does not Work when Attempting to Back Up a VMware VM

Symptom

The HotAdd transport mode is not supported for this VM and as a result the backup is failing over to the NBDSSL (encrypted network block device) mode. (Backup job is running slow). The backup of a VMware VM is not using HotAdd Transport. For more details about HotAdd transport, see link.

Verify the following HotAdd prerequisites:

- The HotAdd backup proxy must be a virtual machine. HotAdd involves attaching a virtual disk to the backup proxy, like attaching disk to a virtual machine.
- The HotAdd proxy must have access to the same datastore as the target virtual machine.
- The VMFS version and data block sizes for the target VM must be the same as the datastore where the HotAdd proxy resides. If the HotAdd proxy is a virtual machine that resides on a VMFS-3 volume, choose a volume with block size appropriate for the maximum virtual disk size of virtual machines that customers want to back up, as shown in VMFS-3 Block Size for HotAdd Backup Proxy. This caveat does not apply to VMFS-5 volumes, which always have 1MB file block size.

  The following table displays the VMFS-3 Block Size for HotAdd Backup Proxy:

VMFS Block Size	Maximum Target Disk Size
1 MB	256 GB
2 MB	512 GB
4 MB	1024 GB
8 MB	2048 GB

- In vSphere 5.1 and older, the maximum supported VMDK size is 1.98 TB.
- The disks that are to be HotAdd must be SCSI. IDE drives are not compatible with HotAdd.
- VMware Tools must be installed and up-to-date on the VM and the backup proxy.
- Datastore needs sufficient space for a VM snapshot.
- HotAdd may fail if any disk was created with a newer hardware version than the VM being backed up. For example, if a disk was moved from a hardware ver-

sion 8 VM to a hardware version 7 VM. To resolve, upgrade the hardware version of the VM.

- A single SCSI controller can have a maximum of 15 disks attached. To run multiple concurrent jobs with more than 15 disks, you need to add more SCSI controllers to your backup proxy machine.
- In case of standalone ESX connection (ESX server is not managed by vCenter), you can only HotAdd disks of VMs which are located on the same ESX as the backup proxy machine.
- HotAdd may fail if you are trying to back up the VM through the ESX added as a standalone server into UDP but actually being managed by vCenter.
- Hot Add may fail if the VM you are trying to back up and the proxy server are in different clusters.

Solution

Disable “automount” on the backup proxy machine using “diskpart” utility.

More Granular Control for VSS over Windows Guest OS Quiescings

Symptom

I want to specify more granular control for VSS when creating quiescing snapshot for Windows guest OS.

Solution

From vSphere 6.5, vSphere web service API allows the following more granular controls for VSS when creating quiescing snapshot for Windows guest OS:

- You can configure anywhere the timeout (default 15 minutes) for quiescing virtual machines from five minutes to four hours.
- VSS backup type – VSS_BT_COPY was previously used as the default when creating snapshot. But now VSS_BT_FULL, VSS_BT_INCREMENTAL, VSS_BT_DIFFERENTIAL, and VSS_BT_LOG are available also. Log truncation is triggered according to application settings.
- VSS backup context is introduced to enforce application (context VSS_CTX_BACKUP) quiescing or file system (context VSS_CTX_FILE_SHARE_BACKUP) quiescing.

Using Arcserve UDP, you may specify the parameters in registry to implement the control.

Note: To specify, you require VMware Tools 10.1.0 or higher installed on the guest OS of VM

Follow these steps:

1. Log onto proxy machine.
2. Create a registry value at proxy server level or VM level.

   Note: If you add the registry value at both the VM level and proxy level, the setting at the VM level registry holds priority over the setting at the Proxy level registry.

   At Proxy Server level (applicable to all backup jobs running in this proxy server)

   a. Open the registry key from the following location:

   [HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll]

   b. Add following DWORD values with proper values:

- VssUseEnhancedSnapshot
- VssTimeoutMinute
- VssBackupType
- VssBackupContext
- VssBootableSystemState
- VssPartialFileSupport

At VM Level

a. Open the registry key from the following location:

[HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll\<vm instance uuid>]

b. Add following DWORD values with proper values:

- VssUseEnhancedSnapshot
- VssTimeoutMinute
- VssBackupType
- VssBackupContext
- VssBootableSystemState
- VssPartialFileSupport

Possible values of the registry values:

VssUseEnhancedSnapshot

- 0 – Do not use enhanced control and the lower registry values are not effective
- 1 - Use enhanced control and the lower registry values are effective

VssTimeoutMinute

Range from 5 to 240

VssBackupType

- 0 - VSS_BT_COPY (default)
- 1 - VSS_BT_FULL
- 2 - VSS_BT_INCREMENTAL
- 3 - VSS_BT_DIFFERENTIAL
- 4 - VSS_BT_LOG

VssBackupContext

- 0 - ctx_auto
- 1 - ctx_backup (default)
- 2 - ctx_file_share_backup

VssBootableSystemState

- 0 – false
- 1 - true (default)

VssPartialFileSupport

- 0 - false (default)
- 1 - true

Unable to configure the Virtual Machine to Enable the "Disk.EnableUUID" Parameter

Symptom

An error occurs during the reconfiguration of the virtual machine to enable the 'disk.EnableUUID' parameter, required for an application-consistent backup. This issue occurs when the virtual environment has recovered from an error.

Solution

This issue is specific to VMware Environment. This issue occurs when the Microsoft VSS attempts to perform an application-quiesced snapshot, but fails to read the disk UUID or Serial Number value for one or more virtual machine disks.

The disk.EnableUUID parameter of the VM must be enabled. VMs created on 4.1 or later have this parameter enabled by default. The backup job automatically configures to avoid data inconsistency and perform application-consistent backup. If the backup job fails to enable disk.EnableUUID, manually configure the parameter using the following steps:

Follow these steps:

1. Power off the virtual machine.
2. Log in to the vCenter Server or the ESXi/ESX host through the vSphere Client.
3. Right-click the virtual machine and click Edit settings.
4. Click the Options tab.
5. Navigate to Advanced > General > Configuration Parameters.
6. Add or modify the row disk.EnableUUID and set it to True.
7. Click OK to save.

8. Click OK to exit.
9. Reboot the virtual machine for changes to take in effect.

   Note: If the parameter is already set to True There is a non disruptive workaround, and that is to edit the vmx and set the disk.enableUUID value to false, then vMotion the vm (to any other host), to reload the vmx in host memory. This effectively disables application level quiescing (file system quiescing is still available though).

Recovery Operations Fail When Recovering Data Using the HOTADD or SAN Transport Mode

Valid on Windows platforms.

Symptom

Recovery operations fail when recovering data using the HOTADD or SAN transport mode. The following message appears in the Activity Log:

An unknown error has occurred. Contact Technical Support.

Solution

Recovery operations fail using the HOTADD transport mode or SAN transport mode when the disk settings are not configured properly.

To configure the disk, perform the following steps:

1. Log in to the backup proxy system using an account with administrative privileges.
2. Open Windows Command Line.
3. From the command line, type the following command:

   diskpart

   Press Enter.
4. Type SAN and then press Enter.

   The current SAN policy displays.
5. Type the following command, and press Enter:

   SAN POLICY = OnlineAll

   The SAN policy is configured as do not automatically mount SAN hosted volumes.
6. To clear the read-only attribute of the specific SAN disk, select the disk from the disk list, type the following command, and press Enter:

   attribute disk clear readonly
7. Type exit and then press Enter.

   The disk is configured and you can resubmit the job.

   If the job fails again, mount the HOTADD disks manually using disk management on the proxy system.

To mount the disks manually, perform the following steps:

1. Log in to the backup proxy system using an account with administrative privileges.
2. Open Windows Control Panel and double-click Administrative Tools.

The Administrative Tools window opens.

3. From the Favorites list, double-click Computer Management.

   The Computer Management opens.

4. Expand Storage and click Disk Management.

   The disks display.

5. Right-click the disk that you want to mount and click Online.

   The disk is mounted and you can resubmit the job.

Recover VM Operation fails when a non-default port is specified

Symptom

Recover VM Operation fails when a non-default port is specified for VMware vCenter Server.

Solution

To resolve, set non-default port number of vCenter to the value of VDDKport in registry of backup proxy machine.

To set the number of VDDKport, perform the following steps on the proxy server:

1. Open the Windows registry.
2. Navigate to HKLM\SOFTWARE\Arcserve\Unified Data Protection\Engine.
3. Right-click VDDKport, select Modify, and specify the value as required.

Scheduled Incremental or Full Backup Job Fails for Hyper-V VM

Symptom

Sometimes the scheduled incremental or full backup job fails for Hyper-V virtual machines and the following errors are displayed in the event viewer on Hyper-V host:

- DM operation add for the virtual machine <vm name> failed with error: Ran out of memory (0x8007000E) (Virtual machine ID <vm ID>)
- Could not create backup checkpoint for virtual machine <vm name>: This operation returned because the timeout period expired. (0x800705B4). (Virtual machine ID <vm ID>)
- Could not create backup checkpoint for virtual machine <vm name>: Element not found. (0x80070490). (Virtual machine ID <vm ID>)
- VSS writers inside virtual machine <vm name> failed to perform BackupComplete to its shadow copy (VSS snapshot) set: A function call was made when the object was in an incorrect state for that function (0x80042301). (Virtual machine ID)
- The Hyper-V VSS writer has encountered an error when processing this virtual machine. (For more information about Hyper-V VSS writer errors, refer to the product documentation).

Solution 1

The resolution is to increase the RAM size on the Hyper-V server and then resubmit the backup job.

Solution 2

If the VSS writer inside the virtual machine does not work properly, then the backup job fails. To resolve the issue, check the event log of both Hyper-V host and the virtual machine. Check the VSS warnings and errors, and take appropriate actions.

Hyper-V VSS NTDS Writer Fails While Taking the VSS snapshot in the VM

Symptom

In a Domain Controller VM, if the AutoMount feature is not enabled, the VSS NTDS writer fails while taking the VSS snapshot in the VM. As a result, the Hyper-V VSS writer fails to take the VSS snapshot on the Hyper-V host.

The Hyper-V HBBU backup job fails with the following activity log:

The Hyper-V VSS writer has encountered an error when processing this virtual machine. (For more information about Hyper-V VSS writer errors, refer to the product documentation).

Solution

Enable the *AutoMount* feature in the VM.

Follow these steps:

1. Open the command prompt window.
2. Open diskpart and execute the following command:

   automount enable

MAC Address Changes are Not Retained After VM Recovery

Valid on Windows platforms and VMware VM

Symptom

The MAC addresses of virtual machines are not retained after recovering virtual machines.

Solution

MAC addresses are not retained during recovery, to prevent duplicates. To retain MAC address information, set the following registry key on the proxy server:

Location: HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine
Key Name: RetainMACForVDDK
Value Type: String

Key Value: 1

VM Recovery Fails with Error - Unable to Open VMDK Files

Symptom

The virtual machine recovery fails and the following error message is displayed in the logs:

Unable to open vmdk file (file name with .vmdk). VMware reported the following error: You do not have access rights to this file. For more information, see the restore debug log. If necessary contact Arcserve support.

The following messages are found in the restore debug logs:

- [VDDKLOG] CnxAuthdConnect: Returning false because SSL verification requested and target authd does not support SSL
- [VDDKLOG] CnxConnectAuthd: Returning false because CnxAuthdConnect failed
- [VDDKLOG] Cnx_Connect: Returning false because CnxConnectAuthd failed
- [VDDKLOG] Cnx_Connect: Error message: SSL required

Solution

The reason could be that the SSL authentication is disabled on the ESX host. To resolve this issue, use one of the following methods:

Using the vSphere Client

1. Log into the vCenter/ESX server.
2. Navigate to the ESX server settings:

   Configuration, Advanced settings, Configuration, Defaults security

3. Enable the following option:

   config.defaults.security.host.ruissl

Using the command line

1. Connect to the ESX host using SSH.
2. Open the following file:

   /etc/vmware/config

3. Set the security.host.ruissl entry to TRUE.
4. Save the file and restart the management agents.

Problems Caused by Duplicate VM UUID

Symptom 1

VM node is overwritten after importing another VM node into Console.

For example, you have 2 virtual machines VM1 and VM2 with the same UUID (which is called Instance UUID for VMware, and VM UUID for Hyper-V) in the ESXi hosts which are managed by different vCenter VC1 and VC2. You import VM1 into console and it shows up in the node list view of console. Later on, you import VM2 into console. In the node list view VM1 is overwritten by VM2 (in other words, VM2 is added but VM1 is gone).

Symptom 2

VM node's information in the Hypervisor column changes back and forth while auto discovery is running.

For example, you have 2 virtual machines VM1 and VM2 with the same UUID (which is called Instance UUID for VMware, and VM UUID for Hyper-V) in the ESXi hosts which are managed by different vCenter VC1 and VC2. You import VM1 into console. And you also import at least one VM from the vCenter VC2 so that both VC1 and VC2 are added into the Node Discovery List (you can check the list at Node Discovery Configuration page which is at Settings tab). While node discovery runs, it firstly connects to VC1 and detects VM1 by the UUID, so that Hypervisor column is updated with information of VC1. But later on, when it connects to VC2, it detects the VM2 by the same UUID, so that Hypervisor column is updated with information of VC2.

Solution

Arcserve UDP uses VM's UUID (which is called Instance UUID for VMware, and VM UUID for Hyper-V) to identify a VM node. Although it is very rare that VMs have duplicated UUID, problematic behaviors may happen in Arcserve UDP in such case.

To solve the problems, refer the following steps to manually change the UUID of the VM (applicable for VMware VM only). After you have manually changed the UUID of the VM, delete the original VM from the Arcserve UDP Console and re-import the VM.

1. Open the following URL:

   https://<vCenter host name>/mob/

2. Log in as an administrator.
3. Search "content" in the NAME column, and click the link in the VALUE column in the same row.

4. Search “rootFolder” in NAME column, and click the link in the VALUE column in the same row.
5. Search “childEntity” in NAME column. In the VALUE column in the same row, find the Datacenter in which the VM reside and click on its link.
6. Search “vmFolder” in NAME column, and click the link in the VALUE column in the same row.
7. Search “childEntity” in NAME column. In VALUE column in the same row, click on “more…” to expand the VM list. Search the VM you are looking for and click the link.
8. Search ReconfigVM_Task in the Methods table, and click the link.
9. In the new browser opened, remove all contents in the VALUE field, and type the following code:

```
<spec>
<instanceUuid>2499952a-6c85-480e-b7df-4cb-
d2137eb69</instanceUuid>
</spec>
```

   Note: The 2499952a-6c85-480e-b7df-4cbd2137eb69 string mentioned above is a sample UUID. You should replace it with the UUID that you want to apply.

10. Click on Invoke Method link to apply the new UUID.
11. To verify that new UUID is applied, close the newly opened browser and go back to the page where you did step 8.
12. Search “config” in NAME column and click on the link in the VALUE column in the same row.
13. Search “instanceUuid” in NAME column. The UUID of the VM is shown in the VALUE column in the same row.

File System Catalog Job Fails or Recovery Point Check Fails for Host-Based Agentless Backup

Valid on Windows platforms. Valid for VMware VM only.

Symptom

- File system catalog job fails for host-based agentless backup recovery points
- Recovery Point Check fails during host-based agentless backup job and the next Incremental backup gets converted to a Verify backup.

Solution

This may be caused by a known issue of VMware (see the VMware KB article). When you quiesce a VMware VM, the snapshot contains corrupted data. The backup reads data from the snapshot and the data that is backed up also becomes corrupted.

Note: This problem occurs with all VMware ESXi versions and on a VM with guest OS running Windows 2008 R2 SP1 and Windows 2012. Arcserve UDP cannot detect the data corruption problem because VMware does not return an error in such cases. You may not be aware of the problem until you try to restore data.

You can follow the Arcserve KB article to check whether the problem is caused by this VMware known issue. The workaround recommended by VMware is to disable the VSS writers, such as MSSearch Service Writer (ignore it, if not installed) and Shadow Copy Optimization Writer (typically present in every Windows VM), in guest OS of VM. You can manually disable the writers per the VMware KB article.

Arcserve UDP also provides a simple way to disable the writers, if VMware Tools snapshot quiescing method is used. Follow these steps to disable the writers.

At Proxy Server level (applicable for VMs protected by this proxy server)

1. Log in to the proxy server.
2. Open the registry key from the following location:

   ```
   [HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll]
   ```

3. Create a Multiple-String value with the name *DisableSpecificVSSwriters*.
4. Enter the names of the VSS writer that is expected to be disabled (each writer name occupies one line).
5. Save the registry key.

At VM level

1. Log in to the proxy server.
2. Open the registry key from the following location:

   ```
   [HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll\<vm instance uuid>]
   ```

3. Create a Multiple-String value with the name *DisableSpecificVSSwriters*.
4. Enter the names of the VSS writer that is expected to be disabled (each writer name occupies one line).
5. Save the registry key.

   Notes:

   - If you add the registry key in the VM and proxy level registry both, then the setting in the VM level registry will have the priority over the setting in the Proxy level registry.
   - This registry setting works only when the VMware Tools snapshot quiescing method is used in the backup plan.
   - If you have ever configured writers manually in the guest OS of the VM per the VMware KB article, your configuration will be overwritten.
   - The name of the writer is case sensitive and must be an exact match as shown in the output of “vssadmin list writers”.
   - If you want to enable all writers again, do not delete the registry value *DisableSpecificVSSwriters*. Instead, retain the registry value but remove the content within the registry. If *DisableSpecificVSSwriters* does not exist, Arcserve UDP assumes that nothing needs to be changed in terms of enabling or disabling VSS writers.

