

Arcserve® Unified Data Protection

Solutions Guide

Version 5.0

arcserve®

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.

Copyright © 2015 Arcserve (USA), LLC and its affiliates and subsidiaries. All rights reserved. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

Arcserve Product References

This document references the following Arcserve products:

- Arcserve® Backup
- Arcserve® Unified Data Protection Agent for Windows
- Arcserve® Unified Data Protection Agent for Linux
- Arcserve® High 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.

<https://www.arcserve.com/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: Understanding Arcserve UDP 15

Introduction	15
How Arcserve UDP Works	17
Task-Based Scenarios	19

Chapter 2: Installing Arcserve UDP 23

How to Install Arcserve UDP	23
Review the Prerequisites and Considerations.....	25
Decide the Installation Type	25
Install Arcserve UDP Using the Setup Wizard	26
Install Arcserve UDP Using the Command Line.....	31
Verify the Installation.....	36
Communication Ports Used by Arcserve UDP	36
How the Installation Process Affects Operating Systems	41
How to Install Arcserve UDP Updates	47
Review the Considerations for Installing Updates	49
Specify Updates Preference	50
Check and Install the Updates.....	53
(Optional) Install Arcserve UDP Updates Silently.....	54
Verify that the Updates are Successfully Installed	55
How to Uninstall Arcserve UDP	56
Standard Uninstall.....	56
Silent Uninstall	57
Remove Components Left Behind by the Uninstaller	58
How to Manage Arcserve UDP Licenses.....	60
Review the Prerequisite	61
Add a License	62
Remove a License.....	63
Verify the License	63
How to Upgrade to Arcserve UDP	64
No Charge Edition (NCE).....	65

Chapter 3: Exploring and Configuring Arcserve UDP 67

Arcserve UDP User Interface	67
Navigating Arcserve UDP	68
Tabs	69

Job Monitor Dialog	80
How to Configure Arcserve UDP	81
Configure Server Communication Protocol	82
Configure Database	83
Configure Arcserve UDP Backup Data Synchronization	86
Configure SRM	87
Node Discovery Configuration	89
Configure Email and Alert	90
Update Configuration	92
Configure Administrator Account	93
Installation Settings	94
Map the Plan to the User Account	94
How to Migrate Arcserve r16.5 Recovery Points to Arcserve UDP	96
Create a Data Store to Replicate Data from a Arcserve r16.5 Recovery Point	96
Replicate Arcserve r16.5 Data to the UDP Data Store	97

Chapter 4: Adding and Managing Source Nodes 99

How to Add Nodes to the Console	99
Review the Prerequisites	101
Add Nodes	102
Discover Nodes	105
Import Nodes	106
How to Manage Nodes	111
Review the Prerequisites	112
Update Hypervisor Information	112
Specify the Hypervisor	113
Update VM Information	115
Update Nodes	116
Export Node	118
Synchronize Data	119
Delete Nodes from the Console	119
Deploy Agent to Nodes	120
Deploy Agent to Nodes	121
Perform Preflight Checks for Your Backup Jobs	121
How to Add and Manage Node Groups	133
Review the Prerequisites	134
Add Node Groups	134
Modify Node Groups	135
Delete Node Groups	136

Chapter 5: Adding and Managing Destinations 137

How to Add a Destination	137
Review the Prerequisites	138
Add a Recovery Point Server	139
(Optional) Deploy the Recovery Point Server	141
Add a Data Store	143
Verify the Destination	145
How to Manage a Data Store	146
Review the Prerequisites	147
Modify a Data Store	148
Delete a Data Store from the Console	156
Stop a Data Store	157
Start a Data Store	158
Browse Recovery Points in a Data Store	158
Delete Node Data from a Data Store	160
Troubleshooting: How to Use a Data Store When One or Multiple folders are Full.....	160
How to Manage a Recovery Point Server	161
Review the Prerequisites	162
Update a Recovery Point Server	162
Delete a Recovery Point Server from the Console	163
Import a Data Store.....	164
Install/Upgrade Recovery Point Server	165

Chapter 6: Creating Plans to Protect Data 167

How to Create a Windows Backup Plan	170
Review the Prerequisites and Considerations.....	171
Create a Backup Plan with a Backup Task	174
(Optional) Perform a Manual Backup	188
Verify the Backup	188
How to Create a Linux Backup Plan	189
Review the Prerequisites and Considerations.....	190
Create a Backup Plan	190
(Optional) Perform a Manual Backup	204
Verify the Backup	205
Troubleshooting	206
How to Create a Host-Based Virtual Machine Backup Plan	207
Review the Prerequisites and Considerations.....	208
Create a Host-Based Backup Plan	210
(Optional) Perform a Manual Backup	234
Verify the Plan	235
Troubleshooting	235

How to Create a Virtual Standby Plan	241
Review the Prerequisites and Considerations.....	243
Create a Plan with a Backup Task.....	244
Add a Virtual Standby Task to the Plan	258
(Optional) Run the Virtual Standby Job Manually.....	269
Pause and Resume Heartbeat	270
Pause and Resume Virtual Standby Job	271
Verify the Plan	272
Applying Best Practices	273
How to Create a Virtual Standby Plan for Arcserve High Availability Nodes	274
Review the Prerequisites for Remote Virtual Standby.....	275
Create a Virtual Standby Plan for HA Nodes	275
Configure Remote Converters.....	285
Verify the Plan	286
How to View Virtual Standby Settings From the Monitor Server.....	287
Review the Prerequisites and Considerations.....	289
Log In To Monitor Servers	289
Understanding the Virtual Standby Summary Screen.....	290
View Activity Log	293
View Virtual Standby Settings	293
View the Email Settings.....	298
How to Protect Virtual Standby Machines	299
Review the Prerequisites and Considerations.....	300
Power On Virtual Standby Machines	301
Protect Virtual Standby Machines After it is Powered On	305
Verify the Virtual Standby Machine is Protected	306
How to Replicate Data Between Data Stores Managed from a UDP Console	307
Review the Prerequisites and Considerations.....	308
Create a Backup Plan with a Backup Task	308
Add a Replicate Task to the Plan.....	322
(Optional) Perform a Manual Replication	324
Verify the Plan	325
How to Replicate Data Between Data Stores Managed From Different UDP Consoles.....	325
Review the Prerequisites	327
Create a User Account for the Source Console	327
Create a Plan to Define the Destination Data Store.....	328
Map the Plan to the User Account.....	330
Send the Plan and User Account Details to the Source Administrator	331
Receive the Plan and User Account Details from the Destination Administrator.....	332
Create a Replication Plan to Send Data to the Destination Console.....	332
Verify the Data is Replicated	337
Applying Best Practices	337

How to Perform an Offline Data Replication Using RPS Jumpstart	339
Review the Prerequisites	341
Create a Temporary Data Store on an External Device.....	342
Replicate Source Data to the Temporary Data Store	343
Delete the Temporary Data Store from the Source Console	344
Send the External Device to the Destination Location.....	344
Receive the External Device	344
Import the Temporary Data Store from the External Device.....	345
Create a Destination Data Store	345
Replicate Data from the Temporary Data Store to the Destination Data Store	346
Verify that the Data is Replicated	346
(Optional) Set the Concurrent Node Count for RPS Jumpstart	347
How to Create a Copy Recovery Points Plan	347
Review the Prerequisites and Considerations.....	348
Create a Plan with a Backup Task.....	349
Add a Copy Recovery Points Task to the Plan	363
Verify the Plan	365
How to Create a File Copy Plan	366
Review the Prerequisites and Considerations.....	367
Create a Plan with a Backup Task.....	368
Add a File Copy Task to the Plan	382
Verify the Plan	392

Chapter 7: Restoring Protected Data 395

How to Restore From a Recovery Point	395
Review the Restore Prerequisites and Considerations	397
Specify the Recovery Point Information to Restore.....	401
Restore the Recovery Point Content.....	408
Verify that Content was Restored	409
How to Restore From a File Copy	409
Review the Restore Prerequisites and Considerations	411
Specify the File Copy Information to Restore	412
Restore the File Copy Content	421
Verify that Content was Restored	422
How to Restore Files/Folders	424
Review the Restore Prerequisites and Considerations	425
Specify the File/Folder Information to Restore.....	429
Restore the File/Folder	441
Verify that the File/Folder was Restored	442
How to Restore a Virtual Machine	443
Review the Restore Prerequisites and Considerations	445

Specify the Virtual Machine Information to Restore	446
Restore the Virtual Machine	460
Verify that the Virtual Machine was Restored	464
How to Restore Microsoft Exchange Mail	464
Review the Restore Prerequisites and Considerations	466
Specify the Microsoft Exchange Mail Information to Restore	468
Restore the Microsoft Exchange Mail	479
Verify that the Microsoft Exchange Mail was Restored	480
How to Restore a Microsoft Exchange Application	481
Review the Restore Prerequisites and Considerations	483
Specify the Microsoft Exchange Information to Restore	485
Restore the Microsoft Exchange Application	490
Verify that the Microsoft Exchange Application was Restored	491
How to Restore Exchange Mail on a VMware Virtual Machine	492
Review the Restore Prerequisites and Considerations	495
Specify the Exchange Mailbox Database	497
Select the Exchange Mail Objects to be Restored	504
Define the Restore Options	505
Restore the Exchange Mail	508
Verify that the Exchange Mail was Restored	510
How to Restore a Microsoft SQL Server Application	511
Review the Restore Prerequisites and Considerations	512
Specify the Microsoft SQL Server Information to Restore	514
Restore the Microsoft SQL Server Application	519
Verify that the Microsoft SQL Server Application was Restored	520
How to Restore a Pervasive PSQL Database	520
Review the Restore Prerequisites and Considerations	522
Determine the Database and Location to Restore	522
Verify that the Pervasive PSQL Database was Restored	534
How to Restore an Oracle Database	534
Review the Prerequisites and Considerations	536
Restore the Server Parameter File	536
Restore the Parameter File	537
Restore the Archived Redo Logs	538
Restore the Tablespaces or Data Files	538
Restore System, or Undo Tablespaces or Data Files	539
Restore All Tablespaces and Data Files	541
Restore Control Files	542
Restore the Entire Database (Tablespaces and Control Files)	544
Recover the Oracle Database Using Bare Metal Recovery	545
How to Perform a File-Level Recovery on Linux Nodes	546
Review the Prerequisites	548

(Optional) Recover Data from the iSCSI Volume to the Target Machine	549
Specify the Recovery Point.....	550
Specify the Target Machine Details.....	554
Specify the Advanced Settings	558
Create and Run the Restore Job.....	562
Verify that Files are Restored.....	563
How to Perform a Bare Metal Recovery (BMR) for Linux Machines	563
Review the BMR Prerequisites.....	565
Get the IP Address of the Target Machine Using the Live CD	566
(Optional) Recover Data to the iSCSI Volume of the Target Machine	567
(Optional) Recover Data from the iSCSI Volume to the Target Machine	568
Review the Backup Server.....	569
Specify the Recovery Points	570
Specify the Target Machine Details.....	572
Specify the Advanced Settings	573
Create and Run the Restore Job.....	577
Verify that the Target Node is Restored	584
How to Perform a BMR Using a Backup	584
Review the BMR Prerequisites and Considerations	586
Define BMR Options.....	587
Verify that the BMR was Successful	605
BMR Reference Information	605
Troubleshooting BMR Issues.....	611
How to Perform a BMR Using a Virtual Standby VM	614
Review the BMR Prerequisites and Considerations	616
Define BMR Options.....	617
Verify that the BMR was Successful	639
BMR Reference Information	640
Troubleshooting BMR Issues.....	645
How to Restore Microsoft Clustered Nodes and Shared Disks	648
Review the Prerequisites	649
Restore Files of a Cluster Shared Disk	650
Restore a Specific Node in a Cluster	650
Restore a Corrupted Cluster Shared Disk.....	651
Restore the Entire Clustered Nodes and Shared Disk	651

Chapter 8: Generating Arcserve UDP Reports 653

How to Generate Arcserve UDP Reports	653
Using Filters and Actions.....	655
Generate a Report.....	656
Schedule Emails.....	657

Send Report by Email	660
Chapter 9: Managing Arcserve High Availability	663
How Arcserve High Availability Works	663
Manage HA Control Services	663
Manage HA Licenses	664
Manage Scenarios	664
Remote Installation	676
High Availability Reports	678
Appendix A: Troubleshooting	679
Add Permissions for VDDK 5.1 and 5.5 at vCenter Server Level	679
Operating System Not Found	679
Virtual Standby Jobs Fail Due to Internal Errors.....	680
Virtual Standby Jobs Fail Using the HotAdd Transport Mode	681
Virtual Standby Jobs to Hyper-V Systems Fail	683
Data Store Name is Already in Use.....	684
Failed to perform backup of a virtual disk. System error=[The device is not ready(21)].	684
The Arcserve UDP Agent Service Runs Slowly	685
Failed to Create a Snapshot for Hyper-V Virtual Machines When Multiple Jobs are Running	687
Convert Incremental Backup to a Verify Backup Because the Virtual Machine Snapshots Either Changed from the Last Backup Job or Needs Consolidation.....	688
Settings disabled when opening agent UI	688
Backup destination settings disabled when opening the Linux agent UI	689
Scheduled Incremental or Full Backup Job Fails for Hyper-V VM	690
Unable to Restore Files	691
Troubleshooting Backup Job Failure after Changing the Console Hostname/IP Address	691
Data Store Switches to the Restore Only Mode.....	693
Backup Job Fails.....	695
Pause or Resume Fails When Agent Is Not Connected to the Network.....	696
Arcserve UDP Cannot Communicate with the Arcserve UDP Agent (Windows) Web Service on Remote Nodes	697
Unable to Apply Backup Settings to Node	698
The Hyper-V VSS NTDS Writer Fails While Taking the VSS snapshot in the VM.....	698
Host-based Backup Does Not Use HotAdd Transport Mode.....	699
MAC Address Changes are Not Retained After VM Recovery	699
Agentless Host-based Backup for Hyper-V VM Fails after Upgrading Arcserve UDP	700
Appendix B: Data Deduplication	701
Types of Data Deduplication	701

How Data Deduplication Works	702
When Should You Use Deduplication.....	703
Configuring Deduplication Data Stores in Arcserve UDP	704
Deduplication, Encryption, and Compression	705
Deduplication Limitations	705

Appendix C: Command Line Data Integrity Tool for Deduplication Data Store 707

Appendix D: Arcserve UDP Terms and Definitions 709

Agent-Based Backup.....	709
Compression.....	709
configuration	710
dashboard	710
Data Store	710
Destination	710
Discovered Nodes.....	710
Encryption	711
Host-Based Agentless Backup	712
HOTADD Transport Mode	712
Job	712
jobs.....	712
NBD Transport Mode	712
NBDSSL Transport Mode	713
Nodes	713
Plan.....	713
Protected Nodes.....	713
Recent Event	713
Recovery Point	713
Recovery Point Server	713
Replicate.....	713
Resources	714
SAN Transport Mode.....	714
Systems	714
Tasks.....	714
Unprotected nodes	714

Chapter 1: Understanding Arcserve UDP

This section contains the following topics:

[Introduction](#) (see page 15)

[How Arcserve UDP Works](#) (see page 17)

[Task-Based Scenarios](#) (see page 19)

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 or Microsoft Hyper-V 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
- Backs up data to recovery point servers
- Replicates backup data to recovery point servers and remote recovery point servers
- Copies selected source files to a secondary backup location
- Copies recovery points to an additional location
- Creates virtual standby machines from backup data
- Restores backup data and performs Bare Metal Recovery (BMR)
- Monitors Arcserve High Availability

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 or Microsoft Hyper-V virtual machine 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.

The following diagram illustrates the major functions that Arcserve UDP lets you perform.



How Arcserve UDP Works

Arcserve UDP is a unified data protection solution that lets you protect your computer systems. The following are the high level steps that you need to follow 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 and Hyper-V 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, and creation of virtual standby machines.
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.



Task-Based Scenarios

Task-based scenarios are information modules that provide everything you want to know about performing a task, all in one easy-to-follow location.

Arcserve Unified Data Protection:

- How to Install Arcserve UDP
- How to Install Arcserve UDP Updates
- How to Uninstall Arcserve UDP
- How to Manage Arcserve UDP Licenses
- How to Add a Destination
- How to Add Nodes
- How to Manage a Data Store
- How to Manage a Recovery Point Server
- How to Manage Nodes
- How to Create a Windows Backup Plan
- How to Create a Host-Based Virtual Machine Backup Plan
- How to Create a Virtual Standby Plan
- How to Add and Manage Node Groups
- How to Create a Linux Backup Plan
- How to Create a Copy Recovery Points Plan
- How to Create a File Copy Plan
- How to Replicate Data Between Data Stores Managed from a UDP Console
- How to Replicate Data Between Data Stores Managed from Different UDP Consoles
- How to Perform an Offline Data Replication Using RPS Jumpstart
- How to Create a Virtual Standby Plan for Arcserve HA Nodes
- How to View Virtual Standby Settings from the Monitor Server
- How to Protect Virtual Standby Machines
- How to Perform a Bare Metal Recovery (BMR) for Linux Machines
- How to Perform a Bare Metal Recovery Using a Backup
- How to Perform a Bare Metal Recovery Using a Virtual Standby VM
- How to Perform a File-Level Recovery on Linux Nodes
- How to Restore a Microsoft SQL Server Application

- How to Restore from a Recovery Point
- How to Restore from a File Copy
- How to Restore Files/Folders
- How to Restore a Virtual Machine
- How to Restore Microsoft Exchange Mail
- How to Restore a Microsoft Exchange Application
- How to Restore Microsoft Clustered Nodes and Shared Disks

Arcserve UDP Agent for Windows:

- How to Install Arcserve UDP Agent (Windows)
- How to Install Arcserve UDP Agent (Windows) Updates
- How to Uninstall Arcserve UDP Agent (Windows)
- How to Navigate the Arcserve UDP Agent (Windows) User Interface
- How to Perform a Backup
- How to Restore from a Recovery Point
- How to Restore from a File Copy
- How to Restore Files/Folders
- How to Restore a File/Folder for an Agentless Virtual Machine
- How to Restore a Virtual Machine
- How to Restore Microsoft Exchange Mail
- How to Restore a Microsoft Exchange Application
- How to Restore a Microsoft SQL Server Application
- How to Restore a Pervasive PSQL Database
- How to Restore an Oracle Database
- How to Restore an Active Directory
- How to Perform an Authoritative Restore of an Active Directory after a BMR
- How to Restore Microsoft Clustered Nodes and Shared Disks
- How to Perform a Bare Metal Recovery Using a Backup
- How to Perform a Bare Metal Recovery Using a Virtual Standby VM
- How to Copy a Recovery Point
- How to Create a Boot Kit

Arcserve UDP Agent for Linux:

- How to Install Arcserve UDP Agent (Linux)

- How to Upgrade Arcserve UDP Agent (Linux)
- How to Uninstall Arcserve UDP Agent (Linux)
- How to Navigate the Arcserve UDP Agent (Linux) User Interface
- How to Manage the Licenses
- How to Manage Jobs
- How to Back Up Linux Nodes
- How to Modify and Rerun a Backup Job
- How to Perform a File Level Recovery
- How to Create a Bootable Live CD
- How to Create a CentOS-Based Live CD
- How to Perform a Bare Metal Recovery (BMR) for Linux Machines
- How to Automatically Recover a Virtual Machine
- How to Restore Volumes on a Target Node
- How to Restore an Oracle Database Using Arcserve UDP Agent (Linux)
- How to Integrate and Automate Arcserve UDP Agent (Linux) with the Existing IT Environment
- How to Manage the Backup Server Settings
- How to Manage the Backup Server from the Command Line
- How to Manage the Non-Root Users

Chapter 2: Installing Arcserve UDP

This section contains the following topics:

- [How to Install Arcserve UDP](#) (see page 23)
- [How to Install Arcserve UDP Updates](#) (see page 47)
- [How to Uninstall Arcserve UDP](#) (see page 56)
- [How to Manage Arcserve UDP Licenses](#) (see page 60)
- [How to Upgrade to Arcserve UDP](#) (see page 64)
- [No Charge Edition \(NCE\)](#) (see page 65)


How to Install Arcserve UDP

As a storage administrator, you manage data backup and restore of machines in your network setup. Using Arcserve UDP, you can centrally manage and protect Windows nodes, Linux nodes, and virtual machines in VMware ESX servers or Microsoft Hyper-V servers. The Arcserve UDP installation provides you with the following options:

- **Arcserve UDP - Full:** Installs all the Arcserve UDP components. You can install this on a system from where you want to manage data protection functions. This system must meet the hardware requirements necessary to install Arcserve UDP. For information about the supported systems, see the Arcserve UDP Release Notes. After the installation, you log in to the Arcserve UDP Console (Console) and perform data management functions. The Console lets you can 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.

The following diagram illustrates how to install Arcserve UDP.

How to Install Arcserve Unified Data Protection



What To Do Next?

1. [Review the Prerequisites and Considerations](#) (see page 25)
2. [Decide the Installation Type](#) (see page 25)
3. [Install Arcserve UDP Using the Setup Wizard](#) (see page 26)
4. [Install Arcserve UDP Using the Command Line](#) (see page 31)
5. [Verify the Installation](#) (see page 36)
6. [\(Optional\) Communication Ports Used](#) (see page 36)
7. [\(Optional\) How the Installation Process Affects Operating Systems](#) (see page 41)

Review the Prerequisites and Considerations

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

Prerequisites

- Review the Arcserve UDP Release Notes. The Release Notes contain 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 UDP 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 UDP components.
- Verify that you have the user names and passwords of the systems where you are installing Arcserve UDP components.
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Considerations

Before the installation, you should determine how to set up your Arcserve UDP 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 UDP 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.

Install Arcserve UDP 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 will be required 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 on the **License Agreement** dialog 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 UDP Agent (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 UDP Agent (Windows) change tracking driver and click **Next**.

By default, this option is selected.

- Without this driver installed, Arcserve UDP Agent (Windows) cannot perform a local backup.
- With this driver installed, you would still need to have a valid Arcserve UDP Agent (Windows) license to perform a local backup.

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

The **Destination Folder** dialog opens.

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. To use the SSL protocol for a hostname that contains an underscore (_) character, you must manually run the following batch file before using the UDP Agent or Console:

UDP Agent: INSTALLDIR \Management\BIN\changeToHttps.bat

UDP Console: INSTALLDIR \Management\BIN\changeToHttps.bat

- b. Enter the port number for the agent.
- c. Enter the port number for the Console.
- 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 2008 R2 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. Do one of the following:

Microsoft SQL Server 2008 R2 Express (included):

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

- a. Specify the location where you want to install Microsoft SQL Server 2008 R2 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 2008 R2 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.
 - 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 **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**.

- c. 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.

13. Click **Finish**.

A dialog informing you that a system restart is required and asking if you want to reboot immediately or later.

When the reboot is finished, Arcserve Unified Data Protection is installed on your computer.

Applicable to Update 2: When installing Update 2 in a multi-node environment, it is important that all associated agent nodes also install Update 2 and you do not have a mix of agent nodes at different update levels trying to provide the same protection.

Install Arcserve UDP Using the Command Line

You can install Arcserve UDP 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 UDP.

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\CA\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 UDP Agent.

Default: 8014

Example:

-AgentPort:8014

Note: Use this option when you want to install the Arcserve UDP Agent.

-Driver:<DRIVER>

(Optional) Specifies whether to install Arcserve UDP 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 UDP 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 CA ARCserve Universal Agent service.

0: Does not stop the CA ARCserve Universal Agent service if it is running during the installation process.

1: Stops the CA 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.

4. Restart the target computer after you complete the silent installation.

Verify the Installation

To verify the installation, confirm the existence of the arcserve UDP Agent and Recovery Point Server services from the Windows Services dialog. You can also confirm the presence of Arcserve UDP by clicking Start, All Programs on the Windows operating system.

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 UDP and you are ready to back up your Windows machine.

Communication Ports Used by Arcserve UDP

The following table lists the TCP and UDP port numbers used by Arcserve UDP.

Components installed on Microsoft Windows

Port Number	Port Type	Initiated by	Listening Process	Description
8015	TCP	UDP Console	Tomcat7.exe	Default HTTP/HTTPS communication port between remote management consoles and the UDP server. Default HTTP/HTTPS communication port between remote management consoles and the UDP Agent. Note: You can modify the default communication port when you install the UDP components.

Port Number	Port Type	Initiated by	Listening Process	Description
8014	TCP	UDP Agent	Tomcat7.exe	Default HTTP/HTTPS communication port between remote management consoles and the UDP server. Default HTTP/HTTPS communication port between remote management consoles and the UDP Agent. Note: You can modify the default communication port when you install the UDP components.
8014	TCP	UDP Server	httpd.exe	Default HTTP/HTTPS communication port between the UDP server and UDP consoles. *Default shared port and the only port you must open when you use the UDP server as the replication destination. Do not open port 5000-5060 used by data stores with global deduplication enabled. Note: You can modify the default communication port when you install the UDP components.

Port Number	Port Type	Initiated by	Listening Process	Description
8016	TCP	UDP Server	Tomcat7.exe	Reserved for UDP Server 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.
1433	TCP	Remote Java	sqlsrvr.exe	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.
5000-5060	TCP	UDP Server	GDDServer.exe	Reserved for UDP RPS Global Deduplication DataStore Service. One UDP GDD data store will use 3 free ports start from 5000. It is needed when data store with GDD enabled for backup or restore task is used.
4090	TCP	UDP Agent	HATransServer.exe	To transfer data for Virtual Standby task in the proxy mode.
8006				To shutdown Tomcat used by the UDP console.
18005			CA.ARCserve.CommunicationFoundation.Windowsservice.exe	To shutdown Tomcat used by the UDP server or Agent.

Port Number	Port Type	Initiated by	Listening Process	Description
6052	TCP	ARCserve Backup GDB	CA.ARCserve.CommunicationFoundation.WindowsService.exe	CA.ARCserve.CommunicationFoundation.WindowsService.exe Communication that lets the Console and the CA ARCserve Backup Global Dashboard Primary server synchronize data.
6054	TCP	ARCserve Backup Primary server		CA.ARCserve.CommunicationFoundation.WindowsService.exe Communication that lets the Console and the CA ARCserve Backup Primary server synchronize data.

Components installed on Linux

Port Number	Port Type	Initiated by	Listening Process	Description
67	UDP	UDP Linux	bootpd	Incoming, used for PXE boot server. Only required if user wants to use the PXE boot feature. Note: The port numbers cannot be customized.
69	UDP	UDP Linux	tffpd	Incoming, used for the PXE boot server. Only required if user wants to use the PXE boot feature. Note: The port numbers cannot be customized.

Port Number	Port Type	Initiated by	Listening Process	Description
8014	TCP	UDP Linux	Java	Both incoming and outgoing. Default HTTP/HTTPS communication port between the remote consoles and the UDP agent for Linux .
18005	TCP	UDP Linux	Java	Used by Tomcat. Ignore this port for the firewall setting.
22	TCP	SSH service		UDP Linux 3rd party dependency. Default for SSH service, however, you can change this port. This port is required for both incoming and outgoing communications.

Production node protected by UDP Linux remotely

Port Number	Port Type	Initiated by	Listening Process	Description
22	Port Type	SSH service		UDP Linux 3rd party dependency. Default for SSH service, however, you can change this port. This port is required for both incoming and outgoing communications.

Ports listed above 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 need to be open.

Similarly, for remote replications, the Remote administrator need 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.

How the Installation Process Affects Operating Systems

The following installation processes update various Windows Operating Systems:

Installation of Unsigned Binary Files

Arcserve UDP installs binary files that are developed by third parties, other Arcserve products, and Arcserve UDP that are not signed. The following table describes these binary files.

Binary Name	Source
ab.exe	Apache
abs.exe	Apache
ApacheMonitor.exe	Apache
apr_dbd_odbc-1.dll	Apache
apr_ldap-1.dll	Apache
htcacheclean.exe	Apache
htdbm.exe	Apache
htdigest.exe	Apache
htpasswd.exe	Apache
httpd.exe	Apache
httxt2dbm.exe	Apache
libapr-1.dll	Apache
libapriconv-1.dll	Apache
libaprutil-1.dll	Apache
libeay32.dll	OpenSSL
libhttpd.dll	Apache

logresolve.exe	Apache
openssl.exe	Apache
rotatelog.exe	Apache
ssleay32.dll	OpenSSL
wintty.exe	Apache
zlib1.dll	Apache
libbind9.dll	ISC bind
libdns.dll	ISC bind
libisc.dll	ISC bind
libiscfg.dll	ISC bind
liblwres.dll	ISC bind
msvcm80.dll	Microsoft
msvcpr80.dll	Microsoft
msvcr80.dll	Microsoft
win_nsupdate.exe	ISC bind
msvcm90.dll	Microsoft
sqlite3.exe	SQLite
zlib10.dll	Zlib Compression Library
tcnative-1.dll	Tomcat
tomcat7.exe	Tomcat
AxShockwaveFlashObjects.dll	Adobe
ShockwaveFlashObjects.dll	Adobe
LogSet_logo-win12r2_20140417_232307.cab	Microsoft
BaseLicInst.exe	CA License
CALicense.msi	CA License
CALLicense.msi	CA License
BaseLicense.exe	CA License

Installation of Binary Files with Incorrect File Version

Arcserve UDP installs binary files that are developed by third parties, other Arcserve products, and Arcserve UDP that contain incorrect file version information. The following table describes these binary files.

Binary Name	Source
apr_dbd_odbc-1.dll	Apache
openssl.exe	Apache
zlib1.dll	Apache
libbind9.dll	ISC bind
libdns.dll	ISC bind
libisc.dll	ISC bind
libiscfg.dll	ISC bind
liblwres.dll	ISC bind
win_nsupdate.exe	ISC bind
decora-d3d.dll	Java Runtime Environment
decora-sse.dll	Java Runtime Environment
fxplugins.dll	Java Runtime Environment
glass.dll	Java Runtime Environment
glib-lite.dll	Java Runtime Environment
gstreamer-lite.dll	Java Runtime Environment
javafx-font.dll	Java Runtime Environment
javafx-iiio.dll	Java Runtime Environment
jfxmedia.dll	Java Runtime Environment
jfxwebkit.dll	Java Runtime Environment
libxml2.dll	Java Runtime Environment
libxslt.dll	Java Runtime Environment
prism-d3d.dll	Java Runtime Environment
libcurl.dll	VMware
liblber.dll	VMware
libldap.dll	VMware
libldap_r.dll	VMware
sqlite3.exe	SQLite
zlib10.dll	Zlib Compression Library

AxShockwaveFlashObjects.dll	Adobe
ShockwaveFlashObjects.dll	Adobe
sqljdbc_auth.dll	Java Runtime Environment
UpdateData.exe	CA License
dc21x4vm.sys	Intel
NETwew00.sys	Intel
NETwew02.sys	Intel
netwlv64.sys	Intel
NETwNs64.sys	Intel
Netwsw00.sys	Intel
CNN08CL1FX.dll	Canon
CNN08CL2FX.dll	Canon
dedrvor.dll	Microsoft
dedrvpj.dll	Microsoft
dedrvsc.dll	Microsoft
dedrvzd.dll	Microsoft
dexpsff1.dll	Microsoft
hpbresw81.dll	Microsoft
hpbx3w81.dll	Microsoft
hpcf1tw8.dll	Microsoft
hpcf1twb.dll	Microsoft
hpcstw81.dll	Microsoft
hpicl3.dll	Microsoft
hpires.dll	Microsoft
LXPTMV.dll	Microsoft
LXPJLMW.dll	Microsoft
sadrvor.dll	Microsoft
sadrvpj.dll	Microsoft
sadrvc.dll	Microsoft
sadrvd.dll	Microsoft
smxpsff1.dll	Microsoft

Installation of Binary Files Without OS in Manifest

Arcserve UDP installs binary files that are developed by third parties, other Arcserve products, and Arcserve UDP that do not contain the operating system in manifest or have executables with manifest but do not support the latest operating system.

Binary Name	Source
openssl.exe	Apache
sqlite3.exe	SQLite
CALicnse.exe	CA License
CAMinfo.exe	CA License
CAregit.exe	CA License
ErrBox.exe	CA License
lic98log.exe	CA License
lic98Service.exe	CA License
lic98version.exe	CA License
LicDebug.exe	CA License
LicRCmd.exe	CA License
LogWatNT.exe	CA License
mergecalic.exe	CA License
mergeolf.exe	CA License
ab.exe	Apache
abs.exe	Apache
ApacheMonitor.exe	Apache
htcacheclean.exe	Apache
htdbm.exe	Apache
htdigest.exe	Apache
htpasswd.exe	Apache
httpd.exe	Apache
httxt2dbm.exe	Apache
logresolve.exe	Apache
rotatelog.exe	Apache
wintty.exe	Apache
win_nsupdate.exe	ISC bind
jabswitch.exe	Java Runtime Environment

unpack200.exe	Java Runtime Environment
vmware-vdiskmanager.exe	VMware
vmware-mount.exe	VMware
tomcat7.exe	Tomcat
vcredis_x64.exe	Microsoft
vcredis_x86.exe	Microsoft
BaseLicInst.exe	CA License
silent.exe	CA License
UpdateData.exe	CA License
DatabaseMail.exe	Microsoft SQL Server
DCEXEC.EXE	Microsoft SQL Server
SQLAGENT.EXE	Microsoft SQL Server
SQLIOSIM.EXE	Microsoft SQL Server
sqlmaint.exe	Microsoft SQL Server
sqlservr.exe	Microsoft SQL Server
sqlsubss.exe	Microsoft SQL Server
xpdsi.exe	Microsoft SQL Server
java.exe	Java Runtime Environment
javacpl.exe	Java Runtime Environment
java-rmi.exe	Java Runtime Environment
javaw.exe	Java Runtime Environment
javaws.exe	Java Runtime Environment
jp2launcher.exe	Java Runtime Environment
keytool.exe	Java Runtime Environment
kinit.exe	Java Runtime Environment
klist.exe	Java Runtime Environment
ktab.exe	Java Runtime Environment
orbd.exe	Java Runtime Environment
pack200.exe	Java Runtime Environment
policytool.exe	Java Runtime Environment
rmid.exe	Java Runtime Environment
rmiregistry.exe	Java Runtime Environment


servertool.exe	Java Runtime Environment
ssvagent.exe	Java Runtime Environment
tnameserv.exe	Java Runtime Environment
jqs.exe	Java Runtime Environment

How to Install Arcserve UDP Updates

The process of getting and installing Arcserve UDP updates is a two-part process: 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 for building another server with the same configuration/patch level.

The following diagram illustrates the process to install Arcserve UDP updates:



Perform the following tasks to install Arcserve UDP updates:

1. [Review the Considerations for Installing Updates](#) (see page 49)
2. [Specify Update Preferences](#) (see page 50)
3. [Check and Install the Updates](#) (see page 53)
4. (Optional) [Install Arcserve UDP Updates Silently](#) (see page 54)
5. [Verify that the Updates are Successfully Installed](#) (see page 55)

Review the Considerations for Installing Updates

Review the following considerations before installing Arcserve UDP updates:

- When installing a Arcserve UDP update or a Arcserve UDP Agent (Windows) update, it is important to maintain optimal performance between the Console, the Recovery Point Server (RPS), and the 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, you can 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, you can 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.
- When updating a Arcserve UDP update, a system reboot may be required.
- If you are installing Arcserve UDP Version 5.0 Update 2, please be aware of the following issue:
 - A problem has recently been noticed that seems to cause unexpected reboots of servers after Update 2 is installed. This reboot occurs without any prior notification or indication to the user that a reboot will occur after the Update 2 installation is complete.
 - To avoid any possible service disruption due to unplanned reboots of production systems, Arcserve has removed Update 2 from being available for download and installation via the online Update system from within the product.
 - However, you can still install Update 2 manually to ensure a safe installation and you will then be provided with the option to reboot your system immediately or at a later (more convenient) point in time.
 - For more information about this issue and the solution, please refer to the corresponding KB article: [Arcserve UDP Update 2 is available exclusively as a manual download & upgrade from the regular direct download link.](#)
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Specify Updates Preference

Arcserve UDP lets you specify the following Updates preference.

Follow these steps:

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

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

The screenshot shows the Arcserve UDP Configuration console. At the top, the 'configuration' tab is selected. On the left sidebar, 'Update Configuration' is highlighted. The main content area is titled 'Updates' and contains three sections: 'Download Server', 'Test Connection', and 'Update Schedule'. In the 'Download Server' section, 'Arcserve Server' is selected with a radio button, and a 'Proxy Settings' button with a green checkmark is visible. The 'Test Connection' section has a 'Test Connection' button. The 'Update Schedule' section has a checkbox for 'Automatically check for updates' which is unchecked, and a time picker set to 'Every Sunday At 3 : 00 AM'. At the bottom right, there are 'Save', 'Reset', and 'Help' buttons.

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.

If you specify more than one staging server, the first listed server will be designated as the primary staging server. Arcserve UDP will initially attempt to connect to the primary staging server. If for any reason the first listed server is not available, then the next listed server will become the primary staging server. The same sequence will be 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 UDP Console can get update from this staging server.
- If the specified staging server does not have any update, then the UDP Console will not be able to download update from this staging server. The log displays the following message: No new update available.


Note: You cannot download Arcserve UDP updates from a staging server if HTTPS is enabled on it for web communication.

■ Proxy Settings

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

Select **Proxy Settings** to specify if you want the Arcserve UDP updates to be downloaded via a proxy server. A proxy server acts as an intermediary between your download server (staging or client) and the Arcserve server to ensure security, increased performance, and administrative control. This will be the connection to the Arcserve server from which your download server will get the updates.