Incremental Backup Converts to Verify Backup or the Backup Size Increases in Hyper-V

Valid on Hyper-V VM

Symptom

- I have performed an incremental change in a Hyper-V virtual machine. When I perform an incremental backup, the entire virtual machine is backed up instead of backing up only the changed data.
- I have a proxy server with Arcserve UDP Version 7.0 that backs up a virtual machine from one Hyper-V host (example HOST1). I have another proxy server with an older version of Arcserve UDP that backs up a virtual machine from the same Hyper-V host (HOST1). In such cases, the CBT is inactive and the Incremental jobs do not run. The Incremental backup converts to the Verify backup.

Solution

The root causes of above symptoms can be one of the following reasons:

- The loss of change block tracking (CBT) data. The following circumstances will result in CBT data loss:
  - The Hyper-V host crashes or is powered off abnormally.
  - The CBT service is stopped or the service abnormally quits.
  - The CBT service did not complete its work while the Hyper-V host was shutting down.
- Different versions of CBT in the Hyper-V server and the proxy server.

  Example: Consider you have two Arcserve UDP environments, one is Arcserve UDP Version 6.0 and another is Arcserve UDP Version 7.0. These two Arcserve UDP environments back up different VMs in the same Hyper-V server. The Arcserve UDP Version 7.0 environment automatically detects the older version of CBT in the Hyper-V server and upgrades it to the latest version. In such cases, the Arcserve UDP Version 5 environment converts the remaining scheduled incremental backup to a full backup.

  If Arcserve UDP detects different CBT versions, the Activity Log displays a warning message.

The solution is to upgrade all the proxy servers that protects virtual machines from one Hyper-V host to the same version of Arcserve UDP.

Host-based Backup Fails for Hyper-V VM That Has a Special Differencing Disk Configuration

Valid for Hyper-V VM

Symptom

If a differencing disk is configured in a Hyper-V virtual machine, the backup job for that virtual machine fails. It displays the following error message in the activity log:

Failed to prepare for backup of the virtual machine

The following error message is displayed in the backup job log file under C:\Program Files\Arcserve\Unified Data Protection\Engine\Logs,

The virtual disk file \\?\UNC\<IP_Address_VM>\HYPERV_HBBU_SNAPSHOT@<snapshot_name>\WIN12-SQL\VIRTUAL HARD DISKS\WIN12-SQL-1.VHDX was not exposed.

The problem occurs only when the virtual machine has the following differencing disk configurations. All the configurations must apply.

- The virtual machine has one regular virtual hard disk (Fixed size or Dynamically expanding) Disk1 that is attached to one IDE or SCSI controller of the virtual machine.
- The virtual machine has one differencing virtual hard disk (Disk2) that is also attached to one IDE or SCSI controller of the virtual machine.
- The parent disk of Disk2 is specified to Disk1.

Solution

This error occurs because of an abnormal or incorrect configuration. To resolve this error, detach either the differencing disk or its parent from the virtual machine. Arcserve UDP does not support such differencing disk configuration.

Backup Job Fails for a VMware Virtual Machine

Valid for VMware VM

Symptom

When I back up a VMware virtual machine, the backup job fails with either of the following error messages in the activity log:

Abort backup because backup job has been configured to use the "Microsoft VSS inside VM" snapshot method. However, only the "VMware Tools" snapshot method is applicable because Host-based VM Backup failed to deploy the necessary tools into the VM.

Or

Abort backup because backup job has been configured to use the "VMware Tools" snapshot method. However, only the "Microsoft VSS inside the VM" snapshot method is applicable because Host-based VM Backup failed to undeploy tools from inside VM.

Solution

The first error can occur because of multiple reasons. You have selected the **Microsoft VSS inside VM** option but:

- You did not update the VM with the required credentials
- The credentials are not correct
- VMware Tools are not installed or updated.

In this case, Arcserve UDP cannot deploy necessary tools to the virtual machine to use the new snapshot method.

To resolve this error, update the virtual machine with correct credentials. Verify that VMware Tools are updated and running in the virtual machine. After the verification, resubmit the backup job.

Solution

The second error may occur in the following scenario. You have used the **Microsoft VSS inside VM** option in the previous backup jobs. Now, you want to use the **VMware Tools** option but the credentials of the virtual machine have changed (for example, you have changed the password of the guest OS but you did not update the virtual machine node in Console), or VMware Tools is not running for some reason. In such cases, Arcserve UDP cannot undeploy the tools (which were deployed by the previous backup job) from the virtual machine to use the new snapshot method.

To resolve this error, perform one of the following steps:

- Update the virtual machine with correct credentials. Verify that VMware Tools are updated and running in the virtual machine guest OS. After the verification, resubmit the backup job.
- Manually undeploy the tools from the virtual machine:
  a. Log in to the virtual machine.
  b. Navigate to the following folder:

     C:\Program Files\ARCServe\ASVMOperationTools\custom-freeze-vmware-snapshot\auto-deploy
  c. Right-click the auto-undeploy.bat batch file and select Run as administrator.
  d. Delete the following folders:

     C:\Program Files\ARCServe\as-hbbu-vmwarebackup

     C:\Program Files\ARCServe\ASVMOperationTools
  e. Resubmit the backup job.

Disable Re-scan of HBA Adapters During Incremental Backup

Applicable to VMware ESX

Applicable only to UDP v6.5, does not apply to v6.5 Updates.

Symptom

When the source node and the proxy server are in different VMware ESX servers, and I run an Incremental backup, my backup takes a longer time to run. I also receive multiple rescan messages. I want to disable the rescanning of all host bus adapter (hba) of other VMware ESX servers during Incremental backups.

Solution

You can disable the scan by creating a registry key in the proxy server and assigning a value to it.

Follow these steps:

1. Log in to the machine.
2. Navigate to the following folder:

   ```
   HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine
   ```

3. Create the following registry key (DWORD), if it is not already created

   ```
   DisableAllESXNodeRescan
   ```

4. Set the value of the registry key to 1.

Disable the Consecutive Snapshot Creation in VMware VM for a Backup

In the previous version of Arcserve UDP, when multiple VMware VM backup jobs start at the same time, the jobs create snapshots on the ESX host in parallel to jobs. Sometimes such parallel snapshot creation operations increase the disk I/O on the ESX host. To avoid such situation, by default Arcserve UDP Version 7.0 serialize the snapshot creation operation when the jobs are running on the same proxy server and the snapshot creation operation is for the same ESX host. In other words, the snapshot creation operation occurs one after another and ideally there is only one snapshot created at a time. But this is not applicable if the backup jobs are running on different proxy servers, or the backup jobs are for different ESX hosts.

To avoid a situation where the snapshot creation of one job hangs (or takes very long time) and the next jobs are blocked, the next job will wait for the previous snapshot creation operation for at most **five minutes**. After five minutes, the next snapshot creation process starts.

You can disable the behavior of waiting for five minutes by following these steps:

1. Log in to the proxy server.
2. Open the registry key from the following location:

   [HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll]

3. Create a DWORD value with the name **CreateSnapshotSequentially**.
4. Set the value to 0.

In addition, you can also change the default timeout value (five minutes) by following these steps:

1. Log in to the proxy server.
2. Open the registry key from the following location:

   [HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll]

3. Create a DWORD value with the name **CreateSnapshotTimeout**.
4. Set the value with a number.

   The unit is in seconds.

When Restored from a Higher Version of ESXi host to a Lower version ESXi host, VM gets stuck at booting stage

When restored from a higher version of ESXi host to a lower version ESXi host, VM gets stuck at booting stage after it is powered on.

Valid on Windows platforms. Valid for VMware VM only.

Symptom

Restore a VM from a higher version ESXi host to a low version ESXi host, power on the VM. The VM gets stuck at the booting stage.

Solution

The guest OS version of this VM may not be supported by the lower version ESXi. As a workaround, you can use one of the following options:

- Restore the VM to a ESXi which supports that guest OS version or
- Upgrade the existing ESXi host to the suitable version.

For example, a Windows Server 2012 R2 VM is backed up from an ESXi 5.5 and restored to ESXi 5.0 update 1. As ESXi 5.0 starts to support Windows Server 2012 R2 guest OS after Update 2, this issue may occur.

You need to upgrade ESXi 5.0 Update 1 to Update 2 so that Windows Server 2012 R2 guest OS is supported. As a workaround, you can change the guest OS version of restored VM to the version that is supported by the current version of ESXi (in above example, change to Windows Sever 2008 R2").

Note: This workaround may not resolve the issue.

RAM Utilization reaches 99% When Backup Jobs are Submitted to VM

RAM utilization reaches 99% when backup jobs are submitted to VMs on Windows 2012 Hyper-V CSV.

Symptom

During backup of VMs that are part of 2012 Hyper-V cluster RAM utilization gradually increases and reaches to 99% on CSV owner Hyper-V host. As a result, Virtual Machines and Hyper-V host stop working during backup.

Solution

This issue is caused due to known bugs in the Microsoft code.

To fix this problem, download and apply hotfix provided by Microsoft on all the Hyper-V host from the following link:

https://support.microsoft.com/en-in/kb/2878635

Hyper-V Restore Job Fails, Cannot Connect to Utility on Host

Hyper-V VM restore job fails and an error message appears.

Symptom

When restoring a Hyper-V VM, restore job fails and the following error appears in activity log:

Failed to connect to the Hyper-V restore utility on host xxxxx

In the restore job debug log, the following error messages appear:

Failed to connect to xxxxx:10218. error:Attempt to connect timed out without establishing a connection (rc=-536805332)

Failed to connect to xxxxx:10218, error -536805332

Solution

Normally, such problem is caused as the connection between backup proxy and Hyper-V host is blocked by firewall. If possible, turn off the firewall and try restore again. If you cannot turn off the firewall, register the port (or port range) used by restore job to the firewall exception.

By default, the restore job randomly chooses an available port in the range of 1024 and 65535. You can manually specify the range by registry values set in Hyper-V host.

Follow these steps:

1. Log into Hyper-V host.
2. Run the command *regedit* to open registry editor.
3. Navigate to the following registry key (create keys if any of them does not exist):

   HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\HyperVRestoreStub
4. Create the following two registry values (DWORD):

   PortMin

   PortMax
5. Specify values with beginning and end number of the port range.

   Notes:

- Mandatory Range of specified values: 1024-65535
- You can specify the same number for both registry values. In this case, the fixed port number is used for the restore job. When multiple restore jobs run simultaneously and as soon as one restore job occupies this port, then other restore jobs fail.

Automatic Protection Fails to Detect and Protect VM

Virtual machine is not detected and automatically protected as part of automatic protection.

Symptom

Although a VM can be seen in hypervisor (vCenter/ESX or Hyper-V) by native client application (for example, vSphere Client or Hyper-V Manager), automatic protection does not detect and as a result the VM cannot be automatically protected.

Reasons

- The VM does not have a valid status in hypervisor. For example, Automatic Protection skips the VMs that have status of Disconnected, Orphaned or Inaccessible in vCenter/ESX .
- The VM is created by Arcserve UDP tasks such as Virtual Standby (VSB), Instant VM (IVM) or Assured Recovery (AR) (or is cloned from VSB/IVM/AR VM). Arcserve UDP intentionally skips such VMs due to following considerations:
  - Arcserve UDP backing up VMs created by VSB/IVM/AR, without limitation may lead to the situation of *infinite backup loop*. For example, you use automatic protection to protect the whole ESX. In the backup plan, after adding a backup task you add a VSB task that creates VSB VM in the same ESX. Thus, after first backup a new VSB VM is created in the ESX. The new VSB VM is discovered by automatic protection and added into backup plan. Then, during next backup a new VSB VM of previous VSB VM is created and added into the backup plan. The process continues until the ESX storage runs out of free space.
  - For Linux IVM/AR VM, unless storage migration converts to a normal VM, Arcserve UDP cannot back up the VM. The backup job can complete but backed up recovery point cannot be restored. This situation happens due to a technical limitation of Linux IVM implementation.

Solution

If you want to back up VSB/IVM/AR VMs, or the VMs are converted to independent VMs by clone or storage migration, follow the steps given below for vSphere VM and Hyper-V VM to let Automatic Protection detect and protect them.

For vSphere VM

1. Log into vSphere web client and locate the VM.
2. Power off the VM.

3. Right click on the VM and select Edit Settings from menu.
4. Select the **VM Options** tab and expand Advanced.
5. Click **Edit Configuration**.
6. On the Configuration Parameters dialog box, locate parameter with any of following names:
    - VCM
    - UDP_IVM
    - UDP_ARVM
    - UDP_IVM_LINUX
    - UDP_VSBVM
    - UDP_ARVM_LINUX
    - UDP_ARIVM_LINUX
7. Clear value of the parameter, and click **OK** to save.

For Hyper-V VM

1. Log into Hyper-V host and open Hyper-V Manager.
2. Locate the VM.
3. Right click on the VM and select Settings from menu.
4. On the Settings dialog box, select Name in the left panel.
5. Remove all lines that start with following strings:
    - VCM
    - UDP_IVM
    - UDP_ARVM
    - UDP_IVM_LINUX
    - UDP_VSBVM
    - UDP_ARVM_LINUX
    - UDP_ARIVM_LINUX
6. Click **OK** to save.

Set Read block size when backing up VMDK file

Symptom

In certain environment, the backup throughput of VMware VM gets affected by read block size when backing up VMDK file.

Solution

By default, the VMDK read block size is 2 MB. You can modify the size.

Follow these steps:

1. Log into Agentless backup proxy.
2. Run the command *regedit* to open registry editor.
3. Navigate to the following registry key:

   HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll
4. Create the registry values (DWORD) using the following name:

   ReadVirtualDiskBlockSizeKB
5. Specify desired value as the read block size (unit is KB).

Lun Space Reservation Inherited while Performing Lun Clone

Symptom

When Lun clone is performed during backup using hardware snapshot, by default space reservation is inherited from source Lun.

Solution

You must use the registry key that Arcserve UDP provides to disable space reservation. The registry key is at the following location:

HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine

DisableLunSpaceReservation =1

Specifies that the Lun space reservations are disabled.

Note: This is applicable only when FlexClone license is applied.

For more information, view Considerations for NetApp iSCSI/FC Support for VMware.

Virtual Standby Related

This section includes the following troubleshooting topics related to virtual standby:

- Operating System Not Found
- Virtual Standby Jobs Fail Due to Internal Errors
- Virtual Standby Jobs Fail Using the HotAdd Transport Mode
- Virtual Standby Jobs to Hyper-V Systems Fail
- Issue Related to Duplicate Agent UUID
- Terminal EC2 Resources option is not shown

Operating System Not Found

Valid on Windows platforms.

Symptom:

The following message appears when the power on Virtual Standby virtual machine operation fails:

Operating System Not Found.

Solution:

The above behavior can occur on virtual machines that contain SCSI and IDE devices. If this problem occurs, examine how disks are configured on your virtual machine and verify that the boot sequence of the recovered virtual machine is the same as the source virtual machine. If the boot sequence is different, update the BIOS on the recovered virtual machine to match that of the source.

Note: Use (0:1) to represent the first IDE disk.

Virtual Standby Jobs Fail Due to Internal Errors

Valid on Windows operating systems.

Symptom 1:

Virtual standby jobs fail. One of the following messages appears in the Activity Log:

Failed to convert virtual disk

An internal error occurred, contact technical support

In addition, VDDK reports the following error message:

Unknown Error.

Solution 1:

To correct this problem, consider the following solutions:

- Conversion operations can fail when there is not enough free disk space on the data store that is specified in the Virtual Standby policy. VDDK returns the message because the VDDK API (currently) does not support the capability to detect the amount of free disk space on the data store. To correct this problem, free the amount of disk space on the original data store that is required to complete the operation and then resubmit the job.
- Network disturbance and high network traffic can cause the conversion operations to fail. To correct this problem, verify that source node and the ESX Server system or the vCenter Server system can communicate with each other though the network, and then resubmit the job.
- Multiple concurrent connections consisting of backup or recover VM jobs to the ESX Server system or the vCenter Server system, which includes vSphere SDK connections through the VMware vSphere Client, can cause the jobs to fail. To correct this problem, close all unnecessary connections and then resubmit the job.

  This problem is the result of a VMware VDDK connection limitation. The following Network File Copy (NFC) protocol limits apply:

  ESXi 5: Limited by a transfer buffer for all NFC connections and enforced by the host; the sum of all NFC connection buffers to an ESXi host cannot exceed 32MB. 52 connections through vCenter Server which includes the per-host limit.

  Note: Connections cannot be shared across disks. The maximum limits do not apply to SAN or HotAdd connections. If the NFC client fails to shut down properly, connections can remain open for ten minutes.

- Examine the Tasks and Events sections of the VMware vSphere Client log to discover internal errors for the specific virtual machine. Correct the internal errors and then resubmit the job.

  Example: Another application or operation is using the VMDK file. To correct this problem, release the file and then resubmit the job.

Symptom 2:

Virtual standby jobs fail. One of the following messages appears in the Activity Log:

Failed to convert virtual disk

An internal error occurred, contact technical support

In addition, VDDK reports the following error message:

Open vmdk failed with error File not found.

Solution 2:

This problem can occur when:

- VDDK did not process a snapshot properly.
- VDDK did not delete a snapshot manually or internal to the virtual machine.