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



The image shows a 'Proxy Settings' dialog box with a title bar and a close button. It contains two radio buttons: 'Use browser proxy settings (for IE and Chrome only)' and 'Configure proxy settings'. The second option is selected. Below the first option is a note: 'Note: Administrator log in credentials will be used as the proxy credentials.' The 'Configure proxy settings' section has a 'Proxy Server' text box with the placeholder '<Proxy Server Name>', a 'Port' text box, and a checked checkbox labeled 'Proxy server requires authentication'. Below the checkbox are 'User Name' and 'Password' text boxes. The 'User Name' box has the placeholder '<domain name>\<user name>'. The 'Password' box is filled with dots. At the bottom are 'OK', 'Cancel', and 'Help' buttons.

- **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 display 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 UDP Console: Update server unavailable.

Note: The test connection is automatically performed when you open the **Update Configuration** page from the **configuration** tab in the 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 will be 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.

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.

(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

Perform *one* of the following to verify that the updates are successfully installed:

- From the Arcserve UDP Console, click **log** and then verify that the installed updates are listed in the activity logs.
- 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 UDP 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 UDP Console
- Arcserve UDP Recovery Point Server
- Arcserve UDP 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 in to the computer to uninstall Arcserve UDP components.
Note: Log in to 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%\CA\SharedComponents\arcserve Unified Data  
Protection\Setup\uninstall /q /ALL
```

To uninstall selected components

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

■ x64 operating system:

To uninstall all components

```
%ProgramFiles(x86)%\CA\SharedComponents\arcserve Unified Data  
Protection\Setup\uninstall /q /ALL
```

To uninstall selected components

```
%ProgramFiles(x86)%\CA\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 UDP component that you want to uninstall.

Example:

The following syntax lets you uninstall Arcserve UDP Recovery Point Server silently.

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

Component	<Product Code>

Arcserve UDP Agent (x86 platforms)	{CAAD8AEA-A455-4A9F-9B48-C3838976646A}
Arcserve UDP Agent (x64 platforms)	{CAAD1E08-FC33-462F-B5F8-DE9B765F2C1E}
Arcserve UDP Recovery Point Server	{CAAD8172-1858-4DC7-AE81-C887FA6AFB19}
Arcserve UDP Console	{CAAD3E40-C804-4FF0-B1C0-26D534D438C0}

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

Remove Components Left Behind by the Uninstaller

Uninstallation of Arcserve UDP leaves behind some components such as the CA Licensing Components, SQL Server Express, Microsoft Visual C++ components installed as dependency, and driver related (wdf) files. These components consist of multiple individual files, which are installed and removed with the corresponding component. The CA Licensing Components file is not automatically removed during uninstallation because it is a shared component with other CA products and numerous components.

Important! The CA licensing is shared by all CA products, please make sure that you do not have any other CA product installed on your machine or else you may lose the licensing for all CA products installed on that machine.

Important! If the components are removed, any programs that are installed after Arcserve UDP Agent (Windows) and depend on these components may not function properly.

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

Remove CA Licensing Component manually

1. Go to *C:\Program Files (x86)\CA\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)\CA
 - C:\Program Files\CA
 - Folder where you extracted the zip file.
6. Remove the registry key for the CA Licensing component.
 - For x64 platform:
HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\ComputerAssociates\Licensing
 - For x86 platform:
HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\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++ 2010 x86 Redistributable - 10.0.40219* and then click Uninstall.
3. Select *Microsoft Visual C++ 2010 x64 Redistributable - 10.0.40219* and then click Uninstall.
4. Select Microsoft SQL Server 2008 R2 and then click Uninstall.
5. To remove only the Arcserve UDP database, select "ARCSERVE_APP" and click Uninstall.

How to Manage Arcserve UDP Licenses

Using Arcserve UDP, you can centrally manage licenses for all the recovery point servers, physical nodes, and virtual nodes that are added to the Console. The licensing model grants a single overall license to the application with a predetermined number of active license rights included in the overall license pool.


Each new user of the application (member server) is granted an active license from the license pool on a first-come, first-served basis until the total number of available licenses is exhausted. If all the active licenses are already in use and you want to license a new server, you have to manually release a license from one of the licensed servers and then apply that license to the new server.

For all scenarios, when there are no licenses available, you will get an error message in the Activity Log.

You can easily remove license rights to allow other member servers to gain license privileges. From the Console, you can access the License Management dialog and view the active license counts for each component. You can also manage which licenses are applied to which servers.

You can use Arcserve UDP for a trial period. At the end of the trial period, if you have not obtained a license, Arcserve UDP will automatically revert to a [No Charge Edition](#) (see page 65) (NCE) with limited capabilities.

The following diagram displays the process to manage licenses:



What To Do Next?

- [Review the Prerequisite](#) (see page 61)
- [Add a License](#) (see page 62)
- [Remove a License](#) (see page 63)
- [Verify the License](#) (see page 63)

Review the Prerequisite

Review the following prerequisites before managing the licenses:

- You have installed Arcserve UDP.
- You have a valid license.
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Add a License

Arcserve UDP protects only the licensed nodes. If there are enough licenses, then the licenses are automatically applied to the nodes. If you do not want to protect any node, you can release the license from that node and use that license to protect any other node.

Follow these steps:

1. Log in to the Console.
2. Click **Help, Manage Licenses**.
The **License Management** dialog opens.
3. Check the license key on your media case or on your license certificate.
4. Enter the license key on the **License Management** dialog and click **Add**.

License Management

All licenses

License Name	Total	Remaining	Unlicensed

Add **Refresh**

Licensed Nodes

☒ Node Name	Node Type

Page **Release**

Close **Help**


5. Close and open the **License Management** dialog.
The license is added and is listed on the **All licenses** area.

Remove a License

If you do not want to protect any node, you can release the license from that node. You can use the released license to protect any other node.

Follow these steps:

1. Log in to the Console.
2. Click **Help, Manage Licenses**.
The **License Management** dialog opens.
3. From the right pane, select **Licensed Nodes** from the drop-down list.



4. Select a node from the displayed list and click **Release**.
5. Close and open the **License Management** dialog.
The license is removed from the node.

Verify the License

Verify that the correct license is applied to a node. Run a backup job for the node. If the backup is successful, the node is licensed.

How to Upgrade to Arcserve UDP

When you install Arcserve Unified Data Protection to a computer that already has a previous version of any Arcserve product, Arcserve UDP will prompt you for a confirmation. The following table describes various types of upgrade scenarios.

Existing arcserve products	Upgrades to	Considerations
Arcserve D2D (r15, r16, r16.5)	Arcserve Unified Data Protection Agent for Windows or Arcserve Unified Data Protection	If you install Arcserve Unified Data Protection - Full, then the setup wizard also installs the Console and Recovery Point Server. All the D2D settings are migrated. Note: If a D2D node is managed by Arcserve Central Protection Manager, then the settings cannot be migrated.
Arcserve Central Protection Manager	Arcserve Unified Data Protection	No settings are migrated. You have to create new plans.
Arcserve Central Host-Based VM Backup, Arcserve Central Reporting, Arcserve Central Virtual Standby	Arcserve Unified Data Protection	No settings are migrated. You have to create new plans. If you use the same backup destination as previously used for Arcserve D2D r16 or r16.5, you must run a backup job before you run the first virtual standby job.

No Charge Edition (NCE)

Starting with Arcserve UDP Version 5.0 Update 2, at the end of the trial period, a free and fully functional No Charge Edition (NCE) is provided for any users who have not yet obtained a proper license. This NCE is for use on workstation-class hardware (Laptops or Desktops running Microsoft Client Operating Systems) and continues to provide full and complete access to all functions and features that were available during the trial period, with some limited capabilities.

Highlights:

- After the trial period expires, the Workstation Edition (trial period edition) will automatically revert to the NCE.
- Your NCE nodes can still be managed from the Arcserve UDP Console.
- Provides a very simple key-based upgrade path to the full Arcserve UDP "Workstation Edition"
- You can perform backup to a local disk, or to a shared folder, or to any other supported destination that is not an RPS without requiring a license key.
- Using NCE, you will not be able to select an RPS as the backup destination. As a result, you will lose the ability to leverage the Global Deduplication feature, which dramatically reduces the amount of data actually transferred during backup cycles. This feature is available after upgrading to the full Workstation Edition.
- You will not have Live Chat capabilities, but will still be able to email Arcserve Support for questions or to resolve issues.

Frequently Asked Questions:

Q. Can I use the trial version to test all features of Arcserve UDP?

A. Yes, you can use the trial version to leverage all the great features of Arcserve UDP until trial period expires. When the trial period expires, the Workstation Edition of Arcserve UDP will automatically revert to the NCE.

Q. What will happen if a Recovery Point Server (RPS) is selected as the destination for an NCE node?

A. You can still select an RPS as your backup destination under certain conditions. If your Arcserve UDP environment has license counts available, they will be consumed on a need-basis.

Q. Does Arcserve UDP know when to consume a license?

A. Arcserve UDP is intelligent enough to determine which nodes need a license, and will only use (consume) a license when required. As a result, if you are performing a backup to a shared folder you will not consume a license. However, if you select an RPS as your destination it will consume a license (if a license is available). You could then leverage (select) an RPS as your backup destination from your NCE node and it would consume one of the available licenses (but no longer be an NCE node).

Q. Does NCE work for server-class operating systems such as Windows 2012?

A. No. NCE is only for use on desktops and laptops running on any of the supported Windows client operating system (such as Windows 7, 8, or 8.1). You should check the [Compatibility Matrix](#) to view a list of all supported operating systems.

Q. What about product support for NCE?

A. You can leverage support for NCE by connecting to the online community based support, directly from within the product. With the full Workstation Edition, you can leverage some of the more enhanced and speedy support offerings such as the "Live Chat" capability, which is unavailable for the NCE.

Chapter 3: Exploring and Configuring Arcserve UDP

This section contains the following topics:

[Arcserve UDP User Interface](#) (see page 67)

[How to Configure Arcserve UDP](#) (see page 81)

[How to Migrate Arcserve r16.5 Recovery Points to Arcserve UDP](#) (see page 96)


Arcserve UDP User Interface

Before you use Arcserve UDP, become familiar with the user interface. The Arcserve UDP 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

After you install Arcserve UDP, you can log into the Console using the username and password that you specified during installation. The Arcserve UDP console lets you manage all Arcserve UDP functions. The following basic UI elements are available throughout the Arcserve UDP user interface.



Tabs

Lets you navigate to the various functions of Arcserve UDP.

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 UDP solution provides the following tabs to perform data protection functions.



dashboard

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

- Customize the position of the four graphs. Place the mouse over the name of any of the four options, and you can drag one graph to other location on the screen.
- Click one of the two options available for every screen if you want to refresh or maximize a screen. Click Refresh icons in any of the four screens to get the latest data size. Click Maximize icon of a screen to view only that screen in the dashboard.
- View the last backup status of nodes or plans according to the filters you select in the chart in the **Last Backup Status**.



Copyright © 2015 Arcserve (USA), LLC and its affiliates and subsidiaries. All rights reserved.

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

Last Backup Status

Last Backup Status refers to the latest backup status and provides you multiple filters to view the status. Based on your selection from the filter option, you can view the last backup status. For example, select **All Nodes** to see the last backup status of all nodes or select any plan to see the last backup status of the nodes protected by this plan. When you select **All Nodes**, you can see the status as **Successful**, **Failed**, **No Backups**, **Canceled**, and **Missed**. **Successful** indicates that the nodes are successfully backed up. **Failed** indicates that the last backup is not successful. **No Backups** indicate that the nodes do not have any plan associated with it. **Canceled** indicates that the last backup was stopped. **Missed** indicates that the last backup was not performed as scheduled.

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 **No Backups** from the pie chart, the **resources** page opens. The **resources** page displays the nodes that do not have any plan. Also, the **No Backups** filter is preselected on the **resources** page.

Raw Data and Actual Data Storage

The graph refers to the raw data and actual data storage. You can click **Raw Data** or **Actual Data Storage** to hide or view information about any of the two options. You can hover to the point to see the detail data size information using the tooltip.

Raw data

Refers to the original data that Arcserve UDP gets from source.

Actual Data

Refers to the data size that is saved on the disk after being compressed or deduplicated by Arcserve UDP.

Restorable Data and Actual Data Storage

The graph refers to the data that you can restore and the actual data storage. You can click **Restorable Data** or **Actual Data Storage** to hide or view information about any of the two options. You can hover to the point to see the detail data size information using the tooltip.

Restorable Data

Refers to the actual data that can be restored.

Actual Data Storage

The graph displays information about the actual data storage. You can click **Actual Data Storage** to hide or view information about the actual data storage. You can hover to the point to see the detail data size information using the tooltip.

resources

The **resources** tab lets you manage the Arcserve UDP resources: Nodes, Destinations, Virtual Standby, and Plans. Use this tab to add resources to Arcserve UDP 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.

The screenshot shows the Arcserve UDP interface with the **resources** tab selected. The interface is divided into several sections:

- Left Navigation Panel:** Contains a tree view with categories: **Nodes** (All Nodes, Nodes without a Plan, Plan Groups), **Virtual Standby** (All Nodes, Action Required, Standby VM Running, Source Running, Source and VM Running), **Destinations** (Recovery Point Servers), and **Plans** (All Plans).
- Top Section:** Displays "Nodes: All Nodes" with a filter dropdown set to "(No filter applied)". Below this is a table with columns: **Node Name**, **Plan**, and **PFC Status**. The table lists three nodes: "<Node 1>", "<Node 2>", and "<Node 2>".
- Right Panel:** Displays details for "<Node 1>". It includes a **Status** section with a green checkmark, a **Last Backup Job (Incremental)** summary, a **Last Virtual Standby Job** summary, and a **Recent Events** section with a "View Logs" link.

Annotations with green arrows point to specific features:

- Navigate Resources Group:** Points to the left navigation panel.
- Perform Actions like Add and Update Nodes:** Points to the "Add Nodes" button in the top section.
- Apply Filters:** Points to the filter dropdown in the top section.
- View Resource Summary:** Points to the "<Node 1>" header in the right panel.
- Corresponding Node Details:** Points to the row for "<Node 1>" in the table.

	Node Name	Plan	PFC Status
☒	<Node 1>	sf	
☐	<Node 2>	sf	
☐	<Node 2>		

<Node 1> Details:

- Status:** ✔
- Last Backup Job (Incremental):** 3/27/2014 11:00:00 AM, Duration: 00:01:29
- Last Virtual Standby Job:** 3/27/2014 11:01:40 AM, Duration: 00:03:41
- More Virtual Standby Information:** [More Virtual Standby Information](#)
- Virtual Standby Status:** ✔ Active
- Source Status:** ✔ Running
- Standby VM Status:** ⏻ Powered Off
- Recent Events:** [View Logs](#)
- Virtual Standby:** ✔ 3/27/2014 11:01:40 AM
- Backup - Incremental:** ✔ 3/27/2014 11:00:00 AM
- Virtual Standby:** ✔ 3/27/2014 9:09:50 AM

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 you select in the center pane.

resources

Nodes

- All Nodes
- Nodes without a Plan
- Plan Groups
 - sf
- Virtual Standby
 - All Nodes
 - Action Required
 - Standby VM Running
 - Source Running
 - Source and VM Running
- Destinations
 - Recovery Point Servers
- Plans
 - All Plans

Nodes: All Nodes

Actions ▾ | Add Nodes | Filter ⌵ (No filter applied) >>

Filter Name

Node Status	Protection Failure	Protection Type	Application
☐ Protected	☐ Backup Failure	☐ Backup	☐ SQL Server
☐ Unprotected	☐ Restore Failure	☐ Virtual Standby	☐ Exchange
	☐ Merge Failure	☐ Replication	
	☐ Catalog Failure		
	☐ Replication Failure		
	☐ Virtual Standby Failure		

Apply Reset Save Delete

	Node Name	Plan	PFC Status
☒	<Node 1>	sf	
☐	<Node 2>	sf	
☐	<Node 2>		

<Node 1>

Actions ▾

Status ✓

- ✓ **Last Backup Job (Incremental)**
3/27/2014 11:00:00 AM Duration: 00:01:29
- ✓ **Last Virtual Standby Job**
3/27/2014 11:01:40 AM Duration: 00:03:41
- [More Virtual Standby Information](#)
- ✓ Virtual Standby Status Active
- ✓ Source Status Running
- ⏻ Standby VM Status Powered Off

Recent Events [View Logs](#)

- ✓ Virtual Standby 3/27/2014 11:01:40 AM
- ✓ Backup - Incremental 3/27/2014 11:00:00 AM
- ✓ Virtual Standby 3/27/2014 9:09:50 AM
- ✓ Backup - Full 3/27/2014 8:00:01 AM

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.

resources

Nodes

All Nodes

Nodes without a Plan

▶ vCenter/ESX Groups

▶ Hyper-V Groups

▶ VM Backup Proxy Groups

▶ Plan Groups

Destinations

Recovery Point Servers

Plans

All Plans

Destinations: Recovery Point Servers

Actions ▾ | [Add a Recovery Point Server](#)

Name	Plan Count	Data Protected
<RPS ID>		
✓ <Data Store Name>	13	20.5 TB

RPS ID > Data Store Name

Actions ▾

▼ Status ✓

Running

Backup Destination 9.5 TB free of 15.5 TB

Data destination 9.5 TB free of 15.5 TB

Index destination 9.5 TB free of 15.5 TB

Hash destination 194.3 GB free of 371.4 GB

Memory allocation 48.7 GB free of 64 GB

▶ Recent Events [View Logs](#)

▼ Settings

Compression Type: Standard
Compression

Backup Destination: D:\15TBcommon_1814

Concurrent Active Nodes: No Limit

Deduplication

Data File Path: D:\15TB_data

Hash File Path: C:\15tb_hash_1814

Index File Path: D:\15tb_index_1814

Hash Memory Size: 8192 MB

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.

resources

- Nodes
 - All Nodes
 - Nodes without a Plan
- Plan Groups
 - d2d-rps
- Virtual Standby
 - All Nodes
 - Action Required
 - Standby VM Running
 - Source Running
 - Source and VM Running
- Destinations
 - Recovery Point Servers
- Plans
 - All Plans

Plans: All Plans

Actions ▾ | [Add a Plan](#)

	Plan Name	Nodes Protected				Status	Active Job Count
		Total	✓	⚠	✗		
☒	<Plan>	1	1	0	0	✓ Deployed: Successful	0

<Plan>

Actions ▾

Settings ▾

Task 1 Backup: Agent-Based Windows

Source

Destination

Schedule

Daily Backup

Backup Type: Incremental
Start Time: 10:00 PM

Start time for scheduled backup: 4/27/2014 9:50 PM

Recovery points
Daily: 7
Custom / Manual: 31

Catalogs
Daily Backups: No
Custom / Manual Backups: No
Exchange Server: No

Advanced

Task 2 Virtual Standby

Source

Virtualization Server

Virtual Machine

jobs

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

jobs

Most Recent Jobs

- All Jobs Completed
- Jobs Successfully Completed
- Jobs Failed
- Jobs Canceled
- Jobs In Progress

Most Recent Jobs : All Jobs Completed

☐ Group Jobs by Plan [Refresh](#)

Status	Task	Node Name	Job Time
✓	Virtual Standby	<Node>	4/28/2014 10:01:45 PM
✓	Backup - Incremental	<Node>	4/28/2014 10:00:05 PM

<Plan>

4/28/2014 10:01:45 PM

Status: Finished

Duration: 00:04:45

Job Details

Job ID: 4

Node Name: <Node>

Task: Virtual Standby

ESX Host/Center: <IP Address>

ESX Node: localhost.ca.com

History: [View Logs](#)

When a job is in progress, the right pane displays the job monitor that displays 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](#) (see page 653).

reports

- Alert
 - Alert Report
- Data Trend
 - Backup Size Trend Report
- Backup
 - Node Backup Status Report
 - Virtualization Protection Status Report
 - Managed Capacity Report
- Data Distribution
 - Data Distribution on Media Report

Filters/Actions

Global Action

Job NodesAll

Last7Days

GroupsAll Nodes

Node TierAll Tiers

Managed Capacity Report

Name of Report

Filters/Actions

Local Action

Job NodesAll

Protected Nodes

GroupsAll Nodes

Node TierAll Tiers

Unit:TB

Job Nodes	Protected Nodes	Status	Plan	Product	Installed Application	Last Successful Backup Time

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.

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 so on and click Search.

Note: You cannot delete the activity logs.

log						
Severity	Warning and Error	Node Name		Job ID		Job Type All
Time	All	Generated From		Message		
Severity	Time	Node Name	Generated From	Job ID	Job Type	Message
ⓘ	3/28/2014 3:01:38 PM	<node 1>	<node 1>	6	Backup	Snapshot storage area for volume does not have enough free space. The snapshot storage area will be switched to another volume to continue the backup. (Storage Area=[\?\Volume{f4145744-50c5-11e3-adc7-806e6f6e6963}], Volume Name=[\?\Volume{f4145744-50c5-11e3-adc7-806e6f6e6963}])
ⓘ	3/28/2014 10:13:02 AM	<node 2>	<node 2>	1251	Backup	The system volume [System Reserved] is configured as a mirrored volume. As a result, the system will fail to boot if you are attempting to restore data using BMR.
ⓘ	3/28/2014 10:01:10 AM	<node 3>	<node 3>	1394	Backup	Snapshot storage area for volume does not have enough free space. The snapshot storage area will be switched to another volume to continue the backup. (Storage Area=[\?\Volume{22cc83ef-48e2-11e3-9712-806e6f6e6963}], Volume Name=[\?\Volume{22cc83ef-48e2-11e3-9712-806e6f6e6963}])
ⓘ	3/28/2014 9:58:45 AM	<vm name>	<node ID>	102	Backup	Virtual machine is powered off.

configuration

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

For more information about the **configuration** tab, see [How to Configure Arcserve UDP](#) (see page 81).

The screenshot displays the 'configuration' tab in the Arcserve UDP user interface. On the left is a vertical sidebar with a list of configuration options: Database Configuration, arcserve Backup Data Synchronization Schedule (highlighted), SRM Configuration, Active Directory Discovery Configuration, Email and Alert Configuration, Update Configuration, Administrator Account, Installation Settings, and Share Plan. The main content area is titled 'configuration' and contains the following settings:

- Enable/Disable:** Radio buttons for 'Enable' (selected) and 'Disable'.
- Repeat Method:** A section with a dropdown menu set to 'Every number of days', followed by the text 'Every', a text input field containing '1', and the text 'day(s) (1-999)'.
- Scheduled Time:** A section with a 'Time' label, two dropdown menus (the first set to '2' and the second to '30'), and the text 'Hour:Minute, e.g. 13:30'.

At the bottom right of the configuration area, there are four buttons: 'Run Now', 'Save', 'Reset', and 'Help'.

high availability

The **high availability** tab lets you manage and control arcserve High Availability functions.

high availability

Control Services and Scenarios

<Control Service ID>

Scenarios

- 111
- 123
- 1231
- 1234567
- 123899
- 234577878
- 8789789978
- FSHA
- Full System HA SC
- Full System Raid 5
- Full System Scenario

Scenarios

Actions | Create New Scenario

Type item filter text

Name	State	Product	Server	Mode	Flag	Comment
111	Editing	DR	FullSystem	Online		
123	Editing	DR	FullSystem	Online		123123
1231	Editing	DR	FileServer	Online		
1234567	Editing	DR	FullSystem	Online		
123899	Editing	DR	FullSystem	Online		
234577878	Editing	DR	FullSystem	Online		
8789789978	Editing	DR	FullSystem	Online		

Recent Critical Events


There is no any recent critical

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 UDP, you can specify the following Arcserve UDP configuration settings.

- [Server Communication Protocol](#) (see page 82)
- [Database Settings](#) (see page 83)
- [arcserve Backup Data Synchronization](#) (see page 86)
- [SRM Configuration](#) (see page 87)
- [Node Discovery Configuration](#) (see page 89)
- [Email and Alert Configuration](#) (see page 90)
- [Update Configuration](#) (see page 92)
- [Administrator Account](#) (see page 93)
- [Installation Settings](#) (see page 94)
- [Share Plan](#) (see page 94)

Configure Server Communication Protocol

The Arcserve UDP 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 in to the computer where the Arcserve UDP 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:
 - a. 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 UDP Console.
C:\Program Files\CA\arcserve Unified Data Protection\Management\BIN

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

The communication protocol was changed to HTTPS.
 - b. 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 UDP Console.
C:\Program Files\CA\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 UDP Console.

Note: To update the communication protocol used by the Arcserve UDP Recovery Point Server and the Arcserve UDP Agent to communicate with the Arcserve UDP Console, you must 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 UDP database using the procedure described in [Re-create the Arcserve UDP Database](#) (see page 84), and then configure the database.

Follow these steps:

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

The screenshot displays the 'configuration' tab in the Arcserve UDP console. On the left, a sidebar lists various configuration options, with 'Database Configuration' highlighted. The main area is titled 'configuration' and contains the following sections:

- SQL Server:**
 - SQL Server Machine Name: shrp001-1104567
 - SQL Server Instance: ARCSERVE_APP
 - SQL Server port: (1025-65535) [checkbox] Auto detect
- Authentication:**
 - ☒ Windows Authentication Mode
 - ☐ SQL Server and Windows Authentication Mode
 - Username: [text box]
 - Password: [text box]
 - [Test button]
- Database Connection Pool:**
 - Max Connection(s): 20 (1-99)
 - Min Connection(s): 3 (1-99)

At the bottom right, there are three buttons: Save, Reset, and Help.

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 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 UDP Database

For various reasons, you may want to re-create the Arcserve UDP 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 UDP 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 UDP database, all current data is lost.


To re-create the Arcserve UDP database

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

Note: If Microsoft SQL Server Management Studio Express is not installed on the Arcserve UDP server, you can 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 UDP database is deleted.

4. Configure the new database. For more information, see [Configure Database](#) (see page 83).

The Arcserve UDP solution re-creates the database. The name of the database instance is **ARCSERVE_APP**.

Configure Arcserve UDP Backup Data Synchronization

The **arcserve Backup Data Synchronization Schedule** page enables you to configure the system to set a scheduled time and repeat method of how many days, which day of the week, or which day of the month the user can synchronize the arcserve Backup database with the Arcserve UDP database.

Follow these steps:

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

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

The screenshot shows the 'configuration' tab in the console. On the left, the 'SRM Configuration' option is highlighted in the navigation pane. The main area displays the SRM configuration settings. The 'Enable' radio button is selected. The 'Repeat Method' is set to 'Every number of days' with a value of 1. The 'Scheduled Time' is set to 14:30. At the bottom right, there are buttons for 'Run Now', 'Save', 'Reset', and 'Help'.

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.



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. By default, **Discovery Configuration** is disabled.

Follow these steps:

1. From the Console, click the **configuration** 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.

The screenshot shows the 'configuration' tab in the Arcserve UDP console. The left sidebar lists various configuration options, with 'Node Discovery Configuration' selected. The main panel displays the configuration for node discovery. The 'Enable' radio button is selected. The 'Repeat Method' is set to 'Every number of days' with a value of '1' day(s). The 'Scheduled Time' is set to '23:59' (Hour:Minute, e.g. 13:30). Below this is a 'Node Discovery List' table with columns: Node Region, Username, Computer Name Filter, Last Job Start Time, Last Job End Time, and Last Disco. The table contains one entry for '<node_region>' with Username 'Administrator', Computer Name Filter '*', Last Job Start Time '11/11/2014 12:02:16 PM', Last Job End Time '11/11/2014 12:02:39 PM', and Last Disco '0 Machines'. The page has a left sidebar with navigation links and a bottom bar with 'Save', 'Reset', and 'Help' buttons.

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 Adobe Flash Player ActiveX (version 10.0 or higher) and Microsoft .NET Framework (version 2.0 or higher) for the Report Chart export feature on the server to export images in a report successfully.

Follow these steps:

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

The screenshot shows the 'configuration' tab selected in the top navigation bar. On the left, a sidebar lists various configuration options, with 'Email and Alert Configuration' highlighted. The main content area is titled 'Email Settings' and contains several input fields and checkboxes. The 'Service' dropdown is set to 'Other'. The 'Email Server' field is empty, and the 'Port' is set to '25'. The 'Requires Authentication' checkbox is checked. The 'Account Name' and 'Password' fields are empty. The 'Subject' field is pre-filled with 'arcserve Unified Data Protection Alert'. The 'From' and 'Recipients' fields are empty. There are three checkboxes at the bottom: 'Use SSL' (unchecked), 'Send STARTTLS' (unchecked), and 'Use HTML Format' (checked). Below the 'Email Settings' section, there is a section for 'Enable Proxy Settings' which is currently disabled. At the bottom of the form, there is a 'Test Email' button and a 'Send Email Alerts' section with a 'Discovered Nodes' checkbox. At the very bottom of the page, there are four buttons: 'Save', 'Delete', 'Reset', and 'Help'.

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 that you enter 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.

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 **configuration** tab.
2. From the left pane, click **Update Configuration**.

The screenshot shows the 'configuration' tab selected in the top navigation bar. On the left, the 'Update Configuration' option is highlighted in the sidebar. The main content area is titled 'Updates' and contains three sections: 'Download Server', 'Test Connection', and 'Update Schedule'. In the 'Download Server' section, 'Arcserve Server' is selected with a radio button, and a 'Proxy Settings' button is visible next to it. The 'Test Connection' section has a 'Test Connection' button. The 'Update Schedule' section has a checkbox for 'Automatically check for updates' which is checked, and a schedule dropdown set to 'Every Sunday At 3 : 00 AM'. At the bottom right, there are 'Save', 'Reset', and 'Help' buttons.


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

Follow these steps:

1. For **Download Server**, select one of the following options:
 - For **Arcserve Server**, click **Proxy Settings** to complete Proxy Setup.

This screenshot shows the 'Proxy Settings' dialog box overlaid on the 'Update Configuration' page. The dialog has two main sections: 'Use browser proxy settings (for IE and Chrome only)' which is selected with a radio button, and 'Configure proxy settings' which is unselected. Below the first section, a note states: 'Note: Administrator log in credentials will be used as the proxy credentials.' The 'Configure proxy settings' section contains fields for 'Proxy Server', 'Port', 'User Name', and 'Password', along with a checkbox for 'Proxy server requires authentication'. At the bottom of the dialog are 'OK', 'Cancel', and 'Help' buttons. A red arrow points from the 'Proxy Settings' button in the background window to the dialog box.

- For **Staging Server**, click **Add Server** to provide staging server details.



You can add multiple staging servers.


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

Configure Administrator Account

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

Follow these steps:

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



Installation Settings

The **Installation Settings** page lets you specify default settings for installing Arcserve UDP Agent and Arcserve UDP 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 **configuration** tab.
2. From the left pane, click **Installation Settings**.

The screenshot shows the 'configuration' tab selected in the top navigation bar. On the left, a sidebar lists various configuration options, with 'Installation Settings' highlighted. The main content area is titled 'Installation Settings' and contains the following fields and options:

- Installation Folder:** A text box containing the path '%ProgramFiles%\CA\arcserve Unified Data Protection'.
- Port:** A text box containing the value '8014'.
- Protocol:** Radio buttons for 'HTTP' (selected) and 'HTTPS'.
- Change Tracking Driver:** A checkbox labeled 'Install Agent Change Tracking Driver' which is checked.

Below these fields, a note states: 'For more secure communication, HTTPS is the recommended communication protocol.' At the bottom right of the configuration area, there are three buttons: 'Save', 'Reset', and 'Help'.

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 **configuration** 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.

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

Arcserve UDP lets you migrate Arcserve r16.5 recovery point data to a Arcserve UDP data store. This lets you use Arcserve r16.5 data in Arcserve UDP

Create a Data Store to Replicate Data from a 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 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.
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 to this data store 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 encrypted, provide the session password. If the session is not 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**.
8. The recovery point data from Arcserve r16.5 is replicated to the UDP data store.

Chapter 4: Adding and Managing Source Nodes

This section contains the following topics:

[How to Add Nodes to the Console](#) (see page 99)

[How to Manage Nodes](#) (see page 111)

[How to Add and Manage Node Groups](#) (see page 133)

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
- Arcserve High Availability

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

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

The following diagram illustrates how to add nodes to the Console.



What To Do Next?

- [Review the Prerequisites](#) (see page 101)
- [Add Nodes](#) (see page 102)
- [Discover Nodes](#) (see page 105)
- [Import Nodes](#) (see page 106)
 - [Import Nodes from a File](#) (see page 106)
 - [Import Nodes from a vCenter/ESX Server](#) (see page 107)
 - [Import Nodes from a Hyper-V Server](#) (see page 108)
 - [Import Nodes from Arcserve High Availability](#) (see page 110)

Review the Prerequisites

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

1. Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.
2. Log in to 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. Arcserve UDP Agent (Windows) is installed on this node.
- **Linux:** Linux source nodes that you want to protect. Arcserve UDP Agent (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 UDP Agent (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

Add nodes by: Adding Windows Node

Node Name/IP Address: xyz123

Username: Administrator

Password:

Add Description:

arcserve Backup: ☒ Installed

Authentication Type: Windows Authenticaio

caroot Username: Administrator

caroot Password:

Port: 1234

Add to List

☒ Node Name	VM Name	Hypervisor
You have not added any node to the list.		

Remove

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

Adding Linux Node

The screenshot shows the 'Add Nodes to Arcserve UDP Console' dialog box. The 'Add nodes by' dropdown is set to 'Adding Linux Node'. The form contains the following fields:

- Node Name/IP Address: test12
- ☐ SSH Key Authentication
- Root Username: test
- Password: [masked]
- ☐ Non-Root Credentials
- Non-Root Username: [empty]
- Password: [empty]
- Add Description: [empty]

At the bottom right of the form is an 'Add to List' button. To the right of the form is a table with columns 'Node Name', 'VM Name', and 'Hypervisor'. The table is currently empty, with a message 'You have not added any node to the list.' below the header. At the bottom right of the table area is a 'Remove' button. At the bottom of the dialog box are 'Help', 'Save', and 'Cancel' buttons.

Notes:

- For Linux, on selecting **SSH Key Authentication**, you do not need to enter username and password details. 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.
- You cannot log in to the Linux nodes or the Linux Backup Server from the Arcserve UDP Console.

Adding Linux Backup Server Node

The screenshot shows the 'Add Nodes to Arcserve UDP Console' dialog box. The 'Add nodes by' dropdown is set to 'Adding Linux Backup Server Node'. The form contains the following fields:

- Node Name/IP Address: xyz123
- Username: root
- Password: [masked]
- Port: 8014
- Protocol: ☐ HTTP ☒ HTTPS
- Add Description: [empty]

At the bottom right of the form is an 'Add to List' button. To the right of the form is a table with columns 'Node Name', 'VM Name', and 'Hypervisor'. The table is currently empty, with a message 'You have not added any node to the list.' below the header. At the bottom right of the table area is a 'Remove' button. At the bottom of the dialog box are 'Help', 'Save', and 'Cancel' buttons.

The details of the selected option are displayed.

2. 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 would be listed on the right pane.

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

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

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 to 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.

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 File](#) (see page 106)
- [Import Nodes from a vCenter/ESX Server](#) (see page 107)
- [Import Nodes from a Hyper-V Server](#) (see page 108)
- [Import Nodes from Arcserve High Availability](#) (see page 110)

Import Nodes from a File

When you have multiple physical nodes to add, you can use this method. You can import in the following scenarios:

- When you have multiple nodes to add, instead of adding one node at a time, you can create a file a CSV or TXT file in the <NodeName>, <UserName>, <Password> format. Now import this file and all the nodes are added to the Console at one go.
- When you saved nodes as a file using the **Export** option.

Follow these steps:

1. From the **Add nodes by** drop-down list, select **Importing from file**.
2. Click **Browse** to select your saved CSV or TXT file.
3. Click **Upload**.
The nodes are displayed at the left pane.
4. Verify the nodes by providing login credentials.
Note: You can only add verified nodes.
5. (Optional) Click **Browse** if you want to import more nodes.
6. Select the check boxes of verified nodes, and click **Add to List**.
The selected nodes move to the right pane.
7. From the right pane, select the nodes that you want to add, and click **Save**.
The nodes are added and displayed on the **Nodes: All Nodes** page.

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 **Browse**.
In the left pane, a node tree is displayed.
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, and then click **Add to List**.
The selected nodes are added to the right pane.
5. 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 **Browse**.

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 UDP 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, and then click **Add to List**.

The selected nodes are added to the right pane.

5. 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, 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 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 Arcserve HA

When you already have a Arcserve High Availability scenario with a arcserve D2D recovery point as a source, then you can import this source as a node to Arcserve UDP. You can then use this node to create a virtual standby machine.

Follow these steps:

1. From the **Add nodes by** drop-down list, select **Importing Nodes from arcserve HA**.
2. Specify the Arcserve High Availability Control Service details and click **Browse**.

Note: If you have already added Hostname/IP Address, you can select them from the Hostname/IP Address drop-down menu. Based on the selected option, the user name and password appear.

3. From the left pane, select the nodes, and then click **Add to List**.

The selected nodes are added to the right pane.

4. Select the nodes and click **Save**.

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

A confirmation message opens asking if you want to configure Converter information.

5. Click **No**, if you do not want to configure.


(Optional) Click **Yes**, if you want to configure. The **Configure Remote Converters** dialog is displayed.

6. (Optional) Specify the connection information for the converters, and click **Update**.

The converters related information of the added nodes is updated and the node is added to the Console.

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. The following diagram illustrates how you can manage a node.



What To Do Next?

- [Review the Prerequisites](#) (see page 112)
- [Update Hypervisor Information](#) (see page 112)
- [Specify the Hypervisor](#) (see page 113)
- [Update VM Information](#) (see page 115)
- [Update Nodes](#) (see page 116)
- [Export Node](#) (see page 118)
- [Synchronize Data](#) (see page 119)
- [Delete Nodes from the Console](#) (see page 119)
- [Deploy Agent to Nodes](#) (see page 120)
- Perform Preflight Checks for Your Backup Jobs

Review the Prerequisites

Before starting to manage the nodes, complete the following prerequisites:

- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.
- Log in to the Console.
- Add a node.