To correct this problem, resubmit the job. If the job fails again, delete the recovered virtual machine and resubmit the job.

Symptom 3:

Virtual standby jobs fail. One of the following messages appears in the Activity Log:

Unable to apply '<Plan Name>' to node '<Node Name>'. The Arcserve UDP Agent web service on the converter "<Converter Name>" is busy. Retry later.

In addition, UDP Console log file (ARCApp.log) reports the following error message:

[ERROR] deployVsbTask: Failed to invoke D2D web service API - timeout. javax.xml.ws.WebServiceException: java.net.SocketTimeoutException: Read timed out

Solution 3:

This problem can occur due to timeout. To correct this problem, perform the following steps:

1. Login to **UDP Console** with appropriate credentials.
2. Open command line interface and run the below command:

   regedit

   The Registry opens.
3. Navigate to ***\SOFTWARE\Arcserve\Unified Data Protection\Engine\WebService***.

4. Verify that the **timeoutValue** key exist. If the key does not exist, create manually.

   Add/modify key as mentioned below:

   - **Key Name:** timeoutValue
   - **Value:** <enter the value in minutes>. For example, if you want to set the timeout value as 20 minutes specify 20 as the value.

5. Exit **regedit**.
6. Navigate to the **UDP Console** installation folder. For example, *C:\Program Files\Arcserve\Unified Data Protection\Management\Configuration*.
7. Open the file **ConsoleConfiguration.xml** using notepad.
8. Find the below text under the section **<TimeoutConf>**:

   <webServiceRequestTimeout>600</webServiceRequestTimeout>

9. Modify the value of **webServiceRequestTimeout** in seconds. For example, if you want to set the timeout value as 20 minutes specify 1200 as the value..
10. **Save** the file and exit.
11. **Restart** the **UDP Console management service** to allow the settings to take effect.
12. Redeploy the plan and check the result.

Virtual Standby Jobs Fail Using the HotAdd Transport Mode

Valid on Windows platforms.

Symptom:

Recovery operations fail when recovering data using the HotAdd transport mode. The following message appears in the Activity Log:

An unknown error has occurred. Contact technical support.

In addition, VDDK reports the following error message:

Unknown Error.

Solution:

Recovery operations fail using the HotAdd transport mode when the disk settings are not configured properly.

To configure the disk, follow these steps:

1. Log in to the backup proxy system using an account with administrative privileges.

   Open Windows Command Line.

2. From the command line, type the following command:

   diskpart

   Press Enter.

   Type SAN and then press Enter.

   The current SAN policy displays.

3. Type the following command:

   SAN POLICY = OnlineAll

   Press Enter.

   The SAN policy is configured as do not automatically mount SAN hosted volumes.

4. To clear the read only attribute of the specific SAN disk, select the disk from the disk list and type the following command:

   attribute disk clear readonly

   Press Enter

5. Type exit and then press Enter.

   The disk is configured and you can resubmit the job. If the job fails again, mount the HotAdd disks manually using disk management on the proxy system.

   To mount the disks manually, follow these steps:

1. Log in to the backup proxy system using an account with administrative privileges.

   Open Windows Control Panel and double-click Administrative Tools.

   The Administrative Tools window opens.
2. From the Favorites list, double-click Computer Management.

   The Computer Management opens.
3. Expand Storage and click Disk Management.

   The disks display.
4. Right-click the disk that you want to mount and click Online.

   The disk is mounted and you can resubmit the job.

Virtual Standby Jobs to Hyper-V Systems Fail

Valid on Windows operating systems.

Symptom:

The following message appears in the Activity Log:

Virtual Standby job failed to get the Hyper-V VM.

Solution:

Virtual Standby jobs fail when:

- The Virtual Standby web service is unable to retrieve information about the virtual machine from the Hyper-V system. Communication problems between the Arcserve UDP and the Hyper-V system occur when the required Hyper-V services are not running on the Hyper-V system.

  Solution: Verify that all of the required Hyper-V services are running on the Hyper-V system.

- The Hyper-V system does not contain sufficient amount of free disk space that is required to create the Virtual Standby virtual machine or to create a snapshot of the Virtual Standby virtual machine.

  Solution: Consider reconfiguring the Hyper-V system to allow sufficient free disk space in the system volume.

Note: If you discover other possible causes, contact Arcserve Support.

Issue Related to Duplicate Agent UUID

Symptom:

A monitor that exists in the Console with the same agent UUID needs to overwrite the agent UUID.

Solution:

1. From your machine, run *regedit*.
2. Navigate to the location: *HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine*
3. Delete the Value Data of NodeID.
4. Restart Arcserve UDP Agent Service.

Modifications to Virtual Private Cloud in Virtual Standby Task Do Not Reflect in Amazon EC2

Symptom:

After performing a Virtual Standby task, if I modify the Virtual Private Cloud (VPC) in Networks Settings of Virtual Standby Task to use a different Amazon VPC for the subsequent tasks, the modifications are not updated in the Amazon EC2.

Solution

Amazon EC2 does not allow modifying the VPC from UDP Console after running a Virtual Standby task at least once. As a workaround, now Arcserve UDP lets you terminate old instance and create a new instance with the existing data from previous sessions.

Follow these steps:

1. Configure a backup plan with VSB to EC2 as the secondary task.
2. Perform some successful VSB to AWS tasks.
3. Modify the VPC setting in the VSB task [network settings].

   Updated VPC settings is reflected on the instance after the next sucessful VSB task.

Notes:

- The old instance created in EC2 is ignored and the new instance with updated VPC is created in EC2 to reflect the network changes.

- New network setting is applied to old snapshots also.

Terminal EC2 Resources option is not shown

Symptom:

When nodes are configured with a plan that has Virtual Standby to EC2 task in UDP 7.0 U1, UDP 7.0 or UDP 6.5 U4, and if you upgrade the console to UDP 7.0 U2, the **Terminate EC2 Resources** option does not show for those specific nodes.

Solution:

Click **modify** option for the related plan and save it.

Copy Recovery Points Related

This section includes the following troubleshooting topics related to Copy to Recovery Points (CRP):

- Bandwidth Congestion with Copy Recovery Point to Cloud Jobs
- Configure the Registry for Copy Recovery Point Job
- Merge Job Skipped

Configure the Registry for Copy Recovery Point Job

Symptom

The Copy Recovery Point job did not run.

Solution

You can control the number of retry jobs and the time interval of retry jobs for the Copy to Recovery task using the registry keys below. The registry key is in the machine where you have installed UDP Agent.

Retry job for CRP

The registry key is located in the Arcserve UDP Agent at the following location:

HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFCopySession\nodename1

Add new key as mentioned below:

Key Name: MaxRetryCount

Key Type: REG_DWORD

Value: 1 as default and maximum (10)

Note: This update is applicable to only one node “nodename1”.

For Copy Recovery Point to Cloud jobs, Arcserve UDP uses a temporary path, which can be configured with the registry keys for Temp path.

Configure Temp Path

The registry key is located in the Arcserve Backup Server at the following location:

HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFCopySession

Key Name: LocalTempPath

Key Type: REG_SZ

Note: By default, it stores in Arcserve product home path.

Bandwidth Congestion with Copy Recovery Point to Cloud Jobs

Symptom

Running Copy Recovery Points (CRP) jobs often result in slowing down of agent machine due to bandwidth congestion.

Solution

Arcserve helps you define a specific bandwidth for CRP jobs. As a result, even if the machine has multiple jobs running, with CRP jobs, bandwidth is controlled automatically on the agent machine only for CRP jobs. Using a separate process AFCRPBackend.exe for CRP jobs, you can configure the Policy based QoS to AFCRPBackend.exe and throttle the outbound bandwidth.

Follow these steps:

1. Type MMC in Start box and press Enter.

   Microsoft Management Console opens.

2. From the MMC console, press CTRL+M.

   Add or Remove Snap-ins dialog appears.

3. From the Add or Remote Snap-ins dialog, select Group Policy Object Editor under Available snap-ins and click Add.

4. In the Select Group Policy Object dialog, leave the default setting of Local Computer and click Finish.

5. Click OK.

   The Add or Remote Snap-ins dialog closes.

6. From the left pane of the MMC console Window, expand Local Computer, Computer Configuration, Windows Settings, and right click Policy-based QoS and select Create new policy from the menu.

7. In the policy-based QOS windows perform the following options, and click Next:

   - Enter a name for the new policy.

   - Set the DSCP value to 0.

   - Select the check box of Specify Outbound Throttle Rate.

   Note: We recommend selecting MBps as option.

8. For QoS policy applies to, select the check box of Only applications with this executable name and enter the following exe path, and click Next:

   C:\Program Files\Arcserve\Unified Data Protection\Engine\BIN\AFCRPBackend.exe

   Note: AFCRPBackend.exe helps ensure that bandwidth restriction is applicable only to CRP jobs even if multiple jobs are running on the agent machine.

9. Proceed with default settings to the last screen and click Finish.

Merge Job Skipped

Symptom

- The merge Job was skipped because the session is currently locked by the copy recovery point job. Verify if any copy recovery point job is pending.
- The merge job was skipped because the session is currently locked by the on-demand copy recovery point job. The on-demand copy recovery point job is either running or waiting to run.

Solution

To continue the merge job we need to remove the lock files from the backup destination of the node.

Run the Delete lock tool from *C:\Program Files\Arcserve\Unified Data Protection\Engine\BIN\AFDeleteLockTool.exe*.

Note: This tool deletes all CRP related locks (CRP scheduled and ad-hoc CRP) and merge starts immediately. Merged sessions will not be available for CRP to copy to destination.

Usage:

AFDeleteLockTool.exe -path <BackupDestinationPath> -user <USERNAME> -password <PASSWORD> -type <crp/adhoccrp>

Remarks:

-path: -specify Backup Destination Path

-type: -specify type of Task to delete lock

crp: Deletes Lock related to Scheduled CRP jobs.

adhoccrp: Deletes Lock related to ad hoc CRP jobs.

-user: -specify Username for backup destination

-password: -specify Password for backup destination

-path and -type is Mandatory

-user and -password is optional. (either both should be supplied or both should be skipped)

Examples:

- If you want to remove only CRP scheduled tasks locks, please run the tool as below:

  AFDeleteLockTool.exe -path I:/Backup/RPS678 -type crp

- If you want to remove only Ad-hoc CRP job locks, please run the tool as below:

  AFDeleteLockTool.exe -path I:/Backup/RPS678 -user User123 -password "**********" -type crp

- If you want to remove all CRP locks, please run the tool as below:

  AFDeleteLockTool.exe -path I:/Backup/RPS678 -user User123 -password "**********" -type crp -type adhoccrp

Arcserve UDP Reports Related

This section includes the following troubleshooting topic related to Arcserve UDP reports:

How to Display Charts to View Arcserve UDP Reports

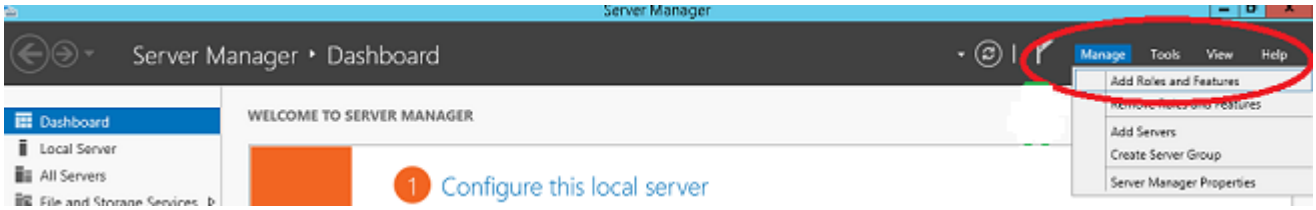
How to Display Charts to View Arcserve UDP Reports

At times when you want to view the reports, the charts may not be displayed. Complete the below prerequisite to troubleshoot the problem for Windows 2012/2012r2.

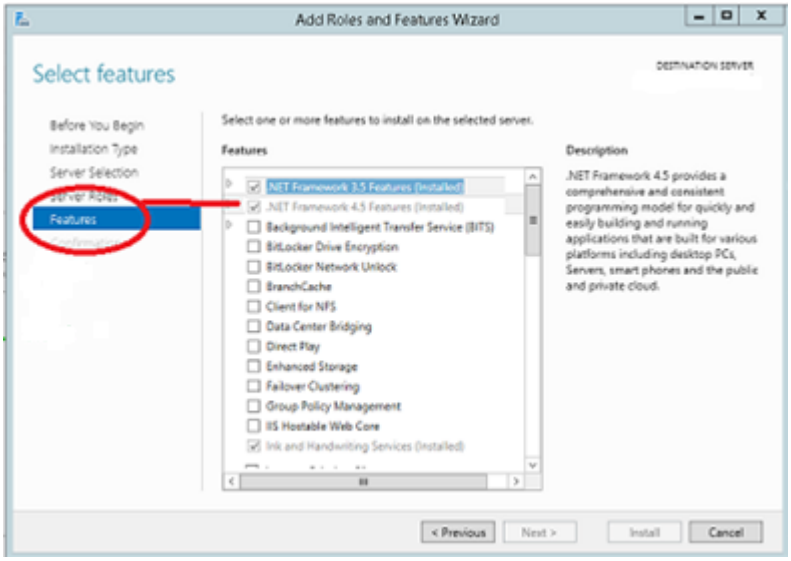
To view the charts, you need to install .NET 3.5 framework or later version.

To install .NET Framework, follow these steps:

1. Open Server Manager.
2. On the Dashboard, from the top-right corner, click **Manage**, and then select the **Add Roles and Features** option.



3. On the Add Roles and Features Wizard dialog, do the following:
   a. From the left pane, click **Features**.
   b. On the **Features** tab, select the **.NET Framework 3.5 Features (Installed)** check box.
   c. Click **Next**.



4. From the Confirm Installation Selections screen, click **Install**.

   The installation of .NET Framework is completed successfully.

Important! Even after implementing the prerequisite, if charts are not displayed using **Print/Email/Save** options, then verify if the Registry Editor has FCImgExportDll. To verify the dll, navigate to the following path:

C:\Program Files\Arcserve\Unified Data Protection\Management\BIN

To unregister and register FCImgExportDLL.dll, follow these steps:

1. Navigate to following path in the command prompt:

   C:\Windows\Microsoft.NET\Framework64\v4.0.30319

2. To unregister the DLL, run the following command:

   C:\Windows\Microsoft.NET\Framework64\v4.0.30319>Regasm.exe /u "C:\Program Files\Arcserve\Unified Data Protection\Management\BIN\FCImgExportDLL.dll"

3. To register the DLL, run the following command:

   C:\Windows\Microsoft.NET\Framework64\v4.0.30319>Regasm.exe "C:\Program Files\Arcserve\Unified Data Protection\Management\BIN\FCImgExportDLL.dll" /codebase

Microsoft Office 365 Exchange Online Related

This section includes the following troubleshooting topics related to Microsoft Office 365 Exchange Online:

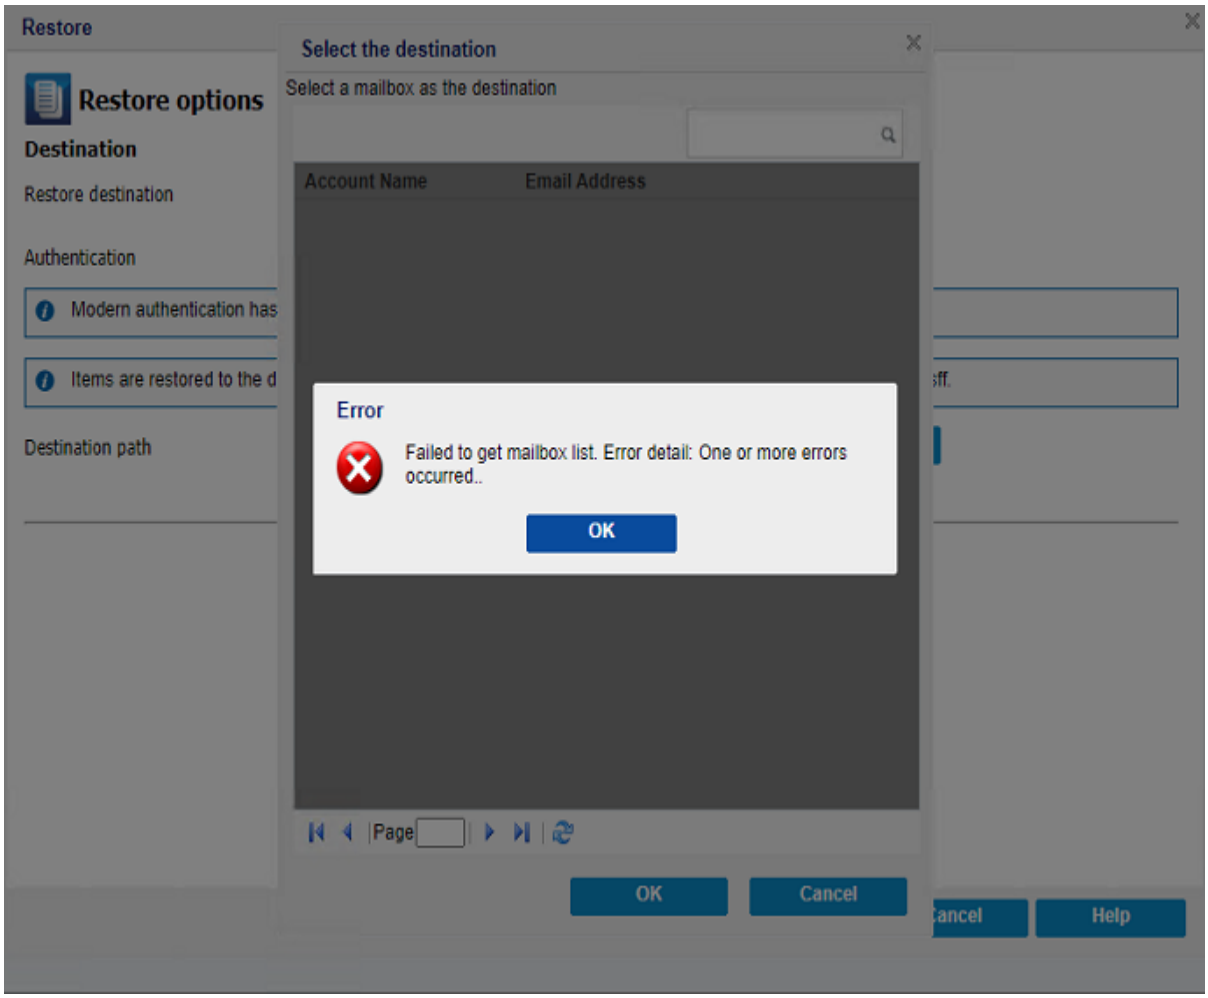
- Unable to View the List of Mailboxes on the *Select the destination* Dialog
- Access or Permission Denied for Exchange Online Group

Unable to View the List of Mailboxes on the *Select the destination* Dialog

Symptom:

While defining the restore options, the Select the destination dialog does not display the list of mailboxes. When you restore the mail data to an alternate location using the existing authentication and click the Browse button to select a destination, the following error message appears:

Failed to get mailbox list. Error detail: One or more errors occurred.