Update Hypervisor Information

After a VM node is added into Arcserve UDP, the connection-related information, such as hostname or credentials of hypervisor of VM may change. In such cases, Arcserve UDP lets you update the hypervisor information.

Follow these steps:

1. Click the **resources** tab.
2. Right-click on 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 (such as Exchange granular restore), 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. You can specify the hypervisor details in such cases and the VM uses the hypervisor host license instead of using another license.

The following examples describe when you can 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:

- You cannot specify the hypervisor for a physical node.
- You cannot specify the hypervisor for a VM node that is imported from vCenter/ESX or Hyper-V.
- You cannot specify the hypervisor for a VM on EC2.
- You can specify the hypervisor for multiple VMs, which 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.
 - Select the node name, and from the right 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**, **ESX**, and **Other** (Xen, Kernel-based Virtual Machine, Red Hat Enterprise Virtualization).

The image shows a 'Specify Hypervisor' dialog box. It has a title bar with the text 'Specify Hypervisor'. Inside the dialog, there are six fields: 'Hypervisor Type' with a dropdown menu showing 'ESX', 'ESX Hostname/IP Address' with a dropdown menu, 'Username' with a text box containing 'root', 'Password' with an empty text box, 'Port' with a text box containing '443', and 'Protocol' with a dropdown menu showing 'HTTPS'. At the bottom of the dialog, there are three buttons: 'Help', 'OK', and 'Cancel'.

Specify Hypervisor	
Hypervisor Type	ESX
ESX Hostname/IP Address	
Username	root
Password	
Port	443
Protocol	HTTPS
Help OK Cancel	

5. Enter the hypervisor details and click **OK**.
The hypervisor information is specified.

Update VM Information

Using Arcserve UDP, 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 UDP.
- The user name or password of the node was updated after the node was registered with Arcserve UDP.

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.
 - Select the node name, and from the right pane click the **Actions** drop-down list.
3. Click **Update**.
The **Update node** dialog opens.
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](#) (see page 109).
2. In VMware, when you update nodes, Arcserve UDP automatically install 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 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](#) (see page 109).

Export Node

You can export the nodes as a CSV (.csv) file. When required, you can import the CSV 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.csv file.
4. Click **Open** or **Save**.
The node list is exported.

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. Perform one of the following actions:
 - (Node Level) Right-click the node name.
 - (Node Level) Select the node name, and from the center pane click the **Actions** drop-down list, and then click **Sync Data**.
 - (Node Level) Select the node name, and from the right pane click the **Actions** drop-down list.
 - (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 UDP.

 - **Full Synchronize arcserve Backup**
 - **Incremental Synchronize arcserve Backup**
 - **Full Synchronize arcserve UDP Agent**

The **Information** dialog explains that the selected synchronization method is submitted.

Delete Nodes from the Console

Using Arcserve UDP, 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.
 - Select the node name, and from the right 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 UDP Agent for a node, use **Install/Upgrade Agent**. If the destination machine contains a prior version of Arcserve UDP agent, then use the upgrade option to get the latest version. Otherwise, use the install option.

Note: You can deploy Arcserve UDP 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, and click **OK**.
The node is installed or upgraded with the latest version of Arcserve UDP Agent.

Deploy Agent to Nodes

To upgrade or install Arcserve UDP Agent for a node, use **Install/Upgrade Agent**. If the destination machine contains a prior version of Arcserve UDP agent, then use the upgrade option to get the latest version. Otherwise, use the install option.

Note: You can deploy Arcserve UDP 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 UDP 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 UDP 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. PFC runs automatically when you perform the following actions:

- Import virtual machines from a [vCenter Server/ESX Server system](#) (see page 107) or [Hyper-V](#) (see page 108).
- [Update a node](#) (see page 116)

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.

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.

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 UDP 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 UDP 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-consistent backup.

Item	Description
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 UDP 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 UDP 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.

Solutions for Preflight Check Items of VMware 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 UDP and use Arcserve UDP Agent (Windows) to back up the VM.

Status	Message	Solution
Warning	Changed Block Tracking is enabled with snapshots present. A full disk backup will be applied.	To apply the used block backup, 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\CA\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 UDP or use Arcserve UDP Agent (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 UDP or use Arcserve UDP Agent (Windows) to back up the VM.

Status	Message	Solution
Warning	The virtual Raw Device Mapping (RDM) disk backs up as a full disk.	Create an agent-based backup plan in Arcserve UDP or use Arcserve UDP Agent (Windows) to back up the VM.
Warning	The independent disk is not backed up.	Create an agent-based backup plan in Arcserve UDP or use Arcserve UDP Agent (Windows) to back up the VM.
Warning	The application backs up the disk on the NFS data store as a full disk.	Create an agent-based backup plan in Arcserve UDP or use Arcserve UDP Agent (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 UDP or use Arcserve UDP Agent (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 UDP or use Arcserve UDP 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 UDP or use Arcserve UDP Agent (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	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.

Status	Message	Solution
Warning	VMware does not support application-consistent quiescing if the guest OS has dynamic 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. 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	Failed to access the virtual machine.	Provide the built-in or domain administrator credentials to log in to the virtual machine guest operating system. 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 2003 and later.
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 UDP or use Arcserve UDP Agent (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 in the VM.
- The VM must run on ESXi 4.0 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

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 <i>Key-Value Pair</i> and <i>Live virtual machine backup</i> 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	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 UDP or use Arcserve UDP 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 UDP or use Arcserve UDP Agent to back up the virtual machine.
Warning	Unable to back up the disk on the remote share.	Create an agent-based backup plan in Arcserve UDP or use Arcserve UDP 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 UDP or use Arcserve UDP Agent to back up the virtual machine.
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 UDP or use Arcserve UDP Agent to back up the virtual machine.
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.

Status	Message	Solution
Warning	An application-consistent snapshot is not supported. The integration service is not operational (Failed state).	Refer to the Integration Services column.
Warning	Failed to get the virtual machine by instance UUID.	Verify if the virtual machine exists on the Hyper-V server.
Warning	The virtual machine is not running.	Refer to the Power State column.
Warning	Failed to access the virtual machine.	Provide credentials with administrative privileges.

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.

How to Add and Manage Node Groups

Using Arcserve UDP, 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 UDP 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.

- **Hyper-V Groups:** Displays the nodes that you add using the **Importing from Hyper-V** option.
- **VM Backup Proxy Groups:** Displays Agentless nodes that are protected by Backup, Host-based Agentless task
- **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](#) (see page 134)
- [Add Node Groups](#) (see page 134)
- [Modify Node Groups](#) (see page 135)
- [Delete Node Groups](#) (see page 136)

Review the Prerequisites

Before working on the node groups, complete the following prerequisites:

- Review the [Compatibility Matrix](#) which 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 **Add** under **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 UDP 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 **Modify**.

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 UDP. 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**.
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.

Chapter 5: Adding and Managing Destinations

This section contains the following topics:

[How to Add a Destination](#) (see page 137)

[How to Manage a Data Store](#) (see page 146)


[How to Manage a Recovery Point Server](#) (see page 161)

How to Add a Destination

A destination is a location where you store your backup data. 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](#) (see page 138)
2. [Add a Recovery Point Server](#) (see page 139)
3. [\(Optional\) Deploy the Recovery Point Server](#) (see page 141)
4. [Add a Data Store](#) (see page 143)
5. [Verify the Destination](#) (see page 145)

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 UDP.
- Verify that you have administrator privileges to install Arcserve UDP.
- Review the [Compatibility Matrix](#) which 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 in to Arcserve UDP, 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 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.

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 Folder

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 UDP, 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 UDP 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.
 - Select a recovery point server, and from the right pane click the **Actions** drop-down list.

A list of options is displayed.

4. Click **Install/Upgrade Recovery Point Server**.

The **Product Install and Upgrade** page is displayed.

Destinations: Recovery Point Servers

Actions ▾ | [Add a Recovery Point Server](#)

Product Install and Upgrade

If the destination machine contains a prior version of Recovery Point Server, it will automatically be upgraded to the latest version. In addition, the selections for the port number, protocol, and installation path will be kept the same as previously specified.

Installation Folder

%ProgramFiles%\CA\arcserve Unified Data Protection

Port

8014

Protocol

☒ HTTP ☐ HTTPS

For more secure communication, HTTPS is the recommended communication protocol.

Change Tracking Driver

☒ Install Agent Change Tracking Driver

Installation or upgrade of this driver will require a reboot of the destination machine.

Start Time to Install or Upgrade

☒ Run immediately

☐ 3/20/14  1 : 00 AM

☐ Run manually at a later time

OK

Cancel

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 a 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.
 - Select a recovery point server, and from the right pane click the **Actions** drop-down list.

A list of options is displayed.

4. Click **Add a Data Store**.

The **Add a Data Store** page is displayed with the name of the specified recovery point server.

5. Specify the following fields:

Data Store Name

Defines the name of the data store.

Backup Destination Folder

Defines the location of the folder where the data store is created.

Note: For non-deduplication and deduplication data store, the backup destination path should be an empty folder.

Compression Type

Specifies whether to use the standard 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](#) (see page 709).

Concurrent Active Nodes

Specifies the maximum concurrent jobs on the data store.

No Limit: The default means that all jobs in this Data Store are started immediately.

Limit to: 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 coming after five backup jobs waits in the queue, but you can submit another job such as File System Catalog.

Notes:

- 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.

Encrypt Data

Specifies whether to enable data encryption. 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.

Deduplicate Data

Specifies whether to enable data deduplication. 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.

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. Selecting a high speed Solid State Drive (SSD) can increase the deduplication capacity with a low memory allocation required.

Note: The **Hash Destination** path should be a blank folder.

Hash Destination is on a Solid State Drive (SSD)

Specifies if the hash folder is on a solid state drive.

Memory Allocation

Specifies the amount of physical memory that you allocate to keep hashes.

Deduplication Block Size

Defines the deduplication block size. The options are 4 KB, 8 KB, 16 KB, and 32 KB. The deduplication block size also impacts the Deduplication capacity estimation. For example, if you change the default 4KB to 8KB, the Deduplication capacity estimations double. Increasing the deduplication block size can decrease the deduplication percentage.

Note: You cannot specify the same path for the following four folders: **Backup Destination folder**, **Data Destination**, **Index Destination**, and **Hash Destination**.

6. (Optional) Set values for the **Estimate Memory and Storage Requirements** options. These options help to estimate the capacity of a Deduplication data store.
7. Click **Save**.

The data store is created and gets displayed on the center pane. Click the data store to view the details in the right pane.

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.

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.

The following diagram illustrates the various operations you can perform on existing data stores:



What To Do Next?

- [Review Prerequisites](#) (see page 147)
- [Modify a Data Store](#) (see page 148)
- [Delete a Data Store from the Console](#) (see page 156)
- [Stop a Data Store](#) (see page 157)
- [Start a Data Store](#) (see page 158)
- [Browse Recovery Points in a Data Store](#) (see page 158)
- [Delete Node Data from a Data Store](#) (see page 160)
- [Troubleshooting: How to Use a Data Store When the Backup Destination Folder is Full](#) (see page 160)

Review the Prerequisites

To manage a data store, complete the following prerequisites:

- You have already added a data store.
- Review the [Compatibility Matrix](#) which 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:

- Type of compression
- Remove data encryption
- Deduplication options: Deduplicate Data and Deduplication Block Size.

Some considerations you must be aware 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.
 - Select the data store, and from the right 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.

Data Store Name

Defines the name of the data store.

Backup Destination Folder

Defines the location of the folder where the data store is created.

Note: For non-deduplication and deduplication data store, the backup destination path should be an empty folder.

Compression Type

Specifies whether to use the standard 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](#) (see page 709).

Concurrent Active Nodes

Specifies the maximum concurrent jobs on the data store.

No Limit: The default means that all jobs in this Data Store are started immediately.

Limit to: 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 coming after five backup jobs waits in the queue, but you can submit another job such as File System Catalog.

Notes:

- 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.

Encrypt Data

Specifies whether to enable data encryption. 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.

Deduplicate Data

Specifies whether to enable data deduplication. 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.

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. Selecting a high speed Solid State Drive (SSD) can increase the deduplication capacity with a low memory allocation required.

Note: The **Hash Destination** path should be a blank folder.

Hash Destination is on a Solid State Drive (SSD)

Specifies if the hash folder is on a solid state drive.

Memory Allocation

Specifies the amount of physical memory that you allocate to keep hashes.

Deduplication Block Size

Defines the deduplication block size. The options are 4 KB, 8 KB, 16 KB, and 32 KB. The deduplication block size also impacts the Deduplication capacity estimation. For example, if you change the default 4KB to 8KB, the Deduplication capacity estimations double. Increasing the deduplication block size can decrease the deduplication percentage.

Note: You cannot specify the same path for the following four folders: **Backup Destination folder**, **Data Destination**, **Index Destination**, and **Hash Destination**.

7. (Optional) Set the values for **Estimate Memory and Storage Requirements**, and click **Save**.

These options help to estimate the capacity of a Deduplication data store.

The data store is updated.

Modify the Data Store Threshold

A data store has default threshold setup in the system and physical memory. To free up space or replace the existing disk with a bigger disk, you can modify the default threshold manually. 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 nondeduplication data store, the threshold monitors the storage space only of the backup destination folder. All the five items that the thresholds monitors have two types of values: warning threshold and error threshold.

Thresholds registry locations and default values

1. Registry location: [HKEY_LOCAL_MACHINE\SOFTWARE\CA\ARCserve Unified Data Protection\Engine\DataStore\XXXXXXX\CommStore]
Threshold values: "WarnPathThreshold"="0.03" and "ErrorPathThreshold"="100"
2. Registry location: [HKEY_LOCAL_MACHINE\SOFTWARE\CA\ARCserve Unified Data Protection\Engine\DataStore\XXXXXXX\GDD\DataRole]
Threshold values: "WarnPathThreshold"="0.03" and "ErrorPathThreshold"="100"
3. Registry location: [HKEY_LOCAL_MACHINE\SOFTWARE\CA\ARCserve Unified Data Protection\Engine\DataStore\XXXXXXX\GDD\HashRole]
Threshold values: "WarnPathThreshold"="0.03" and "ErrorPathThreshold"="100",
"WarnMemThreshold"="0.03" and "ErrorMemThreshold"="10"
Note: Hash role monitors the memory and disk usage both. Path represents the disk usage and Mem represents the memory.
4. Registry location: [HKEY_LOCAL_MACHINE\SOFTWARE\CA\ARCserve Unified Data Protection\Engine\DataStore\XXXXXXX\GDD\IndexRole]
Threshold values: "WarnPathThreshold"="0.03" and "ErrorPathThreshold"="100"
5. For physical memory available in both the system and physical memory allocated to the data store:
Registry location: [HKEY_LOCAL_MACHINE\SOFTWARE\CA\ARCserve Unified Data Protection\Engine\DataStore\XXXXXXX\GDD\HashRole]
Threshold values: "WarnMemThreshold"="0.03" and "ErrorMemThreshold"="10"

Follow these steps to modify the threshold:

1. Navigate to the respective registry location.
2. Manually modify the default value of threshold.

When the data store approaches a threshold, the following warning message is displayed:

Running Low Disk Space: Data Destination.

The screenshot shows the 'resources' console. On the left, there's a sidebar with 'Nodes' and 'Plans'. The main area displays a table titled 'Destinations: Recovery Point Server'. The table has columns: Name, Plan Count, Data Protected, Deduplication, Compression, Overall Data Reduction, and Space Occupied. Below the table, there's a detailed view for a specific destination, showing its status as 'Running' and a warning message: 'Low Disk Space Backup Destination', 'Low Disk Space Data Destination', 'Low Disk Space Index Destination', and 'Low Disk Space Hash Destination'.

Note: You can view the error and warning messages from the **log** tab in the Console.

The data store threshold is modified.

Fix the Data Store Threshold Warning and Error Messages

The **log** tab display errors or warning statuses related to data store threshold. The following diagrams displays the different types of errors or warnings for specific folders:

Example Graphics of Error and Warning Messages for Threshold:

Warning or Error Message for Four Folders

The screenshot shows the 'log' tab with a table of log entries. The first entry has a severity of 'Warning and Error' and a message: 'The backup destination folder,data destination,index destination,hash destination of data store "004" is running out of free space/memory, and may not be able to save the specified backup data.' The message is highlighted with a red box.

Warning or Error Message for One Item

The screenshot shows the 'log' tab with a table of log entries. The first entry has a severity of 'Warning and Error' and a message: 'The data destination of data store "006" has sufficient free space/memory to save the specified backup data, but is close to its maximum capacity.' The message is highlighted with a red box.

Error Message for Only One Item

The screenshot shows the 'log' tab with a table of log entries. The first entry has a severity of 'Warning and Error' and a message: 'The data destination of data store "006" is running out of free space/memory, and may not be able to save the specified backup data.' The message is highlighted with a red box.

When Does Error or Warning Message Appear

If the threshold value is less than 1, then the value is percentage or the value unit is MB. For example, for backup destination folder, WarnPathThreshold="0.03" leads to the following report status:

- Report warning status if the volume free size is less than volume size 3%
- Report error status if the volume free size less than 100MB.

Follow these steps to fix the messages:

1. Navigate to the respective registry location.
2. Manually modify the default value of threshold to change the thresholds or release more space.

Note: If the threshold is reached, you can release the space manually. The updated status is available in 15 minutes.

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.

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.
 - Select the data store, and from the right 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.

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.
 - Select the data store, and from the right 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**.

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.
 - Select the data store, and from the right 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**.

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.

To delete a node from a data store, see [Delete Node Data From a Data Store](#) (see page 160).

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.
 - Select the data store, and from the right 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.

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 it. 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, this job is called the Purge 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 is displayed.
4. Select the node that you want to delete.
5. Click **Actions, Delete**.

Note: To delete multiple nodes from a data store, press the ctrl key and select the nodes, then click **Action, 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.

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:

You can stop the data store, copy the corresponding folder into a large volume, then specify the new path to import data store, and overwrite the existing one to continue using it.


Note: Verify that before copying the folder, you stop the data store. During copy, if some files cannot be copied, you can skip copying those files.

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.

The following diagram illustrates how the storage administrator can manage a recovery point server:



What To Do Next?

- [Review the Prerequisites](#) (see page 162)
- [Update a Recovery Point Server](#) (see page 162)
- [Delete a Recovery Point Server from the Console](#) (see page 163)
- [Import a Data Store](#) (see page 164)
- [Install/Upgrade Recovery Point Server](#) (see page 165)

Review the Prerequisites

To manage a recover point server, complete the following prerequisites:

- Log in to the Console.
- Add a recovery point store.
- Review the [Compatibility Matrix](#) which 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.Select a recovery point server, and from the right pane 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.

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.

Select a recovery point server, and from the right pane 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.

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.

Select a recovery point server, and from the right pane 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.

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.

Select a recovery point server, and from the right pane click the **Actions** drop-down list.

4. Click **Install/Upgrade Recovery Point Server**.

The install path and reboot 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**.

Chapter 6: Creating Plans to Protect Data

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 Task

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

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.

Replicate to a remotely-managed RPS


Lets you create a task to replicate or send data to a remote recovery point server.

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
Backup: Host-Based Agentless	■ Replicate ■ Virtual Standby ■ Copy Recovery Points ■ Replicate to a remotely-managed RPS
Backup: Agent-Based Linux	None
Replicate data from a remote RPS	■ Virtual Standby ■ Replicate
Virtual Standby*	None

* Use this Virtual Standby task to create virtual standby machines for nodes that you import from Arcserve High Availability.

The following diagram illustrates how different tasks form a backup plan. The diagram also shows parameters that you can define in each task.



This section contains the following topics:

[How to Create a Windows Backup Plan](#) (see page 170)

[How to Create a Linux Backup Plan](#) (see page 189)

[How to Create a Host-Based Virtual Machine Backup Plan](#) (see page 207)

[How to Create a Virtual Standby Plan](#) (see page 241)

[How to Create a Virtual Standby Plan for Arcserve High Availability Nodes](#) (see page 274)

[How to View Virtual Standby Settings From the Monitor Server](#) (see page 287)

[How to Protect Virtual Standby Machines](#) (see page 299)

[How to Replicate Data Between Data Stores Managed from a UDP Console](#) (see page 307)

[How to Replicate Data Between Data Stores Managed From Different UDP Consoles](#) (see page 325)

[How to Perform an Offline Data Replication Using RPS Jumpstart](#) (see page 339)

[How to Create a Copy Recovery Points Plan](#) (see page 347)

[How to Create a File Copy Plan](#) (see page 366)

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](#) (see page 172)

To backup Microsoft clustered nodes and shared disks, review the following prerequisites:

- [Review the Prerequisites to Back Up Microsoft Clustered Nodes and Shared Disks](#) (see page 173)

The following diagram illustrates the process to create a Windows backup plan:



What To Do Next?

1. [Review the Prerequisites and Considerations](#) (see page 171)
2. [Create a Backup Plan](#) (see page 174)
3. [\(Optional\) Perform a Manual Backup](#) (see page 188)
4. [Verify the Backup](#) (see page 188)

Review the Prerequisites and Considerations

Verify that you have completed the following prerequisite tasks:

- Log in to the Console.
- (Optional) Create data store to store the backup data.
- [Review the prerequisites to back up an Oracle database](#) (see page 172).
- [Review the Prerequisites to back up Microsoft clustered nodes and shared disk](#) (see page 173).
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Review the Prerequisites for Oracle Database

To back up an Oracle database with consistent data, ensure that the following prerequisites are met:

- The ARCHIVELOG mode is enabled to archive the Redo logs.

Follow these steps to check if the ARCHIVELOG mode is enabled:

- a. Log in to 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.

If the ARCHIVELOG mode is not enabled, you must start the ARCHIVELOG mode to backup 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 will be written to the flash recovery area. If you do not want to write archive logs to the flash recovery area, you can 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: No Archive Mode

Automatic archival: Disabled

Archive destination:

E:\app\oracle\oradata\<oracle_database_name>\arch

Oldest online log sequence: 21

Current log sequence: 23

- Oracle VSS Writer Service is started and 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 datafiles, server parameter file, control files, archived redo logs, and online redo logs for the backup.

- Review the [Compatibility Matrix](#) which 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 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](#) which 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.

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. 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: After a plan is paused, all jobs are paused except the restore job and the copy job. The running jobs are not affected. If you pause a plan that has pending jobs, then those pending jobs will also pause. When you resume the plan, the pending jobs does not resume immediately. After you resume the plan, the pending jobs will run from the next scheduled time. You can find the schedule of the next job form the home page of Arcserve UDP Agent (Windows).

6. From the **Task Type** drop-down list, select **Backup, Agent-Based Windows**.

Add a Plan

New Plan

☐ Pause this plan

Task1: Backup: Agent-Based Windows

+

Add a Task

Product Installation

Task Type

Backup: Agent-Based Windows

Source

Destination

Schedule

Advanced

+

Add Nodes

Remove

☒

Node Name

VM Name

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 and click **Add Node**.
2. Select one of the following options:

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.


Adding Windows Nodes

Opens the **Add Nodes to Arcserve UDP Console** dialog. Select this option if you have not added the nodes and you want to manually add the nodes to protect.

Discovering 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. (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** area.

4. Select the nodes from the **Available Nodes** area and click the **Add all nodes (>>)** or **Add selected nodes (>)** icon.

The selected nodes are displayed on the **Selected Nodes** area.

5. Click **OK** to close the dialog.
6. To choose **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](#) (see page 711).
 - c. Optionally, provide an encryption password.
 - d. Confirm the encryption password.
 - e. Select a type of compression. For more information, see [Compression Type](#) (see page 709).

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](#) (see page 182).

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.

New Backup Schedule

Custom

Backup Type: Incremental

Start Time: 8:00 AM

☐ Sunday ☐ Monday ☐ Tuesday
☐ Wednesday ☐ Thursday ☐ Friday
☐ Saturday

Repeat: ☒

Every: 3 Hours

Until: 6:00 PM

Help Save Cancel

- 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 checkbox 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 it is displayed on the **Schedule** page.

Source **Destination** **Schedule** **Advanced**

+ Add Delete		Type	Description	Su	Mo	Tu	We	Th	Fr	Sa	Time
☒			Custom Incremental Backups Every 3 Hours	✓	✓	✓	✓	✓	✓	✓	8:00 AM - 6:00 PM
☐			Weekly Incremental Backup						✓		8:00 PM

Add Merge Schedule

- Click **Add** and select **Add Merge Schedule**.

The **Add New Merge Schedule** dialog opens.

- Specify the start time to start the merge job.
- Specify **Until** to specify an end time for the merge job.
- Click **Save**.

The Merge Schedule is specified and it is displayed on the **Schedule** page.

Add Throttle Schedule

- Click **Add** and select **Add Throttle Schedule**.

The Add New Throttle Schedule dialog opens.

- Specify the throughput limit in MB per minutes unit.
- Specify the start time to start the backup throughput job.
- Specify **Until** to specify an end time for the throughput job.
- Click **Save**.

The Throttle Schedule is specified and it is displayed on the **Schedule** page.

- Specify the start time for the scheduled backup.

Start time for scheduled backup 10/24/2014 1 : 27 PM

Recovery Point Retention

7	Daily Backups
	Weekly Backups
	Monthly Backups
31	Custom / Manual Backups

- 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.

- Specify the catalog details.

Catalogs

Generate file system catalogs (for faster search) after

☐ Daily Backups

☐ Weekly Backups

☐ Monthly Backups

☐ Custom / Manual Backups

Generate Exchange catalogs (for granular restore) after

☐ All Backups for Nodes with Exchange Installed

Catalogs let you generate the Exchange Granular Restore catalog and the File System catalog. The Exchange Granular Restore catalog is required to restore Exchange mailboxes, mailbox folders, and individual mail objects. The File System catalog is required to perform faster and easier 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: Retention settings must be set 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 have to 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 will run at 6:00 AM, 7:00 AM, 8:00 AM, but NOT 9:00 AM.

Backup Throttle Schedule

Backup throttle schedule lets you control the backup throughput speed, which 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 will be 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 will adjust 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 will be 500 MB/minute and from 8:00 PM to 10:00 PM the throttle limit will be 2500 MB/minute.

If you do not define any backup schedule and backup throughput schedule, the backup will run 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 being 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 will process these sets one by one.
- If a merge job is resumed after a pause, the job detects at which point it was 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**

Truncate log

☐ Truncate SQL Server log

Weekly

☐ Truncate Exchange Server Log

Weekly

Run Commands

☐ Before a backup is started

☐ On exit code

0

☒ Run Job ☐ Fail Job

☐ After a snapshot is taken

☐ After a backup is over

Username for Commands

Password for Commands

Enable Email Alerts



Email Settings

Job Alerts



Missed jobs



Backup, Catalog, Replication, File Copy, Restore or Copy Recovery Point job failed/crashed/canceled



Backup, Catalog, Replication, File Copy, Restore or Copy Recovery Point job successfully completed



Merge job stopped, skipped, failed or crashed



Merge job success

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.

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.

(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 should perform an immediate manual backup without waiting for the next scheduled backup to occur.

Follow these steps:

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 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.

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 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 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 a remote share folder.

Note: You cannot log in to the Linux nodes or the Linux Backup Server from the Arcserve UDP Console.

The following diagram illustrates the process to protect Linux nodes:



What To Do Next?

1. [Review the Prerequisites and Considerations](#) (see page 190)
2. [Create a Backup Plan](#) (see page 190)
3. [\(Optional\) Perform a Manual Backup](#) (see page 204)
4. [Verify the Backup](#) (see page 205)
5. [Troubleshooting](#) (see page 206)

Review the Prerequisites and Considerations

Complete the following prerequisites:

- Log in to the Console.
- Review the [Compatibility Matrix](#) which 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. Click the **resources** tab.
2. From the left pane, navigate to **Plans**, and click **All Plans**.

If you have added any plans, these plans will be displayed in the center pane.

- On the center pane, click **Add a Plan**.

The **Add a Plan** page opens.

- Enter a plan name.

- (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: After a plan is paused, all jobs are paused except the restore job and the copy job. The running jobs are not affected. If you pause a plan that has pending jobs to run, then those pending jobs will also pause. When you resume the plan, the pending jobs will not resume immediately. After you resume the plan, the pending jobs will run from the next scheduled time.

- From the **Task Type** drop-down menu select **Backup, Agent-Based Linux**.

The screenshot shows the 'Add a Plan' page. At the top, there is a 'New Plan' text input field and a 'Pause this plan' checkbox. Below this is a sidebar on the left with a dark blue header 'Task1: Backup: Agent-Based Linux'. The sidebar contains three buttons: 'Add a Task' (with a plus icon), 'Product Installation', and 'Add a Task' (with a plus icon). The main content area has a 'Task Type' dropdown menu set to 'Backup: Agent-Based Linux'. Below this are four tabs: 'Source', 'Destination', 'Schedule', and 'Advanced'. The 'Source' tab is active. It contains a 'Linux Backup Server' dropdown menu with an 'Add' button. Below this are two buttons: '+ Add Nodes' and 'Remove'. At the bottom, there is a table with a checkbox and two columns: 'Node Name' and 'VM Name'.

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 nodes 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. Select the **Linux Backup Server** from the drop-down list.

The screenshot shows the 'Source' tab selected among four tabs: 'Source', 'Destination', 'Schedule', and 'Advanced'. Below the tabs, there is a section for 'Linux Backup Server' with a dropdown menu showing '<server name/IP address>' and an 'Add' button. Below this are two buttons: '+ Add Nodes' and 'Remove'. At the bottom, there is a table with two columns: 'Node Name' and 'VM Name'. The table contains one row with a checkbox in the 'Node Name' column and the text '<linux node>'.

Node Name	VM Name
☐ <linux node>	

3. (Optional) Click **Add** to add a new Linux Backup Server to the list.
4. Click **Add Nodes** and select one of the following options:

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.

Adding Linux Nodes

Opens the **Add Nodes to arcserve UDP Console** dialog. Select this option if you have not added the nodes and you want to manually add the nodes to protect.

5. Select the nodes from the **Available Nodes** column and click the **Add all nodes** or **Add selected nodes** button.

The selected nodes are displayed in the **Selected Nodes** column.

6. Click **OK** to close the dialog.
7. (Optional) Provide the details for the following options:

Volumes to be excluded for all listed nodes

Specify the volume that you do not want to backup. If you do not want to backup multiple volumes, separate each volume using a colon (:).

Files/folders to be excluded for all listed nodes

Specify the files and folders that you do not want to backup. 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.

Volumes to be excluded for all listed nodes

/NFS

Files/folders to be excluded for all listed nodes

/(tmp):/(*.iso)

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 a backup destination and enter the complete path of the storage location.
 - 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
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 UDP Agent (Linux) in the Release Notes.
 - If you have selected **CIFS share**, then type the Backup Destination detail in the following format:
//hostname/share_folder
Note: The shared folder name cannot contain any spaces.
 - If you have selected **Source local**, then provide the path of the local destination.

3. Click the arrow button to validate the Backup Destination information.

If the backup destination is invalid, an error message is displayed.

4. Specify the retention settings.

Note: For more information about the recovery sets, see *Understanding the Recovery Sets* (see page 195).

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 through 30, or the last day of the month.

Note: The 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.

5. 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 that this option provides a good balance between the CPU usage and the disk space usage. This compression is the default setting.

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.

6. Select an algorithm from the **Encryption Algorithm** drop-down list and type the encryption password, if necessary.

- 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 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](#) (see page 711).

- 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.

- b. 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.

The destination 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 UDP Agent (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.

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 per the schedule. You can add multiple schedules and provide retention settings. The default value is to repeat Incremental Backup every 3 hours from 8:00 AM to 6:00 PM.

You can edit or delete a backup job schedule.

Task Type Backup: Agent-Based Linux


Source Destination **Schedule** Advanced

Add		Delete								
☐	Type	Description	Su	Mo	Tu	We	Th	Fr	Sa	Time
☐		Custom Incremental Backups Every 3 Ho...	✓	✓	✓	✓	✓	✓	✓	12:00 AM - 6:0...

Follow these steps:

1. Click the **Schedule** tab and click **Add**.
2. Select **Backup Job Schedule**.

The **New Backup Schedule** dialog opens.



The image shows a 'New Backup Schedule' dialog box with a close button (X) in the top right corner. The dialog is divided into several sections. The 'Backup Type' section has a dropdown menu set to 'Incremental'. The 'Start Time' section has a text field with '8:00 AM' and a calendar icon. Below this is a section for selecting days of the week, with checkboxes for Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, and Saturday, all of which are currently unchecked. The 'Repeat' section has a checked checkbox. Below the checkbox is a section for frequency, with 'Every' followed by a text field containing '3' and a dropdown menu set to 'Hours'. At the bottom of this section is an 'Until' section with a text field containing '6:00 PM' and a calendar icon. At the bottom of the dialog are three buttons: 'Help', 'Save', and 'Cancel'.

New Backup Schedule ✕

Backup Type: Incremental

Start Time: 8:00 AM 

☐ Sunday ☐ Monday ☐ Tuesday
☐ Wednesday ☐ Thursday ☐ Friday
☐ Saturday

Repeat: ☒

Every: 3 Hours 

Until: 6:00 PM 

[Help](#) **Save** [Cancel](#)

3. 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.

4. Specify the backup start time.
5. (Optional) Select the Repeat check box and specify the repeat schedule.
6. Click Save.

The **New Backup Schedule** dialog closes.

The backup 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 backup throughput and pre/post script settings.

The following diagram displays the **Advanced** tab:

Source **Destination** **Schedule** **Advanced**

Throttle

☐ Backup throughput limit: MB/min

Pre/Post script Settings

Run on Linux Backup Server

Before job is started	None ▼
After job is over	None ▼

Run on source node

Before job is started	None ▼
After job is over	None ▼
Before snapshot is taken	None ▼
After snapshot is taken	None ▼

Follow these steps:

1. Click the **Advanced** tab.
2. Specify the throttle backup value.

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/CA/d2dserver/usr/prepost`

Note: For more information about creating the pre/post scripts, see *Manage Pre/Post Scripts for Automation* (see page 201).

4. 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.

(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/CA/d2dserver/usr/prepost

Follow these steps:

1. Place the file in the following location of the Backup Server:
/opt/CA/d2dserver/usr/prepost
2. Provide the execution permission to the script file.
3. Log into the Arcserve UDP Agent (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 should perform an immediate manual backup without waiting for the next scheduled backup to occur.

Follow these steps:

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 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.

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 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.

Troubleshooting

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 in to 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.

resources

Nodes: Linux Backup Server Groups

Actions	Node Name	Plan
☒	Server 1	
☐	Server 2	
☐	Server 3	

Update ...
Delete

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, you have to ensure specific prerequisites (To back up SQL Server and Exchange Server, there are no prerequisites 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](#) (see page 209)

The following diagram illustrates the process to protect host-based virtual machine nodes.



What To Do Next?

1. [Review the Prerequisites and Considerations](#) (see page 208)
2. [Create a Host-Based Backup Plan](#) (see page 210)
3. [\(Optional\) Perform a Manual Backup](#) (see page 234)
4. [Verify the Plan](#) (see page 235)
5. [Troubleshooting](#) (see page 235)

Review the Prerequisites and Considerations

Verify if you have completed the following prerequisites:

- Log in to the Console.
Prepare a host-based backup proxy server where you have installed Arcserve UDP Agent (Windows).
- To run functions such as Preflight Check, pre/post commands, or application log purge, use 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.
- 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](#) (see page 209).
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Consider the following points before you back 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 is an expected behavior.

Review the Prerequisites to Perform Application Consistent Backup for Oracle Database

To back up an Oracle database with consistent data, ensure that the following prerequisites are met:

- The ARCHIVELOG mode is enabled to archive the Redo logs.

Follow these steps to check if the ARCHIVELOG mode is enabled:

- Log in to the Oracle server as an Oracle user with SYSDBA privileges.
- Enter the following command at the SQL*Plus prompt:

```
ARCHIVE LOG LIST;
```

Archive log settings for the current instance is displayed.

- Configure the following settings:

Database log mode: Archive Mode

Automatic archival: Enabled

- Start the ARCHIVELOG mode.

If the ARCHIVELOG mode is not enabled, you must start the ARCHIVELOG mode to backup the database.

Follow these steps to start the ARCHIVELOG mode:

- Shut down the Oracle server.
- Run the following statements in Oracle:

```
CONNECT SYS/SYS_PASSWORD AS SYSDBA
```

```
STARTUP MOUNT;
```

```
ALTER DATABASE ARCHIVELOG;
```

```
ALTER DATABASE OPEN;
```

By default, archive logs will be written to the flash recovery area. If you do not want to write archive logs to the flash recovery area, you can 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: No Archive Mode

Automatic archival: Disabled

Archive destination:

E:\app\oracle\oradata\<oracle_database_name>\arch

Oldest online log sequence: 21

Current log sequence: 23

Oracle VSS Writer Service is started and 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 to pause the plan.

The plan will not run until you clear the check box to resume the plan.

Note: After a plan is paused, all jobs are paused except the restore job and the copy job. The running jobs are not affected. If you pause a plan that has pending jobs, then those pending jobs will also pause. When you resume the plan, the pending jobs does not resume immediately. After you resume the plan, the pending jobs will run from the next scheduled time. You can find the schedule of the next job form the home page of Arcserve UDP Agent (Windows).

6. From the **Task Type** drop-down list, select **Backup, Host-Based Agentless**.

Add a Plan
<Enter Plan Name>
☐ Pause this plan

Task1: Backup: Host-Based Agentless

+ Add a Task

Product Installation

Task Type
Backup: Host-Based Agentless

Source Destination Schedule Advanced

Backup Proxy

☒ Node Name	VM Name	Hypervisor

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 then click **Set Backup Proxy** to provide the proxy server details.

The **Backup Proxy** dialog opens.

The proxy server is a node where you install the Arcserve UDP Agent (Windows). If the Agent is not installed on this proxy server, then, 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.