Solution:

To resolve this error, upload a valid certificate to the Azure Portal under specified Application ID. For more information, see How to Use Self-Signed Certificate.

Access or Permission Denied for Exchange Online Group

Symptom:

When there is no mailbox attached to the Teams Private Group, the following error message appears:

Backup mailbox [abcd-pvt-group] failed:Access is denied

Solution:

To avoid this issue, follow these steps:

1. Create a shared mailbox:
    a. Go to the Microsoft Login page, and then log into the Microsoft Office 365 Admin account.
    b. Navigate to the Exchange Online admin center page.
    c. Under Dashboard, go to **recipients** > **shared**.
    d. Click **+** to add or create a shared mailbox.
2. Set a password and assign the Office 365 license for shared mailbox:
    a. Return to the Office 365 Admin account.
    b. Go to **Users** > **Active users**, and then select shared mailbox.
    c. Set the Password and assign the Office 365 license.
3. Add the shared mailbox as the member of the private group:
    a. Navigate to the Exchange Online admin center page, and then from Dashboard, click **Groups**.
    b. Select the private mailbox, and then click **Edit**.
    c. On the Edit Settings page, add the shared mailbox as the member of the group, and then click **Save**.
4. Add the shared mailbox for Global Administrator:
    a. Go to the Azure Portal.
    b. Navigate to the Roles section, and then select the Global Administrator role.
    c. Click **Add Assignments**, select shared mailbox, and then click **Save**.
5. On the UDP console, update the node with email address of shared mailbox, and then re-run the backup job.
6. Verify if the backup job is successful.

Chapter 24: Appendix

This section contains the following topics:

Command Line Tool for Deduplication Data Store

The command line tool (as_gddmgr.exe) lets you check data integrity at recovery point level and data store level for deduplication data stores.

You can use this tool to regenerate a hash database in case of a hash database failure. You can also use this tool to query and manage the backend purge and disk reclamation for deduplication data stores.

The output of the command is displayed on the Windows command console. The tool also generates a log file in the “Logs” folder. For example, as_gddmgr_2014-9-4_11-14-22-655.log is a log file that contains all details.

Location:

You can find as_gddmgr.exe in the "Bin" folder of the UDP installation path.

Syntax:

as_gddmgr.exe

-Scan CheckRecoveryPoint <data store name> -Node [<All> |<UDP agent node name>] -RecoveryPoint [<Latest>|<recovery point number>] [-LogLevel <n>]

-Scan VerifyRefCount <data store name> [-LogLevel <n>]

-Scan VerifyData <data store name> [-Password <data store password>] [-LogLevel <n>]

-Scan VerifyAll <data store name > [-Password < data store password >] [-LogLevel <n>]

-Scan RebuildHash <data store name> [-NewHashPath <new hash path>] [-LogLevel <n>]

-Scan RebuildHashWithIndexPath <index path> -NewHashPath <new hash path> [-LogLevel <n>]

-Purge Start <data store name>

-Purge Stop <data store name>

-Purge Status <data store name>

-Purge StartToReclaim <data store name>

-Purge StartToIdentifyObsoletedData <data store name>

Options:

CheckRecoveryPoint

Rehydrates the specified recovery point(s) as full, then checks data integrity.

Node <All> | <UDP agent node name>

Specifies the agent node name.

RecoveryPoint <All> | <recovery point number>

Specifies the recovery point to check for integrity.

Password <data store password>

Specifies the data store password.

LogLevel <n>

Specifies the log level number.

VerifyRefCount

Scans index files and reference files to verify reference count recorded in the hash database. Before you specify this option, manually stop the deduplication data store.

VerifyData

Scans data files and then regenerates the hash keys by comparing this with the reference file. Before you specify this option, manually stop the deduplication data store.

VerifyAll

Performs both the VerifyRefCount and VerifyData operations. Before you specify this option, manually stop the deduplication data store.

RebuildHash

Specify data store name and then regenerate the hash database by scanning index and reference files. Before you specify this option, manually stop the deduplication data store.

RebuildHashWithIndexPath

Specify the deduplication index path and then regenerate the hash database by scanning index and reference files. The option is used only when the data store is not present on any recovery point servers.

Start

Enables running the purge and disk reclamation in parallel with other regular Arcserve UDP jobs. Running purge in parallel may cause throughput degradation of regular Arcserve UDP jobs.

Stop

Disables running the purge and disk reclamation in parallel with other regular Arcserve UDP jobs.

Status

Queries the status of purge or disk reclamation.

StartToReclaim

Enables running the purge and disk reclamation in parallel with other regular Arcserve UDP jobs. This option skips the identify obsolete data phase that finds out the obsolete data block, and then directly start disk reclamation phase, which free up the disk space. The side effect of this option is that it might degrade the efficiency of the disk reclamation because the identify obsolete data phase might find out more obsolete data blocks in the data files, but the disk reclaim does not wait for the identify obsolete data phase to complete. In addition, running purge in parallel might cause the through-put degradation of regular Arcserve UDP jobs.

StartToIdentifyObsoletedData

Enables running the purge and disk reclamation in parallel with other regular Arcserve UDP jobs. This option starts identifying obsolete data phase. The option is useful if that user wants to skip the ongoing disk reclamation phase.

Note: Be aware that the following options might run for a long time because the operation scans many files in the deduplication data store.

- VerifyRefCount
- VerifyData
- VerifyAll
- RebuildHash
- RebuildHashWithIndexPath

Examples:

as_gddmgr.exe -Scan CheckRecoveryPoint GDDDataStore1 -Node myComputer -RecoveryPoint 18

as_gddmgr.exe -Scan CheckRecoveryPoint GDDDataStore1 -Node All -RecoveryPoint Latest

as_gddmgr.exe -Scan VerifyRefCount GDDDataStore1

as_gddmgr.exe -Scan VerifyData GDDDataStore1 -Password 123

as_gddmgr.exe -Scan VerifyAll GDDDataStore1

as_gddmgr.exe -Scan RebuildHash GDDDataStore1

as_gddmgr.exe -Scan RebuildHash GDDDataStore1 -NewHashPath C:\NewHashPath

as_gddmgr.exe -Scan RebuildHashWithIndexPath D:\GDDDataStore\Index -NewHashPath D:\NewHashPath

as_gddmgr.exe -Purge Start GDDDataStore1

as_gddmgr.exe -Purge Stop GDDDataStore1

as_gddmgr.exe -Purge Status GDDDataStore1

as_gddmgr.exe -Purge StartToReclaim GDDDataStore1

as_gddmgr.exe -Purge StartToIdentifyObsoletedData GDDDataStore1

How to Show Recovery Point Check Option

From Arcserve UDP 7.0 version, recovery point check function is hidden by default from plan wizard. We recommend to use Assure Recovery task to detect possible data problem. We recommend to use Assure Recovery task to detect possible data problem. For details, refer to How to configure Assure Recovery. If required, you can still use this option to perform data check and view in the Plan wizard.

Follow these steps:

1. Log into UDP Console.
2. Navigate to the UDP Console installation folder. For example, *C:\Program Files\Arcserve\Unified Data Protection\Management\Configuration*.
3. Open the file ConsoleConfiguration.xml using any text editor.
4. Find the below text under the section <TimeoutConf>:

   <recoveryPointCheck>false</recoveryPointCheck>
5. Modify the value from false to true.
6. Save the file and exit.
7. Restart the UDP Console management service to allow the setting to take effect.

   Note: The Arcserve UDP version is upgraded from previous version and this option is already selected in a plan, and is not hidden by default.

   This option lets you detect data corruption issues by verifying the file system of the volumes. When the backup job completes, Arcserve UDP mounts the recovery point and runs the chkdsk Windows command. If the chkdsk command detects an error, the next backup job is converted to a Verify backup job. This option is applicable for both VMware and Hyper-V virtual machines with the Windows guest OS. Review the following considerations before enabling this option:

   - The following types of volume are not supported and they are skipped by **Recovery Point Check**:
     - The volume whose file system type is not NTFS
     - The volume whose type is striped with parity
     - The volume that is in that storage pool
   - The chkdsk command cannot detect all file system problems. The recovery point check may pass but the recovery point can still be corrupted.
   - Depending on the size of the file system of the guest OS, the chkdsk command may take a longer time to run. The chkdsk uses a large amount of system memory on the Backup Proxy server and affects the performance of the

proxy server. This result in the backup job taking a longer time to complete. *In the worst case, the system memory of the Backup Proxy server may get exhausted and the server may become non-responsive, especially when there are numerous concurrent backup jobs or huge volumes are being checked.* Check recovery point itself can monitor the system memory usage and, if the memory usage reaches to a threshold, check recovery point will suspend itself for some time and release some the system memory. However, as a best practice, disable this option unless it is necessary or you have a powerful Backup Proxy server. Alternatively, you can distribute the load to multiple proxy servers by creating multiple plans and specifying different proxy serves in each of the plan.

- If the backup is crash consistent, there are high chances that chkdsk will detect problems (due to the nature of a crash consistent backup). As a best practice, do not enable this option for a crash consistent backup.
- If you want to enable the Recovery Point Check option but you do not want the next backup job to be converted to a Verify backup job, create a DWORD value named CheckRecoveryPointIgnoreError in the registry of the proxy server and set the DWORD value to 1. Create the DWORD value at the following location:

  HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll

  The DWORD is applicable to all the backup jobs that are running on the current proxy server. If you want to control the behavior of a specific virtual machine, you can set the value at the following location:

  HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll\<VM GUID>

  Note: If you add the registry key in both the VM and proxy level registry, then the setting in the VM level registry will have the priority over the setting in the Proxy level registry.
- If you want to fail the backup job after the Recovery Point Check detects problem (so that you can be aware of data problem promptly), create a DWORD value named CheckRecoveryPointDontFailJob in the registry of the proxy server and set the DWORD value to 0. Create the DWORD value at the following location:

  HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll

The DWORD is applicable for all the backup jobs that are running on the current proxy server. If you want to control the behavior of a specific virtual machine, you can set the value at the following location:

HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AFBackupDll\<VM GUID>.

Note: If you add the registry key in both the VM and proxy level registry, then the setting in the VM level registry will have the priority over the setting in the Proxy level registry.

How to apply different version of VDDK other than the built-in Version in Arcserve UDP

VMware Virtual Disk Development Kit (VDDK) 6.7.1 is built-in with Arcserve UDP 7.0. At the same time, Arcserve UDP 7.0 is compatible with VDDK 5.5. If you want to perform VSB/Instant VM/Assured Recovery to VMware 5.0 or 5.1, use VDDK 5.5 instead of VDDK 6.7.1.

Important! Arcserve certifies only VDDK 6.7.1 and 5.5 . Other versions of VDDK are not guaranteed to work with Arcserve UDP Version 7.0.

To modify the VDDK manually, follow these procedures:

- Modify VDDK for Virtual Standby Backup (VSB)
- Modify VDDK for Host Based Backup (HBBU)

Modifying the VDDK Manually for Virtual Standby Backup (VSB)

You can modify VDDK manually or by using Windows batch file. This topic explains procedures for both the options.

Considerations:

- For Instant VM and Assured Recovery Test, you need to perform this action on Recovery Server/Proxy Server.
- For VSB to VMware, use a dedicated Agent server as VSB monitor and configure the VSB task in proxy mode. Then, manually switch the VDDK to version 5.5 on the VSB monitor server.
- If the VSB task is configured in an MSP plan or a cross-site plan switch the VDDK on the RPS server as the monitor server is not present. As a result, this RPS server is not available to use as a proxy server for Host-Based Agentless backup task.

To modify the VDDK manually, follow these steps:

1. Log into the computer where the Arcserve UDP Agent is installed using an administrative account or an account that has administrative privileges.
2. Rename the folder *VDDK* to VDDK6.7.
3. Rename the folder *VDDK5.5* to VDDK from the following default location:

   C:\Program Files\Arcserve\Unified Data Protection\engine\BIN

Note: The location of the BIN folder can vary depending upon the location where you installed the Arcserve UDP Agent.

4. Run the following command to verify the state of the service:

   Note: You need to delete the service, only when not in running status.

   sc query vstor2-mntapi20-shared

   If the service does not exist, delete using the next step. If the service is running, before deleting using the next step run the following command to stop:

   sc stop vstor2-mntapi20-shared

5. Run the following command to delete the service:

   sc delete vstor2-mntapi20-shared

6. Run the following command to verify if the service is deleted successfully:

   sc query vstor2-mntapi20-shared

   Note: If the state of the service is STOP_PENDING, restart the machine.

   To modify the VDDK using a Windows batch file, follow these steps:

1. Log into the computer where the Arcserve UDP Agent is installed using an administrative account or an account that has administrative privileges.
2. Launch the *ChangeToVDDK55.bat* utility tool from the following default location:

   C:\Program Files\Arcserve\Unified Data Protection\engine\BIN

   Note: The location of the BIN folder can vary depending upon the location where you installed the Arcserve UDP Agent.

   VDDK 5.5 is installed after a new job is launched.

Modifying the VDDK Manually for Host Based Backup (HBBU)

You can modify the VDDK for host based backup (HBBU) manually.

Modify the VDDK 6.7 for backup of vSphere 5.5 or higher:

Follow these steps:

1. Download VDDK from the VMware website.
2. Extract all files into a temporary folder.
3. Locate the folder named ***bin*** that is available under the below path and rename the folder (for example, *bin_old*):

   <Arcserve UDP installation path>\Engine\BIN\VDDK\BIN\VDDK64\

Example: C:\Program Files\Arcserve\Unified Data Protection\Engine\BIN\VDDK\BIN\VDDK64\

4. Locate the folder named ***bin*** from the extracted files in the temporary folder and copy to the below path:

   <Arcserve UDP installation path>\Engine\BIN\VDDK\BIN\VDDK64\.

   The VDDK is successfully modified.

Modify the VDDK 5.5 for backup of vSphere 5.0 and 5.1:

Follow these steps:

1. Download VDDK from the VMware website.
2. Extract all files into a temporary folder.
3. Locate the folder named ***bin*** that is available under the below path and rename the folder (for example, *bin_old*):

   <Arcserve UDP installation path>\Engine\BIN\VDDK5.5\BIN\VDDK64\

   Example: C:\Program Files\Arcserve\Unified Data Protection\Engine\BIN\VDDK5.5\BIN\VDDK64\

4. Locate the folder named ***bin*** from the extracted files in the temporary folder and copy to the below path:

   <Arcserve UDP installation path>\Engine\BIN\VDDK5.5\BIN\VDDK64\.

   The VDDK is successfully modified.

How to Migrate Arcserve D2D r16.5 Backup Data of Two Servers Having Same Host Name to RPS Data Store

You have two servers with the same host name but different FQDN DNS name, and you use Arcserve D2D r16.5 to protect them. Now, if you want to migrate the data to Arcserve UDP RPS data store, follow these steps:

1. Stop Arcserve D2D service on Server 1.
2. Using the Arcserve UDP Jumpstart job, migrate data of Server 1 to the RPS data store.
3. Upgrade the Arcserve UDP Agent in Server 1.
4. Create a plan (or using an existing plan) to protect Server 1 and choose the same RPS data store as the backup destination.
5. Deploy the plan.
6. Stop Arcserve D2D service on Server 2.
7. Using the Arcserve UDP Jumpstart job, migrate data of Server 2 to RPS data store.
8. On the RPS server, navigate to the data store backup destination folder and locate the folder for Server 2 by searching the host name of Server 2.

   For example, if the host name of Server 2 is “MyServer2”, then the folder name is “MyServer2”.
9. Upgrade the Arcserve UDP Agent in Server 2.
10. On Server 2, start the UDP Agent if it is not started.
11. Open the registry (on any specific server?) and find the following key:

    HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine”
12. Note the value of “NodeID”.