Note: If the proxy server for a host-based agentless plan is an x86 OS and the virtual machine node resides on ESXi 5.5, the plan deployment will fail for the virtual machine. The reason for the deployment failure is that VMware VDDK 5.5.x is required to interact with ESXi 5.5. If the proxy server is an x86 OS, then Arcserve UDP Agent (Windows) uses VMware VDDK 5.1.2. The x86 OS does not support VMware VDDK 5.5.x.

2. Select one of the following options to locate the backup proxy:



Backup Proxy

How do you want to locate a backup proxy?

Select from Node List

Select from Node List

Enter Hostname/IP Address

Node Name	VM Name	Description

Select from Node List

Lets you select the node from the displayed list to assign it as a proxy server. The nodes that you have already added are displayed on the list.

Enter Hostname/IP Address

Lets you add a node as a proxy server using the node name or IP address, username, and password. You do not have to add the port number and protocol. The port number and protocol are configured at the **configuration** tab.

3. Click **Save**.

The backup proxy is selected and the dialog closes.

4. Click **Add Nodes** and select one of the following options:

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.


Importing from Hyper-V

Opens the **Add Nodes to Arcserve UDP Console** dialog. Select this option if you have not added the nodes and you want to import the nodes from a Hyper-V server.

Importing from vCenter/ESX

Opens the **Add Nodes to Arcserve UDP Console** dialog. Select this option if you have not added the nodes and you want to import the nodes from a vCenter/ESX server.

5. (Optional) Select a filter from the **Groups** drop-down list to filter nodes. Enter keywords to further filter your nodes.



The nodes are displayed on the **Available Nodes** area.

6. Select the nodes from the **Available Nodes** area and click the **Add all nodes (>>)** or **Add selected nodes (>)** icon.

The selected nodes are displayed on the **Selected Nodes** area.

7. (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 the 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. For VM which resides ESXi 4.x, VIX must be installed on proxy machine. 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](#) (see page 116).

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 job 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](#) (see page 238) topic.
8. (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](#) (see page 712)

- [NBD transport mode](#) (see page 712)
- [NBDSSL transport mode](#) (see page 713)
- [SAN transport mode](#) (see page 714)

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](#) (see page 216).

9. (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 checkbox 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. 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.

10. (Optional) Select the snapshot separation option for Hyper-V. This option is applicable for Hyper-V only.

Backup each VM individually using separate snapshot

Indicates that when you select this option Arcserve UDP captures separate snapshot for each virtual machine that is specified in the current plan. However, it increases the workload of the Hyper-V host when multiple snapshots are captured. If you do not select this option, Arcserve UDP captures one VSS snapshot for all the virtual machines if the backup jobs start at the same time. It is recommended to disable this option.

For more information about the separate snapshot state of a virtual machine, see the [troubleshooting](#) (see page 240) topic.

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) to use 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.

Host-based VM backup can execute backups using the following transport modes:

- [HOTADD transport mode](#) (see page 712)
- [NBD transport mode](#) (see page 712)
- [NBDSSL transport mode](#) (see page 713)
- [SAN transport mode](#) (see page 714)

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 UDP Agent (Windows) backup proxy server.
2. Open Windows Registry Editor and browse to the following key:
[HKEY_LOCAL_MACHINE\SOFTWARE\CA\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\CA\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\CA\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 UDP Agent (Windows) backup proxy server for the virtual machines.
2. Open Windows Registry Editor and browse to the following key:
[HKEY_LOCAL_MACHINE\SOFTWARE\CA\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 do not have the option to 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](#) (see page 711).
 - c. Optionally, provide an encryption password.
 - d. Confirm the encryption password.
 - e. Select a type of compression. For more information, see [Compression Type](#) (see page 709).

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](#) (see page 182).

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.

New Backup Schedule

Custom

Backup Type: Incremental

Start Time: 8:00 AM

☐ Sunday ☐ Monday ☐ Tuesday
☐ Wednesday ☐ Thursday ☐ Friday
☐ Saturday

Repeat: ☒

Every: 3 Hours

Until: 6:00 PM

Help Save Cancel

- 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 checkbox 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.

Source **Destination** **Schedule** **Advanced**

+ Add Delete		Type	Description	Su	Mo	Tu	We	Th	Fr	Sa	Time
☒			Custom Incremental Backups Every 3 Hours	✓	✓	✓	✓	✓	✓	✓	8:00 AM - 6:00 PM
☐			Weekly Incremental Backup						✓		8:00 PM

Add Merge Schedule

- Click **Add** and select **Add Merge Schedule**.
- The **Add New Merge Schedule** dialog opens.
- Specify the start time to start the merge job.
- Specify **Until** to specify an end time for the merge job.
- Click **Save**.

The Merge Schedule is specified and it is displayed on the **Schedule** page.

Add Throttle Schedule

- Click **Add** and select **Add Throttle Schedule**.
- The Add New Throttle Schedule dialog opens.
- Specify the throughput limit in MB per minute unit.
- Specify the start time to start the backup throughput job.
- Specify **Until** to specify an end time for the throughput job.
- Click **Save**.

The Throughput Schedule is specified and it is displayed on the **Schedule** page.

- Specify the start time for the scheduled backup.

Start time for scheduled backup 10/24/2014  1 : 27 PM

Recovery Point Retention

7	Daily Backups
	Weekly Backups
	Monthly Backups
31	Custom / Manual Backups

- 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.

- Specify the catalog details.

Catalogs	Generate file system catalogs (for faster search) after
	☐ Daily Backups
	☐ Weekly Backups
	☐ Monthly Backups
	☐ Custom / Manual Backups
	Generate Exchange catalogs (for granular restore) after
	☐ All Backups for Nodes with Exchange Installed

Catalogs let you generate the Exchange Granular Restore catalog and the File System catalog. The Exchange Granular Restore catalog is required to restore Exchange mailboxes, mailbox folders, and individual mail objects. 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.

- (Optional) Select one of the backup options in **Recovery Point Check**.

Recovery Point Check	Tests for data corruption by mounting the recovery point and running chkdsk.
	☐ Daily Backups
	☐ Weekly Backups
	☐ Monthly Backups
	☐ Custom / Manual Backups

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 backup job fails. 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 RAM 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. As a best practice, disable this option for frequent backup jobs such as custom or daily backup job.
- The recovery points check is done at the last stage of the backup job. If the chkdsk command detects error, the backup job fails and the associated recovery point is cleaned. In some cases the backup job takes longer time to complete and still you don't get a recovery point if an error is detected. If you want the backup to be successful even if the recovery point check fails, create a DWORD value named *CheckRecoveryPointDontFailJob* in the registry of the proxy machine and set the DWORD value to 1. Create the DWORD value at the following location:

HKEY_LOCAL_MACHINE\SOFTWARE\CA\ARCserve Unified Data Protection\Engine\AFBackupDll

The DWORD is applicable for all the backup jobs running on the current proxy server. If you want to control the behavior of a specific virtual machine, you can set that value under *HKEY_LOCAL_MACHINE\SOFTWARE\CA\ARCserve Unified Data Protection\Engine\AFBackupDll\<VM GUID>*.
- If the backup is crash consistent, there are high chances that chkdsk will detect problems (due to the nature of a crash consistent backup). There may be frequent backup job failure. As a best practice, do not enable this option for a crash consistent backup.

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: Retention settings must be set 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 have to 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 will run at 6:00 AM, 7:00 AM, 8:00 AM, but NOT 9:00 AM.

Backup Throttle Schedule

Backup throttle schedule lets you control the backup throughput speed, which 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 will be 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 will adjust 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 will be 500 MB/minute and from 8:00 PM to 10:00 PM the throttle limit will be 2500 MB/minute.

If you do not define any backup schedule and backup throughput schedule, the backup will run 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 being 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 will process these sets one by one.
- If a merge job is resumed after a pause, the job detects at which point it was 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.

Note: For VMware ESX server older than version 5.0, VMware VIX API is required on the proxy server to perform the operations such as application log truncation, pre-post command. It is also recommended to install VMware VIX API on the machine where you have installed the Arcserve UDP Console for operations such as PFC.

The following image displays the Advanced tab:

The screenshot shows the 'Advanced' tab of a backup configuration window. At the top, there are four tabs: 'Source', 'Destination', 'Schedule', and 'Advanced', with 'Advanced' being the active tab. Below the tabs, there are two main sections: 'Truncate log' and 'Run Commands'. The 'Truncate log' section has a note: 'Only VMware Windows Virtual Machines support Truncate log. For VMware ESX server before version 5, the proxy server must have VIX installed.' It contains two checkboxes: 'Truncate SQL Server log' and 'Truncate Exchange Server Log'. Each checkbox is followed by a dropdown menu set to 'Weekly'. The 'Run Commands' section has a note: 'Only Windows Virtual Machines support Run Commands. For VMware ESX server before version 5, the proxy server must have VIX installed.' It contains three checkboxes: 'Before a backup is started', 'On exit code', and 'After a snapshot is taken'. The 'Before a backup is started' checkbox is followed by a text input field. The 'On exit code' checkbox is followed by a text input field containing '0' and two radio buttons: 'Run Job' (selected) and 'Fail Job'. The 'After a snapshot is taken' checkbox is followed by a text input field. The 'After a backup is over' checkbox is followed by a text input field.

Source **Destination** **Schedule** **Advanced**

Only VMware Windows Virtual Machines support Truncate log. For VMware ESX server before version 5, the proxy server must have VIX installed.

Truncate log ☐ Truncate SQL Server log

Weekly

☐ Truncate Exchange Server Log

Weekly

Only Windows Virtual Machines support Run Commands. For VMware ESX server before version 5, the proxy server must have VIX installed.

Run Commands ☐ Before a backup is started

☐ On exit code 0 ☒ Run Job ☐ Fail Job

☐ After a snapshot is taken

☐ After a backup is over

Username for Commands	
Password for Commands	
Enable Email Alerts	☒ Email Settings
Job Alerts	☐ Missed jobs
	☒ Backup, Catalog, Restore, or Copy Recovery Point job failed/crashed/canceled
	☐ Backup, Catalog, Restore, or Copy Recovery Point job successfully completed
	☒ Merge job stopped, skipped, failed or crashed
	☐ Merge job success
	☐ Skip/Merge job waiting in the job queue
	☐ Hypervisor could not be reached (before backup)
Backup destination free space is less than	☐ 5 % v

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**. This is applicable only for VMware.

User Name

Lets you specify the user who is authorized to run a script.

Note: Use only the default administrator or domain administrator user credentials. If the user name does not belong to the default administrator or domain administrator, then you must disable User Access Control (UAC).

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. This is applicable only for Windows VM.

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. This is applicable only for Windows VM.

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. This is applicable only for Windows VM.

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**.

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 plan is automatically deployed to the source virtual machine node.

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. Use the vSphere SDK or VIX , so that the network access is not required to log in to the VM. 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 default administrator account or default 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](#) (see page 117) 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. Use the remote Windows Management Instrumentation (WMI) to log in to the virtual machine.

If either of the account is not set, use the additional administrator account to connect to the virtual machine and launch the command or script. Use a network to access the virtual machine.

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](#) (see page 117) 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.

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 virtual machine proxy server in your backup environment. By default, Host-Based VM Backup can run up to ten backup jobs concurrently. In environments that contain many virtual machines that are associated with a virtual machine proxy system, 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 http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1022543 (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\CA\ARCserve Unified Data Protection\Engine
3. Locate the following key:
VMMMaxJobNum

Note: The VMMMaxJobNum key is already created and the default value is 10.

4. Right-click VMMMaxJobNum and 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

6. Click OK.

The limit is defined.

7. Restart the Arcserve D2D service.

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 should perform an immediate manual backup without waiting for the next scheduled backup to occur.

Follow these steps:

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 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.

The manual backup is successfully performed.

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.

Troubleshooting

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 Update 2 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 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 5 and another is Arcserve UDP Version 5 Update 2. These two Arcserve UDP environments back up different VMs in the same Hyper-V server. The Arcserve UDP Version 5 Update 2 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.

- 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\CA\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.

The 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 occurs because of the following 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:\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:\as-hbbu-vmwarebackup`
`C:\ASVMOperationTools.`
 - e. Resubmit the backup job.

The backup job is complete but the VM is in the Backing up status

Valid for Hyper-V VM

Symptom:

In Hyper-V 2012 or later, the virtual machine stays in the *Backing up* status although the agent-less host-based backup job of this virtual machine has already finished. I cannot perform some operations such as power-on or power-off the virtual machine during that time in the Hyper-V manager. If the virtual machine is a Hyper-V cluster, I cannot perform live migration for it. In addition, if another backup job for this VM starts at the same time, the backup job fails with the following error:

The Hyper-V VSS writer has encountered an error when processing this virtual machine.

This problem occurs in the following situations:

- There are several backup jobs starting at the same time or at times closer to each other (within one minute).
- One or more backup jobs are complete, but there is still at least one backup job in progress.

Solution:

If the backup jobs start at the same time or at times closer to each other, Arcserve UDP takes one VSS snapshot for all virtual machines instead of taking one VSS snapshot for each virtual machine. This avoids unnecessary workload to the Hyper-V host. After the VSS snapshot is taken, all the virtual machines inside this VSS snapshot instance are locked (in the *Backing up* status). Arcserve UDP cannot release the snapshot until all backup jobs are finished, even if the backup job of a virtual machine is already completed.

The VSS snapshot has a limitation. Only one snapshot can be taken for a virtual machine at a time. If another backup job of the same virtual machine starts at this time, it fails and provides the error message. This error does not happen in Hyper-V 2008R2 because Hyper-V 2008R2 has a different VSS snapshot mechanism.

While the virtual machine is locked, you can still use the guest OS. The lock has no impact on the usage or availability of the guest OS. However, to avoid this situation, you can perform either of the following tasks:


- Enable the **Hyper-V Snapshot Separation** option on the resource tab of the agent-less host-based backup. Then, Arcserve UDP takes a separate snapshot for each virtual machine specified in the plan. The virtual machine is released after the backup is complete.
- Use different plans to protect virtual machines that have different storage sizes. Include the virtual machines with similar storage size in one plan. It will ensure that the backup jobs take a similar amount of time to complete. Also, set different schedules different plans.

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 gets skipped.

The following diagram illustrates the process to create a virtual standby plan.



What To Do Next?

1. [Review the Prerequisites and Considerations](#) (see page 243)
2. [Create a Plan with a Backup Task](#) (see page 244)
3. [Add a Virtual Standby Task to the Plan](#) (see page 258)
4. [\(Optional\) Run the Virtual Standby Job Manually](#) (see page 269)
5. [Pause and Resume the Heartbeat](#) (see page 270)
6. [Pause and Resume the Virtual Standby Job](#) (see page 271)
7. [Verify the Plan](#) (see page 272)
8. [Applying Best Practices](#) (see page 273)

Review the Prerequisites and Considerations

Verify if you have completed the following prerequisites:

- Log in to 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

Note: If you import nodes from Arcserve High Availability, you can create a directly virtual standby task without creating any other task. For more information on creating virtual standby task from Arcserve High Availability nodes, see [How to Create a Virtual Standby Plan for Arcserve High Availability Nodes](#) (see page 274).

- Review the [Compatibility Matrix](#) which 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.

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.

Note: For more information on host-based agentless backup, see [How to Create a Host-Based Virtual Machine Backup Plan](#).

Note: For more information on replicating a backup data, see [How to Create a Recovery Point Server Replication Plan](#) (see page 307).

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: After a plan is paused, all jobs are paused except the restore job and the copy job. The running jobs are not affected. If you pause a plan that has pending jobs, then those pending jobs will also pause. When you resume the plan, the pending jobs does not resume immediately. After you resume the plan, the pending jobs will run from the next scheduled time. You can find the schedule of the next job form the home page of Arcserve UDP Agent (Windows).

6. From the **Task Type** drop-down list, select **Backup, Agent-Based Windows**.

Add a Plan

New Plan

☐ Pause this plan

Task1: Backup: Agent-Based Windows

+

Add a Task

Product Installation

Task Type

Backup: Agent-Based Windows

Source

Destination

Schedule

Advanced

+

Add Nodes

Remove

☒ Node Name

VM Name

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 and click **Add Node**.
2. Select one of the following options:

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.


Adding Windows Nodes

Opens the **Add Nodes to Arcserve UDP Console** dialog. Select this option if you have not added the nodes and you want to manually add the nodes to protect.

Discovering 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. (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** area.

4. Select the nodes from the **Available Nodes** area and click the **Add all nodes (>>)** or **Add selected nodes (>)** icon.

The selected nodes are displayed on the **Selected Nodes** area.

5. Click **OK** to close the dialog.
6. To choose **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](#) (see page 711).
 - c. Optionally, provide an encryption password.

- d. Confirm the encryption password.
- e. Select a type of compression. For more information, see [Compression Type](#) (see page 709).

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](#) (see page 182).

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.

The screenshot shows the 'New Backup Schedule' dialog box. At the top, there is a dropdown menu set to 'Custom'. Below this, the 'Backup Type' is set to 'Incremental'. The 'Start Time' is '8:00 AM'. A section for selecting days of the week has checkboxes for Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, and Saturday, all of which are currently unchecked. The 'Repeat' checkbox is checked. Below the repeat checkbox, there is a section for frequency: 'Every 3 Hours' and 'Until 6:00 PM'. At the bottom, there are three buttons: 'Help', 'Save', and 'Cancel'.

- 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 checkbox 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 it is displayed on the **Schedule** page.

		+ Add	Delete							
☐	Type	Description	Su	Mo	Tu	We	Th	Fr	Sa	Time
☒		Custom Incremental Backups Every 3 Hours	☒	☒	☒	☒	☒	☒	☒	8:00 AM - 6:00 PM
☐		Weekly Incremental Backup						☒		8:00 PM

Add Merge Schedule

- a. Click **Add** and select **Add Merge Schedule**.

The **Add New Merge Schedule** dialog opens.

- Specify the start time to start the merge job.
- Specify **Until** to specify an end time for the merge job.
- Click **Save**.

The Merge Schedule is specified and it is displayed on the **Schedule** page.

Add Throttle Schedule

- Click **Add** and select **Add Throttle Schedule**.

The Add New Throttle Schedule dialog opens.

- Specify the throughput limit in MB per minutes unit.
- Specify the start time to start the backup throughput job.
- Specify **Until** to specify an end time for the throughput job.
- Click **Save**.

The Throttle Schedule is specified and it is displayed on the **Schedule** page.

3. Specify the start time for the scheduled backup.

Start time for scheduled backup 

Recovery Point Retention

7	Daily Backups
	Weekly Backups
	Monthly Backups
31	Custom / Manual Backups

- 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.

- Specify the catalog details.

Catalogs

Generate file system catalogs (for faster search) after

☐ Daily Backups

☐ Weekly Backups

☐ Monthly Backups

☐ Custom / Manual Backups

Generate Exchange catalogs (for granular restore) after

☐ All Backups for Nodes with Exchange Installed

Catalogs let you generate the Exchange Granular Restore catalog and the File System catalog. The Exchange Granular Restore catalog is required to restore Exchange mailboxes, mailbox folders, and individual mail objects. The File System catalog is required to perform faster and easier 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: Retention settings must be set 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 have to 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 will run at 6:00 AM, 7:00 AM, 8:00 AM, but NOT 9:00 AM.

Backup Throttle Schedule

Backup throttle schedule lets you control the backup throughput speed, which 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 will be 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 will adjust 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 will be 500 MB/minute and from 8:00 PM to 10:00 PM the throttle limit will be 2500 MB/minute.

If you do not define any backup schedule and backup throughput schedule, the backup will run 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 being 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 will process these sets one by one.
- If a merge job is resumed after a pause, the job detects at which point it was 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

Truncate log

☐ Truncate SQL Server log

Weekly

☐ Truncate Exchange Server Log

Weekly

Run Commands

☐ Before a backup is started

☐ On exit code

0

☒ Run Job ☐ Fail Job

☐ After a snapshot is taken

☐ After a backup is over

Username for Commands

Password for Commands

Enable Email Alerts



Email Settings

Job Alerts



Missed jobs



Backup, Catalog, Replication, File Copy, Restore or Copy Recovery Point job failed/crashed/canceled



Backup, Catalog, Replication, File Copy, Restore or Copy Recovery Point job successfully completed



Merge job stopped, skipped, failed or crashed



Merge job success

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.

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.

Note: 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 nodes imported from Arcserve High Availability. You have to manually power on recovery point snapshots for such nodes.

Note: 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 UDP 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 you want to monitor the status of the source server.

Note: The monitor server can be any physical computer or virtual machine provided that the server is not the backup source.

User Name

Specify the user name that is required to log in to the monitoring system.

Password

Specify the password for the user name that is required to log in to the monitoring system.

Protocol

Specify HTTP or HTTPS as the protocol that you want to use for communication between the Arcserve UDP and the ESX Server system (monitoring server).

Port

Specify the port that you want to use for data transfer between the Arcserve UDP 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 UDP 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 UDP 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 UDP server and the Hyper-V Server system (monitoring server).

Port

Specify the port that you want to use for data transfer between the Arcserve UDP server and the Hyper-V Server system (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.

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 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 checkbox 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, nodes replicated from a remote recovery point server, and nodes imported from Arcserve High Availability.

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 and nodes imported from Arcserve High Availability, 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 and nodes imported from Arcserve High Availability, 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 <i>is not</i> specified	The Power on the standby virtual machine with customized network configurations option <i>is</i> 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 configurations are

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**:

Network Adapter Settings **DNS Update Settings**

Specify the virtual network, virtual NIC, and TCP/IP settings for each network adapter.

Source Machine Network Adapter	Standby VM - Virtual Network	Standby VM - NIC Type
Adapter1	Use plan settings - VM Network	Use plan settings - E1000
	Use plan settings - VM Network	
	VM Network	
	VM NIC Performance	

TCP/IP Settings

Source Machine	Standby VM
Adapter: Adapter1 IP addresses/Subnet masks DHCP Enabled Gateways Automatic DNS servers Automatic WINS servers Automatic	☐ Retain the network settings from the backup. ☒ Customize the TCP/IP settings.
	IP Addresses Automatic(DHCP Enabled) Gateway Addresses Automatic DNS Addresses Automatic WINS Addresses Automatic

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


When you submit the backup job, the password for the backup is stored on the Arcserve UDP Agent (Windows) node that you are protecting. The Arcserve UDP solution then replicates the recovery points to a remote recovery point server. The converter on the remote server then converts the replicated data to virtual machine data and stores the data on the remote destination. However, the converter cannot convert the replicated recovery points because the backup passwords reside on the Arcserve UDP Agent (Windows) node.

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.



The dialog box is titled "Set Backup Passwords for Node <Node Name>". It contains a blue header bar with the text: "Enter one or more backup encryption passwords. During the conversion process, all passwords will be used in succession to try to decrypt the session. Note: When all of the passwords are not valid, conversion jobs will fail." Below the header bar, there are two buttons: "+ Add" and "Delete". Below these buttons is a table with four columns: "Password", "Confirm Password", "Comment", and "Create Time". The table currently contains one row with the text "No password" in the "Password" column. At the bottom of the dialog box, there are three buttons: "Save", "Cancel", and "Help".

✓ Password	Confirm Password	Comment	Create Time
No password			

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, 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 job runs immediately after the backup job is over.

The virtual standby job is manually run.

Pause and Resume Heartbeat

The Arcserve UDP 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 individual node level.
- You can pause and resume heartbeats for one or more nodes in one step.
- The Arcserve UDP 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 UDP 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 UDP.
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 UDP 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

Follow these steps:

1. Log in to Arcserve UDP.
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

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 particular processes, 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 avoid the VM from entering the save mode, exclude the following processes, folders, and files from the antivirus scanning:

- Process list
 - C:\Program Files\CA\ARCserve Unified Data Protection\Engine\TOMCAT\bin\tomcat7.exe
 - C:\Program Files\CA\ARCserve Unified Data Protection\Engine\TOMCAT\JRE\bin
 - java.exe
 - java-rmi.exe
 - javaw.exe
 - keytool.exe
 - rmid.exe
 - rmiregistry.exe

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 targets Hyper-V virtual machines and Hyper-V processes:


- 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
- Hyper-V process:
 - %windows%\system32\Vmms.exe
 - %windows%\system32\Vmwp.exe

How to Create a Virtual Standby Plan for Arcserve High Availability Nodes

The Arcserve UDP solution integrates with Arcserve High Availability (HA) to create virtual standby machines from the recovery points that are replicated by HA.

The following diagram illustrates the process to create a virtual standby plan for HA nodes:



What To Do Next?

- [Review the Prerequisites for Remote Virtual Standby](#) (see page 275)
- [Create a Virtual Standby Plan for HA Nodes](#) (see page 275)
- [Configure Remote Converters](#) (see page 285)
- [Verify the Plan](#) (see page 286)

Review the Prerequisites for Remote Virtual Standby

Verify that you have completed the following prerequisite tasks:

- Created Arcserve UDP Agent (Windows) or host-based product scenario from Arcserve High Availability control service
Note: For more information on creating a Arcserve High Availability scenario, see the Arcserve High Availability documentation.
- Added the HA nodes using **Importing Nodes from arcserve HA** in **Add Nodes**.
Note: For more information on adding HA nodes, see Import Nodes from arcserve HA.
- Imported nodes from Arcserve High Availability to the remote Arcserve UDP.
Note: The virtual standby task uses the recovery points replicated by Arcserve High Availability, and not the Arcserve High Availability nodes, to create virtual standby machines.
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Create a Virtual Standby Plan for HA Nodes

After you import the nodes from Arcserve High Availability to Arcserve UDP, you can create a virtual standby plan to create the standby virtual machines. The virtual machines can be either VMware or Hyper-V.

Note: The **Virtual Standby** task in Task 1 is meant for nodes imported from Arcserve High Availability only. If you add **Virtual Standby** in Task 1, you cannot add another task to this 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 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.

6. From the **Task Type** drop-down list, select **Virtual Standby**.

Add a Plan

Sample Plan ☐ Pause this plan

Task1: Virtual Standby

Task Type: Virtual Standby

Source Virtualization Server Virtual Machine Advanced

☐ Add Nodes Remove

☒ Node Name VM Name

Now specify the Source, Virtualization Server, Virtual Machine, 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 but the plan will not be deployed unless you add any nodes.

Follow these steps:

1. Click the **Source** tab and then click **Add Nodes**.
2. Select one of the following options:


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.

Importing Nodes from Arcserve HA

Opens the **Add Nodes to arcserve UDP Console** dialog. Select this option if you have not added the nodes and you want to manually add the nodes to protect.

3. (Optional) Select a filter from the **Groups** drop-down list to filter nodes. Enter keywords to further filter your nodes.



The nodes are displayed on the **Available Nodes** area.

4. Select the nodes from the **Available Nodes** area and click the **Add all nodes (>>)** or **Add selected nodes (>)** icon.

The selected nodes are displayed on the **Selected Nodes** area.

5. Click **OK** to close the dialog.

The source is specified.

Specify the Virtualization Server

You can provide the virtualization server details.

Follow these steps:

1. Enter the following details if the **Virtualization Type** is **VMware**.

Source **Virtualization Server** **Virtual Machine** **Advanced**

Virtualization Type

☒ VMware ☐ Hyper-V

ESX Host/vCenter

machine name/IP

Username

root

Password

.....

Protocol

☐ HTTP ☒ HTTPS

For more secure communication, HTTPS is the recommended protocol.

Port

port number

ESX Node

node name ▼

ESX Host/vCenter

Specify the host name of the ESX or vCenter Server system.

Username

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 UDP 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.

2. Enter the following details if the **Virtualization Type** is **Hyper-V**.

Source	Virtualization Server	Virtual Machine
Virtualization Type	☐ VMware	☒ Hyper-V
Hyper-V Host Name		
Username	Administrator	
Password	••••••••	
Protocol	☒ HTTP ☐ HTTPS	
For more secure communication, HTTPS is the recommended protocol.		
Port	8014	
Connect		

HyperV Host Name

Specify the host name of the Hyper-V system.

Username

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 UDP server and the Hyper-V Server system (monitoring server).

Port

Specify the port that you want to use for data transfer between the Arcserve UDP server and the Hyper-V Server system (monitoring server).

The virtualization server details are specified.

Specify the Virtual Machine

Specify the virtual machine details for VMware or Hyper-V virtualization servers.

Follow these steps:

1. Enter the following details for VMware virtual machine if you have selected VMware as the **Virtualization Server**.

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.

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.

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.

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 checkbox 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.

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.

2. Enter the following details for Hyper-V virtual machine if you have selected Hyper-V as the **Virtualization Server**.

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 count 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 checkbox to specify the location on the Hyper-V server where you want to store the conversion data for each virtual disk.

Note: Arcserve UDP does not support creating virtual disk images (VHD 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 check box 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.

The virtual machine is specified.

Specify the Advanced Settings

The Advanced page lets you specify advanced settings such as alerts.

Task Type	Virtual Standby ▼		

Source	Virtualization Server	Virtual Machine	Advanced

Enable Email Alerts	☒	Email Settings	
Job Alerts	☒	VM powered on for source machine configured with manual power ON	
	☒	Virtual Standby errors/failure/crash	
	☒	Virtual Standby success	
	☒	The Virtual Standby did not start successfully from the Recovery Point Snapshot	
Hypervisor is not reachable	☒		
VM storage free space less than	☒	5	% ▼

Follow these steps:

1. Specify the following details on the Advanced page.

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.

- **VM powered on for source machine configures 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.

2. Click **Save**.

The advanced settings are specified.

The changes are saved and the virtual standby task is automatically deployed to the virtual standby server.

Configure Remote Converters


Virtual standby lets you convert the recovery points replicated by arcserve Replication and High Availability. The recovery points are converted to a virtual machine format that is compatible with Microsoft Hyper-V, VMware vCenter, or ESXi.

When the nodes are imported from Arcserve High Availability to a remote Arcserve UDP server, the nodes can then be converted to a virtual machine format. The nodes are converted from the replica folder of Arcserve High Availability. By default, the converter is the location where the recovery points are replicated. You have to specify the node name and login credentials of the converter.

Follow these steps:

1. Log in to the Console.
2. Click the **resources** tab.
3. From the left pane, click **All Nodes**.
4. Select the **Converter** option as shown in the following diagram.

Nodes: All Nodes



The **Converter** column is added to the Console.

5. Click the converter that you want to configure from the **Converter** column.

The **Configure Remote Converters** dialog opens.

6. Specify the Port, Protocol, Username, and Password for the selected converter and click Update to save the information.

The converter is configured.

Note: When you import nodes from arcserve Replication and High Availability for the first time, a dialog pops-up automatically that asks you to configure the converter information. If you click Yes on the dialog, the Configure Remote Converters dialog opens.

Host Name	Port	Protocol	Username	Password	Information
<machine_name>	8014	HTTP			Need more information

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 steps 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 machines for Arcserve High Availability nodes are created.

How to View Virtual Standby Settings From the Monitor Server

After you have created and deployed the virtual standby plan, you can view the virtual standby settings from the monitor server.

The following diagram illustrates the process to view the virtual standby settings from the monitor server:



What To Do Next?

- [Review the Prerequisites and Considerations](#) (see page 289)
- [Log in to Monitor Servers](#) (see page 289)
- [Understanding the Virtual Standby Summary Screen](#) (see page 290)
- [View Activity Log](#) (see page 293)
- [View Virtual Standby Settings](#) (see page 293)
- [View the Email Settings](#) (see page 298)

Review the Prerequisites and Considerations

Verify that you have complete the following prerequisite tasks:

- Logged in to the Console
- Created and deployed a virtual standby plan
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Log In To Monitor Servers

Arcserve UDP lets you log in directly to the server that is monitoring the Arcserve UDP Agent (Windows) source nodes. From the monitor server, you can perform maintenance tasks and view information about the health of the source nodes.

Follow these steps:

1. Log in to the Console.
2. Click the **resource** tab.
3. Click **All Nodes**.
4. On the center pane, select the node that has the virtual standby task.
5. Right-click the node and select **Log in to Monitor Server**.

The monitor server interface opens in a new window.



Note: If a new browser window does not open, verify that the pop-up options for your browser allow all pop-ups or pop-ups only for this website.

6. Click the **Virtual Standby** tab.

The **Virtual Standby** page opens.

You are logged in to the monitor server.

Understanding the Virtual Standby Summary Screen

The **Virtual Standby Summary** screen displays icons that provide a quick visual indication of the current status, along with guidance for the urgency of any actions that you need to take.

The following icons appear on the home page:



Successful
(No action is necessary)



Caution
(Action may be necessary soon)



Warning
(Immediate action is necessary)


The **Virtual Standby Summary** screen displays the following information:

- **Servers list**--Displays a list of source servers (source nodes) that this monitoring server is protecting. The list sorts source servers by their current status. For example, **All**, **Action Required**, **Server Running**, and so on.
Note: The Servers list appears only when you are logged in to the monitoring server. For more information, see [How to Use the Servers List](#) (see page 290).
- **Virtual Standby Summary**--Displays summary information for the selected source server. For more information, see [Monitor the Status of Virtual Conversion Jobs](#) (see page 292).
- **Virtual Standby Settings**--Displays summary information about virtual conversion settings for the selected source server. For more information, see [View Virtual Standby Settings for Source Servers](#) (see page 291).
- **Recovery Point Snapshots**--Displays a list of recovery point snapshots that are available for the selected source server. For more information, see [View the Recovery Point Snapshots List](#) (see page 292).
- **Tasks**--Displays a list of tasks that you can perform for the selected source server.
- **Support and Community Access**--Provides a mechanism that lets you initiate various support-related functions.

Understanding the Servers List

The Servers list on the **Virtual Standby Summary** screen displays a list of source servers that a monitoring server is protecting. The list sorts servers by their current status. For example, **All**, **Action Required**, **Source Running**, and so on.

To perform maintenance tasks or to view information about a Arcserve UDP Agent (Windows) node, click the **Virtual Standby** tab and then click the server as illustrated by the following screen:



View Virtual Standby Settings for Source Servers

The **Virtual Standby Summary** screen displays information about the virtual machines that are protecting source servers.

Virtual Machine Information	
Type:	VMware ESX
ESX Host Name:	<host_name>
Version:	5.5.0
Virtual Machine Name:	<VM_Name>
Processor:	1
Memory:	4096 MB
Data Store:	DatastoreISCSI-2
Network Adapter:	
▼ Adapter1	
Adapter Type:	E1000
Network Connection:	VM Network


Monitor the Status of Virtual Conversion Jobs

Virtual standby lets you monitor the status of in-progress virtual conversion jobs. In addition, virtual standby lets you view the summary about the virtual conversion data and the virtual machines that are protecting your Arcserve UDP Agent (Windows) source servers.

Follow these steps:

1. Log in to the **Monitor Server**.
2. Click the **Virtual Standby** tab.

Virtual Standby Summary displays information about in-progress virtual conversion jobs and displays a summary about virtual conversion jobs and the virtual machine that is protecting the source server.



View the Recovery Point Snapshots List

The **Virtual Standby** screen displays a list of the most recent recovery point snapshots. The list box displays the date and time the backup of the Arcserve UDP Agent (Windows) source nodes when completed.

From the list of recovery point snapshots list, you can power on virtual machines. For more information, see [Power On Virtual Standby Machines from Recovery Point Snapshots](#) (see page 301).

Recovery Point Snapshots - Ready to Power on	
Time of backup	Action
1/9/2014 8:00:03 AM	Power On VM from this Snapshot
1/9/2014 6:35:21 AM	Power On VM from this Snapshot
1/9/2014 6:17:06 AM	Power On VM from this Snapshot
1/9/2014 4:39:27 AM	Power On VM from this Snapshot

Note: If the Virtual Standby destination is a VMware ESX server, the maximum number of recovery point snapshots that are displayed is 29. If the Virtual Standby destination is a Microsoft Hyper-V server, the maximum number of recovery point snapshots that are displayed is 24.

View Activity Log

Virtual standby lets you view Activity Log information about virtual conversion jobs. The Activity Log contains virtual conversion job records for the Arcserve UDP Agent (Windows) source node that you are protecting.

Note: The Activity Log (activity.log) is stored in the following directory on the node where Arcserve UDP Agent (Windows) is installed:

C:\Program Files\CA\arcserve Unified Data Protection\Engine\Logs

Follow these steps:

1. Login to the monitor server, and click the **Virtual Standby** tab.
2. Expand the servers from the **Sources** pane, and click a node to view its activity log.
3. From the **Virtual Standby Tasks** pane, click **View Logs**.

The **Activity Log** dialog opens.

View Virtual Standby Settings

The **Virtual Standby Settings** dialog contains information about the plan assigned to the node. You can view information about the Virtualization Server, the Virtual Machine, the Stand-in Server, and the Preferences defined in the plan that is assigned to the Arcserve UDP Agent (Windows) source node. You cannot edit settings from this dialog.

Follow these steps:

1. Click the **resource** tab on the Console.
2. From the left pane, navigate to **Nodes**, and click **All Nodes**.
3. On the center pane, select the node that you specified as **Monitor** in the **Virtual Standby** task, and click **Login Agent**.

A new browser opens and you are automatically logged in to the monitor server.

Note: If a new browser window does not open, verify that the pop-up options for your browser allow all pop-ups or pop-ups only for this website.

4. Click the **Virtual Standby** tab.

The **Virtual Standby Summary** screen appears.

5. From **Expand All** or **Server Running from the Servers** list, and click the node whose Virtual Standby settings you want to view.
6. From the **Virtual Conversion Tasks** list located on the right side of the **Virtual Standby Summary** screen, click **Virtual Standby Settings**.

The **Virtual Standby Settings** dialog opens.

Virtualization Server Options

■ **VMware Systems:**

The following options apply to VMware systems:

- **Virtualization type**--VMware.
- **ESX Host/vCenter**--Identifies the host name of the ESX or vCenter Server system.
- **User Name**--Identifies the user name that is required to log in to the VMware system.
- **Password**--Identifies that a password for User Name is required to log in to the VMware system.
- **Protocol**--Displays the communication protocol used between the source arcserve UDP Agent node and the monitoring server.
- **Port**--Identifies the port used for data transfer between the source server and the monitoring server.

■ **Monitoring:**

The following options apply to VMware systems.

- **Monitor Server**--Identifies the host name of the server that monitors the source server.
- **User Name**--Identifies the user name that is required to log in to the monitoring server.
- **Password**--Identifies that a password for User Name is required to log in to the monitor server.
- **Protocol**--Identifies the communication protocol used between the arcserve Central Virtual Standby server and the ESX Server system (monitoring server).

- **Port**--Identifies the port used for data transfer between the arcserve Central Virtual Standby server and the ESX Server system (monitoring server).
- **Use monitor server as proxy for data transfer**--Identifies that the monitor server copies the conversion data from the arcserve UDP Agent source server to the ESX Server data store.

Note: The Use monitor server as proxy for data transfer option is enabled by default. You can disable this option to allow the arcserve UDP Agent source server to copy the conversion data directly to the ESX Server data store.

- **Hyper-V Systems:**

The following options apply to Hyper-V systems:

- **Virtualization type**--Hyper-V.
- **Hyper-V Host Name**--Identifies the host name of the Hyper-V system.
- **User Name**--Identifies the user name that is required to log in to the Hyper-V system.
- **Password**--Identifies that a password for User Name is required to log in to the Hyper-V system.
- **Port**--Identifies the port used for data transfer between the source server and the monitoring server.

Virtual Machine Options

VMware Systems:

- **VM Name Prefix**--Identifies the prefix added to the display name for the virtual machine on the ESX Server system.
Default: UDPVM_
- **VM Resource Pool**--Identifies the name of resource pool where the standby virtual machine is grouped.
- **Datastore**--Identifies the location where you want to store the conversion data.

- **Use one datastore for all virtual machine source disks**--Indicates that the application copies all of the disks related to the virtual machine to one data store.
- **Choose a datastore for each VM source disk**--Indicates that the application copies disk-related information for the virtual machine to the corresponding data store.
- **Networks**--Identifies the NICs, virtual networks, and paths that the ESX Server system uses to communicate with the virtual machines.

Connect all virtual NICs to the following virtual network--Identifies the virtual NICs that are mapped to the virtual network. This option is specified when the virtual machine contains virtual NICs and a virtual network.

Choose a virtual network for each virtual NIC--Identifies the name of the virtual network that you want the NIC to use to communicate.
- **CPU Count**--Identifies the minimum and maximum CPU count supported by the standby virtual machine.
- **Memory**--Identifies the total amount of RAM in MB allocated for the standby virtual machine.
- **Recovery Point Snapshot**--Specify the number of recovery points for the standby virtual machine. The maximum number of recovery points is 24 for Hyper-V virtualization servers

Hyper-V Systems:

- **VM Name Prefix**--Identifies the prefix added to the display name for the virtual machine on the Hyper-V system.

Default: UDPVM_
- **Path**--Identifies the location on the Hyper-v Server where the conversion data is stored.

- **Networks**--Identifies the NICs, virtual networks, and paths that the Hyper-V server uses to communicate with the virtual machines.
- **CPU Count**--Identifies the minimum and maximum CPU count supported by the standby virtual machine.
- **Memory**--Identifies the total amount of RAM in MB allocated to the standby virtual machine.
- **Recovery Point Snapshot**--Specify the number of recovery points for the standby virtual machine. The maximum number of recovery points is 24 for Hyper-V virtualization servers

Stand-in Settings

Recovery:

- **Manually start the Virtual Machine**--Indicates that the virtual machines are powered on and provisioned manually when the source server fails or stops communicating.
- **Automatically start the Virtual Machine**--Indicates that the virtual machines are powered on and provisioned automatically when the source server fails or stops communicating.
- Heartbeat Properties:
 - Timeout**--Identifies the length time that the monitor server must wait for a heartbeat before it powers on a recovery point snapshot.
 - Frequency**--Identifies the frequency that the source server communicates heartbeats to the monitor server.

7. Click **Cancel** to close the **Virtual Standby Settings** dialog.

The virtual standby settings are viewed.

View the Email Settings

Configure the email settings to receive email alerts. You can configure the following email alert settings:

Follow these steps:

1. Click the **resource** tab on the Console.
2. From the left pane, navigate to **Nodes**, and click **All Nodes**.
3. On the center pane, select the node that you specified as **Monitor** in the **Virtual Standby** task, and click **Login Agent**.

A new browser opens and you are automatically logged in to the Monitor server.

Note: If a new browser window does not open, verify that the pop-up options for your browser allow all pop-ups or pop-ups only for this website.

4. Click the **Virtual Standby** tab.

The **Virtual Standby Summary** screen appears.

5. From **Expand All** or **Server Running from the Servers** list, and click a node to view its Virtual Standby settings.
6. From the Navigation pane, expand **Virtual Standby Tasks** and click **Virtual Standby Settings**.

The **Virtual Standby Settings** dialog opens.

7. Click the **Preferences** tab.
 - **Missing heartbeat for source machine**--Indicates that Virtual Standby sends alert notifications when the monitor server does not detect a heartbeat from the source server.
 - **VM powered on for source machine configured with auto power ON**--Indicates that 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.
 - **VM powered on for source machine configured with manual power ON**--Indicates that Virtual Standby sends alert notifications when it manually powers on a virtual machine.

- **VM storage free space less than**--Indicates that 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.
- **Virtual Standby errors/failure/crash**--Indicates that Virtual Standby sends alert notifications when it detects an error that occurred during the conversion process.
- **Virtual Standby success**--Indicates that the process of creating a virtual standby virtual machine completed successfully.
- **Hypervisor is not reachable**--Indicates that Virtual Standby sends alert notifications when it detects that it cannot communicate with the ESX Server system or the Hyper-V system.
- **The Virtual Standby did not start successfully from the Recovery Point Snapshot**--Indicates that the process of creating a virtual standby virtual machine from a recovery point snapshot did not complete successfully.


The email settings are viewed.

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:

How to Protect Virtual Standby Machines



What To Do Next?

- [Review the Prerequisites and Considerations](#) (see page 300)
- [Power On Virtual Standby Machines](#) (see page 301)
- [Protect Virtual Standby Machines After Power On](#) (see page 305)
- [Verify the Virtual Standby Machine Is Protected](#) (see page 306)


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](#) which 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 Task to the Plan](#) (see page 258).

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 VMware installs VMware Tools on the virtual machine or Windows Hyper-V installs Integration Services 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 UDP Agent (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 UDP server. For more information, see [Power on Virtual Standby Virtual Machines from Recovery Point Snapshots](#) (see page 301). 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 in to the Hyper-V server that is monitoring the nodes that you are protecting.
2. Start Hyper-V Manager by doing the following:
 - 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](#) (see page 301). 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 UDP Agent (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](#) (see page 271).

3. After you deploy the plan, log in to Arcserve UDP Agent (Windows) on the virtual standby machine and schedule a repeat method for the Arcserve UDP Agent (Windows) backup job.

For more information, see the Arcserve UDP Agent (Windows) User Guide.

Note: Arcserve UDP lets you automatically resynchronize the plans to the managed Arcserve UDP Agent (Windows) nodes on a weekly basis. This mechanism lets Arcserve UDP restart the backup jobs on the virtual standby machines by redeploying the plan that was in effect on the Arcserve UDP Agent (Windows) node to the virtual standby machine. The plan deployment process behaves in this manner because the source node and the virtual standby machine have the same host name, which lets Arcserve UDP resynchronize the plan. The only limitation to this behavior is Arcserve UDP and the virtual standby machine must be able to communicate with each other through the network. After Arcserve UDP resynchronizes and deploys the plan to the virtual standby machine, you then resume the Virtual Standby job on the virtual standby machine. For more information, see [Pause and Resume Virtual Standby Jobs](#).

Verify the Virtual Standby Machine is Protected

Verify the virtual standby machines are protected by confirming 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 and it is protected.


The virtual standby machines are 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.

The following diagram illustrates how to replicate data between data stores managed from a UDP Console.



What To Do Next?

1. [Review the Prerequisites and Considerations](#) (see page 308)

2. [Create a Plan with a Backup Task](#) (see page 174)
3. [Add a Replicate Task to the Plan](#) (see page 322)
4. [\(Optional\) Perform a Manual Replication](#) (see page 324)
5. [Verify the Plan](#) (see page 325)

Review the Prerequisites and Considerations

Verify if you have completed the following prerequisites:

- Log in to the Console.
- Install the server component and create Data Stores.
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

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. 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: After a plan is paused, all jobs are paused except the restore job and the copy job. The running jobs are not affected. If you pause a plan that has pending jobs, then those pending jobs will also pause. When you resume the plan, the pending jobs does not resume immediately. After you resume the plan, the pending jobs will run from the next scheduled time. You can find the schedule of the next job form the home page of Arcserve UDP Agent (Windows).

6. From the **Task Type** drop-down list, select **Backup, Agent-Based Windows**.

Add a Plan

New Plan ☐ Pause this plan

Task1: Backup: Agent-Based Windows

Task Type: Backup: Agent-Based Windows

Source Destination Schedule Advanced

☒ Node Name	VM Name

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 and click **Add Node**.
2. Select one of the following options:

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.


Adding Windows Nodes

Opens the **Add Nodes to Arcserve UDP Console** dialog. Select this option if you have not added the nodes and you want to manually add the nodes to protect.

Discovering 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. (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** area.

4. Select the nodes from the **Available Nodes** area and click the **Add all nodes (>>)** or **Add selected nodes (>)** icon.

The selected nodes are displayed on the **Selected Nodes** area.

5. Click **OK** to close the dialog.
6. To choose **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](#) (see page 711).
 - c. Optionally, provide an encryption password.

- d. Confirm the encryption password.
- e. Select a type of compression. For more information, see [Compression Type](#) (see page 709).

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](#) (see page 182).

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.

New Backup Schedule

Custom

Backup Type: Incremental

Start Time: 8:00 AM

☐ Sunday ☐ Monday ☐ Tuesday
☐ Wednesday ☐ Thursday ☐ Friday
☐ Saturday

Repeat: ☒

Every: 3 Hours

Until: 6:00 PM

Help Save Cancel

- 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 checkbox 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 it is displayed on the **Schedule** page.

Source **Destination** **Schedule** **Advanced**

+ Add ▼ Delete		Type	Description	Su	Mo	Tu	We	Th	Fr	Sa	Time
☒			Custom Incremental Backups Every 3 Hours	✓	✓	✓	✓	✓	✓	✓	8:00 AM - 6:00 PM
☐			Weekly Incremental Backup						✓		8:00 PM

Add Merge Schedule

- Click **Add** and select **Add Merge Schedule**.

The **Add New Merge Schedule** dialog opens.

- Specify the start time to start the merge job.
- Specify **Until** to specify an end time for the merge job.
- Click **Save**.

The Merge Schedule is specified and it is displayed on the **Schedule** page.

Add Throttle Schedule

- Click **Add** and select **Add Throttle Schedule**.

The Add New Throttle Schedule dialog opens.

- Specify the throughput limit in MB per minutes unit.
- Specify the start time to start the backup throughput job.
- Specify **Until** to specify an end time for the throughput job.
- Click **Save**.

The Throttle Schedule is specified and it is displayed on the **Schedule** page.

- Specify the start time for the scheduled backup.

Start time for scheduled backup  :

Recovery Point Retention

7	Daily Backups
	Weekly Backups
	Monthly Backups
31	Custom / Manual Backups

- 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.

- Specify the catalog details.

Catalogs

Generate file system catalogs (for faster search) after

☐ Daily Backups

☐ Weekly Backups

☐ Monthly Backups

☐ Custom / Manual Backups

Generate Exchange catalogs (for granular restore) after

☐ All Backups for Nodes with Exchange Installed

Catalogs let you generate the Exchange Granular Restore catalog and the File System catalog. The Exchange Granular Restore catalog is required to restore Exchange mailboxes, mailbox folders, and individual mail objects. The File System catalog is required to perform faster and easier 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: Retention settings must be set 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 have to 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 will run at 6:00 AM, 7:00 AM, 8:00 AM, but NOT 9:00 AM.

Backup Throttle Schedule

Backup throttle schedule lets you control the backup throughput speed, which 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 will be 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 will adjust 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 will be 500 MB/minute and from 8:00 PM to 10:00 PM the throttle limit will be 2500 MB/minute.

If you do not define any backup schedule and backup throughput schedule, the backup will run 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 being 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 will process these sets one by one.
- If a merge job is resumed after a pause, the job detects at which point it was 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:

The screenshot shows the 'Advanced' tab of a configuration window. It features four main sections: 'Truncate log', 'Run Commands', 'Username for Commands', and 'Password for Commands'. The 'Truncate log' section has two options: 'Truncate SQL Server log' and 'Truncate Exchange Server Log', each with a 'Weekly' frequency dropdown. The 'Run Commands' section includes checkboxes for 'Before a backup is started', 'After a snapshot is taken', and 'After a backup is over', each followed by a text input field. There is also an 'On exit code' checkbox with a value of '0' and radio buttons for 'Run Job' (selected) and 'Fail Job'. The 'Username for Commands' and 'Password for Commands' sections each have a single text input field.

Source	Destination	Schedule	Advanced
Truncate log			
		☐ Truncate SQL Server log	Weekly
		☐ Truncate Exchange Server Log	Weekly
Run Commands			
		☐ Before a backup is started	
		☐ On exit code	0
		☐ After a snapshot is taken	
		☐ After a backup is over	
Username for Commands			
Password for Commands			

Enable Email Alerts ☒ **Email Settings**

Job Alerts ☐ Missed jobs

☐ Backup, Catalog, Replication, File Copy, Restore or Copy Recovery Point job failed/crashed/canceled

☐ Backup, Catalog, Replication, File Copy, Restore or Copy Recovery Point job successfully completed

☐ Merge job stopped, skipped, failed or crashed

☐ Merge job success

Enable Resource Alerts ☒

CPU Usage	Memory Usage
Alert Threshold: 85 %	Alert Threshold: 85 %

Disk Throughput	Network I/O
Alert Threshold: 50 MB/s	Alert Threshold: 60 %

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 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. The replication destination must be a data store in the recovery point server. You can create multiple replicate task to perform multiple replications.

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.
3. Click the **Destination** tab and enter the recovery point server details and retry schedule details.

Task Type: Replicate ▼

[Source](#) [Destination](#) [Schedule](#) [Advanced](#)

Recovery Point Server:

Data Store:

When replication job fails:

Start retry

minutes later(1~60)

Retry

times(1~99)

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

- 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.

Source Destination Schedule Advanced

Add	Delete									
☒	Type	Description	Su	Mo	Tu	We	Th	Fr	Sa	Time

Numbers of Recovery Points to Retain at Replication Destination

Custom, Daily, Weekly, and Monthly Backups are defined by their corresponding schedules in the Backup task.

Custom	31
Daily	
Weekly	
Monthly	

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 have to first perform a manual backup. The replicate task is associated with a backup task. You cannot run the replicate task as a standalone job. If a plan includes a backup task and a replicate, then when you manually run the backup job, the replication job runs automatically after the completion of the backup job. So, you perform a manual replication by performing a manual backup.

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, 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 replication job runs immediately after the backup job is over.

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:

How to Replicate Data Between Data Stores Managed From Different UDP Consoles



What To Do Next?

1. [Review the Prerequisites](#) (see page 327)
2. [Create a User Account for the Source Console](#) (see page 327)
3. [Create a Plan to define the Destination Data Store](#) (see page 328)
4. [Map the Plan to the User Account](#) (see page 94)
5. [Send the Plan and User Account Details to the Source Administrator](#) (see page 331)
6. [Receive the Plan and User Account Details from the Destination Administrator](#) (see page 332)
7. [Create a Replication Plan to Send Data to the destination Console](#) (see page 332)
8. [Verify the Data is Replicated](#) (see page 337)

Review the Prerequisites

Review the following prerequisites before replicating data:

- Review the [Compatibility Matrix](#) which 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.

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 remote 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.

Add a Plan

Task1: Replicate from a remote RPS

+

Add a Task

Product Installation

Task Type

Replicate from a remote RPS

Source


Destination

Schedule

Advanced

Replicate recovery point from a remote Recovery Point Server.

6. Click the **Destination** tab and specify the recovery point server and the data store.
7. (Optional) Select the **Server is behind NAT router** checkbox 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](#) (see page 182).

11. Click **Save**.

The **Add New Merge Schedule** dialog closes.

12. Enter the recovery points retention details.

Number of Recovery Points to Retain at Replication Destination

Custom, Daily, Weekly, and Monthly Backups are defined by their corresponding schedules in the Backup task.

Daily	
Weekly	
Monthly	
Custom / Manual	31

13. Click the **Advanced** tab and provide 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.

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 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 **configuration** tab.
2. From the left pane, click **Share Plan**.



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

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

- Select the **User Account**.
- 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.
- Click **Add all plans** or **Add selected plans** to add the plans in the **Selected Plans** column.
- 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, you can 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**

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

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.

The screenshot shows a user interface for configuring tasks. On the left, there is a list of tasks: 'Task1: Backup: Agent-Based Windows' (with a green checkmark), 'Task2: Replicate to a remotely-managed RPS' (highlighted in dark blue), and 'Add a Task' (with a plus icon). Below this list is a 'Product Installation' button. On the right, the 'Task Type' is set to 'Replicate to a remotely-managed RPS'. Below this, there are three tabs: 'Source' (active), 'Destination', and 'Schedule'. Under the 'Source' tab, the text 'Replicate Recovery Point from Task1: Backup:' is visible.

7. Click the **Destination** tab and enter the following details.

Source	Destination	Schedule
Remote Console:	<Remote Console IP Address>	
Username:	Administrator	
Password:	••••••••	
Port:	8015	
Protocol:	☒ HTTP ☐ HTTPS	
Enable Proxy:	☐	
Proxy Server:		
Port:	8014	
Proxy server requires authentication	☐	
Username:		
Password:		
Connect		
Plan	▼	

Remote Console

Specify the IP address of the destination Console. The destination administrator provides you the destination Console address.

Username

Specify the username 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 checkbox 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 checkbox 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

Replication over WAN related settings are saved at the following registry key:

[HKEY_LOCAL_MACHINE\SOFTWARE\CA\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:00000001
- "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_Threshold 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 one.

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.

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.

How to Perform an Offline Data Replication Using RPS Jumpstart



What To Do Next?

- [Review the Prerequisites](#) (see page 341)
- [Create a Temporary Data Store on an External Device](#) (see page 342)
- [Replicate Source Data to the Temporary Data Store](#) (see page 343)
- [Delete the Temporary Data Store from the Source Console](#) (see page 344)
- [Send the External Device to the Destination Location](#) (see page 344)
- [Receive the External Device](#) (see page 344)
- [Import the Temporary Data from the External Device](#) (see page 345)
- [Create a Destination Data Store](#) (see page 345)
- [Replicate Data from the Temporary Data Store to the Destination Data Store](#) (see page 346)
- [Verify that the Data is Replicated](#) (see page 346)
- [\(Optional\) Set the Concurrent Node Count for RPS Jumpstart](#) (see page 347)

Review the Prerequisites

Review the following prerequisites before you perform an offline data replication:

- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

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](#) (see page 347).

Administrator—Destination Console

- Verify that there is enough 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 recovery point server and 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.

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.

Note: You can also use an existing data store as a destination data store.

To create a data store, follow the same procedure as described in [Create a Temporary Data Store on an External Device](#) (see page 342). Make sure that you create the data store on the same recovery point server where you imported the temporary 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 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 recovery point server and 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.

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, you cannot set the concurrent node count from the Console. 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\CA\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 Create a Copy Recovery Points Plan

Using Arcserve UDP, you can copy the recovery points to 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 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.

The following diagram illustrates the process to copy recovery points:

How to Create a Copy Recovery Points Plan



What To Do Next?

- [Review the Prerequisites and Considerations](#) (see page 348)
- [Create a Plan with a Backup Task](#) (see page 349)
- [Add a Copy Recovery Points Task to the Plan](#) (see page 363)
- [Verify the Plan](#) (see page 365)

Review the Prerequisites and Considerations

Verify if you have completed the following prerequisites:

- Log in to 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](#) which 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 a non-Windows VM.

Note: For more information on host-based agentless backup, see [How to Create a Host-Based Virtual Machine 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: After a plan is paused, all jobs are paused except the restore job and the copy job. The running jobs are not affected. If you pause a plan that has pending jobs, then those pending jobs will also pause. When you resume the plan, the pending jobs does not resume immediately. After you resume the plan, the pending jobs will run from the next scheduled time. You can find the schedule of the next job form the home page of Arcserve UDP Agent (Windows).

- From the **Task Type** drop-down list, select **Backup, Agent-Based Windows**.

Add a Plan

New Plan ☐ Pause this plan

Task1: Backup: Agent-Based Windows

+

Add a Task

Product Installation

Task Type

Backup: Agent-Based Windows

Source

Destination

Schedule

Advanced

+

Add Nodes

Remove

☒ Node Name

VM Name

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 and click **Add Node**.
2. Select one of the following options:

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.


Adding Windows Nodes

Opens the **Add Nodes to Arcserve UDP Console** dialog. Select this option if you have not added the nodes and you want to manually add the nodes to protect.

Discovering 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. (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** area.

4. Select the nodes from the **Available Nodes** area and click the **Add all nodes (>>)** or **Add selected nodes (>)** icon.

The selected nodes are displayed on the **Selected Nodes** area.

5. Click **OK** to close the dialog.
6. To choose **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](#) (see page 711).
 - c. Optionally, provide an encryption password.

- d. Confirm the encryption password.
- e. Select a type of compression. For more information, see [Compression Type](#) (see page 709).

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](#) (see page 182).

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.



The screenshot shows the 'New Backup Schedule' dialog box. At the top, there is a dropdown menu set to 'Custom'. Below this, the 'Backup Type' is set to 'Incremental'. The 'Start Time' is '8:00 AM'. A section for selecting days of the week has checkboxes for Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, and Saturday, all of which are currently unchecked. The 'Repeat' checkbox is checked. Below it, the frequency is set to 'Every 3 Hours'. The 'Until' time is '6:00 PM'. At the bottom, there are three buttons: 'Help', 'Save', and 'Cancel'.

- 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 checkbox 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 it is displayed on the **Schedule** page.

Source **Destination** **Schedule** **Advanced**

+ Add Delete		Type	Description	Su	Mo	Tu	We	Th	Fr	Sa	Time
☒			Custom Incremental Backups Every 3 Hours	✓	✓	✓	✓	✓	✓	✓	8:00 AM - 6:00 PM
☐			Weekly Incremental Backup						✓		8:00 PM

Add Merge Schedule

- Click **Add** and select **Add Merge Schedule**.

The **Add New Merge Schedule** dialog opens.

- Specify the start time to start the merge job.
- Specify **Until** to specify an end time for the merge job.
- Click **Save**.

The Merge Schedule is specified and it is displayed on the **Schedule** page.

Add Throttle Schedule

- Click **Add** and select **Add Throttle Schedule**.

The Add New Throttle Schedule dialog opens.

- Specify the throughput limit in MB per minutes unit.
- Specify the start time to start the backup throughput job.
- Specify **Until** to specify an end time for the throughput job.
- Click **Save**.

The Throttle Schedule is specified and it is displayed on the **Schedule** page.

- Specify the start time for the scheduled backup.

Start time for scheduled backup  :

Recovery Point Retention

7	Daily Backups
	Weekly Backups
	Monthly Backups
31	Custom / Manual Backups

- 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.

- Specify the catalog details.

Catalogs

Generate file system catalogs (for faster search) after

☐ Daily Backups

☐ Weekly Backups

☐ Monthly Backups

☐ Custom / Manual Backups

Generate Exchange catalogs (for granular restore) after

☐ All Backups for Nodes with Exchange Installed

Catalogs let you generate the Exchange Granular Restore catalog and the File System catalog. The Exchange Granular Restore catalog is required to restore Exchange mailboxes, mailbox folders, and individual mail objects. The File System catalog is required to perform faster and easier 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: Retention settings must be set 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 have to 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 will run at 6:00 AM, 7:00 AM, 8:00 AM, but NOT 9:00 AM.

Backup Throttle Schedule

Backup throttle schedule lets you control the backup throughput speed, which 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 will be 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 will adjust 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 will be 500 MB/minute and from 8:00 PM to 10:00 PM the throttle limit will be 2500 MB/minute.

If you do not define any backup schedule and backup throughput schedule, the backup will run 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 being 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 will process these sets one by one.
- If a merge job is resumed after a pause, the job detects at which point it was 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
Truncate log	☐	Truncate SQL Server log	Weekly
	☐	Truncate Exchange Server Log	Weekly
Run Commands	☐	Before a backup is started	
	☐	On exit code	0 ☒ Run Job ☐ Fail Job
	☐	After a snapshot is taken	
	☐	After a backup is over	
Username for Commands			
Password for Commands			

Enable Email Alerts

**Email Settings**

Job Alerts



Missed jobs



Backup, Catalog, Replication, File Copy, Restore or Copy Recovery Point job failed/crashed/canceled



Backup, Catalog, Replication, File Copy, Restore or Copy Recovery Point job successfully completed



Merge job stopped, skipped, failed or crashed



Merge job success

Enable Resource Alerts

**CPU Usage**Alert Threshold: %**Memory Usage**Alert Threshold: %**Disk Throughput**Alert Threshold: MB/s**Network I/O**Alert Threshold: %**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 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 have to manually run another backup job to start the copy recovery points 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 **Copy Recovery Points**.

The Copy Recovery Points task is added. You do not have to configure the **Source** tab in the Copy Recovery Points task because it displays the backup destination from the Backup task.

3. Click the **Copy Settings** tab and enter the details.

Source	Copy Settings	Schedule
Destination		
Specify the number of Copy Recovery Points to keep.	1	
Compression	Standard Compression	
Encryption Algorithm	No Encryption	
Encryption Password		
Confirm Encryption Password		

Destination

Specifies the destination where you want to keep the copy recovery points. This destination can only be a shared folder. It cannot be a recovery point server.

Specify the number of Copy Recovery Points to keep

Specifies the number of recovery points copy that you want to retain. When the number exceeds, the oldest recovery point is merged until the specified number of recovery points copy is remained.

Default: 1

Maximum: 1344

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 uncompressible 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 uncompressible 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.

4. Click the **Schedule** tab and specify the job schedule.

The Copy Recovery Points job runs after the specified number of successful backups is performed. For example, you have specified that the Copy Recovery Points job will run after five successful backups. If you have four successful backups and one unsuccessful backup, the copy recovery Points job will not start. It will wait until there are five successful backups. After the five successful backups, the recovery point from the fifth backup is copied to the specified destination.

5. 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.

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 UDP, you can copy or move selected source files to a destination (disk or shared folder). 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. After you copy the file, you want to delete that file from the source node. You can create a file copy plan 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. File Copy 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 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.

The following diagram illustrates the process to create a file copy:



What To Do Next?

- [Review the Prerequisites](#) (see page 367)
- [Create a Plan with a Backup Task](#) (see page 368)
- [Add a File Copy Task to the Plan](#) (see page 382)
- [Verify the Plan](#) (see page 392)

Review the Prerequisites and Considerations

Verify if you have completed the following prerequisites:

- Log in to the Console.
- Review the [Compatibility Matrix](#) which 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 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.

Note: When you resume the plan, the file copy job is not resumed automatically. You have to run another backup job to trigger a file copy job.

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: After a plan is paused, all jobs are paused except the restore job and the copy job. The running jobs are not affected. If you pause a plan that has pending jobs, then those pending jobs will also pause. When you resume the plan, the pending jobs does not resume immediately. After you resume the plan, the pending jobs will run from the next scheduled time. You can find the schedule of the next job form the home page of Arcserve UDP Agent (Windows).

6. From the **Task Type** drop-down list, select **Backup, Agent-Based Windows**.

Add a Plan

New Plan ☐ Pause this plan

Task1: Backup: Agent-Based Windows

Task Type: Backup: Agent-Based Windows

Source Destination Schedule Advanced

☒ Node Name	VM Name

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. Click the **Source** tab and click **Add Node**.
2. Select one of the following options:

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.


Adding Windows Nodes

Opens the **Add Nodes to Arcserve UDP Console** dialog. Select this option if you have not added the nodes and you want to manually add the nodes to protect.

Discovering 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. (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** area.

4. Select the nodes from the **Available Nodes** area and click the **Add all nodes (>>)** or **Add selected nodes (>)** icon.

The selected nodes are displayed on the **Selected Nodes** area.

5. Click **OK** to close the dialog.
6. To choose **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](#) (see page 711).
 - c. Optionally, provide an encryption password.

- d. Confirm the encryption password.
- e. Select a type of compression. For more information, see [Compression Type](#) (see page 709).

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](#) (see page 182).

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.

The screenshot shows the 'New Backup Schedule' dialog box. At the top, there is a dropdown menu set to 'Custom'. Below this, the 'Backup Type' is set to 'Incremental'. The 'Start Time' is '8:00 AM'. A section for selecting days of the week has checkboxes for Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, and Saturday, all of which are currently unchecked. The 'Repeat' checkbox is checked. Below the 'Repeat' checkbox, there is a section for frequency: 'Every 3 Hours' and 'Until 6:00 PM'. At the bottom, there are three buttons: 'Help', 'Save', and 'Cancel'.

- 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 checkbox 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 it is displayed on the **Schedule** page.

Source Destination Schedule Advanced

+ Add ▼ Delete		Type	Description	Su	Mo	Tu	We	Th	Fr	Sa	Time
☒			<u>Custom Incremental Backups Every 3 Hours</u>	✓	✓	✓	✓	✓	✓	✓	8:00 AM - 6:00 PM
☐			<u>Weekly Incremental Backup</u>						✓		8:00 PM

Add Merge Schedule

- Click **Add** and select **Add Merge Schedule**.

The **Add New Merge Schedule** dialog opens.

- Specify the start time to start the merge job.
- Specify **Until** to specify an end time for the merge job.
- Click **Save**.

The Merge Schedule is specified and it is displayed on the **Schedule** page.

Add Throttle Schedule

- Click **Add** and select **Add Throttle Schedule**.

The Add New Throttle Schedule dialog opens.

- Specify the throughput limit in MB per minutes unit.
- Specify the start time to start the backup throughput job.
- Specify **Until** to specify an end time for the throughput job.
- Click **Save**.

The Throttle Schedule is specified and it is displayed on the **Schedule** page.

- Specify the start time for the scheduled backup.

Start time for scheduled backup 10/24/2014  1 : 27 PM

Recovery Point Retention

7	Daily Backups
	Weekly Backups
	Monthly Backups
31	Custom / Manual Backups

- 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.

- Specify the catalog details.

Catalogs

Generate file system catalogs (for faster search) after

☐ Daily Backups

☐ Weekly Backups

☐ Monthly Backups

☐ Custom / Manual Backups

Generate Exchange catalogs (for granular restore) after

☐ All Backups for Nodes with Exchange Installed

Catalogs let you generate the Exchange Granular Restore catalog and the File System catalog. The Exchange Granular Restore catalog is required to restore Exchange mailboxes, mailbox folders, and individual mail objects. The File System catalog is required to perform faster and easier 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: Retention settings must be set 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 have to 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 will run at 6:00 AM, 7:00 AM, 8:00 AM, but NOT 9:00 AM.

Backup Throttle Schedule

Backup throttle schedule lets you control the backup throughput speed, which 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 will be 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 will adjust 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 will be 500 MB/minute and from 8:00 PM to 10:00 PM the throttle limit will be 2500 MB/minute.

If you do not define any backup schedule and backup throughput schedule, the backup will run 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 being 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 will process these sets one by one.
- If a merge job is resumed after a pause, the job detects at which point it was 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**

Truncate log

☐ Truncate SQL Server log

Weekly

☐ Truncate Exchange Server Log

Weekly

Run Commands

☐ Before a backup is started

☐ On exit code

0

☒ Run Job ☐ Fail Job

☐ After a snapshot is taken

☐ After a backup is over

Username for Commands

Password for Commands

Enable Email Alerts



Email Settings

Job Alerts



Missed jobs



Backup, Catalog, Replication, File Copy, Restore or Copy Recovery Point job failed/crashed/canceled



Backup, Catalog, Replication, File Copy, Restore or Copy Recovery Point job successfully completed



Merge job stopped, skipped, failed or crashed



Merge job success

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.

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. You can either retain the original copy or delete the original copy after you have copied the files to the specified destination.

Note: You can create a File Copy task if you have selected the following two options in the backup task:

- The backup destination is a local disk or a shared folder. It should not be Arcserve UDP Recovery Point Server.
- The **Generate File System catalog for faster search after each backup** check box is selected.

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 click **Add** from the following options:

Add

Lets you add a plan and select the plan type. When you click **Add**, the **Plan type** dialog opens.

Remove

Lets you remove the selected source and plan type.

Modify

Lets you modify the selected source and plan type.

Note: Arcserve UDP does not copy application files, files with system attributes, and files with temporary attributes.

Note: File copy does not support mounted volumes as the source. If you attempt to select a mounted volume as the source, then no files will be copied.

Note: If a symbolic link is selected when specifying the File Copy source folder, it is replaced with the actual path it points to when saving the settings. When you restore file copies, the actual path is displayed instead of the symbolic link.

4. Select one of the following **Plan types**:

File Copy

Data is copied from the source to the destination (remains on source location) and provides multiple stored versions on the destination.

File Copy - Delete Source

Data is moved from the source to the destination (deleted from source location) and provides more available free space at your source.

When you select **File Copy - Delete Source**, a warning message is immediately displayed alerting you that your specified file copied data will be moved (deleted) from and no longer available at the original source location. Click **OK** to continue to the **File Copy Policies** dialog.

For files copied using the **File Copy - Delete Source** option, Arcserve UDP leaves a stub file with the "D2DARC" extension. The stub file contains information about the destination where the files were moved. 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).

If you want to disable the stub file creation, access or create the registry key "HKLM\SOFTWARE\CA\arcserve Unified Data Protection\Engine\AfArchiveDll" and then create a DWORD named "CreateStubFile" with the value set to 0.

Note: If you disable or delete the file copy stub files, you can no longer track the status and location of moved files.

Depending upon the plan type that is selected, a different **File Copy Plan** dialog opens; however, the selections are similar.

5. Provide the source path and filter options in the **Plan** dialog.

Note: For more information on the **File Copy Plans** or **File Copy - Delete Source Plans** dialog, see [Specify File Copy Plan](#) (see page 386).

6. Click the **Destination** tab and specify the destination settings.

Specify the location to where the data will be Copied

Specifies the destination location for the file copy job. You can only select one destination.

Arcserve UDP lets you specify the settings for file copying your backed-up files to a disk or to the cloud. For file copying, you can specify to perform a copy and retain or a copy and move of your backed-up data. The two processes are similar, with the exception that when you perform a copy and move, the data is moved from the source to the destination (deleted from source location) and provides more available free space at your source. When you perform a copy and retain, the data is copied from the source to the destination (remains on source destination) and provides multiple stored versions.

File Copy to a local or network drive

When selected, lets you specify the full path of the location where you want to move or copy the source files/folders. The destination can be any local volume or folder or a file share accessible by any uniform naming convention (UNC) path. You can browse to this destination location. Clicking the green arrow icon lets you validate the connection to the specified destination.

File Copy to Cloud

When selected, lets you specify the cloud location where you want to move or copy the source files/folders. Arcserve UDP currently supports file copying to multiple cloud vendors, such as Amazon S3 (Simple Storage Service), Windows Azure, Fujitsu Cloud (Windows Azure), and Eucalyptus-Walrus. These cloud vendors are publicly available web services which let you safely and securely store and retrieve any amount of data, at any time, from anywhere on the Web.

You can click the Configure button to display the Cloud Configuration dialog. For more information, see [Specify Cloud Configuration for File Copy](#) (see page 390).

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.

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).

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:

No Compression

No compression is performed. This option has the lowest CPU usage (fastest speed), but also has the largest storage space requirement for your file copy.

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.

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.

Retention Time

This setting only applies to file copied data that is moved (not file copy data that is retained).

Specifies the amount of time (years, months, weeks, 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! The retention time setting only applies to data that has been copied and moved from the source to the destination (and not copied and retained). 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 Versions

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.

7. Click the **Schedule** tab and specify the number of backup.

The File Copy job runs after the specified number of backup is fulfilled.

8. Click **Save**.

The changes are saved and the file copy task is automatically deployed to the node.

Specify File Copy Plan

Depending upon the plan type that is selected, a different File Copy Plan dialog opens; however, the selections are similar.

File Copy Selected:

File Copy Source

Each File Copy plan has a source folder and optional file/folder filters. The file/folder filters determine what information will be copied. A file will be copied to the destination if it satisfies at least one plan.

Source Filters

Source filters enable you to specify and limit what is being copied. These filters are only applied to the corresponding source that is specified.

Include ▾	File Pattern ▾	

Type	Variable	Value

You can use wildcard characters "*" and "?" in File/Folder Patterns

File Copy - Delete Source Selected:

File Copy - Delete Source

Each File Copy plan has a source folder and optional file/folder filters. The file/folder filters determine what information will be copied. A file will be copied to the destination if it satisfies at least one plan.

Source Filters

Source filters enable you to specify and limit what is being copied. These filters are only applied to the corresponding source that is specified.

Type	Variable	Value

You can use wildcard characters "*" and "?" in File/Folder Patterns

File Size Filter

The file size filter lets you specify and limit the source data to be Copied based on the size of the file.

☐ Filter by File Sizes

		MB v

File Age Filter

The file age filters let you specify and limit the source data to be Copied based on the age of the file.

☐ Files not accessed in		month(s) v
☐ Files not modified in		month(s) v
☐ Files not created in		month(s) v


File Copy - Source Selection

Lets you specify the file copy source. You can specify or browse to a source volume or folder.

Source Filters

Filters let you limit the objects to be file copied by certain specified types and values.

For more information about these filters, see [How File Copy Source Filters Work](#).



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 Agent (Windows) 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 Copy - Delete Source Jobs Only)

This filter only applies to File Copy - Delete Source jobs (not File Copy jobs).