    The value is a unique ID in format of GUID, such as “e856e0ba-66d7-4da5-8b98-2250173e349a”.
13. On the RPS server, locate the backup destination folder of Server 2 and update the node ID as **<folder_name>[NodeID value]**.

    Example: MyServer2[e856e0ba-66d7-4da5-8b98-2250173e349a]
14. Create a plan (or use an existing) to protect Server 2 and choose the same RPS data store as the backup destination.
15. Deploy the plan.

How to Deploy Arcserve UDP in Microsoft Azure

You can deploy Arcserve UDP in Microsoft Azure when using Arcserve Unified Data Protection V6.5 Update 2 or higher. The feature helps you deploy the Virtual Standby machines in Microsoft Azure and protect the data. A virtual standby machine is created on Microsoft Azure and related data based on the recovery points from the On-Premise backup.

The virtual standby converts the recovery points to virtual machine formats in Microsoft Azure and prepares a snapshot to easily recover the data when needed.

This feature provides the high availability capability ensuring that the virtual machine takes over immediately when the source machine fails. The standby virtual machine is created by converting the recovery points to Azure virtual machine format.

What To Do Next?

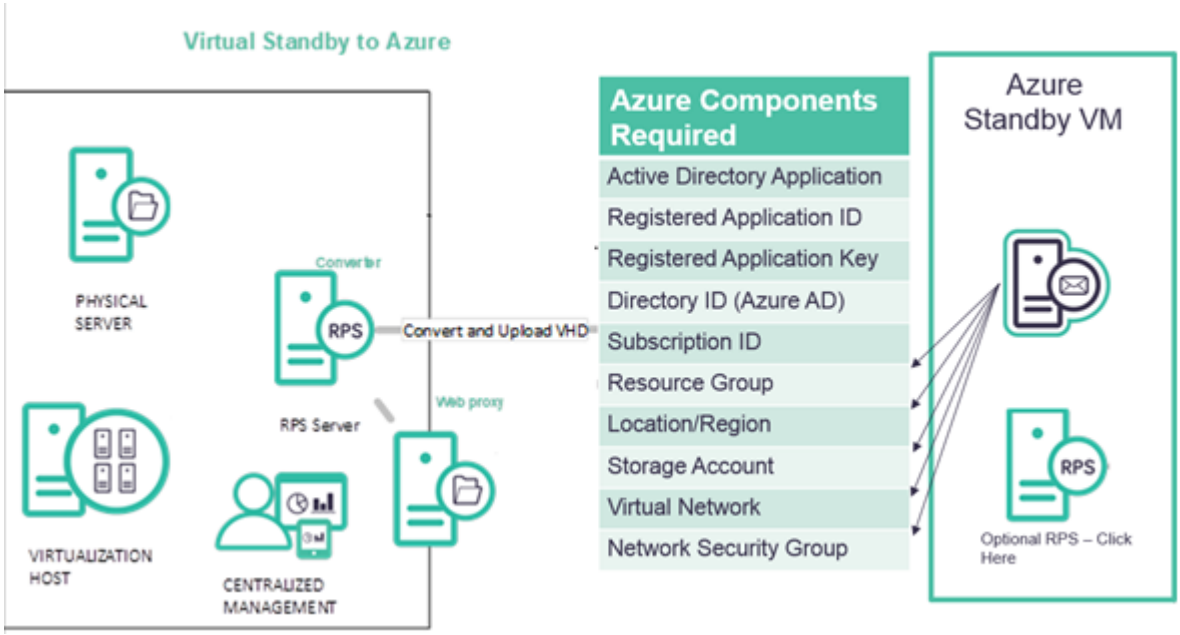
1. Understanding Process Flow
2. Best Practices
3. Considerations
4. Planning Deployment
5. Deploy Arcserve UDP in Microsoft Azure

Understanding Process Flow

Arcserve UDP supports the virtual standby Virtual Machines in Microsoft Azure. If you have already deployed the data protection solution on-premise, you can deploy Arcserve UDP in Microsoft Azure.

The on-premise solution is composed of Arcserve UDP Console and Arcserve UDP Reocvery Point Server (RPS). You can backup Windows systems to local RPS. You can create a plan in Arcserve UDP Console to backup the data in local RPS and then perform the virtual standby task. The virtual standby feature monitors the heartbeat of the source node. If the source node is down the virtual machine on Azure takes over.

The below image shows the process flow of Arcserve UDP for Microsoft Azure:



Best Practices

Review the following best practices before deploying Arcserve UDP to Microsoft Azure:

1. You must install Arcserve UDP Console on-premise to protect the nodes that are in on-premise network.
2. Select HTTPS protocol while installing Arcserve UDP components.
3. The easiest way to create Azure resources is to create at least one test virtual machine that you can use as your standby VM later.
4. Create Arcserve UDP Recovery Point Server in Azure.
    - Keep the TCP ports 8014 and 8015 open for inbound connections.
    - You must resolve the name of the Recovery Point Server to the public IP, if you access the Recovery Point Server from a remote web browser.
    - Use the shared plan task **Replicate to remotely managed RPS**, to replicate tasks.

Considerations

Review the following considerations before deploying Arcserve UDP to Microsoft Azure:

- In Microsoft Azure any single VHD file cannot have a system disk that is larger than 2048 GB and data disk that is larger than 4095 GB. In Arcserve UDP Virtual Standby the source cannot have a system disk that is larger than 2048 GB and data disk that is larger than 4095 GB.
- Microsoft Azure VM and Arcserve UDP VSB support only Windows OS 2008 R2 or higher.
- Arcserve UDP does not support creating the classic virtual machine as Standby VM.
- Microsoft Azure VM requires the system volume and boot volume on the same disk.
- Microsoft Azure does not support machine boot from the UEFI system.
- Microsoft Azure does not support the source machine as Hyper-V server.
- Arcserve UDP Virtual Standby does not support the dynamic system disk for source machine.
- Arcserve UDP in Microsoft Azure supports only Windows nodes.

Planning the Arcserve UDP Deployment in Microsoft Azure

You must complete the following tasks before deploying Arcserve UDP in Microsoft Azure:

1. Review Prerequisites
2. Prepare a Cloud Account in Microsoft Azure
3. Prepare Microsoft Azure Resources

Prerequisites

Review the following prerequisites for Arcserve UDP and Recovery Point Server before deploying Arcserve UDP in Microsoft Azure:

- Review Compatibility Matrix for supported operating systems, databases, and browsers.
- Prepare servers in advance to deploy as Arcserve UDP Console and Recovery Point Server.
- The servers for each component must meet the below system requirements:

- **Server:** Windows 2008 R2 Server or higher
- **CPU:** Minimum 4 Core 2.7 GHz
- **Disk Space:** 40 GB for Operating System installation
- **RAM:** Minimum 8 GB
- **Backup Storage for Recovery Point Server:** Based upon source data size

Prepare a Cloud Account in Microsoft Azure

Before you can create a cloud account for Microsoft Azure, you must complete the following mandatory prerequisite tasks in the given order:

1. Perform the following steps to register the Azure Active Directory application that Arcserve UDP uses to communicate with Microsoft Azure:
    a. Log into the Azure portal using valid credentials.
    b. Select the **Azure Active Directory** option.
    c. Select the **App registrations** option.
    d. Select the **New application registration** option.
    e. Specify details for the following fields and click **Create**:

        Name

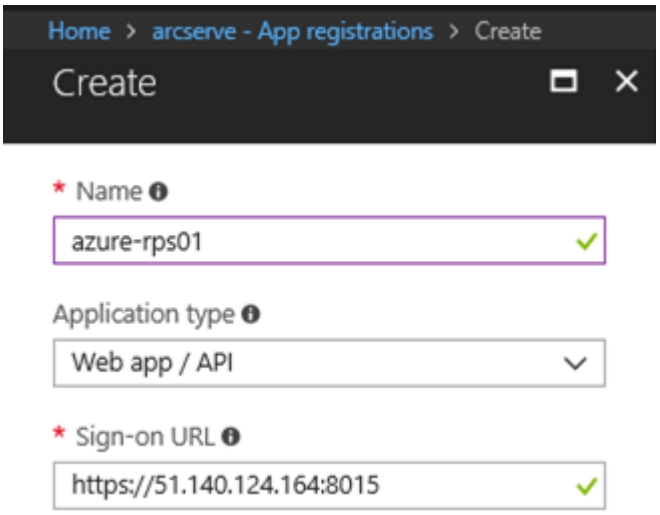
        Refers to the name of **Arcserve UDP Console** server.

        Application type

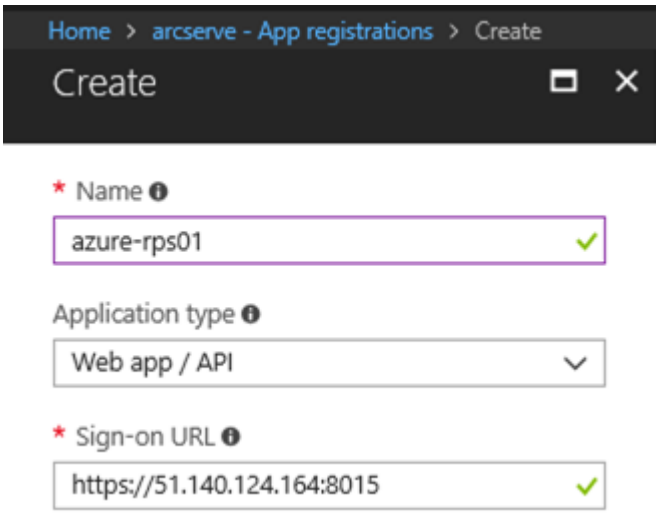
        Specify **Web app / API** as Application type.

        Sign-on URL

        Refers to the URL of the **Arcserve UDP** application.

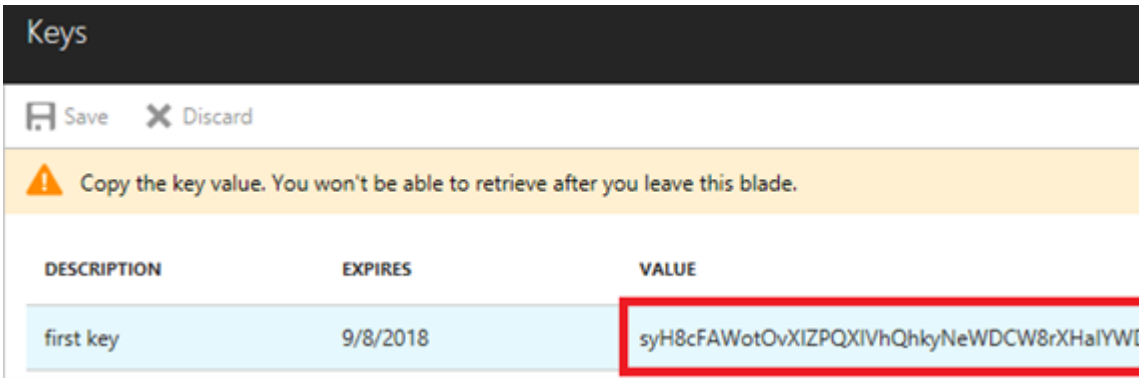


2. Perform the following steps to obtain the Application ID from Microsoft Azure that is used to communicate with Arcserve UDP:
    a. Navigate to **App registrations** in Azure Active Directory.
    b. Copy the **Application ID** that is used to communicate with Arcserve UDP.
    c. Store the Application ID as Client ID.

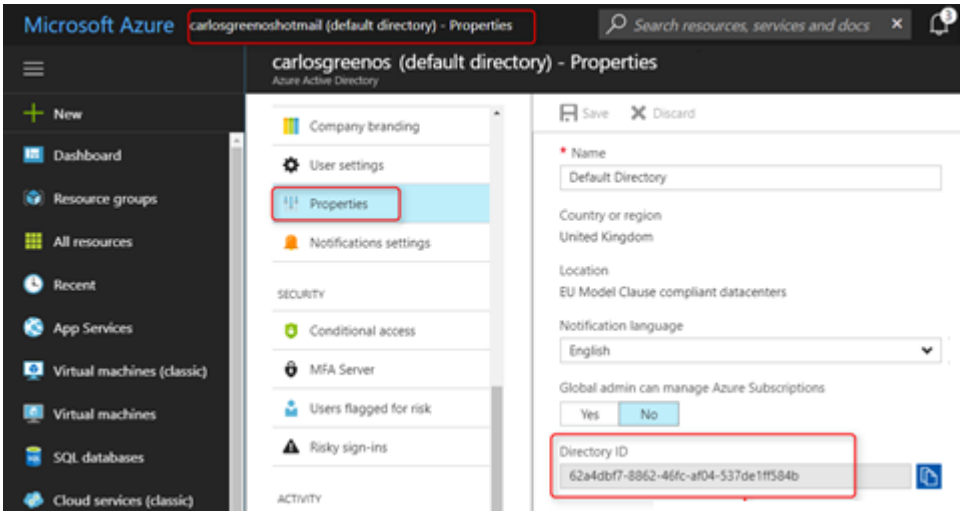


3. Perform the following steps to generate the client secret key for the application:
    a. Navigate to the application settings and click **Keys**.
    b. Enter description and duration for the key and click **Save**.
    c. Copy the displayed key value as you cannot retrieve the key

later. The saved key value is the Client Secret key.

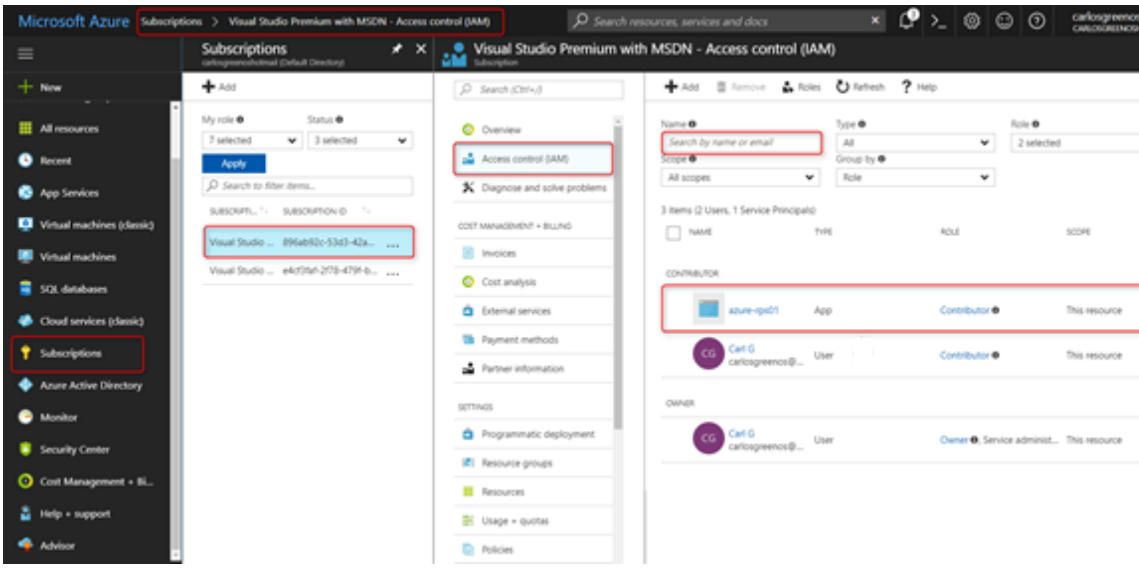


4. Perform the following steps to obtain the Tenant ID for the application:
    a. Navigate to **Azure Active Directory** and select **Properties**.
    b. Note the **Directory ID** available in the **Properties**. Directory ID is used as Tenant ID in Arcserve UDP.



5. Perform the following steps to assign Contributor role to the application:
    a. From the Microsoft Azure portal menu, click **Subscriptions**.
    b. Select your subscription.
    c. Click the Access Control (IAM) tab.
    d. Add your application.

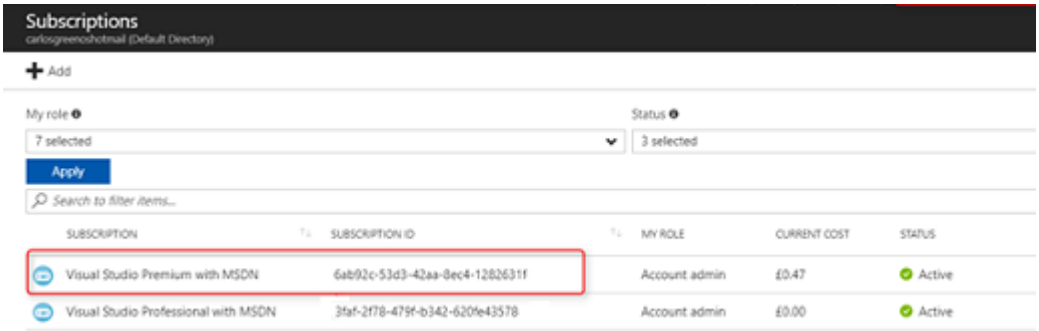
e. Assign the Contributor role to the application.



6. Perform the following steps to get Azure Subscription ID:

   a. Navigate to the Microsoft Azure portal menu and select **Subscriptions**.

   b. Note the value of subscription ID that appears in the **SUBSCRIPTION ID** field. You must use the same Subscription ID when you add Microsoft Azure Cloud account in Arcserve UDP



Prepare Microsoft Azure Resources

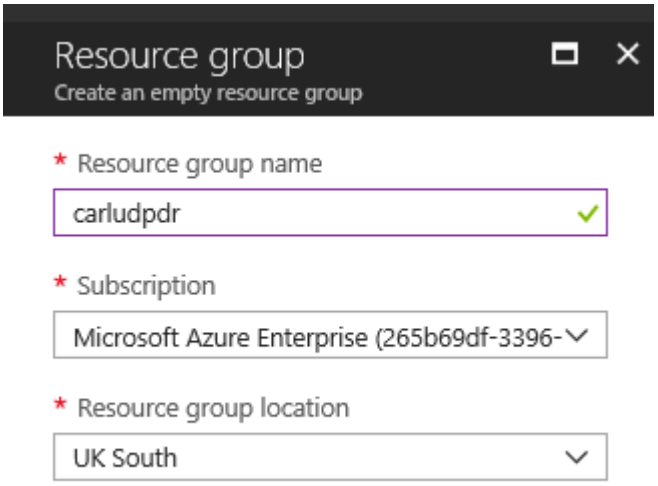
Before you can create a cloud account for Microsoft Azure, you must create the following Microsoft Azure resources:

1. Perform the following steps to create a Resource group:

   a. Log into the Azure portal using valid credentials.

   b. Click **Add** to create an empty resource group.

c. Enter a name and location for the new resource group and click **Create**.



2. Perform the following steps to create a storage account:

   a. Navigate to **Storage accounts** and click **Add**.

   b. Specify the following:

      Note: Ensure that your storage account, virtual network and network security group of the standby VM are at the same location.

      Name

      Specifies name of the storage account.

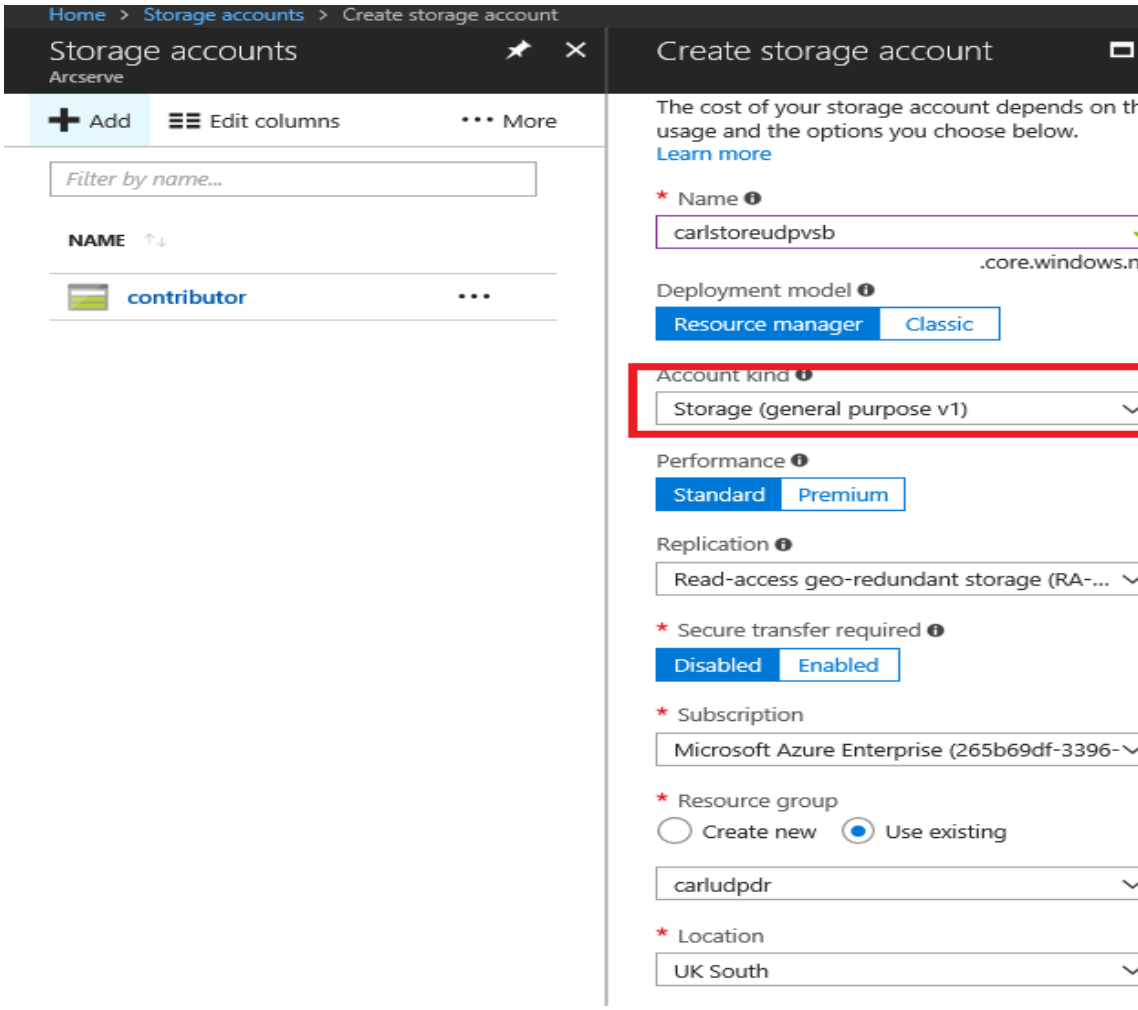
      Deployment model

      Select deployment model based on your requirement.

      Account kind

      Specify **Storage (general purpose v1)** or **StorageV2 (general purpose v2)** as Account kind.

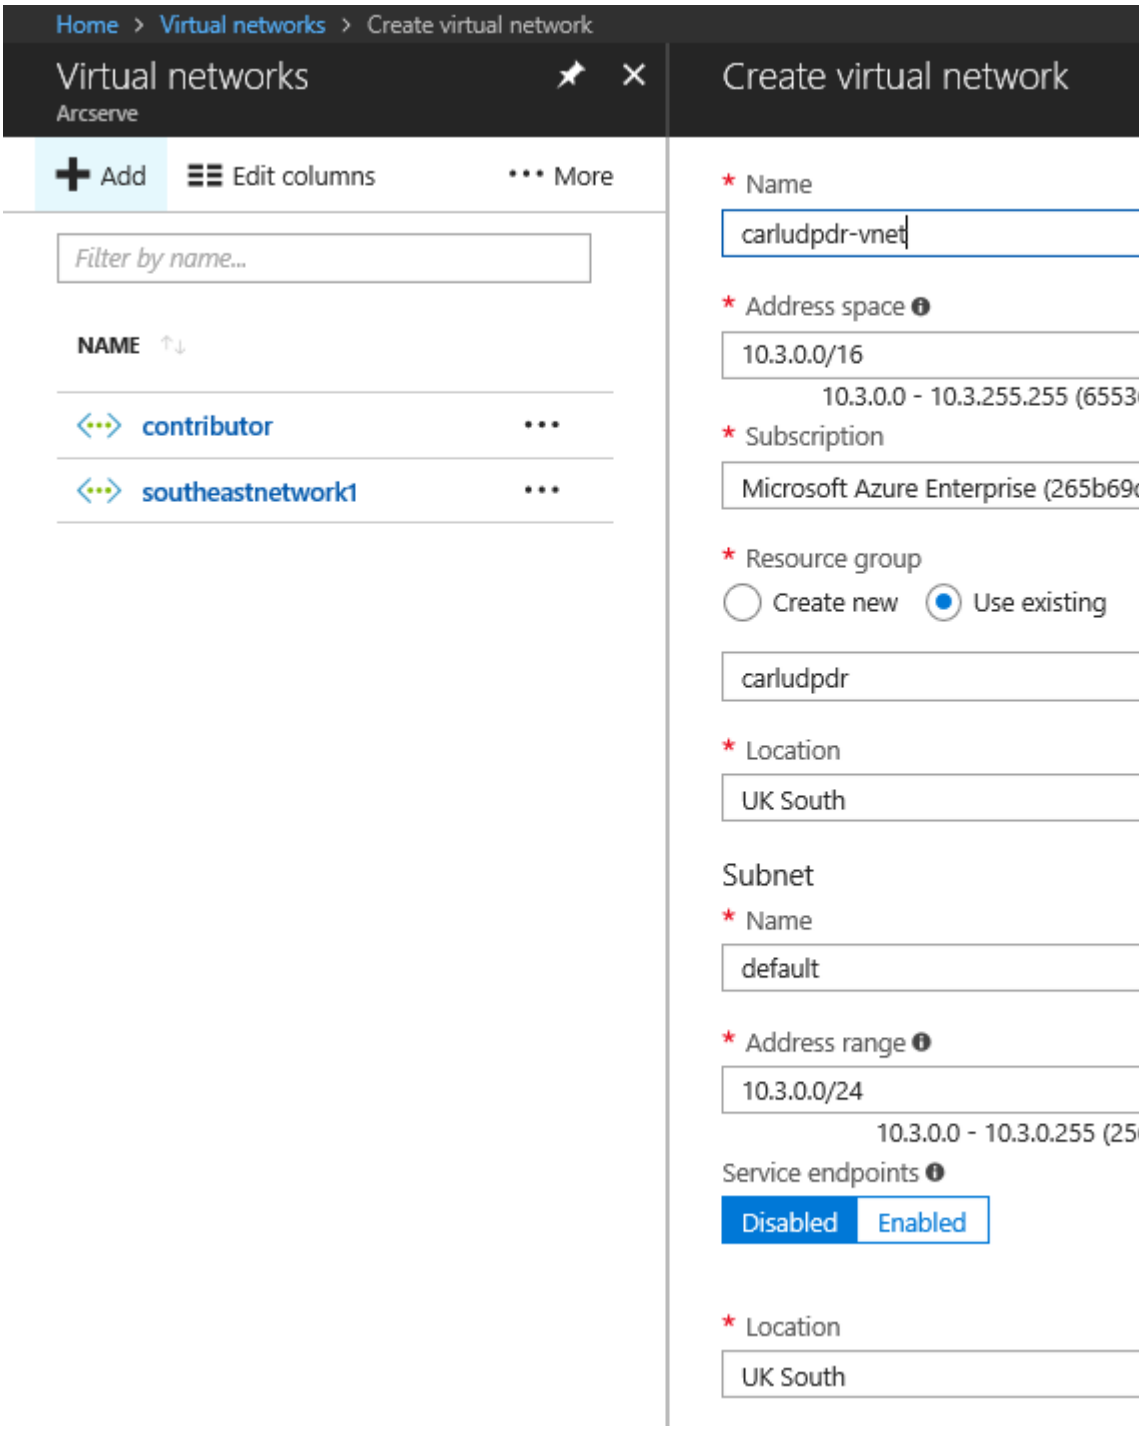
c. Specify other details as required and click **Create**:



3. Perform the following steps to create a Virtual Network and Subnet:

   a. Navigate to Microsoft Azure home, **Virtual networks**, and click **Add**.

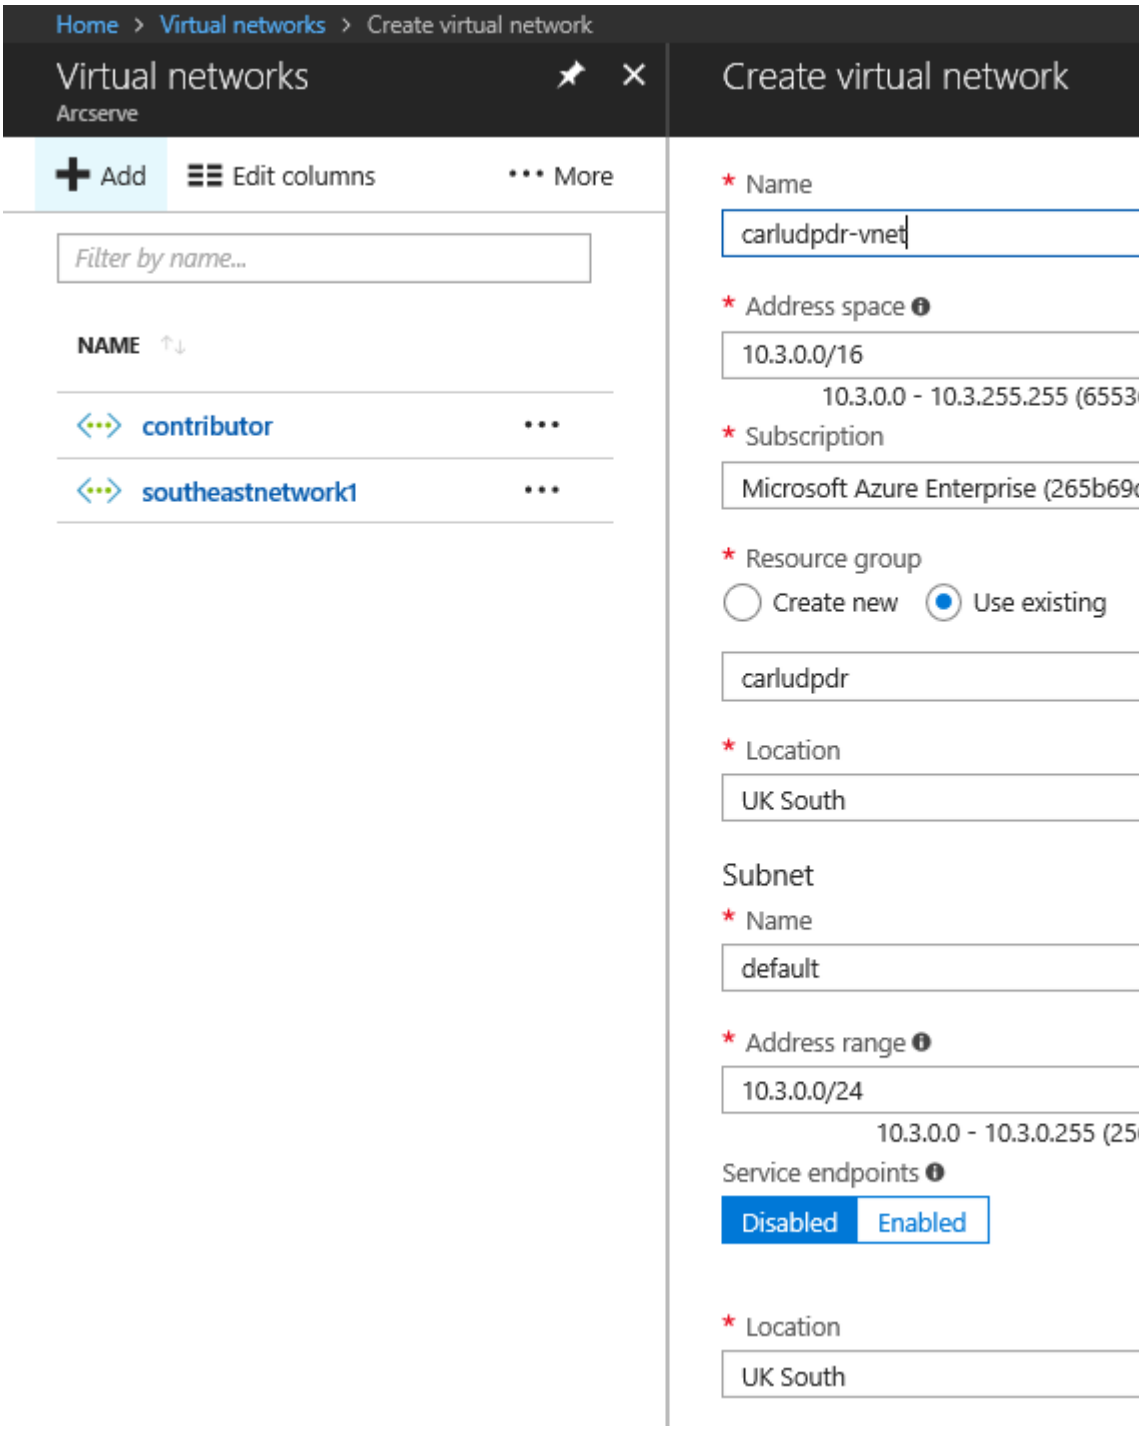
b. Enter the required details and click **Create**.



4. Perform the following steps to create a Network Security group:

   a. Navigate to Microsoft Azure home, **Network security groups**, and click **Add**.

b. Enter the required details and click **Create**.



Microsoft Azure resources are successfully created.

Deploy Arcserve UDP in Microsoft Azure

After completing the prerequisites, you can start deploying Arcserve UDP in Microsoft Azure.