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 Agent (Windows) 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 Copy - Delete Source Jobs Only)

This filter only applies to File Copy - Delete Source jobs (not File Copy jobs).


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 D2D 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.

Specify Cloud Configuration for File Copy

From the **File Copy Settings Destination** dialog, you can click the **Configure** button to display the **Cloud Configuration** dialog.



The **Cloud Configuration** dialog box is used to configure cloud storage settings for file copy jobs. It includes a note about performance, sections for Vendor Type, Connection Settings, and Advanced settings, and buttons for testing the connection and confirming the configuration.

Cloud Configuration [X]

Note: File Copy jobs to/from cloud locations are generally slower than File Copy jobs to/from disks or network shares.

Vendor Type: Amazon S3

Connection Settings

Vendor URL: s3.amazonaws.com

Access Key ID: [Text Field]

Secret Access Key: [Text Field]

☐ Enable Proxy

Advanced

Bucket Name: [Text Field]

Note: Bucket name will be prefixed with 'arcserve-[agent hostname]'

Bucket Region: [Dropdown Menu]

☐ Enable Reduced Redundancy Storage

Test Connection OK Cancel Help

From this dialog you can use the drop-down menu to select which cloud vendor type you want to use for storage of your file copies. The available options are Amazon S3, Windows Azure, Fujitsu Cloud (Windows Azure), and Eucalyptus-Walrus. (Amazon S3 is the default vendor). For more information about Fujitsu Cloud (Windows Azure), see the [Overview](#) and [Registration](#).

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. Specify the **Connection Settings**:

Vendor URL

Identifies the URL address of the cloud provider.

(For Amazon S3, Windows Azure, and Fujitsu Cloud (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 Cloud (Windows Azure) 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 Cloud (Windows Azure), 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 (Domain Name\Username and Password) that is required to use the proxy server.

(Proxy capability is not available for Eucalyptus-Walrus).

2. 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.

(For this field, Amazon S3 and Eucalyptus-Walrus use Bucket Name. Windows Azure and Fujitsu Cloud (Windows Azure) use Container).

Note: For the remainder of this step, all references to Buckets can also be applied to Containers unless specified.

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).

3. Click **Test Connection** to verify the connection to the specified cloud location.
4. Click **OK** to exit the **Cloud Configuration** dialog.

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. 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.

Chapter 7: Restoring Protected Data

This section contains the following topics:


- [How to Restore From a Recovery Point](#) (see page 395)
- [How to Restore From a File Copy](#) (see page 409)
- [How to Restore Files/Folders](#) (see page 424)
- [How to Restore a Virtual Machine](#) (see page 443)
- [How to Restore Microsoft Exchange Mail](#) (see page 464)
- [How to Restore a Microsoft Exchange Application](#) (see page 481)
- [How to Restore Exchange Mail on a VMware Virtual Machine](#) (see page 492)
- [How to Restore a Microsoft SQL Server Application](#) (see page 511)
- [How to Restore a Pervasive PSQL Database](#) (see page 520)
- [How to Restore an Oracle Database](#) (see page 534)
- [How to Perform a File-Level Recovery on Linux Nodes](#) (see page 546)
- [How to Perform a Bare Metal Recovery \(BMR\) for Linux Machines](#) (see page 563)
- [How to Perform a BMR Using a Backup](#) (see page 584)
- [How to Perform a BMR Using a Virtual Standby VM](#) (see page 614)
- [How to Restore Microsoft Clustered Nodes and Shared Disks](#) (see page 648)

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:

How to Restore from a Recovery Point



Perform the following tasks to restore from a recovery point:

1. [Review the Restore Prerequisites and Considerations](#) (see page 397)
2. [Specify the Recovery Point Information to Restore](#) (see page 401)
 - a. [Specify the Recovery Point and Content to Restore](#) (see page 402)
 - b. [Define the Restore Options](#) (see page 405)
3. [Restore the Recovery Point Content](#) (see page 408)
4. [Verify that Content was Restored](#) (see page 409)

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](#) which 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 you another job is running and requests you 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.
- (Optional) Understand how the restore process works. For more information, see [How File Level Restores Work](#) (see page 398).
- (Optional) Review the files skipped during restore. For more information, see [Files Skipped During Restore](#) (see page 399).
- 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 a session:
"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	Remark
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 location.

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	Performance data used by the Windows performance counter.
		perf?00?.bak	
		CATROOT*	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%\inetrv\	metabase.bin	Metabase binary file of earlier IIS versions before 6.0.
XP W2003	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 http://msdn.microsoft.com/en-us/library/windows/desktop/bb891959(v=vs.85).aspx#filesnottobackup
	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\RegBackup*	Backup of current registry table.
Win8 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. Which files or directories were skipped can be found in: C:\Program Files\CA\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](#) (see page 402)
2. [Define the Restore Options](#) (see page 405)

Specify the Recovery Point 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:

- 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** dropdown 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. Click the **Browse Recovery Points** option.

The **Browse Recovery Points** dialog opens. You can see the **Recovery Point Server** details in the **Backup Location**.



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

- 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 **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: A clock icon with a lock symbol indicates the recovery point contains encrypted information and may require a password for restore.

7. 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.

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.

Restore Options

Destination
Select the restore destination

☒ Restore to original location

☐ Restore to

Resolving Conflicts
Specify how to resolve conflicts

☐ Overwrite existing files

☐ Replace active files

☐ Rename files

☒ Skip existing files

Directory Structure
Whether to create root directory during restore

☐ Create root directory

Backup Encryption or Protection Password
The data that you are attempting to restore is encrypted or password protected. Specify the password that is required to restore the data.

Password

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 Username 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:\Folder1\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:\Folder1\SubFolder2\A.txt", "C:\Folder1\SubFolder2\B.txt", and also E:\Folder3\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. 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: A clock icon with a lock symbol indicates the recovery point contains encrypted information and may require a password for restore.

5. Click **Next**.

The **Restore Summary** dialog 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** 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.

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 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:

How to Restore from a File Copy



Perform the following tasks to restore from a File Copy:

1. [Review the Restore Prerequisites and Considerations](#) (see page 411)
2. [Specify the File Copy Information to Restore](#) (see page 412)
 - a. [Specify the File Copy and Content to Restore](#) (see page 412)
 - [Specify Cloud Configuration for Restore](#) (see page 416)
 - b. [Define the Restore Options](#) (see page 418)
3. [Restore the Recovery Point Content](#) (see page 421)
4. [Verify that Content was Restored](#) (see page 422)

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](#) which 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 you another job is running and requests you 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 online help.

- (Optional) Understand how the restore process works. For more information, see [How File Level Restores Work](#) (see page 411).

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:

1. [Specify the File Copy and Content to Restore](#) (see page 412)
2. [Define the Restore Options](#) (see page 418)

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 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** dropdown 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. 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 **Change** to browse to an alternate location where your file copy images are stored.

A 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](#) (see page 416).

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 it is the first time you are performing a catalog synchronization, the **Browse Existing** button will be disabled because there is no existing file copy catalog 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 there were 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 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.

Cloud Configuration

Note: File Copy jobs to/from cloud locations are generally slower than File Copy jobs to/from disks or network shares.

Vendor Type: Amazon S3

Connection Settings

Vendor URL: s3.amazonaws.com

Access Key ID: <Access Key>

Secret Access Key:

☐ Enable Proxy

Proxy Server: <proxy server> Port: 80

☒ Proxy server requires authentication

Username: <domain name>\<user name>

Password:

Advanced

Bucket Name: [dropdown] Click 'Refresh' to load existing buckets

Bucket Region: [text box]

☐ Enable Reduced Redundancy Storage

Test Connection OK Cancel Help

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 Cloud (Windows Azure)**, and **Eucalyptus-Walrus**. (**Amazon S3** is the default vendor). For more information about Fujitsu Cloud (Windows Azure), 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, Windows Azure, and Fujitsu Cloud (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 Cloud (Windows Azure) 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 Cloud (Windows Azure), 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 Cloud (Windows Azure) use Container).

Bucket Region

For Amazon S3 only, the available region for the specified bucket is displayed in this field.

(For Windows Azure, Fujitsu Cloud (Windows Azure), and Eucalyptus-Walrus, the region is not displayed).

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).

3. Click **Test Connection** to verify the connection to the specified cloud location.
4. Click **OK** to exit the **Cloud Configuration** dialog.

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.

Restore Options

Destination
Select the restore destination

☒ Restore to original location

☐ Restore to

Resolving Conflicts
Specify how to resolve conflicts

☒ Overwrite existing files

☐ Replace active files

☐ Rename files

☒ Skip existing files

Directory Structure
Whether to create root directory during restore

☐ Create root directory

Encryption Password
The data that you are attempting to restore is encrypted or password protected. Specify the password that is required to restore the data.

Time	Name	Password
4/19/2014 2:31:30 AM	Customized Incremental Backup	Passed

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 Username 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:\Folder1\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:\Folder1\SubFolder2\A.txt", "C:\Folder1\SubFolder2\B.txt", and also E:\Folder3\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 encryption password manually.
 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 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](#) (see page 425)
2. [Specify the File/Folder Information to Restore](#) (see page 429)
 - a. [Specify the File/Folder Location](#) (see page 430)
 - [Specify Cloud Configuration for Restore](#) (see page 416)
 - b. [Specify the File/Folder to Restore](#) (see page 437)
 - c. [Define the Restore Options](#) (see page 438)
3. [Restore the File/Folder](#) (see page 441)
4. [Verify that the File/Folder was Restored](#) (see page 442)

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](#) which 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 you another job is running and requests you to try again later.
- 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](#) (see page 426).

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](#) (see page 399).

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.



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	Remark
All	Root folder of each volume	CAVolTrc.dat	Used by the Arcserve UDP tracking Driver.
		cavoltrcsnap shot.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 location.

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	Performance data used by the Windows performance counter.
		perf?00?.bak	
		CATROOT*	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%\inetrv\	metabase.bin	Metabase binary file of earlier IIS versions before 6.0.
XP W2003	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 http://msdn.microsoft.com/en-us/library/windows/desktop/bb891959(v=vs.85).aspx#filesnottobackup
		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\RegBackup*	Backup of current registry table.
Win8 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. Which files or directories were skipped can be found in: C:\Program Files\CA\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](#) (see page 430)
 - [Specify Cloud Configuration for Restore](#) (see page 416)
2. [Specify the File/Folder to Restore](#) (see page 437)
3. [Define the Restore Options](#) (see page 438)

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 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** dropdown 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. Click the **Find Files/Folders to Restore** option.

The **Find Files/Folders to Restore** dialog opens.

Restore

Find Files/Folders to Restore

The locations from where you search

☒ File Copy Location
 <file copy location> Change

☒ Backup Location
 <backup location> Browse

☐ Search all recovery points
☒ Select recovery points to search

Select recovery points

Start Time: 11/10/12 End Time: 11/10/12 Filter

Time	Backup Type	Name	Catalog
☒ 11/10/2012 3:23:02 AM	Incremental Backup		Created
☒ 11/10/2012 3:08:02 AM	Incremental Backup		Created
☒ 11/10/2012 2:53:02 AM	Incremental Backup		Created
☒ 11/10/2012 2:38:02 AM	Incremental Backup		Created
☒ 11/10/2012 2:23:02 AM	Incremental Backup		Created
☒ 11/10/2012 2:08:02 AM	Incremental Backup		Created

Page 1 of 1 Displaying 1 - 14 of 14

Previous Next Cancel Help

3. Select **File Copy Location** checkbox and click **Change** to change 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.

Destination


☒ Local or network drive ☐ Cloud

<destination> Browse

OK Cancel

- If you select **Local or network drive**, you can either specify a location or browse to the location where your file copy images are stored.
- You can click green arrow validate icon to verify proper access to the source location.
- If you select **Cloud**, you can 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](#) (see page 416).

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 it is the first time you are performing a catalog synchronization, the **Browse Existing** button will be disabled because there is no existing file copy catalog 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 there were 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. Select the **Backup Location** checkbox and click **Change** to change the Backup Location.

The **Source** dialog opens where you can select the backup location.

Source

☐ Select local disk or shared folder
☒ Select Recovery Point Server

Recovery Point Server setting

Hostname:
 Username:
 Password:
 Port:
 Protocol: ☐ Http ☒ Https
 Data Store:

Node	User Name	Dest Plan Name
abc-1-2012		

Page 1 of 1 | Displaying 1 - 1 of 1

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 selected recovery points are encrypted or password protected. As a result, you must provide the proper encryption password or session password.

Time ▾	Name	Password
9/28/2013 7:45:08 PM		

< >

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.

Cloud Configuration

Note: File Copy jobs to/from cloud locations are generally slower than File Copy jobs to/from disks or network shares.

Vendor Type: Amazon S3

Connection Settings

Vendor URL: s3.amazonaws.com

Access Key ID: <Access Key>

Secret Access Key:

☐ **Enable Proxy**

Proxy Server: <proxy server> Port: 80

☒ Proxy server requires authentication

Username: <domain name>\<user name>

Password:

Advanced

Bucket Name: [dropdown] [refresh icon]
Click 'Refresh' to load existing buckets

Bucket Region: [text field]

☐ Enable Reduced Redundancy Storage

Test Connection OK Cancel Help

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 Cloud (Windows Azure)**, and **Eucalyptus-Walrus**. (**Amazon S3** is the default vendor). For more information about Fujitsu Cloud (Windows Azure), 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, Windows Azure, and Fujitsu Cloud (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 Cloud (Windows Azure) 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 Cloud (Windows Azure), 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 Cloud (Windows Azure) use Container).

Bucket Region

For Amazon S3 only, the available region for the specified bucket is displayed in this field.

(For Windows Azure, Fujitsu Cloud (Windows Azure), and Eucalyptus-Walrus, the region is not displayed).

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).

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.

Restore Options

Destination
Select the restore destination

☒ Restore to original location

☐ Restore to 

Resolving Conflicts
Specify how to resolve conflicts

☐ Overwrite existing files

☐ Replace active files

☐ Rename files

☒ Skip existing files

Directory Structure
Whether to create root directory during restore

☐ Create root directory

Encryption Password
The data that you are attempting to restore is encrypted or password protected. Specify the password that is required to restore the data.

Time	Name	Password
4/19/2014 2:31:30 AM	Customized Incremental Backup	Passed

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 Username 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:\Folder1\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:\Folder1\SubFolder2\A.txt", "C:\Folder1\SubFolder2\B.txt", and also E:\Folder3\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 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:

How to Restore a Virtual Machine



Perform the following tasks to restore a virtual machine:

1. [Review the Restore Prerequisites and Considerations](#) (see page 445)
2. [Specify the Virtual Machine Information to Restore](#) (see page 446)
 - a. [Specify the Virtual Machine and the Recovery Point to Restore](#) (see page 446)
 - b. [Define the Restore Options](#) (see page 449)
 - [Define the Original Location Restore Options](#) (see page 450)
 - [Define the Alternate Location Restore Options](#) (see page 453)
3. [Restore the Virtual Machine](#) (see page 460)
4. [Verify that the Virtual Machine was Restored](#) (see page 464)

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](#) which provides the supported operating systems, databases, and browsers.

Review the following restore considerations:

- Arcserve UDP Agent (Windows) only allows one restore job to run at the same time, while Arcserve UDP allows multiple restore jobs 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 another job is running and requests you to try again later.
- 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).
- You cannot restore a VM to an ESXi 5.5 machine if an x86 OS machine is used as a proxy server. In this case, you will get an error message on the recovery dialog. The reason is that, VMware VDDK 5.5.x is required to interact with ESXi 5.5, but, if the proxy server is an x86 OS, VDDK 5.1.2 is used by Arcserve UDP Agent (Windows) because VDDK 5.5.x is not supported on an x86 OS.

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](#) (see page 446)
2. [Define the Restore Options](#) (see page 449)
 - [Define the Original Location Restore Options](#) (see page 450)
 - [Define the Alternate Location Restore Options](#) (see page 453)

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 a 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 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** dropdown 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.

- Click the **Recover VM** option.

The **Recover VM** dialog opens.

Restore

Recover VM

Backup Location

Recovery Point Server: <recovery_point_server_name> **Change**

Data Store: Data Store 1

Node: <virtual_machine_name>

Node

Select Node: <virtual_machine_name>

Recovery Point Date

October 2014

S	M	T	W	T	F	S
28	29	30	1	2	3	4
5	6	7	8	9	10	11
12	13	14	15	16	17	18
19	20	21	22	23	24	25
26	27	28	29	30	31	1
2	3	4	5	6	7	8

Time Range

12:00 AM - 6:00 AM (1)

6:00 AM - 12:00 PM

12:00 PM - 6:00 PM

6:00 PM - 12:00 AM

Time	Schedule Type	Backup Type	Name	Catalog Status
4:52:19 AM	Custom / Manual	Full	automat...	Disabled

Name	Date Modified	Size
C:		49.66 GB
E:		1000.00 MB
F:		1000.00 MB
G:		500.00 MB
K:		545.00 MB
Volume{73ae4dec-297d-11e4-80be-8}		350.00 MB

Previous **Next** **Cancel** **Help**

- Click **Change** to change the Backup Location.

The **Source** dialog opens. You can select the backup location in this dialog.

Source

☐ Select local disk or shared folder

☒ Select Recovery Point Server

Recovery Point Server setting

Hostname: xxx.xxx.xxx.xxx

Username: Administrator

Password:

Port: 8014

Protocol: ☒ Http ☐ Https

Data Store: New Data Store **Refresh**

Node

abc-1-2012

Page 1 of 1 | Displaying

OK **Cancel**

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

Restore

Restore Options

Destination
Select the restore destination

☒ Restore to original location

☐ Restore to an alternative location

Resolve Conflicts
Specify how to resolve conflicts

☐ Overwrite existing Virtual Machine

Post Recovery

☐ Power on Virtual Machine

Backup Encryption or Protection Password
The data that you are attempting to restore is encrypted or password protected. Specify the password that is required to restore the data.

Password

Previous Next Cancel Help

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](#) (see page 450).

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](#) (see page 453).

2. Specify the **Resolving Conflicts** options that Arcserve UDP will perform if conflicts are encountered during the restore process

The available option is whether to overwrite the existing virtual machine. By default, this overwrite option is not selected.

- 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, and if you restore to the original location, the VM recovery job will fail if the VM still exists on the original location; and if you restore to alternative location, the restore process creates a separate image of this virtual machine and the restore process does not overwrite any existing images that are at the specified restore destination.

3. Specify the **Post Recovery** option.

Select whether power is applied to the virtual machine at the end of the restore process. By default, this option is not selected.

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.



The dialog box titled "Set Credential for Source vCenter/ESX Server" contains the following fields and controls:

- vCenter/ESX Server:** Text box with placeholder "<Server Hostname/IP Address>".
- VM Name:** Text box with placeholder "<Virtual Machine Name>".
- Protocol:** Radio buttons for **HTTP** and **HTTPS**, with **HTTPS** selected.
- Port Number:** Text box with placeholder "<Port Number>".
- Username:** Text box with placeholder "<User Name>".
- Password:** Password field with masked characters "••••••••••".
- Buttons:** **OK** and **Cancel** buttons at the bottom right.

- For Hyper-V the **Set the credentials for the source Hyper-V Server** dialog is displayed.



The dialog box titled "Set the credentials for the source Hyper-V Server" contains the following fields and controls:

- Hyper-V/Hyper-V Cluster Server:** Text box with placeholder "<Server Hostname/IP Address>".
- VM Name:** Text box with placeholder "<Virtual Machine Name>".
- Username:** Text box with placeholder "<User Name>".
- Password:** Text box with placeholder "<Password>".
- Buttons:** **OK** and **Cancel** buttons at the bottom.

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.

Restore

Restore Options

Destination
Select the restore destination

☐ Restore to original location

☒ Restore to an alternative location

vCenter/ESX Server Information

vCenter/ESX Server: <vCenter/ESX Server name> Protocol: ☐ HTTP ☒ HTTPS

User Name: administrator Port Number: 123

Password: ***** **Connect to this vCenter/ESX Server**

VM Settings

VM Name: <VM_Name>

ESX Server: <Server>

Resource Pool: <Resource Pool> **Browse Resource Pools**

VM Datastore: datastore1_230(27.24 GB Fre)

Disk Datastore: Total Source Disk Size: 55.95 GB

Source Disk	Size	Source Volumes	Virtual Disk Type	Target Datastore
Disk0	50.00 GB	V? \Volume{73ae... 297d-11e4-80be-806e6f6e696...	Thin	datastore1_230(27.24 GB)
Disk1	2.00 GB	E-15161814	Thin	datastore1_230(27.24 GB)

Network:

Label	Configuration
Network adapter 1	VM Network

Resolve Conflicts
Specify how to resolve conflicts

☐ Overwrite existing Virtual Machine

Post Recovery

☐ Power on Virtual Machine

Backup Encryption or Protection Password
The data that you are attempting to restore is encrypted or password protected. Specify the password that is required to restore the data.

Password: *****

Previous Next Cancel Help

- For Hyper-V, the **Restore Options** dialog expands to display additional restore to alternative options.

Restore

Restore Options

Destination
Select the restore destination

☐ Restore to original location

☒ Restore to an alternative location

Hyper-V Server Information

Hyper-V/Hyper-V Cluster Server:

Username:

Password:

☐ Add virtual machine to the cluster.

VM Settings

VM Name:

VM Path:

☒ Specify the same virtual disk path for all virtual disks

Path:

Virtual Disk Type:

☐ Specify a virtual disk path for each virtual disk

Network:

Adapter	Connected To
Network Adapter 1	Intel(R) 82579LM Gigabit Network Coi

Resolve Conflicts
Specify how to resolve conflicts

☐ Overwrite existing Virtual Machine

Post Recovery

☐ Power on Virtual Machine

Backup Encryption or Protection Password
The data that you are attempting to restore is encrypted or password protected. Specify the password that is required to restore the data.

Password:

If you select the **Specify a virtual disk path for each virtual disk** option, the following dialog appears:

Restore

Username: Administrator

Password: **Connect**

☐ Add virtual machine to the cluster.

VM Settings

VM Name: <Virtual Machine Name>

VM Path: <Virtual Machine Path> **Browse**

☐ Specify the same virtual disk path for all virtual disks

☒ Specify a virtual disk path for each virtual disk

Source Disk	Size	Source Volumes	Virtual Disk Type	Path
Disk0	60.00 GB	W? Volume{3... e14d-11e3-93e8-806e6f6e6...	Fixed Size	D:\VMs\Virtual Hard Disks
Disk1	1.00 GB	J:\K\	Fixed Size(Quick)	D:\VMs\Virtual Hard Disks
Disk2	10.00 GB	E-1	Dynamically Expand	D:\VMs\Virtual Hard Disks

Network:

Previous Next Cancel Help

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. 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.

Protocol

Specifies the protocol that you want to use for communication with the destination server. The available selections are HTTP and HTTPS.

Default: HTTP.

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.



VM DataStore

Specify the destination VM DataStore for the virtual machine recovery or each virtual disk within the virtual machine.

A virtual machine can have multiple virtual disks and you can specify a different data store for each virtual disk.

For example:

- Disk0 can be restored to Datastore1.
- Disk1 can be restored to Datastore1.
- Disk2 can be restored to Datastore2.

Important! For VM DataStore, this field only populates if the user has full VMware system administrator permissions. If the user does not have proper administrator permissions, Arcserve UDP Agent (Windows) will not continue the restore process after you connect to the vCenter/ESX Server.

Disk Datastore

Specifies the data store (on ESX server) for each of the virtual disks of the VM respectively. The default data store of the VM disk file for the ESX server is shown by default. To assign the virtual disk type, you can select one of the following options: Thin, Thick Lazy Zeroed, or Thick Eager Zeroed.

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.

Follow these steps:

1. On the **Restore Summary** dialog, review the displayed information to verify that all the restore options and settings are correct.

Restore Summary for Original Location:

Restore Summary for Alternate Location (VMware):

Restore

Restore Summary

Verify your settings are correct and then click Finish to start the restore process.

Destination

Restore to an alternative location

vCenter/ESX Server Information

vCenter/ESX Server:

<Server IP Address>

Protocol:

<HTTP or HTTPS>

User Name:

<User Name>

Port Number:

<Port Number>

Password:

VM Name:

<VM Name>

VM Settings

ESX Server:

<ESX Server Name>

Resource Pool:

N/A

VM Datastore:

datastore1_230

Disk Datastore:

Source Disk	Size	Source Volumes	Virtual Disk Type	Target Datastore
Disk0	50.00 GB	\\? \Volume{73ae... 297d-11e4-	Thin	datastore1_230

Network:

Label	Configuration
Network adapter 1	VM Network

Resolve Conflicts

Do not overwrite existing Virtual Machine

Post Recovery

Do not power on Virtual Machine

Previous

Finish

Cancel

Help

Restore Summary for Alternate Location (Hyper-V):

Restore

Restore Summary

Verify your settings are correct and then click Finish to start the restore process.

Destination

Restore to an alternative location

Hyper-V Server Information

Hyper-V/Hyper-V/
Cluster Server: <Hyper-V Server Name>

User Name: <User Name>

Password: *****

Add virtual
machine to the
cluster: No

VM Settings

VM Name: <Virtual Machine Name>

VM Path: <Virtual Machine Path>

Virtual Disks:

Source Disk	Size	Source Volumes	Virtual Disk Type	Path
Disk0	60.00 GB	\\?.\Volume{350b... e14d-11e3-	Dynamically Expanding	D:\VMs\Virtual Hard Disks

Network:

Adapter	Connected To
Network Adapter 1	Intel(R) 82579LM Gigabit Network Connection - Virtual Switch

Resolve Conflicts

Do not overwrite existing Virtual Machine

Post Recovery

Do not power on Virtual Machine

Previous Finish Cancel Help

- 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 Restore Microsoft Exchange Mail

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. For Microsoft Exchange Mail you can then browse these recovery points to locate the individual objects (mailboxes, mailbox folders, or mail) that you want to restore.

The following diagram illustrates the process to restore Microsoft Exchange Mail:



Perform the following tasks to restore a virtual machine:

1. [Review the Restore Prerequisites and Considerations](#) (see page 466)
2. [Specify the Microsoft Exchange Mail Information to Restore](#) (see page 468)
 - a. [Specify the Microsoft Exchange Mailbox Database](#) (see page 469)
 - [Generate a Microsoft Exchange Granular Restore Catalog Now](#) (see page 472)
 - b. [Select the Level of Microsoft Exchange Objects to be Restored](#) (see page 476)
 - c. [Define the Restore Options](#) (see page 477)
3. [Restore the Microsoft Exchange Mail](#) (see page 479)
4. [Verify that the Microsoft Exchange Mail was Restored](#) (see page 480)

Review the Restore Prerequisites and Considerations

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

- If the restore is set to original location, the mailbox whose contents are to be restored should be available in Microsoft Exchange.
- If the restore is set to alternate location, the mailbox to which restore is targeted should be available in Microsoft Exchange.
- Sufficient space available in the target mailbox to restore the specified Exchange objects.
- To perform a Microsoft Exchange Granular Restore, the account specified in the Arcserve UDP Agent (Windows) user interface for restore must have sufficient restore permissions. The following permissions are required for the account:
 - The account must have a mailbox on the Exchange Server system.
 - The account must be a domain account.
 - The account must be a member of the Administrator group.
 - The account must be a member of the Backup Operators group.
 - There must be a mailbox associated with the account and the mailbox must be initialized.

For Microsoft Exchange Server 2007, Exchange Server 2010, and Exchange Server 2013 this mailbox must be in same Organization (Microsoft Exchange Organization) of the Exchange server to which you plan to restore to (restore destination).

- The name of the mailbox must be unique.

A unique name is a name that does not exist in the organization as a subset of characters in another mailbox name.

For example, if there is a mailbox named Administrator in your organization, you cannot use the name Admin.

- The account user must have the proper role assignments:
 - For Microsoft Exchange Server 2007 systems, the User Name must be a domain account with either the Microsoft Exchange Organization administrator role or the Exchange server administrator role.
 - For Microsoft Exchange Server 2010 systems, the User Name must be a domain account with Microsoft Exchange Organization management role.
 - For Microsoft Exchange Server 2013 systems, the User Name must be a domain account with Microsoft Exchange Organization management role.

The mailbox must be operational, unhidden, and initialized. If the mailbox has never received an email, then it is not initialized. To initialize, send an email to the mailbox.

- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Review the following restore considerations:

- The Arcserve UDP solution 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 another job is running and requests you to try again later.

Note: On Microsoft Exchange 2007 and later, Microsoft Exchange Server MAPI Client and Collaboration Data Objects 1.2.1 (and later) are not provided as part of the base product installation. Messaging API (MAPI) is a prerequisite for Microsoft Exchange Granular Restore. If MAPI is not installed on your Exchange server, mailbox or mail level granular restores may fail. To remedy any potential problem, Microsoft provides a download package that contains Microsoft Exchange MAPI and the latest version of Collaboration Data Objects (CDO). To download and install the latest version of this package, see the Microsoft Download Center.

- If you have two or more databases in the storage group and the option “Dismount the database before restore and mount the database after restore” is checked, then any attempt to restore one single database will force the dismount of all the other databases within that same storage group. This dismount of all databases will occur before the restore is launched and will then mount them after the restore is complete.
- To avoid potential security risks and possible restore failures through another agent server, the agent server must install the certificate on the Exchange Server. To install the certificate, see the steps about How to Install the Certificate on the Exchange Server to the agent machine.

Specify the Microsoft Exchange Mail Information to Restore

The Arcserve UDP solution provides granular mailbox recovery capabilities for Microsoft Exchange Server. You can browse and select the recovery points that were captured on the production Microsoft Exchange Server and recover specified messages, folders, and mailboxes that were in the Exchange database at the time of the selected recovery point.

The process involved in restoring Microsoft Exchange Mail is as follows:

1. [Specify the Microsoft Exchange Mailbox Database](#) (see page 469)
 - [Generate an Microsoft Exchange Granular Restore Catalog Now](#) (see page 472)
2. [Select the Level of Microsoft Exchange Objects to be Restored](#) (see page 476)
3. [Define the Restore Options](#) (see page 477)

Specify the Microsoft Exchange Mailbox Database

The Arcserve UDP solution lets you browse recovery points to locate individual objects within a Microsoft Exchange database to perform granular restore. You can either select restore to the original location or restore to a different location option. The Microsoft Exchange Server recovery can only be made using the Restore by Recovery Point method.

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** dropdown 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.

- Click the **Restore Exchange Mail** option.
The **Restore Exchange Mails** dialog opens.

Restore

Restore Exchange Mail

Backup Location

Recovery Point Server: xxx.xxx.xxx.xxx Change

Data Store: New Data Store

Node: abc-1-2012

Recovery Point Date

March 2014

S	M	T	W	T	F	S
23	24	25	26	27	28	1
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22
23	24	25	26	27	28	29
30	31	1	2	3	4	5

Today

Time Range

12:00 AM - 6:00 AM

6:00 AM - 12:00 PM

12:00 PM - 6:00 PM

6:00 PM - 12:00 AM (1)

Time	Schedule Type	Backup Type	Name
11:45:20 PM	Regular	Full	Customized Incremental Backup

Select a mailbox database, then click the Next button

Mailbox Database	Path	Catalog Status
ArchiveMDB	Microsoft Exchange Writer 2013	Not Created
Mailbox Database 1665	Microsoft Exchange Writer 2013	Not Created
MDB_01	Microsoft Exchange Writer 2013	Not Created

Previous Next Cancel Help

- Click **Change** to change the Backup Location.
The **Source** dialog opens. You can select the backup location in this dialog.

Source

☐ Select local disk or shared folder

☒ Select Recovery Point Server

Recovery Point Server setting

Hostname: xxx.xxx.xxx.xxx

Username: Administrator

Password:

Port: 8014

Protocol: ☐ Http ☒ Https

Data Store: DS Refresh

Node	User Name	Dest Plan Name
abc-1-2012		

OK Cancel

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 **Restore Exchange Mail** 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 **Restore Exchange Mail** 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 Microsoft Exchange mailbox databases 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. Specify a Microsoft Exchange mailbox database that you want to restore and click **Next**.

Note: If you did not enable the Exchange Granular Restore option during backup (no catalog generated), a notification message will be displayed asking you if you want to generate an Exchange Granular Restore catalog at this time. If you select No for generating a catalog now, you will not be able to browse to or select a granular recovery point. As a result you will only be able to perform full database restore from the **Browse Recovery Points Restore** dialog.

The **Restore Exchange Mails** dialog is updated to display a listing of the mailbox content for the selected database.

Note: Exchange Granular Restore supports email restores only. Calendar, Contacts, Notes and Tasks restores are not supported.



The Microsoft Exchange mailbox database is specified.

Generate a Microsoft Exchange Granular Restore Catalog Now

If you enabled the Exchange Granular Restore option at backup, the corresponding Exchange Granular Restore catalogs are generated during each backup. These catalogs help you to perform granular recoveries (at mailbox, mailbox folder, and or mail item levels) by letting you browse and select exactly which Exchange object you want to recover. However, if you disabled the Exchange Granular Restore option at backup, the Exchange mailboxes may not be displayed when you are trying to browse the mailbox database in a recovery point because the corresponding catalog was not generated.

Before you perform the restore, you can generate the catalog manually from the Restore Exchange Mails dialog.

Note: The cache file, which is used to record the data change when mounting a writable volume from the backup session, must be on a non-4k sector size disk.

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** dropdown 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. Click the **Restore Exchange Mails** option.

The **Restore Exchange Mail** dialog opens.

Restore

Restore Exchange Mail

Backup Location

Recovery Point Server: xxx.xxx.xxx.xxx Change

Data Store: New Data Store

Node: abc-1-2012

Recovery Point Date

March 2014

S	M	T	W	T	F	S
23	24	25	26	27	28	1
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22
23	24	25	26	27	28	29
30	31	1	2	3	4	5

Today

Time Range

- 12:00 AM - 6:00 AM
- 6:00 AM - 12:00 PM
- 12:00 PM - 6:00 PM
- 6:00 PM - 12:00 AM (1)

Time	Schedule Type	Backup Type	Name
11:45:20 PM	Regular	Full	Customized Incremental Backup

Select a mailbox database, then click the Next button

Mailbox Database	Path	Catalog Status
ArchiveMDB	Microsoft Exchange Writer 2013	Not Created
Mailbox Database 1665	Microsoft Exchange Writer 2013	Not Created
MDB_01	Microsoft Exchange Writer 2013	Not Created

Previous Next Cancel Help


3. Navigate to the Microsoft Exchange mailbox database that you want to restore:
 - a. Specify the backup location. You can either specify a location or browse to the location where your backup images are stored. If necessary, enter the User name and Password credentials to gain access to that location. You can click green arrow validate icon to verify proper access to the source location.

The calendar view will highlight (in green) all dates during the displayed time period that contain recovery points for that backup source.
 - b. Select the calendar date for the backup image you want to restore.

The corresponding Microsoft Exchange mailbox databases for that date are displayed, along with the time of the backup, the type of backup that was performed, and the name of the backup.

4. Select a Microsoft Exchange mailbox database that you want to restore and click **Next**.

The **Submit Exchange Granular Restore Catalog Job** dialog displays informing you that the Exchange Granular Restore catalog has not been generated for the selected database and asking you if you want to generate the catalog now.



5. Click **OK** to launch the process of generating an Exchange Granular Restore catalog.

The Exchange Granular Restore catalog for the selected database is generated. The process of generating a catalog could be time-consuming depending upon the size of the database.

During the catalog generating process, the Job Monitor displays information about the ongoing event, with the estimated time remaining to complete the catalog job.

Note: If you select **Cancel** for generating a catalog now, you cannot browse to or cannot select a granular recovery point. As a result you can only perform a full database restore.

6. When the catalog generating process is finished, click **Next** to continue the Exchange Granular Restore of the selected database.

The granular restore can now be continued.

Note: If you attempt to click **Next** while the catalog is still being generated, a pop-up message appears informing you of this condition.

The Microsoft Exchange Granular Restore catalog is generated.

Select the Level of Microsoft Exchange Objects to be Restored

After generating the Microsoft Exchange Granular Restore catalog, specify the level of Exchange objects to be restored.

Note: Arcserve UDP does not support granular recovery of Exchange public folder objects. You need to use Application Restore to recover the entire public folder database and then extract the specific Exchange object you need.

Follow these steps:

1. From the **Exchange Database** dialog, select the level of Exchange object(s) to be restored (mailbox, folder, or individual mail).

You can select the entire content or partial content of the Exchange object to restore. You can select multiple Exchange objects to restore.

Note: When using Arcserve UDP to restore individual mailbox/mail objects from the Exchange mailbox database, the operating system used for the restore must be same as used when it was backed up (including the same Windows Version number and Service Pack level and also the related version of the visual C++ redistributable package required to support it).

The available Microsoft Exchange objects are:

a. **Mailbox database**

If you select a mailbox database, all of the mailboxes in that database will be restored.

b. **Mailbox (or mailboxes)**

If you select a mailbox level, all corresponding content (folders and individual mail) within that mailbox will be restored.

c. **Folder(s)**

If you select the mailbox folder level, all corresponding mail content within that folder will be restored.

d. **Individual mail object(s)**

If you select the individual mail level, only the selected mail object(s) will be restored.

2. Click **Next**.

The **Restore Options** dialog Opens.

The Microsoft Exchange objects to be restored are specified.

Define the Restore Options

After you select the level of Microsoft Exchange objects, specify the destination for your backup. You can either select to restore to the original location or restore to a different location option.

Note: For Microsoft Exchange Server 2010 and Exchange Server 2013, archived mailbox items cannot be restored to the original location. Archived mailbox items can only be restored to an alternate location or to a local disk. In addition, regular mailbox items cannot be restored to archive mailboxes.