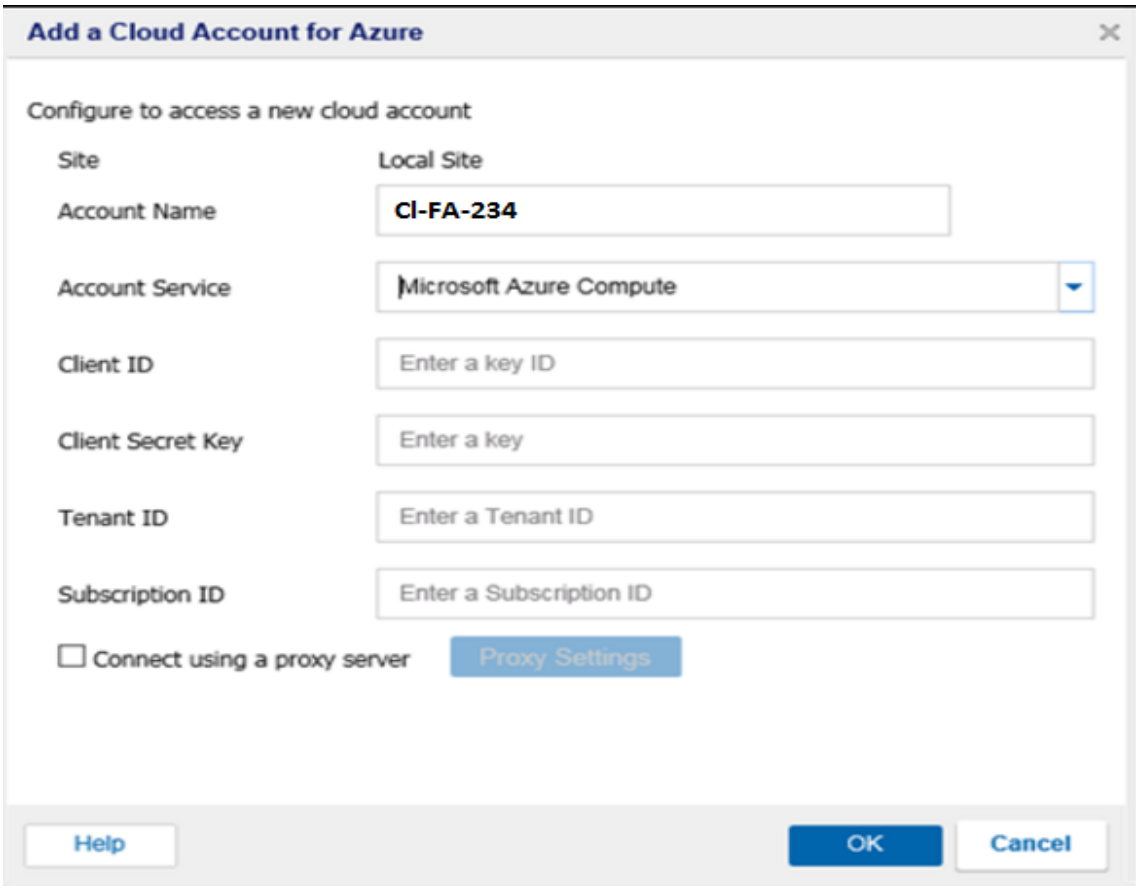
What To Do Next?

1. Add Azure Cloud Account in Arcserve UDP
2. Create a Plan with a Backup Task
3. Add a Virtual Standby Task to the Plan
4. Run the Virtual Standby Job Manually
5. Power-on Virtual Standby VM in Azure
6. Verify Virtual Standby VM Status

Add Azure Cloud Account in Arcserve UDP

Add a Microsoft Azure Compute cloud account to copy Files or recovery points to cloud storage. You can use this account while creating tasks for Virtual Standby to Cloud or Instant Virtual Machine on Microsoft Azure plans.

Note: To add a Cloud Account for Microsoft Azure, you must meet the pre-requisites. For details, view Prerequisites.



Follow these steps:

1. Log into Arcserve Unified Data Protection, and click the **resources** tab.
2. From the left pane, navigate to **Destinations**, and click **Cloud Accounts**.

The **Destinations: Cloud Accounts** page is displayed in the center pane.

3. Click **Add a Cloud Account**.

   The **Add a Cloud Account** page is displayed.

4. For **Account Name**, provide a unique name.

   Account Name specifies the name of the cloud storage. This name will be added to Console for identifying the cloud account. Each cloud account must have a unique storage name.

5. Select the option from the **Account Service** drop-down list.

   Multiple fields appear for configuration.

6. Enter details in the following fields to configure and click **OK**:

   Client ID

   Refers to the Application ID of the Azure Active Directory application. Copy your Client ID prepared in the text editor.

   Client Secret Key

   Refers to the authentication key generated for the Azure Active Directory application that you enter as Client ID. Copy your Client Secret Key prepared in the text editor.

   Important! This Secret Key is crucial for maintaining the security of your accounts. You should keep your keys and your account credentials in a secure location. Do not embed your Secret Key in a web page or other publicly accessible source code and do not transmit over insecure channels.

   Tenant ID

   Refers to the ID of the Azure Active Directory where you created the Azure Active Directory application. Copy your Tenant ID prepared in the text editor.

   Subscription ID

   Refers to a Globally Unique Identifier (GUID) that uniquely identifies your subscription to use Azure services. Copy your Subscription ID prepared in the text editor.

   Proxy Settings

   Specifies the proxy server settings. Select **Connect using a proxy server** to enable this option. If you select this option, you must also include the IP address (or machine name) of the proxy server and the corresponding port number that is used by the proxy server for internet connections. You can also select this option if your proxy server requires authentication. You then must

provide the corresponding authentication information that is required to use the proxy server.

The cloud account is added to the Console.

Create a Plan with a Backup Task

A plan includes different types of tasks that you want to perform. To create a virtual standby machine, you create a plan that includes a backup task and a virtual standby task. A backup task performs a backup of the source nodes and stores the data to the specified destination. The virtual standby feature uses the backup data and converts to a virtual machine format.

Follow these steps:

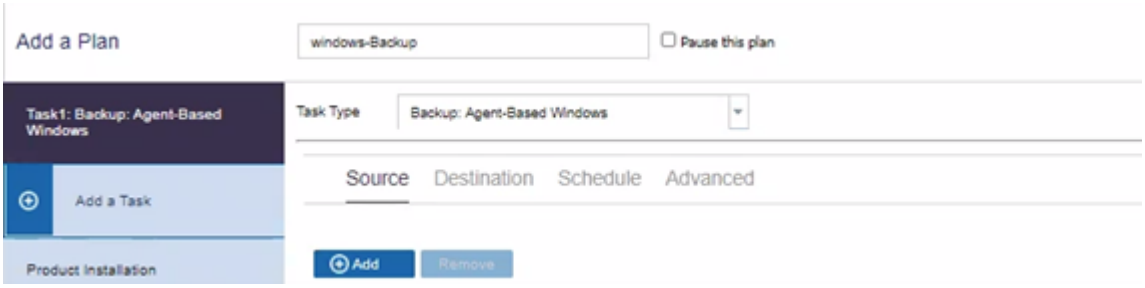
1. Click the **resources** tab on the Console.
2. From the left pane, navigate to **Plans**, and click **All Plans**.

   If you have created plans earlier, those plans are displayed on the center pane.
3. On the center pane, click **Add a Plan**.

   Add a Plan opens.
4. Enter a plan name.
5. (Optional) Select **Pause this plan** check box.

   The plan will not run until you clear the check box to resume the plan.

   Note: If a plan is paused, then any in-progress job is not paused but all corresponding scheduled jobs associated with that plan are paused. However, you can manually run a job. For example, you can manually run backup job and replication job for a node even if the respective plan is paused. In such case, the following task to the on-demand (manual) job does not run. For example, there is a replication task after an on-demand backup job, the replication job does not run for the on-demand backup job. You need to manually run the replication job. When you resume the plan, the pending jobs do not resume immediately. After you resume the plan, the pending jobs run from the next scheduled time.
6. From the **Task Type** drop-down list, select **Backup, Agent-Based Windows**.



Now, specify the Source, Destination, Schedule, and Advanced details.

Add a Virtual Standby Task to the Plan

Create a Virtual Standby to Azure task so that the backup data is converted to a virtual machine format and a virtual machine is created. The virtual standby feature also monitors the heartbeat of the source node so that when the source node is down, the virtual machine immediately takes over as the source node.

Notes:

- Virtual standby cannot automatically power on the recovery point snapshots that are taken from host-based virtual machine nodes, nodes replicated from a remote recovery point server, and the Source of the Virtual Standby task is the one replicated to a different Site. You have to manually power on recovery point snapshots for such nodes.
- When you pause the plan and resume again, the Virtual Standby job does not resume automatically. You must manually run another backup job to start the Virtual Standby job. Also, when you pause the plan, the Pause/Resume Virtual Standby option becomes unavailable. If you do not want the virtual machine to start automatically after you pause the plan, you must manually pause the heartbeat for the nodes.

Follow these steps:

1. Click **Add a Task** from the left pane.

   A new task is added to the left pane.

2. From the **Task Type** drop-down menu, select **Virtual Standby**.

   The Virtual Standby task is added.

3. From the **Source** tab select one source for the virtual standby task.

4. Click the **Virtualization Server** tab and enter the virtualization server and monitoring server details.

   Virtualization Type

   Specify Azure as the Virtualization Type.

Account Name

Select an existing Azure account or click **Add** to create an account.

For more information, see how to add a cloud account.

Resource Group

Select an existing resource group or click **Add** to create a resource group.

For more information, see Resource group in Azure.

Region

Select the region where you want the standby VM to operate in Azure. For more information, see Regions in Azure.

Monitor

Specify the host name of the server that monitors the status of the source server.

Notes:

- You can use any physical computer or virtual machine as the monitor server .
- You cannot use the backup source server as the monitor server.
- Monitor server configuration is not required if the nodes are replicated from a remote recovery point server or the Source of the Virtual Standby task is the one replicated to a different Site.
- Monitor server configuration is not required if the Virtual Standby Source is the replicate task and the replication target RPS server is inside Azure.

User Name

Specify the user name to log into the monitoring system.

Password

Specify the password for the user name to log into the monitoring system.

Protocol

Specify HTTP or HTTPS as the protocol that you want to use for communication between the Arcserve UDP and the monitoring server.

Port

Specify the port that you want to use for data transfer between Arcserve UDP and monitoring server.

5. Click the **Virtual Machine** tab and enter the details for VM Basic Settings, VM DataStore for VMware, VM path for Hyper-V, and VM Network.

VM Name Prefix

Specify a prefix that you want to add to the display name of virtual machine on Azure.

Default value: UDPVM_

Recovery Point Snapshots

Specify the number of recovery point snapshots (recovery points) for the standby virtual machine. The maximum number of recovery point snapshots count is 29.

Default value: 5

Combine all unconverted sessions into a single recovery point snapshot

Select if you want to combine all unconverted sessions into a single recovery point snapshot when next scheduled VSB job takes place.

Default: Selected

Virtual Machine Size

Microsoft Azure provides a wide selection of Virtual Machine Sizes optimized to suit different use cases. They have varying combinations of CPU, memory, storage, and networking capacity. For more information about Virtual Machine Size and how they meet your computing needs, view Sizes of Windows virtual machine in Azure.

Storage Account Name

Select an existing Storage Account Name or create a Storage Account in Azure. When you create Storage Account in Azure, you must specify one of the following as Account Kind

- Storage (general purpose v1)
- StorageV2 (general purpose v2)

For more information, see Storage account in Azure.

Virtual Network

Select an existing Virtual Network or create a Virtual Network in Azure. For more information, see Virtual Network in Azure.

Subnet

Select an existing Virtual Network Subnet based on the selected Virtual Network or add a Subnet in Azure. For more information, see Add Subnet in Azure.

Network Security Group

Select an existing Network Security Group or create a Network Security Group in Azure. Configure the security group rules to open the related ports, including

3389 for remote desktop, 8014, 8015 for Arcserve UDP communication. For more information, see Network Security Group.

Enable auto assign Public IP

Select to assign the public IP to Standby VM automatically when the standby VM starts in Azure.

6. Click the **Advanced** tab and provide the following details:

Automatically start the Virtual Machine

Specify if you want to start the virtual machine automatically.

Note: This option is not available for host-based virtual machine nodes and nodes replicated from a remote recovery point server and the Source of the Virtual Standby task is the one replicated to a different Site. The Virtual Standby Source is the replicate task and the replication target RPS server inside Azure.

Timeout

Specify the time that the monitor server must wait for a heartbeat before a recovery point snapshot is powered on .

Frequency

Specify the frequency that the source server communicates the heartbeats to the monitor server.

Example: The Timeout value specified is 60. The Frequency value specified is 10. The source server communicates heartbeats in 10-second intervals. If the monitoring server does not detect a heartbeat within 60 seconds of the last heartbeat that was detected, the monitor server powers on a virtual machine using the latest recovery point snapshot.

Customize job parameters

You can customize job parameters for the following options:

- *Number of threads uploading for each job*: Default Value: 4
- *Buffer size for each thread*: Default Value: 4096 KB

Enable Email Alerts

Select to enable email alerts. You receive email alerts based on the settings that you provide.

- **Missing heartbeat for source machine**--Virtual standby sends alert notifications when the monitor server does not detect a heartbeat from the source server.

Note: This option is not available for the nodes from Replicate from a remote Recovery Point Server or if the source of the Virtual Standby task is the one that is replicated to a different site.

- **VM powered on for source machine configured with auto power ON**--Virtual Standby sends alert notifications when a virtual machine that was configured to power on automatically when a heartbeat is not detected is powered on .

  Note: This option is not available for host-based virtual machine nodes and the nodes from Replicate from a remote Recovery Point Server or if the source of the Virtual Standby task is the one that is replicated to a different site.

- **VM powered on for source machine configured with manual power ON**--Virtual Standby sends alert notifications when a virtual machine is manually powered on.
- **Virtual Standby errors/failure/crash**--Virtual Standby sends alert notifications if an error is detected during the conversion process.
- **Virtual Standby success**--Virtual Standby sends alert notifications when a virtual machine is powered on successfully.
- **The Virtual Standby did not start successfully from the Recovery Point Snapshot**--Virtual Standby sends alert notifications if the Automatically start the Virtual Machine Stand-in Recovery option is enabled, but a virtual machine is not powered on automatically.

7. Click **Save**.

   The changes are saved and the virtual standby task is automatically deployed to the virtual standby server.

   Note: When the virtual standby task is complete, the virtual machine standby volume is created. The standby virtual machine is created only after the virtual machine is powered on from Arcserve UDP.

   You have successfully created and deployed the Virtual Standby to Azure plan.

Run the Virtual Standby Job Manually

To manually run a virtual standby job, you have to first perform a manual backup. The virtual standby task is associated with a backup task. If a plan includes a backup task and a virtual standby task and you manually run the backup job, the virtual standby job runs automatically after the completion of the backup job.

Follow these steps:

1. Click the **Resources** tab.
2. From the left pane, navigate to **Nodes**, and click **All Nodes**.

   Displays the plans that you added.
3. Select the nodes that you want to backup. The selected node must have assigned a plan.
4. On the center pane, click **Actions**, **Backup Now**.

   The **Run a backup now** dialog opens.
5. Select the backup type and provide a name for the backup job.
6. Click **OK**.

   The backup job runs.

   The virtual standby job runs immediately after the backup job is over.

   The virtual standby job is manually run.

Power-on Virtual Standby VM in Azure

The standby VM volumes are created in Azure after completion of Virtual Standby job. The standby VM is created only when powered on from Arcserve UDP.

Virtual standby can be configured to power on virtual standby machines from recovery point snapshots automatically when the monitoring server does not detect a heartbeat from the source server. Optionally, you can power on virtual standby machines from recovery point snapshots manually in the event a source server fails, an emergency occurs, or you want to offline a source node for maintenance.

Follow these steps:

1. From the **Resources** tab, navigate to the **Virtual Standby** node group.

   The virtual standby nodes are displayed on the center pane.
2. On the center pane, select the node and click **Standby VM**.

   The **Standby VM** dialog opens.
3. On the **Standby VM** dialog, perform the following tasks:
   - Select a date and time snapshot of the recovery point snapshot to power on the virtual machine.
   - Click **Power On VM**.

     The virtual machine is powered on using the data contained in the recovery point snapshot.

You can now verify the status or shut down the Virtual Standby VM. For more information, view the following:

- Verify Virtual Standby VM Status
- Shut Down Virtual Standby VM in Azure

Shut Down Virtual Standby VM in Azure

You can shut down Virtual Standby VM in Azure using Arcserve UDP Console.

Follow these steps:

1. From the Console, click **Resources**.
2. Navigate to **Virtual Standby**.
3. Select and right click the **Virtual Standby VM** from the center pane.

   The Standby VM dialog opens and displays the active snapshots.
4. Select an Active Snapshot and click **Shutdown VM**.
5. Click **Delete**.

   A confirmation dialog opens.
6. If you want to delete the disks attached to Virtual Standby VM, select the option **Delete the attached disk(s)**.
7. Click **OK**.

   The Virtual Standby VM is successfully shut down.

Verify Virtual Standby VM Status

You can verify the Virtual Standby VM status using Arcserve UDP Console.

Follow these steps:

1. From the Console, click **Resources**.
2. Navigate to **Virtual Standby**.
3. Select the **Virtual Standby VM** from the center pane.
4. The Configuration Wizard panel displays Virtual Standby VM status under the **Virtual Standby Status** group.

   The status shows as **Running** if the Virtual Standby VM is powered-on. If not powered-on, the status shows as **N/A** or **Powered Off**.

Arcserve UDP Terms and Definitions

Agent-Based Backup

An Agent-Based backup is a method to back up data using an agent component. The agent is installed on the source node.

Compression

Compression is used for backups. Compression is often selected to decrease disk space usage, but also has an inverse impact on your backup speed due to the increased CPU usage.

The available options are:

No Compression

This option has the lowest CPU usage (fastest speed), but also has the highest disk space usage for your backup image.

Standard Compression

Some compression is performed. This option provides a good balance between CPU usage and disk space usage. This is the default setting.

Maximum Compression

Maximum compression is performed. This option provides the highest CPU usage (lowest speed), but also has the lowest disk space usage for your backup image.

Notes:

- If your backup image contains uncompressible data (such as JPG images, ZIP files, and so on), you may need to allocate additional storage space to handle such data. As a result, if you select any compression option and have uncompressible data in your backup, it could result in an increase in disk space usage.
- If you change the compression level from No Compression to either Standard Compression or Maximum Compression, or if you change from either Standard Compression or Maximum Compression to No Compression, the first backup performed after this compression level change is automatically a Full Backup. After the Full Backup is performed, all future backups (Full, Incremental, or Verify) are performed as scheduled.

  This option is available only for the local or remote share destinations. You cannot change the compression setting if the Arcserve Unified Data Protection agent is backed up to data store.

- If your destination does not have sufficient free space, you may consider increasing the Compression setting of the backup. This option is available only for the local or remote share destinations. You cannot change the compression setting if the Arcserve Unified Data Protection agent is backed up to data store.

configuration

A tab on the Arcserve UDP Console to define configuration parameters such as email alerts, database settings, and installation preferences.

dashboard

A tab on the Arcserve UDP Console that lets you view the last Backup status and storage status. You can view the the latest Actual, Raw and Restorable Data storage.

Data Store

A data store is a physical storage area on a disk. You can create a data store on any Windows system where the recovery point server is installed. Data stores can be local or on a remote share that the Windows system can access.

Destination

Destination is a computer or server where you store backup data. A destination can be a local folder on the protected node, a remote shared folder, or a Recovery Point Server (RPS).

Discovered Nodes

Discovered nodes are physical or virtual systems that are added to the Arcserve UDP Console by discovering them from active directory or vCenter/ESX server, importing from a file, or manually adding them using its IP address.

Encryption

The Arcserve Unified Data Protection solution provides encryption feature for data.

When the backup destination is a recovery point server, the available encryptions are No Encryption and Encrypt data with AES-256. You can set this to create a data store. When the backup destination is the local or remote share, the available encrypt format options are No Encryption, AES-128, AES-192, and AES-256. You can set the option while creating a plan to backup to local or share folder, or set this from backup setting for standalone Arcserve Unified Data Protection Agent.

Key Features of Encryption

1. AES256 encryption method applies to:
   - Data Store
   - Password saved (in protection plan, registry, configuration file, and so on)
2. *For Backup job*: If encryption is enabled, data is encrypted before sending out of the server.
3. *For Replication job:* If replication destination has encryption enabled, data is encrypted before sending out of the server.

Encryption settings

a. Select the type of encryption algorithm that you want to use for backups.

   Data encryption is the translation of data into a form that is unintelligible without a deciphering mechanism. The Arcserve Unified Data Protection solution uses secure, AES (Advanced Encryption Standard) encryption algorithms to achieve maximum security and privacy of your specified data.

b. When an encryption algorithm is selected, provide (and confirm) an encryption password.
   - The encryption password is limited to a maximum of 23 characters.
   - A full backup and all its related incremental and verify backups must use same password to encrypt data.
   - If the encryption password for an incremental or verify backup is changed, a full backup must be performed. This means after changing encryption password, the first backup will be full, despite the original backup type.

   For example, if you change the encryption password and submit a customized incremental or verify backup manually, it automatically converts to a full backup.

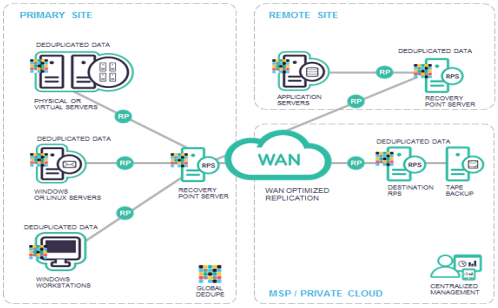
   Note: This option is available only for the local or remote share destinations. You cannot disable the encryption setting if the Arcserve Unified Data Protection agent is backed up to data store.

c. The Arcserve Unified Data Protection solution has encryption password and session password.
   - The encryption password is required for data store.
   - The session password is required for node.
   - If the data store is encrypted, then session password is mandatory. If the data store is not encrypted, the session password is optional.

A password is not required when you are attempting to restore to the computer from which the backup was performed. However, when you attempt to restore to a different computer, a password is required. By default, only for the first login password is required. To enter password even after the first login, the administrator needs to manually stop Arcserve UDP Agent Explorer Extension Service.

Global Deduplication

Arcserve UDP's global deduplication dramatically reduces the amount of data actually transferred during backup cycles. The ability to deduplicate across all the clients in the environment limits the unnecessary storage and transfer of existing data, making it global as data is deduplicated across nodes, jobs and across sites.



Host-Based Agentless Backup

A Host-Based Agentless backup is a method to back up data without using an agent component on the source machine.

HOTADD Transport Mode

The HOTADD transport mode is a data transport method that lets you back up virtual machines configured with SCSI disks. For more information, see the Virtual Disk API Programming Guide on the VMware website.

Job

A job is an Arcserve UDP action to back up, restore, create virtual standby, or replicate nodes.

jobs

A tab on the Arcserve UDP Console where you can monitor the status of all jobs such as backup, replication, and restore. The details include jobs, task types, node IDs, recovery points, and plan names.

NBD Transport Mode

Network Block Device (NBD) transport mode, also referred to as LAN transport mode, uses the Network File Copy (NFC) protocol to communicate. Various VDDK

and VCB operations use one connection for each virtual disk that it accesses on each ESX/ESXi Server host when using NBD.

NBDSSL Transport Mode

Network Block Device Secure Sockets Layer (NBDSSL) transport mode uses the Network File Copy (NFC) protocol to communicate. NBDSSL transfers encrypted data using TCP/IP communication networks.

Nodes

A node is a physical or virtual system that Arcserve UDP protects. Arcserve UDP can protect physical nodes and virtual machines in a vCenter/ESX or Microsoft Hyper-V server.

Plan

A plan is a group of tasks to manage backup, replication, and creation of virtual standby machines. A plan consists of a single or multiple tasks. Tasks are a set of activities to define the source, destination, schedule, and advanced parameters.

Protected Nodes

Protected nodes are the nodes that have scheduled backup plans to back up data on regular intervals.

Recent Event

Recent Events are the jobs that are still running or jobs that were recently completed.

Recovery Point

A recovery point is a point in time backup snapshot of a node. A recovery point is created when you back up a node. Recovery points are stored on the backup destination.

Recovery Point Server

A recovery point server is a destination node where you install the server. You can create data stores in a recovery point server. The Recovery Point Server (RPS) acts as a backup repository for disk images and offers a unique set of technologies that provide the fundamental building blocks of the Arcserve UDP solution. Among the

key features of RPS are True Source Side Global Deduplication, Proven Built-in Replication of Disk Images, RPS Jumpstart or “Offline” Synchronization, and Multi-Tenant Storage

Replicate

Replicate is a task that duplicates the recovery points from one server to another server.

Resources

resources is a tab on the Arcserve UDP Console. From the **resources** tab, you can manage source nodes, destinations, and plans.

SAN Transport Mode

The SAN (Storage Area Network) transport mode lets you transfer backup data from proxy systems connected to the SAN to storage devices.

Systems

Systems are all type of nodes, devices, and virtual machines that can be managed by Arcserve Unified Data Protection. This includes physical, virtual, Linux, and standby virtual machines.

Tasks

A task is a set of activities to define various parameters to back up, replicate, and create virtual standby machines. These parameters include source, destination, schedule, and some advanced parameters. Each task is associated with a plan. You can have more than one task in a plan.

Unprotected nodes

Unprotected nodes are the nodes that are added to Arcserve Unified Data Protection but a plan is not assigned. When a plan is not assigned, you cannot back up data and the node remains unprotected.

Data Deduplication

Data deduplication is technology that eliminates duplicate copies of the same data, thereby reducing storage space. In an organization, there could be various reasons for duplicate data such as a specific email attachment forwarded to multiple users. When you back up this data, you end up saving multiple copies of the same data on the backup storage media.

Data deduplication eliminates redundant data and saves only one instance of the data. All other instances are replaced with a reference to that instance. This method can considerably reduce the storage space that is required to store backup data.

For example, there could be a same 10 MB file that 100 users have stored in their local systems. When you back up all these local systems or nodes, you would need 1000 MB of storage space. With Data Deduplication, you can reduce the storage space to approximately 10 MB because only one instance of the file is stored on the disk. The remaining 99 instances refer to that one instance.

Benefits of Data Deduplication

- Stores more backup data in a storage space
- Reduces the amount of data that is sent over the network
- Performs speedy backup as reference information is stored rather than the actual data
- Reduces cost of network bandwidth and storage media

Types of Data Deduplication

Arcserve UDP supports the following two types of data deduplication.

Source-side Data Deduplication

Ensures that only unique data from the agent is sent to a recovery point server for data backup.

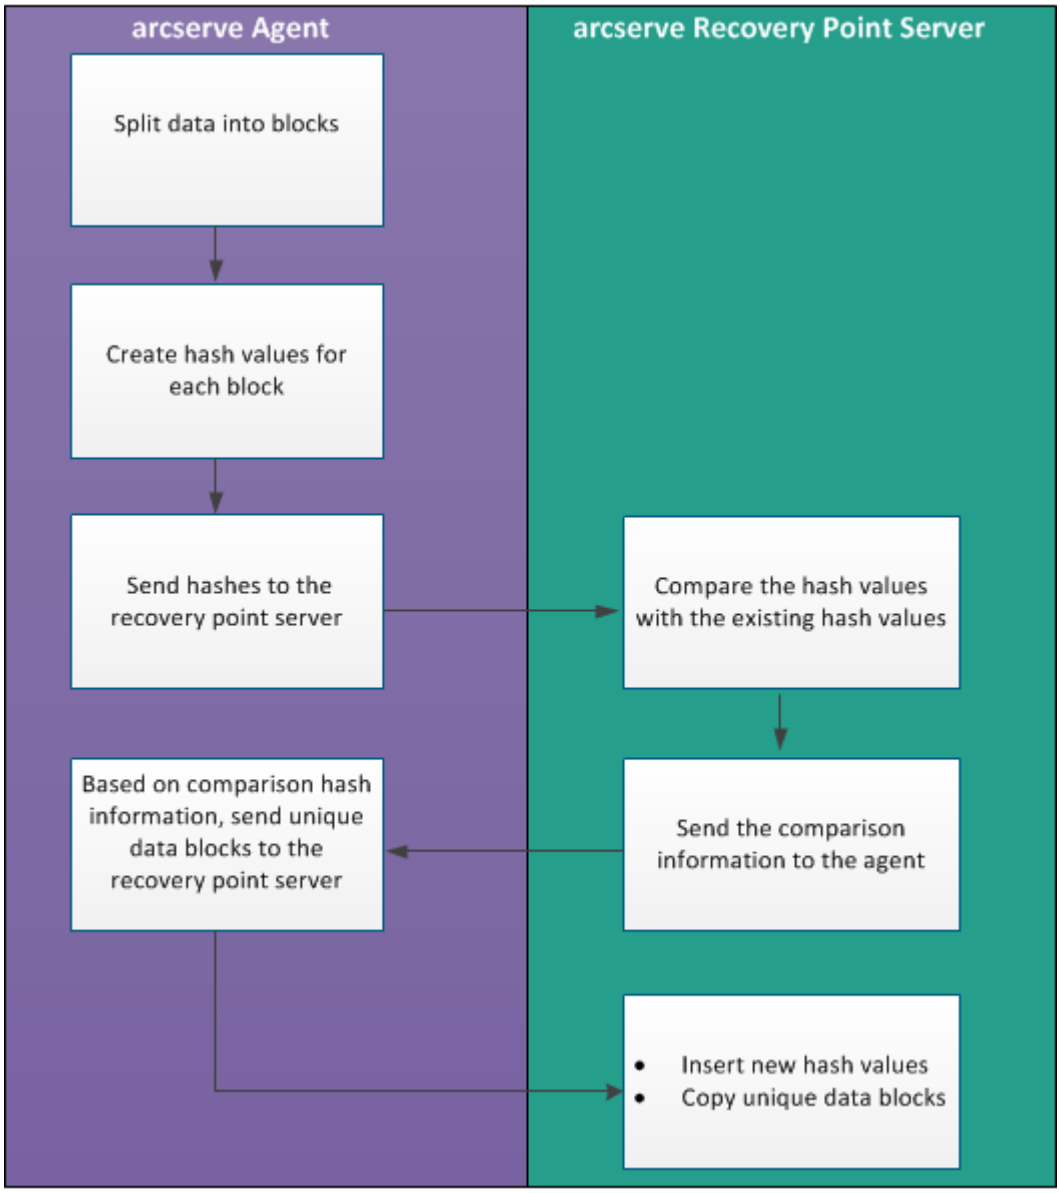
Global Deduplication

Ensures that only unique data from multiple agents are backed up to a recovery point server. If similar data blocks are present on multiple nodes, only one copy is backed up to the recovery point server.

How Data Deduplication Works

Arcserve UDP deduplication process splits data into data blocks and each block is assigned a unique identifier called hash. Hash is calculated based on the volume cluster. The default deduplication block size is 4KB (the default volume cluster size is 4KB for most of the nodes). These hash values are compared with the hash values of the existing backup data and if duplicate references are found, those data blocks are not backed up. Only data blocks with unique references are backed up.

The following diagram illustrates how deduplication works in Arcserve UDP:



When a backup is triggered, the deduplication process on the agent first splits data into blocks and assigns a unique hash key or value to each block. The hash values are then sent to the recovery point server. At the recovery point server, these hash

values are compared with the existing hash values and the duplicate hashes are filtered out. The comparison results are then sent back to the agent. Based on this duplicate hash information, the agent sends the unique data blocks to the recovery point server for backup. The new hash values of these data blocks are also inserted to its existing hash list on the recovery point server.

When there are multiple agents, the deduplication process remains the same, however, duplicate data from multiple agents are filtered out. This eliminates any duplication of data even from multiple agents.

The following are the benefits of using a Data Deduplication in Arcserve UDP:

- **Faster Full Backup**
- **Faster Merge Job**
- **Global Deduplication Support**
- **Optimized Replication**

When Should You Use Deduplication

The following are some of the scenarios where using a deduplication data store could be more effective:

- When you have multiple nodes with similar data. In this scenario if you back up data from all nodes to a data store, you would get a good reduction in the amount of data that is actually stored on the recovery point server. The storage space required could be considerably less.
- When you have to frequently take a full backup of a node. In this scenario, most of your backup data already exists so your backup time could be very less.
- When the network bandwidth is precious. As only unique data blocks travel across the network, you can reduce the network usage.
- When backed-up data frequently moves from one node to another. In this scenario, when you try to back up the new node (where the data moves from its original node), the destination already contains the copy and only the reference information is backed up.

Configuring Deduplication Data Stores in Arcserve UDP

The following are the important parameters to configure for a deduplication data store:

Data destination

Data destination is used to store the protected data. It is better to use larger disk for the data destination because it contains the original data blocks of the source.

Index Destination

Index destination is used to store the index files and it is better to use a different disk to improve the deduplication processing throughput.

Hash destination

Hash destination is used to store the hash files and it is better to use to high speed SSD drive which can improve the deduplication capacity with a low memory allocation required.

If hash destination is configured on a high speed SSD, it could be used to enlarge deduplication capacity with low memory allocation requirement.

Backup destination folder

The destination folder where .D2D files and catalog files reside.

Block size

The “deduplication block size” also impacts the “deduplication capacity estimation”. The default “deduplication block size” is 16 KB. If you set it to 32 KB, then the “deduplication capacity estimation” is doubled. The impact of increasing the deduplication block size is that it can decrease the deduplication percentage and at the same time the memory requirement decreases.

Memory Allocation

To estimate the memory requirement, use the “Estimate Memory and Storage Requirements” tool. If the Memory Allocated is not enough and when the memory is fully used, the new data cannot insert new hash into hash DB. So, any data that are backed up after that cannot be Deduplicated, causing the Dedupe ratio to go down. If you cannot increase the memory for some reason, then try increasing the deduplication block size as it would decrease the memory requirement.

Note: Block Size cannot be changed for an existing data store.

Be aware that a new backup job is not allowed to launch once hash memory is full. But for the ongoing backup job (which was launched before the hash memory is full), it is allowed to continue and get completed. In this case, it would not insert new hash keys to hash database. As a result, impacting the dedupe percentage.

The reason is that all data blocks in the ongoing backup job are still compared with the existing hash keys in the hash database,

• If it is duplicated with the existing hash key, it is not written to the disk any more.

• If it is not duplicated with the existing hash key, it is written to disk. But the new hash key would not be inserted into hash database because hash database is full. As a result, the consequent data blocks could not compare against these new hash keys.

Deduplication, Encryption, and Compression

In addition to data deduplication, we can also apply compression and encryption on a data store.

If you enable encryption, the Arcserve Unified Data Protection Agent for Windows consumes the CPU resource to encrypt the data. As encryption is applied only to the unique data, the CPU resource needed for encryption could be minimum where the deduplication percentage is high.

- With no compression and deduplication, the CPU usage is less for the compression task and the data stored is in the non-compressed format.
- With standard compression and deduplication, the CPU usage is optimal for the compression task and the data stored is in a compressed format and the requirement for storage space is less.
- With maximum compression and deduplication, CPU usage is maximum for the compression task and the data stored is 2-3% more and the requirement for storage space is less.

Deduplication Limitations

You cannot modify compression type, Encryption Setting, and Deduplication Block Size once you create a deduplication data store.