Follow these steps:

1. From the **Restore Options** dialog, select the restore destination.

Restore Options

Destination
Select the restore destination

☒ Restore to original location

User Name: e15\Administrator
Password:
CAS Name: [] Browse

☐ Dump email items [] Browse

How should arcserve UDP Agent resolve duplicate mails
☒ Rename ☐ Overwrite

☐ Restore to alternate location

Enter the account, then click the Browse button to choose the destination.

User Name: e15\Administrator
Password:
Destination: [] Browse
CAS Name: [] Browse

Backup Encryption or Protection Password
The data that you are attempting to restore is encrypted or password protected. Specify the password that is required to restore the data.
Password:

Notes:
For Microsoft Exchange Server 2010 or 2013 systems, the user name must be a domain account with an Exchange Organization Management role.

Previous Next Cancel Help

The available destination options are:

Restore to Original Location

Restores the mails to the original location from where the backup image was captured. Mails will retain the same hierarchy and be restored to its original mailbox and original folder.

- If current machine is not the active Microsoft Exchange server, Arcserve UDP Agent (Windows) will detect the location of the active server and then restore the mails to that active server.
- If mailbox has been moved to another Microsoft Exchange server, but still in the same organization, Arcserve UDP Agent (Windows) will detect the new Exchange server where the original mailbox resides, and then restore to that new server.
- If the display name of the mailbox is changed, any attempt to restore the mailbox (from an earlier backup session) to its original location will fail because Arcserve UDP Agent (Windows) will not be able to find the changed name. To solve this problem, you can specify to restore this mailbox to an alternate location.

Note: When restoring a mailbox or mail to the original location, make sure the destination mailbox is available, or else the restore will fail. Arcserve UDP Agent (Windows) only validates the destination when the restore job is submitted.

Note: For a granular restore catalog job in Microsoft Exchange 2013 only, enter the CAS (Client Access Server) Name. The CAS is a thin, stateless server that serves as a proxy for client connections to the Mailbox servers. To do this, from the Restore Options dialog, click **Browse**. From the Select the Client Access Server dialog, select one of the CAS items and click **OK**.

Dump File Only

Restores the mails to a disk. This disk location can be local or a remote machine. The restored mails will maintain the same hierarchy as they had in the corresponding Microsoft Exchange mailbox. The file name will become the subject of mail.

Note: If the mail subject, folder name, or mailbox name includes any of the following characters, the character will be replaced by hyphen (-) in file name: \ / : * ? " < > |

For this option, you also need to specify what you want Arcserve UDP Agent (Windows) to do to resolve a conflict situation. In Microsoft Exchange, you can have multiple mail objects with the same name under the same folder. However in a File System, two files with the same name cannot co-exist under the same folder.

There are two options to resolve this conflict situation:

Rename

If on the disk, there is a file with the same name as the mail subject, Arcserve UDP will still name the mail subject, but will append a number at the end of the mail subject.

Overwrite

If on the disk, there is a file with the same name as the mail subject, Arcserve UDP Agent (Windows) will overwrite that file.

Note: When you select individual mail objects to restore to the disk (dump), by default the format of the restored mail object will be an Outlook Message (.MSG) file and not a Personal Storage Table (.PST) file.

Restore to Alternate Location

Restores the mails to a specified location or lets you browse to the location where your backup images will be restored. The destination must be a mailbox in the same Microsoft Exchange organization, and a new folder name is required. (If you are attempting to restore mails to an alternate location, the destination cannot be a public folder).

Note: When restoring mail to an alternate location, if the specified destination folder already exists, the restore will continue. However, if the specified folder does not exist, then Arcserve UDP Agent (Windows) will create the folder first and then continue the restore.

- a. Specify the User Name and Password credentials, click the **Browse** button to navigate through a list of all Microsoft Exchange Servers, Storage Groups, Exchange Databases, and Mailboxes in the current organization.
- b. Select a mailbox as the destination.
- c. For a granular restore catalog job in Microsoft Exchange 2013 only, enter the **CAS (Client Access Server) Name**. The CAS is a thin, stateless server that serves as a proxy for client connections to the Mailbox servers. To do this, from the **Restore Options** dialog, click **Browse**. From the **Select the Client Access Server** dialog, select one of the CAS items and click **OK**.

2. Click **Next**.


The **Restore Summary** dialog opens.

Restore the Microsoft Exchange Mail

The **Restore Summary** dialog 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 Microsoft Exchange Mail is restored.

Verify that the Microsoft Exchange Mail was Restored

After the completion of the restore process, verify that the Microsoft Exchange Mail 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 Microsoft Exchange Mail to the restore destination as Original Location or Alternate Location, then login to the user's mailbox to check the restored email.

If you select to restore the Microsoft Exchange Mail to Dump email items, then navigate to the dump folder to check the restored email.

For example, if you select to restore the Microsoft Exchange Mail to C:\dump_folder1, then after restore navigate to this location to verify email

2. Verify emails of the restored Exchange Mail.


The Microsoft Exchange Mail is restored successfully.

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:

How to Restore a Microsoft Exchange Application



Perform the following tasks to restore a Microsoft Exchange Application:

1. [Review the Restore Prerequisites and Considerations](#) (see page 483)
2. [Specify the Microsoft Exchange Information to Restore](#) (see page 485)
 - a. [Specify the Recovery Point and Microsoft Exchange Database](#) (see page 486)
 - b. [Define the Restore Options](#) (see page 487)
3. [Restore the Microsoft Exchange Application](#) (see page 490)
4. [Verify that the Microsoft Exchange Application was Restored](#) (see page 491)

Review the Restore Prerequisites and Considerations

Arcserve UDP Agent (Windows) supports the following versions of Microsoft Exchange Server:

- Microsoft Exchange 2007 - Single Server Environment, Local Continuous Replication (LCR), and Cluster Continuous Replication (CCR) environment.

For Microsoft Exchange 2007 CCR environment, Arcserve UDP Agent (Windows) must be installed on both the active node and passive node of Microsoft Cluster. Backup can be performed from an active node and passive node, but restore can only be performed to an active node.

- Microsoft Exchange 2010 - Single Server Environment and Database Availability Group (DAG) environment.
- Microsoft Exchange 2013 - Single Server Environment and Database Availability Group (DAG) environment.

For Microsoft Exchange Server 2010 DAG and Exchange Server 2013 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.

Note: Microsoft Exchange Server 2007 Single Copy Cluster (SCC) environment is not supported by Arcserve UDP Agent (Windows)

Microsoft Exchange Server can be restored 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 and Microsoft Exchange Server 2013.

Mailbox Database Level (Microsoft Exchange 2007, 2010, and 2013)

Specifies if you want to restore a specific Mailbox Database, you can perform a restore at this level.

Mailbox Level (Microsoft Exchange 2007, 2010, and 2013)

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

- If the restore is set to original location, the mailbox whose contents are to be restored should be available in Microsoft Exchange.
- If the restore is set to alternate location, the mailbox to which restore is targeted should be available in Microsoft Exchange.
- Sufficient space available in the target mailbox to restore the specified Exchange objects.
- To perform a Microsoft Exchange Granular Restore, the account specified in the Arcserve UDP Agent (Windows) user interface for restore must have sufficient restore permissions. The following permissions are required for the account:
 - The account must have a mailbox on the Exchange Server system.
 - The account must be a domain account.
 - The account must be a member of the Administrator group.
 - The account must be a member of the Backup Operators group.
 - There must be a mailbox associated with the account and the mailbox must be initialized.

For Microsoft Exchange Server 2007, Exchange Server 2010, and Exchange Server 2013 this mailbox must be in same Organization (Microsoft Exchange Organization) of the Exchange server to which you plan to restore to (restore destination).

- The account user must have the proper role assignments:
 - For Microsoft Exchange Server 2007 systems, the User Name must be a domain account with either the Microsoft Exchange Organization administrator role or the Exchange server administrator role.
 - For Microsoft Exchange Server 2010 systems, the User Name must be a domain account with Microsoft Exchange Organization management role.
 - For Microsoft Exchange Server 2013 systems, the User Name must be a domain account with Microsoft Exchange Organization management role.
- The mailbox must be operational, unhidden, and initialized. If the mailbox has never received an email, then it is not initialized. To initialize, send an email to the mailbox.
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Review the following restore considerations:

- The Arcserve UDP solution 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 another job is running and requests you to try again later.

Note: On Microsoft Exchange 2007 and later, Microsoft Exchange Server MAPI Client and Collaboration Data Objects 1.2.1 (and later) are not provided as part of the base product installation. Messaging API (MAPI) is a prerequisite for Microsoft Exchange Granular Restore. If MAPI is not installed on your Exchange server, mailbox or mail level granular restores may fail. To remedy any potential problem, Microsoft provides a download package that contains Microsoft Exchange MAPI and the latest version of Collaboration Data Objects (CDO). To download and install the latest version of this package, see the [Microsoft Download Center](#).

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](#) (see page 486)
2. [Define the Restore Options](#) (see page 487)

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 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** dropdown 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. 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.

Restore

Browse Recovery Points Change

Backup Location

Recovery Point Server: <Recovery Point Server>
 Data Store: <Data Store>
 Node: <Node Name>

Recovery Point Date

April 2014

S	M	T	W	T	F	S
30	31	1	2	3	4	5
6	7	8	9	10	11	12
13	14	15	16	17	18	19
20	21	22	23	24	25	26
27	28	29	30	1	2	3
4	5	6	7	8	9	10

Today

Time Range

12:00 AM - 6:00 AM (1)
 6:00 AM - 12:00 PM
 12:00 PM - 6:00 PM
 6:00 PM - 12:00 AM (3)

Time	Type	Backup Type	Name
10:20:20 PM	Custom / Manual	Incremental	Customized Incremental Backup
10:00:04 PM	Daily	Incremental	
9:02:41 PM	Custom / Manual	Incremental	Customized Incremental Backup

Name	Date Modified	Size
C:		77.19 GB
System Reserved		241.66 MB
Microsoft Exchange Writer 2013		5.18 GB
Microsoft Exchange Server		
Microsoft Information Store		
<Recovery Point Server>		
Mailbox Database 072		
mydb1		

Previous Next Cancel Help

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.



Restore

Restore Options

Destination
Select the restore destination

☒ Restore to original location

☐ Dump file only **Browse**

☐ Replay log on database

☐ Restore to Recovery Database

Backup Encryption or Protection Password
The data that you are attempting to restore is encrypted or password protected. Specify the password that is required to restore the data.

Password

☒ Dismount the database before restore and mount the database after restore.

Previous **Next** **Cancel** **Help**

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.

Restore to Recovery Storage Group (Microsoft Exchange 2007)

Restores the database to a Recovery Storage Group (RSG).

An RSG is a storage group that can be used for recovery purposes. You can restore a Microsoft Exchange Mailbox Database from a backup in a Recovery Storage Group and then recover and extract data from it, without affecting the production database that is being accessed by end users.

- If single storage group or database (except a public folder database) from the same storage group are selected to restore, the default restore destination is "Restore to Recovery Storage Group" (or "Restore to Recovery Database").
- If multiple storage groups or databases from multiple storage groups are selected to restore, Microsoft Exchange can only be restored to the original location or restore with "Dump file only" option. The default restore destination is "Restore to original location".

Before restoring Microsoft Exchange 2007 database to a Recovery Storage Group, you must create a Recovery Storage Group and Mailbox Database with the same name.

For example, if you want to restore MailboxDatabase1 from the First Storage Group to a Recovery Storage Group, create a Recovery Storage Group and add the database "MailboxDatabase1" to the Recovery Storage Group.

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.

Note: This option is not applicable for Microsoft Exchange Server 2007.

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.

Restore

Restore Summary

Verify your settings are correct and then click Finish to start the restore process.

Component(s) to be restored

Name	Path
mydb1	Microsoft Exchange Writer\Microsoft Exchange Server\Microsoft Information Store\<server name>

Destination

Restore to original location

Restore Options

Dismount the database before restore and mount the database after restore.: Yes

Previous Finish Cancel Help

- 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 Mail on a VMware Virtual Machine

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 a 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. For Exchange Mail backup you can browse the recovery points and locate the individual objects (mailboxes, mailbox folders, or mail) that you want to restore.

This scenario describes how to restore VMware virtual machine (VM) Exchange Mail objects that were backed up using the Host-based agentless plan. The following diagram illustrates the process to restore Exchange Mail installed on a VMware virtual machine:

How to Restore Exchange Server Mail on a VMware Virtual Machine



Perform the following tasks to restore a virtual machine:

1. [Review the Restore Prerequisites and Considerations](#) (see page 495)
2. [Specify the Exchange Mailbox Database](#) (see page 497)
 - [Generate an Exchange Granular Restore Catalog Now](#) (see page 501)
 - [Assign a Separate Folder for Temporary Files](#) (see page 504)
3. [Select the Exchange Mail Objects to be Restored](#) (see page 504)
4. [Define the Restore Options](#) (see page 505)
5. [Restore the Exchange Mail](#) (see page 508)
6. [Verify that the Exchange Mail was Restored](#) (see page 510)

Review the Restore Prerequisites and Considerations

Verify that you have completed the following tasks before you restore an Exchange Mail:

- If the restore is set to original location, the proxy server is in the same domain with the Exchange VMs.
- If the restore is set to an alternate location, the user name to which restore is targeted is available in a different location.
- Sufficient space is available in the target mailbox to restore the specified Exchange objects.
- If the Exchange VM is using Refs/NTFS- Deduplication volume, the proxy server should support it.
- The Exchange VM is installed with VM tools.
- The proxy server must be a 64-bit machine, and installed with the latest Messaging API (MAPI).
- To perform an Exchange Granular Restore, the account specified in the Arcserve UDP user interface for restore has sufficient restore permissions. The following permissions are required for the account:
 - The account has a mailbox on the Exchange Server system.
 - The account is a domain account.
 - The account is a member of the Administrator group.
 - The account is a member of the Backup Operators group.
 - A mailbox is associated with the account and the mailbox is initialized.
 - The name of the mailbox is unique.

A unique name is a name that does not exist in the organization as a subset of characters in another mailbox name.

For example, if there is a mailbox named Administrator in your organization, you cannot use the name Admin.

The mailbox is operational, unhidden, and initialized. If the mailbox has never received an email, then it is not initialized. To initialize, send an email to the mailbox.
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Review the following restore considerations:

- The Arcserve UDP solution 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.
- The Host-based agentless Exchange Granular Restore supports only the following three Exchange Server versions in the virtual machine:
 - Exchange Server 2007
 - Exchange Server 2010
 - Exchange Server 2013
- The Host-based agentless Exchange Granular Restore supports email restores only. Calendar, Contacts, Notes and Tasks restores are not supported.
- The Host-based agentless Exchange Granular Restore supports only the virtual machines in the VMware Esx/vCenter Hypervisor.
- The Backup proxy server must install the certificate on the Exchange Server. The certificate helps to avoid potential security risks and possible restore failures through the proxy server. If the certificate is not installed, see [Failed to Restore Exchange Object \(Message, Folder, or Mailbox\) to the Original or Alternate Location](#) (see page 692).
- The Exchange catalog is not generated automatically in the virtual machine, unless you have enabled the **Generate Exchange catalogs (for Granular Restore)** option while adding schedule for the Host-based agentless backup plan.

Source Destination **Schedule** Advanced

Type	Description	Su	Mo	Tu	We	Th	Fr	Sa	Time
	Daily Incremental Backup	✓	✓	✓	✓	✓	✓	✓	10:00 PM

Start time for scheduled backup: 10/20/2014 2 : 15 PM

Recovery Point Retention

Daily Backups
 Weekly Backups
 Monthly Backups
 Custom / Manual Backups

Catalogs (For Windows virtual machines only)

Generate file system catalogs (for faster search) after

☐ Daily Backups
☐ Weekly Backups
☐ Monthly Backups
☐ Custom / Manual Backups

Generate Exchange catalogs (for granular restore) after

☐ All Backups for Nodes with Exchange Installed

The Backup Proxy Server must be on a supported 64-bit platform.

Specify the Exchange Mailbox Database

The Arcserve UDP solution lets you browse recovery points to locate individual objects within an Exchange database to perform granular restore. You can either select restore to the original location or restore to a different location. The Exchange Mail recovery can only be made using the Restore by Recovery Point method.

Follow these steps:

1. From Arcserve UDP:

- a. Log in to Arcserve UDP.
- b. Click the **resources** tab.
- c. From the left pane, select **All Nodes**.

All the added nodes are displayed on the center pane.

d. Perform one of the following steps to reach the Restore dialog:

- From the All Nodes page, right click the virtual machine node name, and select **Restore** from the options.
- From the center pane, select the virtual machine node using which you have created the host-based plan, click **Actions**, and from the **Actions** drop-down menu click **Restore**.
- From the All Nodes page, click the virtual machine node name and from the virtual machine node name page that is displayed click the Restore link.

The restore method selection dialog opens.

2. Click the **Restore Exchange Mail** option.

The **Restore Exchange Mail** dialog opens.

Restore

Restore Exchange Mail

Backup Location
e:\Dest\ XYZ123-ex13-1@100.13.15.293 Change

Recovery Point Date

October 2014

S	M	T	W	T	F	S
28	29	30	1	2	3	4
5	6	7	8	9	10	11
12	13	14	15	16	17	18
19	20	21	22	23	24	25
26	27	28	29	30	31	1
2	3	4	5	6	7	8

Today

Time Range

12:00 AM - 6:00 AM
6:00 AM - 12:00 PM
12:00 PM - 6:00 PM
6:00 PM - 12:00 AM (1)

Time	Schedule Type	Backup Type	Name
10:00:39 PM	Daily	Incremental	

Select a mailbox database, then click the Next button

Mailbox Database	Path	Catalog Status
Mailbox Database 0374:	Microsoft Exchange Writer 2013	Created
mdb1	Microsoft Exchange Writer 2013	Created

Previous Next Cancel Help

3. Click **Change** to modify the Backup location.

The **Source** dialog opens. You can select the backup location.

Source

☐ Select local disk or shared folder
☒ Select Recovery Point Server

Recovery Point Server setting

Hostname:
 Username:
 Password:
 Port:
 Protocol: ☐ Http ☒ Https
 Data Store:

Node	User Name	Dest Plan Name
abc-1-2012		

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.

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. You can view the backup location in the **Source** dialog.

- c. Click **OK**.

The recovery points are listed in the **Restore Exchange Mail** dialog.

Select Recovery Point Server

- a. Specify the **Recovery Point Server setting** details and click **Refresh**.

All the agents are listed in the right column under **Node**.

- b. Select the agent from the displayed list and click **OK**.

The recovery points are listed in the **Restore Exchange Mail** dialog.

- From the Recovery Points Date calendar, select the date for the backup image to restore.


All the dates containing recovery points for the specified backup source are highlighted in green.

The corresponding Exchange mailbox databases 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.

- Specify an Exchange mailbox database that you want to restore and click **Next**.

Note: The catalog status of a mailbox database is either Created or Not Created. You can proceed only if the status is Created. If the status of catalog shows Not Created, to use that Exchange mailbox you need to generate catalog now.

The **Restore Exchange Mail** dialog is updated to display a listing of the mailbox content for the selected database.



The Exchange mailbox database is specified.

Generate Exchange Granular Restore Catalog Now

Generating exchange catalog is necessary to restore the Exchange Mail. The catalog is generated automatically if you enabled the Generate Exchange catalogs (for Granular Restore) option while adding schedule for the Host-based agentless backup plan. The corresponding Exchange Granular Restore catalogs are generated during each backup. These catalogs help you to perform granular recoveries (at mailbox, mailbox folder, and or mail item levels) by letting you browse and select exactly which Exchange object you want to recover.

If you have not enabled the Generate Exchange catalogs (for Granular Restore) option when you created the plan, then you must generate now.

Depending upon your action when you add schedule to the Host-based agentless plan, the catalog status in the Mailbox Database displays the following two options:

- **Created:** This status appears for a Mailbox Database when you select the checkbox of the Generate Exchange catalogs and the catalog is generated successfully.
- **Not Created:** This status appears for a Mailbox Database when you do not select the checkbox of the Generate Exchange catalogs or the catalog is not generated successfully.

Before you perform the restore, you can generate the catalog manually from the Restore Exchange Mails dialog.

Follow these steps:

1. 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 virtual machine node that you have used for the host-based plan, and click **Actions**.
 - e. Click **Restore** from the **Actions** drop-down menu.

The restore method selection dialog opens.

2. Click the **Restore Exchange Mails** option.

The **Restore Exchange Mail** dialog opens.

Restore

Restore Exchange Mail

Backup Location

Recovery Point Server: xxx.xxx.xxx.xxx Change

Data Store: New Data Store

Node: abc-1-2012

Recovery Point Date

March 2014

S	M	T	W	T	F	S
23	24	25	26	27	28	1
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22
23	24	25	26	27	28	29
30	31	1	2	3	4	5

Today

Time Range

12:00 AM - 6:00 AM

6:00 AM - 12:00 PM

12:00 PM - 6:00 PM

6:00 PM - 12:00 AM (1)

Time	Schedule Type	Backup Type	Name
11:45:20 PM	Regular	Full	Customized Incremental Backup

Select a mailbox database, then click the Next button


Mailbox Database	Path	Catalog Status
ArchiveMDB	Microsoft Exchange Writer 2013	Not Created
Mailbox Database 1665	Microsoft Exchange Writer 2013	Not Created
MDB_01	Microsoft Exchange Writer 2013	Not Created

Previous Next Cancel Help

3. Navigate to the Exchange mailbox database that you want to restore:
 - a. Specify the backup location. You can either specify a location or browse to the location where your backup images are stored. If necessary, enter the User name and Password credentials to gain access to that location.
The calendar view will highlight (in green) all dates during the displayed time period that contain recovery points for that backup source.
 - b. Select the calendar date for the backup image you want to restore.
The corresponding Exchange mailbox databases for that date are displayed, along with the time of the backup, the type of backup that was performed, and the name of the backup.

4. Select an Exchange mailbox database that you want to restore and click **Next**.

The **Submit Exchange Granular Restore Catalog Job** dialog displays informing you that the Exchange Granular Restore catalog has not been generated for the selected database and asking you if you want to generate the catalog now.



5. Enter the encryption password if the recovery point is encrypted.
6. Click **OK** to launch the process of generating an Exchange Granular Restore catalog.

The Exchange Granular Restore catalog for the selected database is generated. The process of generating a catalog could be time-consuming depending upon the size of the database.

During the catalog generating process, the Job Monitor displays information about the ongoing event, with the estimated time remaining to complete the catalog job.

Note: If you select **Cancel** for generating a catalog now, you cannot browse to or cannot select a granular recovery point. As a result you can only perform a full database restore.

7. When the catalog generating process is finished, click **Next** to continue the Exchange Granular Restore of the selected database.

The granular restore can now be continued.

Note: If you attempt to click **Next** while the catalog is still being generated, a pop-up message appears informing you of this condition.

The Exchange Granular Restore catalog is generated.

Assign a Separate Folder for Temporary Files

When performing Exchange Mail object restore or Exchange catalog job, temporary files are generated in the Arcserve UDP installation volume. When the Exchange database size is large, the temporary files take more space. You can assign the temporary folder for this process by adding the following registry configuration setting in the Agentless backup proxy server:

Location: HKEY_LOCAL_MACHINE\SOFTWARE\CA\ARCserve Unified Data Protection\Engine\Exchange

Name: "DefragTempPath"

Type: REG_SZ

Value: Valid path for the temporary folder.

Note: Configure the registry in the proxy server as this option is applicable to the Exchange Mail granular restore.

Select the Exchange Mail Objects to be Restored

After generating the Exchange Granular Restore catalog, specify the level of Exchange objects to be restored.

Note: Arcserve UDP does not support granular recovery of public folder objects in an Exchange Mail restore. You need to use Application Restore to recover the entire public folder database and then extract the specific Exchange object.

Follow these steps:

1. From the **Exchange Database** dialog, select the level of Exchange object(s) to be restored (mailbox, folder, or individual mail).

You can select the entire content or partial content of the Exchange object to restore. You can select multiple Exchange components to restore.

Note: When using Arcserve UDP to restore individual mailbox/folders/mail objects from the Exchange mailbox database, the operating system of proxy server must be greater or equal than Exchange VMs.

The available Exchange objects are:

a. **Mailbox** (or mailboxes)

If you select a mailbox level, all corresponding content (folders and individual mail) within that mailbox will be restored.

b. **Folder(s)**

If you select the mailbox folder level, all corresponding mail content within that folder will be restored.

c. **Individual mail object(s)**

If you select the individual mail level, only the selected mail object(s) will be restored.

2. Click **Next**.

The **Restore Options** dialog Opens.

The Exchange objects to be restored are specified.


Define the Restore Options

After you select the level of Microsoft Exchange objects, specify the destination for your backup. You can either select to restore to the original location or restore to a different location option.

Note: For Microsoft Exchange Server 2010 and Exchange Server 2013, archived mailbox items cannot be restored to the original location. Archived mailbox items can only be restored to an alternate location or to a local disk. In addition, regular mailbox items cannot be restored to archive mailboxes.

Follow these steps:

1. From the **Restore Options** dialog, select the restore destination.



The available destination options are:

Restore to Original Location

Restores the mails to the original location from where the backup image was captured. Mails will retain the same hierarchy and be restored to its original mailbox and original folder.

- If current machine is not the active Microsoft Exchange server, Arcserve UDP will detect the location of the active server and then restore the mails to that active server.
- If mailbox has been moved to another Microsoft Exchange server, but still in the same organization, Arcserve UDP will detect the new Exchange server where the original mailbox resides, and then restore to that new server.

- If the display name of the mailbox is changed, any attempt to restore the mailbox (from an earlier backup session) to its original location will fail because Arcserve UDP will not be able to find the changed name. To solve this problem, you can specify to restore this mailbox to an alternate location.

Notes:

- When restoring a mailbox or mail to the original location, make sure the destination mailbox is available, or else the restore will fail. Arcserve UDP Agent (Windows) only validates the destination when the restore job is submitted.
- For a granular restore catalog job in Microsoft Exchange 2013 only, enter the CAS (Client Access Server) Name. The CAS is a thin, stateless server that serves as a proxy for client connections to the Mailbox servers. To do this, from the Restore Options dialog, click **Browse**. From the Select the Client Access Server dialog, select one of the CAS items and click **OK**.

Dump email items

Restores the mails to a disk. This disk location can be a local or a remote machine. The restored mails will maintain the same hierarchy as they had in the corresponding Exchange mailbox. The item name will become the subject of mail.

Note: If the mail subject, folder name, or mailbox name includes any of the following characters, the character will be replaced by hyphen (-) in file name: \ / : * ? " < > |

For this option, you also need to specify what you want Arcserve UDP to do to resolve a conflict situation.

There are two options to resolve this conflict situation:

Rename

If on the disk, there is a file with the same name as the mail subject, Arcserve UDP will still name the mail subject, but will append a number at the end of the mail subject.

Overwrite

If on the disk, there is a file with the same name as the mail subject, the Arcserve UDP Agent (Windows) will overwrite that file. This option is not applicable to exchange email for host-based plan option.

Note: When you select individual mail objects to restore to the disk (dump), by default the format of the restored mail object will be an Outlook Message (.MSG) file and not a Personal Storage Table (.PST) file.

Restore to Alternate Location

Restores the mails to a specified location or lets you browse to the location where your backup images will be restored. The destination must be a mailbox in the same Exchange organization, and a new folder name is required. (If you are attempting to restore mails to an alternate location, the destination cannot be a public folder).

Note: When restoring mail to an alternate location, if the specified destination folder already exists, the restore will continue. However, if the specified folder does not exist, then Arcserve UDP will create the folder first and then continue the restore.

- a. Specify the User Name and Password credentials, click the **Browse** button to navigate through a list of all Exchange Servers, Storage Groups, Exchange Databases, and Mailboxes in the current organization.
- b. Select a mailbox as the destination.
- c. For a granular restore catalog job in Microsoft Exchange 2013 only, enter the **CAS (Client Access Server) Name**. The CAS is a thin, stateless server that serves as a proxy for client connections to the Mailbox servers. To do this, from the **Restore Options** dialog, click **Browse**. From the **Select the Client Access Server** dialog, select one of the CAS items and click **OK**.

2. Click **Next**.


The **Restore Summary** dialog opens. You can review the restore settings and options.

Restore the Exchange Mail

The **Restore Summary** dialog 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 confirmation message informs that the Exchange Mail restore is complete.



Verify that the Exchange Mail was Restored

To verify that the restore process is successful, you can check whether all the restored objects are accurately restored on the destination location.

Follow these steps:

1. Navigate to the restore destination you specified.

For example, if you select to restore the Exchange Mail to the restore destination as Original Location or Alternate Location, then login to the user's mailbox to check the restored email.

If you select to restore the Exchange Mail to Dump email items, then navigate to the dump folder to check the restored email.

For example, if you select to restore the Exchange Mail to C:\dump_folder1, then after restore navigate to this location to verify email


2. Verify emails of the restored Exchange Mail.

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:

How to Restore an MS SQL Server Application



Perform the following tasks to restore a Microsoft SQL Server Application:

1. [Review the Restore Prerequisites and Considerations](#) (see page 512)
2. [Specify the Microsoft SQL Server Information to Restore](#) (see page 514)
 - a. [Specify the Recovery Point and Microsoft SQL Server Database](#) (see page 514)
 - b. [Define the Restore Options](#) (see page 516)
3. [Restore the Microsoft SQL Server Application](#) (see page 519)
4. [Verify that the Microsoft SQL Server Application was Restored](#) (see page 520)

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](#) which 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.
- Arcserve UDP Agent (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 a 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.

While you restore to an alternate location, the Instance Name section is unavailable. The Instance name should always be the same. Therefore, you cannot restore a database to an alternate instance present 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 2005 machine to a SQL Server 2008 machine; however, you cannot restore a SQL Server 2008 machine to a SQL Server 2005 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 <http://msdn.microsoft.com/en-us/library/hh213326.aspx>.
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 <http://msdn.microsoft.com/en-us/library/hh213078.aspx>.

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](#) (see page 514)
2. [Define the Restore Options](#) (see page 516)

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 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 server name dropdown 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. 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 SQL Server database to be restored.
4. 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.

Restore

 **Browse Recovery Points**

Backup Location
<Backup Location> Change

Recovery Point Date

March 2014

S	M	T	W	T	F	S
23	24	25	26	27	28	1
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22
23	24	25	26	27	28	29
30	31	1	2	3	4	5

Today

Time Range
12:00 AM - 6:00 AM
6:00 AM - 12:00 PM
12:00 PM - 6:00 PM
6:00 PM - 12:00 AM (1)

Time	Type	Backup Type	Name
10:51:03 PM	Regular	Full	Customized Incremental Backup

Name	Date Modified	Size
C:		9.53 GB
SqlServerWriter		25.00 MB
X XI-01		
MSSQLSERVER		

Previous Next Cancel Help

5. 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.

Restore

Restore Options

Destination
Select the restore destination

☐ Restore to original location

☐ Dump file only

☒ Restore to alternative location

Instance Name	Database Name	New Database Name	Alternative file location
MSSQLSERVER	master*		Browse
MSSQLSERVER	model	model_copy	Browse
MSSQLSERVER	msdb	msdb	Browse

For SQL writer, the "master" database is not allowed to be renamed.

Backup Encryption or Protection Password
The data that you are attempting to restore is encrypted or password protected. Specify the password that is required to restore the data.

Password

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).

Instance Name	Database Name	New Database Name	Alternative file location	
MSSQLSERVER	master*			Browse
MSSQLSERVER	model	new_databasename	c:\newlocation	Browse
MSSQLSERVER	msdb	msdb		Browse

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](#) (see page 512).

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.



Restore

Restore Summary

Verify your settings are correct and then click Finish to start the restore process.

Component(s) to be restored

Name	Path
SqlServerWriter	

Destination

Dump file to: C:\Dump

Previous **Finish** **Cancel** **Help**

- 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 a Pervasive PSQL Database

Arcserve UDP supports restoring a Pervasive PSQL database. This scenario describes how to restore the Pervasive PSQL database on a Microsoft Windows platform for a Pervasive PSQL 11.3 database.

Note: A Pervasive PSQL VSS Writer is not available on the UI, so you must perform the steps manually, as explained in this procedure.

The following diagram illustrates the process to restore a Pervasive PSQL database:



Perform the following tasks to restore a Pervasive PSQL database:

1. [Review the Restore Prerequisites and Considerations](#) (see page 522)
2. [Determine the Database and Location to Restore](#) (see page 522)
 - [Restore User Database to the Original Location](#) (see page 522)
 - [Restore User Databases and System Databases to the Original Location](#) (see page 528)
 - [Restore User Database to an Alternate Location](#) (see page 533)
3. [Verify that the Pervasive PSQL Database was Restored](#) (see page 534)

Review the Restore Prerequisites and Considerations

Review the following general restore considerations:

- Protects the Pervasive PSQL 11.3 database on Windows x86 and Windows x64 systems.
- All Pervasive PSQL consoles, for example the Pervasive Control Center, must be closed before a restore job starts.
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Determine the Database and Location to Restore

During the restore process, you have the option of where you want to restore the database to. The available selections are

- [Restore User Database to the Original Location](#) (see page 522)
- [Restore User Databases and System Databases to the Original Location](#) (see page 528)
- [Restore User Database to an Alternate Location](#) (see page 533)

Restore User Database to the Original Location

During the restore process, you have the option of where you want to restore the database to. The available selections are **Restore User Database to the Original Location**, **Restore User Databases and System Databases to the Original Location**, and **Restore User Database to an Alternate Location**.

This procedure explains how to restore a Pervasive PSQL user database to the original location.

Follow these steps:

1. Access the **Pervasive Control Center** (PCC) from the operating system **Start** menu.


Note: On Windows platforms, the Pervasive Control Center is installed by default when you install a database engine or a client.

2. Locate the database from the **Pervasive Control Center**, and note down the database name being recovered.

Note: In this example, the database being restored is DB03.


3. Delete the database from the **Pervasive Control Center**.

- a. Right-click the database name.
- b. From the pop-up menu, select **Delete**.




4. From the **Delete Item** dialog, select the option **Yes, database name and ddfs**.

This option ensures the database files are deleted.



5. Restore the database files (*.mkd, *.ddf) to the original location.
 - a. From the **Browse Recovery Points** dialog, select the recovery point data to restore and click **Next**.



6. From the **Restore Options** dialog, select **Restore to original location** and **Skip existing files** and click **Next**.

Restore

Restore Options

Destination
Select the restore destination

☒ Restore to original location

☐ Restore to


Resolving Conflicts
Specify how to resolve conflicts

☐ Overwrite existing files
☐ Replace active files
☐ Rename files
☒ Skip existing files

Directory Structure
Whether to create root directory during restore

☐ Create root directory

7. Create a new database with the name noted down from Step 2, and set the location in Step 3 as the new database location.
 - a. Right-click **Databases**.
 - b. From the pop-up menu, select **New, Database**.



8. From the **New database** dialog, enter the following fields and click **Finish** to complete the database creation.

New database

Database
Create a new database.

Database Name:

Location: ...

Database Options


- ☐ Bound
- ☒ Create dictionary files (if they do not exist)
- ☒ Relational integrity enforced
- ☐ Long metadata (V2 metadata)

Database code page: Server Default [Change code page](#)

DSN Options

- ☒ Create 32-bit Engine DSN

The database is recovered.



The user database is restored to the original location.

Restore User Databases and System Databases to the Original Location

During the restore process, you have the option of where you want to restore the database to. The available selections are **Restore User Database to the Original Location**, **Restore User Databases and System Databases to the Original Location**, and **Restore User Database to an Alternate Location**.

This procedure explains how to restore Pervasive PSQL user databases and system databases to the original location.

Follow these steps:


1. Access the **Pervasive Control Center** (PCC) from the operating system **Start** menu.

Note: On Windows platforms, the Pervasive Control Center is installed by default when you install a database engine or a client.

2. Locate the databases from the **Pervasive Control Center**.
3. Delete the databases from the **Pervasive Control Center**.


a. Right-click the database name.

b. From the pop-up menu, select **Delete**.




- c. From the **Delete Item** dialog, select the option **Yes, database name and ddfs**.

This option ensures the database files are deleted.



- d. Repeat for each database that needs to be deleted.
4. Stop the **Pervasive PSQL Relational Engine** service.
 - a. Right-click **Pervasive PSQL Relational Engine**.
 - b. From the pop-up menu, select **Stop Service**.



5. Delete all the database files (*.mdk, *.ddf) for the system databases such as DEFAULTDB, TEMPDB.


By default the database files are located in "C:\ProgramData\Pervasive Software\PSQL\defaultdb" and "C:\ProgramData\Pervasive Software\PSQL\tempdb".

6. Delete the transaction log files in the log folder.

By default the log folder is located in "C:\ProgramData\Pervasive Software\PSQL\Transaction Logs".

7. Restore the whole Pervasive Software data folder to the original location by using the overwrite option.

By default the folder path is "C:\ProgramData\Pervasive Software\PSQL".



8. From the **Restore Options** dialog, select **Restore to original location** and **Overwrite existing files** and click **Next**.

The screenshot shows a 'Restore' dialog box with a title bar and a close button. The main content area is titled 'Restore Options' and contains two sections: 'Destination' and 'Resolving Conflicts'. The 'Destination' section has a label 'Select the restore destination' and two radio button options: 'Restore to original location' (which is selected) and 'Restore to' followed by a text input field, a 'Browse' button, and a small icon. The 'Resolving Conflicts' section has a label 'Specify how to resolve conflicts' and four radio button options: 'Overwrite existing files' (selected), 'Replace active files' (checked with a checkbox), 'Rename files', and 'Skip existing files'. To the right of this section is a 'Directory Structure' section with the label 'Whether to create root directory during restore' and a checkbox labeled 'Create root directory'. At the bottom of the dialog are four buttons: 'Previous', 'Next', 'Cancel', and 'Help'.

Restore

Restore Options

Destination
Select the restore destination

☒ Restore to original location

☐ Restore to 


Resolving Conflicts
Specify how to resolve conflicts

☒ Overwrite existing files
☒ Replace active files
☐ Rename files
☐ Skip existing files

Directory Structure
Whether to create root directory during restore

☐ Create root directory

9. Restore the database files (*.mkd, *.ddf) to original location and overwrite the existing files.



10. Start the **Pervasive PSQL Relational Engine** service.
 - a. Right-click **Pervasive PSQL Relational Engine**.
 - b. From the pop-up menu, select **Start Service**.

The user databases and system databases are restored to the original location.

Restore User Database to an Alternate Location

During the restore process, you have the option of where you want to restore the database to. The available selections are **Restore User Database to the Original Location**, **Restore User Databases and System Databases to the Original Location**, and **Restore User Database to an Alternate Location**.

This procedure explains how to restore a Pervasive PSQL user database to an alternate location.

Follow these steps:

1. Restore the database file to a directory.
2. From the **Pervasive Control Center**, create a new database with a name you want to use, and select the directory used in Step 1 as the new database location.
3. Complete the database creation.

The database will be brought online with the name specified in Step 2.

The user database is restored to an alternate location.

Verify that the Pervasive PSQL Database was Restored

After the completion of the restore process, verify that the Pervasive PSQL database 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 user database to the restore destination as original location, then log in to the original server and check if the Pervasive PSQL database exists.

If you select to restore the user database to the alternate location, then log in to the alternate server provided and check if the Pervasive PSQL database exists.


2. Verify the Pervasive PSQL database was restored.

The Pervasive PSQL database is restored successfully.

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](#) (see page 536)
- [Restore the Server Parameter File](#) (see page 536)
- [Restore the Parameter File](#) (see page 537)
- [Restore the Archived Redo Logs](#) (see page 538)
- [Restore the Tablespaces or Data Files](#) (see page 538)
- [Restore System, Undo Tablespaces, Data Files](#) (see page 539)
- [Restore All Tablespaces and Data Files](#) (see page 541)
- [Restore Control Files](#) (see page 542)
- [Restore the Entire Database \(Tablespaces and Control Files\)](#) (see page 544)
- [Recover the Oracle Database Using Bare Metal Recovery](#) (see page 545)

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](#) which 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 in to 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 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 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 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 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 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 control files are restored.
8. Mount the database to begin the database recovery:

```
SQL> START 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\
01_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 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 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\
01_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;
```

The entire database is restored.

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:

How to Perform a File-Level Recovery



Perform these tasks for a file-level recovery:

- [Review the Restore Prerequisites](#) (see page 548)
- [\(Optional\) Recover Data from the iSCSI Volume to the Target Machine](#) (see page 549)
- [Specify the Recovery Point](#) (see page 550)
- [Specify the Target Machine Details](#) (see page 554)
- [Specify the Advanced Settings](#) (see page 558)
- [\(Optional\) Manage Pre/Post Scripts for Automation](#) (see page 559)
- [Create and Run the Restore Job](#) (see page 562)
- [Verify that Files are Restored](#) (see page 563)

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 5.x does not support the *reiserfs* file system. If the operating system of the Backup Server is RedHat 5.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 UDP Agent (Linux) Live CD to perform the file-level restore because Live CD supports all types of file system.

- Review the [Compatibility Matrix](#) which 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

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.

Note: If you have selected **Source local** as your backup destination, the Backup Server cannot connect to the Source local directly. To access the Source local, you have to perform additional configurations.

To restore files from Source local, follow these steps:

- a. Share the backup destination (Source local) and ensure that the Backup server can connect to the backup destination.
- b. Add the shared destination as the backup storage location to the Backup server.

Now, Source local behaves as an NFS backup storage location and you can restore files from the share.

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** dropdown menu.
The Arcserve UDP Agent (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 UDP Agent (Linux):

- a. Open the Arcserve UDP Agent (Linux) web interface.

Note: During the installation of Arcserve UDP Agent (Linux), you received the URL to access and manage the server. Log in to Arcserve UDP Agent (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. The recent recovery point is selected.

Restore Wizard - File Restore

Select the recovery point you want to recover.

Session Location: NFS share <NFS Share Full Path> **Connect**

Machine: <Machine Name/IP Address>

Date filter: Start 11/16/13 End 11/30/13 **Search**

	Time	Type	Name	Encryption Algorithm	Encryption Password
	11/30/2013 1:27:31 AM	BACKUP_INCREMENTAL	S0000000140	AES 128	
	11/29/2013 10:27:30 PM	BACKUP_FULL	S0000000139	AES 128	
	11/29/2013 7:27:28 AM	BACKUP_INCREMENTAL	S0000000138	AES 128	
	11/29/2013 4:27:29 AM	BACKUP_INCREMENTAL	S0000000137	AES 128	

Files/Folders to be restored **Add** **Remove**

File/Folder Name	Date Modified	Size

<Previous Next> Cancel Help

3. Select a session from the **Session Location** drop-down list, if you want to restore another session, 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 backend up data is stored in Source local, then you must first convert the source node to an NFS server, and then share the session location.

4. Click **Connect**.

All the nodes that have been backed up to this location get listed in the **Machine** drop-down list.

5. Select the node that you want to restore from the **Machine** drop-down list.

All the recovery points of the selected node get listed.


6. 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.

7. 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.



8. 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**.

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.

Follow these steps:

- 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.



Backup Server



Recovery Points



Target Machine



Advanced



Summary

Specify the target machine information for the File Restore.

☒ Restore to original location
 ☐ Restore to alternative location

Target Machine Settings

Host Name/IP	Machine Name/IP Address
Username	
Password	

Resolving Conflicts

How should arcserve UDP Agent(Linux) resolve conflicting files

☒ Overwrite existing files
☐ Rename files
☐ Skip existing files

Directory Structure

Whether to create root directory during restore

☐ Create root directory

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.



Backup Server



Recovery Points



Target Machine



Advanced



Summary

Specify the target machine information for the File Restore.

☐ Restore to original location
 ☒ Restore to alternative location

Target Machine Settings

Host Name/IP	<Machine Name/IP Address>
Username	
Password	
Destination	Browse

Resolving Conflicts

How should arcserve UDP Agent(Linux) resolve conflicting files

☒ Overwrite existing files
☐ Rename files
☐ Skip existing files

Directory Structure

Whether to create root directory during restore

☐ Create root directory

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/CA/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/CA/d2dserver/usr/prepost`

Follow these steps:

1. Place the file in the following location of the Backup Server:
`/opt/CA/d2dserver/usr/prepost`
2. Provide the execution permission to the script file.
3. Log into the Arcserve UDP Agent (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 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 UDP Agent (Linux) Live CD, you can get the IP address of the target machine.

The following diagram displays the process to perform a BMR:



Complete the following tasks to perform a BMR:

- [Review the BMR Prerequisites](#) (see page 565)
- [Get the IP Address of the Target Machine Using the Live CD](#) (see page 566)
- [\(Optional\) Recover Data to the iSCSI Volume of the Target Machine](#) (see page 567)
- [\(Optional\) Recover Data from the iSCSI Volume to the Target Machine](#) (see page 568)
- [Review the Backup Server](#) (see page 569)
- [Specify the Recovery Points](#) (see page 570)
- [Specify the Target Machine Details](#) (see page 572)
- [Specify the Advanced Settings](#) (see page 573)
- [\(Optional\) Manage Pre/Posts Scripts for Automation](#) (see page 574)
- [Create and Run the Restore Job](#) (see page 577)
- [\(Optional\) Perform Post-BMR Operations](#) (see page 578)
- [Verify that the Target Machine is Restored](#) (see page 584)

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 UDP Agent (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.
- Review the [Compatibility Matrix](#) which 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 UDP Agent (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 UDP Agent (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:
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:

```
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.

```
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 UDP Agent (Linux) Live CD or mount the iso file of the Arcserve UDP Agent (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 UDP Agent (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 UDP Agent (Linux) Live CD or mount the iso file of the Arcserve UDP Agent (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 UDP Agent (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** dropdown menu.

The Arcserve UDP Agent (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 UDP Agent (Linux):

- a. Open the Arcserve UDP Agent (Linux) web interface.

Note: During the installation of Arcserve UDP Agent (Linux), you received the URL to access and manage the server. Log in to Arcserve UDP Agent (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 UDP Agent (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 session location.



Backup Server



Recovery Points



Target Machine



Advanced



Summary

Select the recovery point you want to recover.

Session Location: NFS share <NFS Share Full Path> Connect

Machine: <Machine Name/IP Address>

Date filter: Start 11/16/13 End 11/30/13 Search

	Time	Type	Name	Encryption Algorithm	Encryption Password
	11/19/2013 2:40:27 AM	BACKUP_INCREMENTAL	S0000000105		
	11/18/2013 11:40:25 AM	BACKUP_INCREMENTAL	S0000000104		
	11/18/2013 8:40:24 AM	BACKUP_INCREMENTAL	S0000000103		
	11/18/2013 5:40:25 AM	BACKUP_FULL	S0000000102		

Disk Name	Disk Size
/dev/vda	15.00 GB
/dev/vdb	5.00 GB
/dev/vdc	5.00 GB

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 **Reboot** option to automatically restart the target node after the BMR is complete.

5. 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/CA/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. 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/CA/d2dserver/usr/prepost

Follow these steps:

1. Place the file in the following location of the Backup Server:
/opt/CA/d2dserver/usr/prepost
2. Provide the execution permission to the script file.
3. Log into the Arcserve UDP Agent (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.

```
# lvsdisplay -m /dev/mapper/VGTest-LVTest
--- Logical volume ---
LV Name                /dev/VGTest/LVTest
VG Name                VGTest
LV UUID                udoBIx-XKBS-1Wky-3FVQ-mxMf-Fay0-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  
# lvsdisplay -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 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:

How to Perform a Bare Metal Recovery Using a Backup



Complete the following tasks to perform a BMR using a backup:

1. [Review the BMR Prerequisites and Considerations](#) (see page 586)
2. [Define BMR Options](#) (see page 587)
 - [Perform BMR in Express Mode](#) (see page 596)
 - [Perform BMR in Advanced Mode](#) (see page 599)
3. [Verify that the BMR was Successful](#) (see page 605)
4. [BMR Reference Information](#) (see page 605)
5. [Troubleshooting BMR Issues](#) (see page 611)

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 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](#) which 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](#) (see page 606)
 - [Operating Systems that Support UEFI/BIOS Conversion](#) (see page 607)
 - [Managing the BMR Operations Menu](#) (see page 608)

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 UDP 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 Standby VM**

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 [assign the egvcm variable for your book] application.

Note: For this option, you can only recover data if the virtual conversion to a VMDK file (for VMware) was performed using [assign the egvcm variable for your book].

If you select this option, see [Recover using a VMware Virtual Standby VM](#) (see page 626) to continue this procedure.

For more information, see [Recover using a VMware Virtual Standby VM](#) (see page 626) in the 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 [assign the egvcm variable for your book] application.


Note: For this option, you can only recover data if the virtual conversion to a VHD file (for Hyper-V) was performed using [assign the egvcm variable for your book].

If you select this option, see [Recover using a Hyper-V Virtual Standby VM](#) (see page 620) to continue this procedure.

For more information, see [Recover using a Hyper-V Virtual Standby VM](#) (see page 620) in the online help.

4. Select **Restore from an Arcserve UDP 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 <http://www.microsoft.com/en-us/download/details.aspx?id=6408>.

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](#) (see page 607) 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.

arcserve® bare metal recovery

Bare Metal Recovery(BMR)
- Select a Recovery Point

The top pane displays all backed up machines and their backup destinations. If you click on a machine, you can then see the associated recovery points on the bottom pane. Select a recovery point to continue.

Note: By default, only the backed up machines that are detected from local volumes are listed here. After a new removable disk is attached or detached, you can click "Refresh" to refresh the machine list.

You can also click "Browse" to add any backed up machines from the remote shared folder or the data store.

If you fail to browse the remote shared folder, it may be because the NIC driver is not installed or the IP address is incorrect. If necessary, you can perform the following:
[Click here](#) to launch the load driver utility
[Click here](#) to launch the network configuration utility

1 Network Adapter(s) Detected
 Microsoft Hyper-V Network Adapter
 - IP Address: 155.35.70.125
 - Status: Connected

The following backed up machines are detected:

<Machine Name>

Backed up Information

Host Name: appsvr1master

Operating System: Windows Server 2012 R2

Platform: X64

Source: Recovery Point Server

Refresh Browse

The following recovery points are detected for the specified machine. Select a recovery point and continue:

9/28/2014 2:53:43 AM

Boot Firmware:
- UEFI

Backed up Volumes (Used Size/Total Size/Minimum Size Required):

Volume	Used Size	Total Size	Minimum Size Required
\\P\\Volume{48b00294-3b04-4e6f-9716-7ba44a91}	240 MB	300 MB	800 MB
EFI System Partition	260 MB	260 MB	850 MB
C:\	14.56 GB	925.67 GB	465.77 GB
E:\	212 MB	97.66 GB	3.01 GB

Is Data Encrypted?
- Yes

Encryption Library Type
- Microsoft Cryptography Library

Encryption Algorithm Type
- AES-256


Utilities Back Next Abort

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**.

A screenshot of a Windows-style dialog box titled "Enter Session Password". It features a text input field containing eight dots, representing a password. To the right of the input field are two buttons: "OK" and "Cancel". Below the input field, there are two lines of text: "Current password length: 8 characters" and "Maximum password length: 23 characters".


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 nonauthoritative 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](#) (see page 596) if you want minimal interaction during the recovery process.
- Select [Advanced Mode](#) (see page 599) 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 Preinstallation 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.

arcserve® bare metal recovery

Bare Metal Recovery(BMR)
- Start Restore Process

This page displays a summary of the disk/volume restore settings you have made.

Note: After the BMR process is complete and server has been rebooted, you may not want to perform backup jobs from this server. If you are just testing the BMR functionality, it is recommended that you select the "Do not start Agent service automatically after reboot" option.
When you select this option, you can manually start the Agent service (and the Recovery Point Server service, if installed) after reboot if you want to perform backup jobs.

Summary of Restore Settings

Restore Item	Status	Progress	Throughput
Restore source volume 'System Res...	Completed	100.0%	1122.14 MB/Minute
Restore source volume 'C:\' to curre...	Restoring	1.5%	2797.09 MB/Minute

☒ Automatically reboot your system after recovery.
☐ Do not start Agent service automatically after reboot.

Elapsed Time: 00 : 00 : 09
 Estimated Time Remaining: 00 : 07 : 45
 [1.5%] [224MB/14737MB] Restoring basic source volume 'C:\' to current destination disk 0

! Boot volume was restored to current destination disk 0. Please boot your system from this disk.

Utilities Back Next Abort

- 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** 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](#) (see page 608).
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 Preinstallation 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.

arcserve® bare metal recovery

Bare Metal Recovery(BMR)
- Start Restore Process

This page displays a summary of the disk/volume restore settings you have made.

Note: After the BMR process is complete and server has been rebooted, you may not want to perform backup jobs from this server. If you are just testing the BMR functionality, it is recommended that you select the "Do not start Agent service automatically after reboot" option.
When you select this option, you can manually start the Agent service (and the Recovery Point Server service, if installed) after reboot if you want to perform backup jobs.

Summary of Restore Settings

Restore Item	Status	Progress	Throughput
Restore source volume 'System Res...	Completed	100.0%	1122.14 MB/Minute
Restore source volume 'C:\' to curre...	Restoring	1.5%	2797.09 MB/Minute

☒ Automatically reboot your system after recovery.
☐ Do not start Agent service automatically after reboot.

Elapsed Time: 00 : 00 : 09
 Estimated Time Remaining: 00 : 07 : 45
 [1.5%] [224MB/14737MB] Restoring basic source volume 'C:\' to current destination disk 0

! Boot volume was restored to current destination disk 0. Please boot your system from this disk.

Utilities Back Next Abort

- 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

This section contains the following topics:

[How Bare Metal Recovery Works](#) (see page 606)


[Operating Systems that Support UEFI/BIOS Conversion](#) (see page 607)

[Managing the BMR Operations Menu](#) (see page 608)

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 Server 2003	x86	No	No
Windows Server 2003	x64	No	No
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 Server 2012 R2	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 Operations

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 of 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.

This section contains the following topics:

[Slow throughput performance during BMR](#) (see page 611)

[After BMR, dynamic volumes are not recognized by the operating system](#) (see page 612)

[Unable to Reboot Hyper-V VM After BMR](#) (see page 612)

[Unable to Reboot VMware VM After BMR](#) (see page 613)

[Unable to boot the server after performing a BMR](#) (see page 613)

[Failed to submit BMR job to Recovery Point Server](#) (see page 614)

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 remedy 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 remedy 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 Perform a BMR Using a Virtual Standby 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.

The following diagram illustrates the process for how to perform a BMR using a virtual standby VM:



Complete the following tasks to perform a BMR using a backup:

1. [Review the BMR Prerequisites and Considerations](#) (see page 616)
2. [Define BMR Options](#) (see page 617)
 - [Recover Using a Hyper-V Virtual Standby VM](#) (see page 620)
 - [Recover Using a VMware Virtual Standby VM](#) (see page 626)
 - [Perform BMR in Express Mode](#) (see page 632)
 - [Perform BMR in Advanced Mode](#) (see page 634)
3. [Verify that the BMR was Successful](#) (see page 639)
4. [BMR Reference Information](#) (see page 605)
5. [Troubleshooting BMR Issues](#) (see page 611)

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](#) which 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](#) (see page 606)
 - [Operating Systems that Support UEFI/BIOS Conversion](#) (see page 607)
 - [Managing the BMR Operations Menu](#) (see page 608)

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 UDP 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. 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 [assign the egvcm variable for your book] application.

Note: For this option, you can only recover data if the virtual conversion to a VMDK file (for VMware) was performed using [assign the egvcm variable for your book].

If you select this option, see [Recover using a VMware Virtual Standby VM](#) (see page 626) 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 [assign the egvcm variable for your book] application.

Note: For this option, you can only recover data if the virtual conversion to a VHD file (for Hyper-V) was performed using [assign the egvcm variable for your book].

If you select this option, see [Recover using a Hyper-V Virtual Standby VM](#) (see page 620) 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](#) (see page 626) 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](#) (see page 620) to continue this procedure.

Recover using a Hyper-V Virtual Standby 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 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. 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.

arcserve® bare metal recovery

Bare Metal Recovery(BMR)
- Select the type of backup for BMR

Select type of restore source:

☐ Restore from a arcserve Unified Data Protection backup
Use this option to perform a restore from either a backup destination folder or a data store

☒ Recover from a Virtual Standby VM
Use this option to perform a virtual-to-physical(V2P) restore from a virtual standby VM


☐ Source is on a VMware machine

☒ Source is on a Hyper-V machine

▲ Utilities Back Next Abort


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 UDP Agent (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 [assign the egvcm variable for your book].



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.


If you select the **Current State** recovery point, verify the Arcserve UDP Agent service is started on the standby virtual machine.

- If the virtual machine that you are recovering from is powered off, the **Latest State** recovery point is displayed.

If you select the **Latest State** recovery point, an error message is displayed to inform you that the recovery point you are recovering from is the Latest (and not the Current) state and requesting that you start the virtual machine before continuing with the recovery process.

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](#) (see page 632).
- Select **Advanced Mode** if you want to customize the recovery process. For more information, see [Perform BMR in Advanced Mode](#) (see page 634).

Default: Express Mode.

Recover using a VMware Virtual Standby 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 Standby VM** 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. 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.

arcserve® bare metal recovery

Bare Metal Recovery(BMR)
- Select the type of backup for BMR

Select type of restore source:

☐ Restore from a arcserve Unified Data Protection backup
Use this option to perform a restore from either a backup destination folder or a data store

☒ Recover from a Virtual Standby VM
Use this option to perform a virtual-to-physical(V2P) restore from a virtual standby VM

☒ Source is on a VMware machine
☐ Source is on a Hyper-V machine

Utilities Back Next Abort

2. Click **Next**.

The **Select a Recovery Point** screen is displayed with the **ESX/VC Credentials** dialog.

The screenshot displays the Arcserve Bare Metal Recovery (BMR) application window. The title bar reads "arcserve® bare metal recovery". The main window has a dark blue header. Below the header, the text "Bare Metal Recovery(BMR) - Select a Recovery Point" is visible. A message states: "On this page you can select a VM snapshot and recover data from it to your computer." Below this, instructions for connecting to an ESX Server are provided, including links for the load driver utility and network configuration utility. A "Connect to a ESX Server:" section contains "Connect" and "Refresh" buttons. A modal dialog box titled "Please enter ESX/vCenter credentials" is open, featuring input fields for "ESX Server/vCenter:", "User Name:", "Password:", and "VI Port:" (pre-filled with "443"). It also has radio buttons for "HTTP" and "HTTPS" (selected), and "OK" and "Cancel" buttons. At the bottom of the main window, there are "Utilities", "Back", "Next", and "Abort" buttons.


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 UDP Agent (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.


If you select the **Current State** recovery point, verify the Arcserve UDP Agent service is started on the standby virtual machine.

- If the virtual machine that you are recovering from is powered off, the **Latest State** recovery point is displayed.

If you select the **Latest State** recovery point, an error message will be displayed to inform you that the recovery point you are recovering from is the Latest (and not the Current) state and requesting that you start the virtual machine before continuing with the recovery process.

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](#) (see page 632).
- Select **Advanced Mode** if you want to customize the recovery process. For more information, see [Perform BMR in Advanced Mode](#) (see page 634).

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 Preinstallation 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.



The dialog box titled "Summary of Restore Settings" contains two sections. The first section, "Summary of Disk Restore Settings:", has two empty text boxes labeled "Source Disk" and "Destination Disk". The second section, "Summary of Volume Restore Settings:", contains a table with three columns: "Source Volume", "Destination Volume", and "On Disk".

Source Volume	Destination Volume	On Disk
System Reserved (100 MB)	\\?\Volume{f44042c2-bc04-19e3-8c03-005...}	Disk 0
C:\ (39.90 GB)	\\?\Volume{f44042c5-bc04-19e3-8c03-005...}	Disk 0

At the bottom of the dialog, there is a message: "Confirm your restore settings and click OK to continue." and two buttons: "OK" and "Cancel".

- 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.

arcserve® bare metal recovery

Bare Metal Recovery(BMR)
- Start Restore Process

This page displays a summary of the disk/volume restore settings you have made.

Note: After the BMR process is complete and server has been rebooted, you may not want to perform backup jobs from this server. If you are just testing the BMR functionality, it is recommended that you select the "Do not start Agent service automatically after reboot" option.
When you select this option, you can manually start the Agent service (and the Recovery Point Server service, if installed) after reboot if you want to perform backup jobs.

Summary of Restore Settings

Restore Item	Status	Progress	Throughput
Restore source volume 'System Res...	Completed	100.0%	1122.14 MB/Minute
Restore source volume 'C:\' to curre...	Restoring	1.5%	2797.09 MB/Minute

☒ Automatically reboot your system after recovery.
☐ Do not start Agent service automatically after reboot.

Elapsed Time: 00 : 00 : 09
 Estimated Time Remaining: 00 : 07 : 45
 [1.5%] [224MB/14737MB] Restoring basic source volume 'C:\' to current destination disk 0

Boot volume was restored to current destination disk 0. Please boot your system from this disk.

Utilities Back Next Abort

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](#) (see page 608).


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 Preinstallation 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.



The dialog box titled "Summary of Restore Settings" contains two sections. The first section, "Summary of Disk Restore Settings:", has two tabs: "Source Disk" and "Destination Disk". The second section, "Summary of Volume Restore Settings:", contains a table with three columns: "Source Volume", "Destination Volume", and "On Disk".

Source Volume	Destination Volume	On Disk
System Reserved (100 MB)	C:\ (100 MB)	Disk 0
C:\ (23.90 GB)	D:\ (39.90 GB)	Disk 0

At the bottom of the dialog, there is a message: "Confirm your restore settings and click OK to continue." and two buttons: "OK" and "Cancel".

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.

arcserve® bare metal recovery

Bare Metal Recovery(BMR)
- Start Restore Process

This page displays a summary of the disk/volume restore settings you have made.

Note: After the BMR process is complete and server has been rebooted, you may not want to perform backup jobs from this server. If you are just testing the BMR functionality, it is recommended that you select the "Do not start Agent service automatically after reboot" option. When you select this option, you can manually start the Agent service (and the Recovery Point Server service, if installed) after reboot if you want to perform backup jobs.

Summary of Restore Settings

Restore Item	Status	Progress	Throughput
Restore source volume 'C:\' to curre...	Restoring	23.1%	1931.36 MB/Minute

☒ Automatically reboot your system after recovery.
☐ Do not start Agent service automatically after reboot.

Elapsed Time: 00 : 01 : 09
Estimated Time Remaining: 00 : 03 : 54
[23.1%] [2208MB/9564MB] Restoring basic source volume 'C:\' to current destination disk 0

 Boot volume was restored to current destination disk 0. Please boot your system from this disk.

Utilities

BackNextAbort

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

This section contains the following topics:

[How Bare Metal Recovery Works](#) (see page 640)


[Operating Systems that Support UEFI/BIOS Conversion](#) (see page 641)

[Managing the BMR Operations Menu](#) (see page 642)

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 Server 2003	x86	No	No
Windows Server 2003	x64	No	No
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 Server 2012 R2	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 Operations

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 of 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.

This section contains the following topics:

[Slow throughput performance during BMR](#) (see page 645)

[After BMR, dynamic volumes are not recognized by the operating system](#) (see page 646)

[Unable to Reboot Hyper-V VM After BMR](#) (see page 646)

[Unable to Reboot VMware VM After BMR](#) (see page 647)

[Unable to boot the server after performing a BMR](#) (see page 647)

[Failed to submit BMR job to Recovery Point Server](#) (see page 648)

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 remedy 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 remedy 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 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](#) (see page 649)
- [Restore Files of a Cluster Shared Disk](#) (see page 650)
- [Restore a Specific Node in a Cluster](#) (see page 650)
- [Restore a Corrupted Cluster Shared Disk](#) (see page 651)
- [Restore the Entire Clustered Nodes and Shared Disk](#) (see page 651)

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](#) which 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.

Chapter 8: Generating Arcserve UDP Reports

This section contains the following topics:

[How to Generate Arcserve UDP Reports](#) (see page 653)

How to Generate Arcserve UDP Reports

The **reports** tab provides access to various types of reports such as Alerts and Backup Status. 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. The drill-down report includes the following items:

Job Nodes

Displays the node name where backup jobs of Agent Backup or Host-Based VM Backup run.

Protected Nodes

Displays the name of agent node and node protected by the Arcserve UDP 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 UDP Agent, Arcserve UDP Recovery Point Server, 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](#) (see page 655).

The following reports are available in Arcserve UDP:

Alert Report

Displays the alert information for nodes.

Backup Size Trend Report

Displays the backup data size of Arcserve Backup and Arcserve UDP 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.

Node Backup Status Report

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. 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.
- **Not Attempted:** Provides a list of jobs that are not attempted.

Virtualization Protection Status Report

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.

Managed Capacity Report

Displays the raw data size of the last successful full backup for each node that Arcserve Backup, Arcserve UDP Agent, and Host-based VM Backup protect.

Data Distribution on Media Report

Displays the compressed and actual (raw) backup data size at various storage locations during the specified time period. This report allows you to drill down to display more detailed information about disk and deduplication categories.


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.

Notes:

- As a prerequisite, install [Adobe Flash Player ActiveX](#) (version 10.0 or higher) and [Microsoft .NET Framework](#) (version 2.0 or higher) for the Report Chart export feature on the server to export images in a report successfully.
- You cannot install Adobe Flash Player in Windows Server 2012 and 2012 R2. To generate report chart, install Desktop Experience feature in Windows Server 2012 or 2012 R2.

The following screenshot displays the two types of Filters/Actions available on a report page:



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](#) (see page 657).
- **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.
- **Refresh:** Click to update the report related information.
- **Email:** You can email the report. For more information, see [Send Report by Email](#) (see page 660).
- **Save:** You can use the option to export a report. Select one of the formats from **CSV**, **PDF**, and **HTML** and click **Open** or one of the options of **Save** from the dialog displayed at the bottom of the page to export the report.

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.

- 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.

Schedule Emails

Using Arcserve UDP, 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](#) (see page 90).

You can [create a schedule](#) (see page 657), and [edit the schedule](#) (see page 659).

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 contents, 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:

- Log in to the Arcserve UDP.
- Click **reports** on the Navigation bar.
- From the upper right corner of any report, click the global **Filters/Actions** section.
- 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 screenshot shows the 'New Schedule' dialog box. At the top, there is a title bar with the text 'New Schedule' and a close button. Below the title bar, there is a message: 'From this screen, you can edit the schedule, specify the email contents and settings, and specify which reports to be included. After specifying your schedule options, click OK to save your changes, or click Cancel to cancel without the changes.' Below this message, there are four tabs: 'General', 'Email', 'Reports', and 'Schedule'. The 'General' tab is currently selected. In the 'General' tab, there is a text box labeled '* Schedule Name:' with the value 'New Schedule' entered. Below this, there is a text box labeled 'Description:' which is empty. At the bottom of the dialog, there are four buttons: 'Run Now', 'OK', 'Cancel', and 'Help'.

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](#) (see page 657).

Follow these steps:

1. Log in to the Arcserve UDP.
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 UDP, 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](#) (see page 90).

Follow these steps:

1. Log in to the Arcserve UDP.
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](#) (see page 90).

The image shows a 'Send Report by Email' dialog box with a close button (X) in the top right corner. It contains several input fields and options: 'To' with a text box containing 'UserID@companyname.com'; 'CC' with an empty text box and a note 'Use semicolon to separate multiple email recipients.'; 'Priority' with three radio buttons: 'High', 'Normal' (selected), and 'Low'; 'Subject' with a text box containing 'Arcserve Unified Data Protection Alert'; 'Comment' with a text area and a note 'The comments will be inserted at beginning of the email, before all reports.'; and 'Attachment' with three checkboxes: 'PDF', 'HTML', and 'CSV', followed by the text 'Attach the report data to this email with the specified file type(s)'. At the bottom are three buttons: 'OK', 'Cancel', and 'Help'.

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.

Chapter 9: Managing Arcserve High Availability

This section contains the following topics:

[How Arcserve High Availability Works](#) (see page 663)

How Arcserve High Availability Works

Using Arcserve Unified Data Protection, you can monitor and manage Arcserve 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 HA Control Services

To manage Arcserve High Availability features from Arcserve UDP, you must add all the Control Services that you want to manage. After you add the Control Services, you can create Full System scenarios and manage scenarios created in Arcserve High Availability.

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 HA Licenses

Using Arcserve UDP, you can manage Arcserve High Availability licenses from the Console. If you already have the Arcserve High Availability Control Service license, you can select the Control and register the license.

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.
Enter the registration key
4. Click **OK**.

The license is registered.

Manage Scenarios

Arcserve UDP 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](#) (see page 665)
- [Create Full System Scenarios](#) (see page 666)
- [Manage Scenarios](#) (see page 668)
- [Edit Scenarios](#) (see page 669)
- [Manage Scenario Hosts](#) (see page 671)
- [Operations on Scenarios](#) (see page 672)
- [BMR and Reverse Replication](#) (see page 674)
- [Monitor Scenarios](#) (see page 675)

Manage Scenario Group

Arcserve UDP 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

In addition to performing various operations on Arcserve High Availability scenarios, you can create Full System scenarios in Arcserve UDP. You can create scenarios using the scenario creation wizard.

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 Replication and High Availability Administrator Guide.

The **Master and Replica Properties** dialog opens.

15. Review the master and replica properties and then click **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:

Replica Network Adapter

Click to choose the adapter to map to the adapter listed in the Master Network Adapter column.

Apply master adapter information

(Default) Select if the Master Adapter is in the DHCP mode.

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 UDP 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 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 for Full System scenarios.

5. Click **Save** from the **Action** drop-down menu.

The scenario properties are updated.

For a stopped Full System scenario, you could also edit the virtual platform settings.

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**.

To edit network adapter mapping for High Availability or Assured Recovery, perform the following 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:

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 UDP 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 RHA 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 RHA 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 UDP lets you monitor high availability 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 UDP lets you deploy the RHA 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](#) (see page 676) 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 UDP 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.

Appendix A: Troubleshooting

Add Permissions for VDDK 5.1 and 5.5 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](#).

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 through 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 under the following conditions:

- 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.

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, do the following:

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, do the following:

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:

Virtual Standby jobs fail to Hyper-V systems. The following message appears in the Activity Log:

Virtual Standby job failed to get the Hyper-V VM.

Solution:

Virtual Standby jobs fail under the following conditions:

- 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 a 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 more free disk space in the system volume.

Note: If you discover other possible causes, contact Arcserve Support.

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.

The solution is to specify a new name.

Failed to perform backup of a virtual disk. 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.

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\casad2websvc-stdout.*.log  
<D2D_home>\TOMCAT\logs\casad2websvc-stderr.*.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, do the following:

1. Open Registry Editor and access the following key:
 - x86 Operating Systems:
HKEY_LOCAL_MACHINE\SOFTWARE\Apache Software Foundation\Procrun
2.0\CASAD2WebSvc\Parameters\Java
 - x64 Operating Systems:
HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Apache Software Foundation\Procrun
2.0\CASAD2WebSvc\Parameters\Java

2. Do one of the following:

- 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, do the following:

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, do the following:

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.

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:

This may happen because you can only run one snapshot creation at a time.

To solve 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 cluster nodes:

1. Open the Windows registry.
2. Navigate to HKLM\SOFTWARE\CA\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 cluster nodes:

1. Open the Windows registry.
2. Navigate to HKLM\SOFTWARE\CA\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\CA\ARCserve Unified Data Protection\Engine.
3. Right-click **VMMMaxJobNumber**, select **Modify**, and specify the value as required.

Convert Incremental Backup to a Verify Backup Because the Virtual Machine Snapshots Either Changed from the Last Backup Job or Needs Consolidation

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, click on the following VMware Knowledge Base article:

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2003638

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.

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.

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 in to the Backup Server and run the following command:

```
# /opt/CA/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.

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.

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 corrupted and unreadable.

Solution:

When this problem occurs, start the restore process from a UDP agent installed on a Windows 2012 R2 system.

Troubleshooting 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.

Failed to Restore Exchange Object (Message, Folder, or Mailbox) to the Original or Alternate Location

Symptom:

Restoring a host-based backup Exchange object to the original location or alternate location may fail in the following situations:

- If the object is on a virtual machine through a proxy server.
- If the proxy server does not trust the Exchange Server certificate, the restore fails and the following message is displayed:

Failed to communicate with exchange server. It is possible the certificate is not installed on the proxy server. Please ask the administrator to check.

Solution:

Install the certificate on the Exchange Server.

Follow these steps:

1. On the Backup proxy server, open the web browser, and access the Outlook Web App (OWA) URL on the Exchange CAS server selected for restore .
2. Click **Certificate error**.
3. From the **Certificate Invalid** dialog, click **View certificates**.
4. From the **General** tab of the **Certificate Information** dialog, click **Install Certificate**.
5. From the **Certificate Wizard Import** dialog, choose **Local Machine** and then click **Next**.
6. Choose **Place all certificates in the following store**, and then click **Browse**.
7. Choose **Trusted Root Certification Authorities**, click **OK** and then click **Next**.
8. Click **Finish** to import the certificate to the Trusted Root Store on the Backup Proxy Server.
9. When you get the notification that **The import was successful**, click **OK**.

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.

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.Select a recovery point server, and from the right pane 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.

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\CA\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\CA\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.

Pause or 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.

Arcserve UDP Cannot Communicate with the Arcserve UDP Agent (Windows) Web Service on Remote Nodes

Valid on Windows operating systems.

Symptom:

Arcserve UDP cannot communicate with the Arcserve UDP Agent (Windows) web service on remote nodes.

Solution:

The following table describes reasons why Arcserve UDP cannot communicate with the Arcserve UDP Agent (Windows) web service on remote nodes and the corresponding corrective action:

Cause	Corrective Action
The network was not available or not stable when applying plans.	Verify that the network is available and stable and then try again.
The Arcserve UDP Agent (Windows) node could not handle the load when the application tried to communicate with the node.	Verify that the CPU on the remote Arcserve UDP Agent (Windows) node is in a normal state and then try again.
The Arcserve UDP Agent (Windows) service on the remote node was not running when applying plans.	Verify that the Arcserve UDP Agent (Windows) on the remote node is running and then try again.
The Arcserve UDP Agent (Windows) service was not communicating properly.	Restart the Arcserve UDP Agent (Windows) service on the remote node and then try again.

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 centre pane, click **Actions**, and then select **Deploy Now**.

The plan gets redeployed and the backup settings are applied to the node.

The 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`

Host-based 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.

MAC Address Changes are Not Retained After VM Recovery

Valid on Windows platforms.

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\CA\ARCserve Unified Data Protection\Engine
Key Name: RetainMACForVDDK
Value Type: String
Key Value: 1

On virtual machines with two NIC cards, set the RetainMACForVDDK registry key if you wish to set one as Manual. Otherwise, all cards are set to Automatic after recovery.

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 cancelled. 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 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 cancelled, 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.

Appendix B: Data Deduplication

Related Topics:

[Compression](#) (see page 709)

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 4 KB. If you set it to 8 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 UDP Agent (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

The following are the limitations of a deduplication data store:

- You cannot modify the compression type once you create a deduplication data store.

Appendix C: Command Line Data Integrity Tool for Deduplication Data Store

The command line tool (ca_gddmgr.exe) lets you check data integrity at recovery point level and data store level for deduplication data stores. You can run this utility after the backup is complete.

You can also use this tool to regenerate a hash database in case of a hash database failure.

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, ca_gddmgr_2014-9-4_11-14-22-655.log is a log file that contains all details.

Location:

You can find ca_gddmgr.exe in the "Bin" folder of the UDP installation path.

Syntax:

ca_gddmgr.exe

-Scan ChecksumUDPSession <data store name> -Node [<All> |<UDP agent node name>]
-RecoveryPoint [<All>|<recovery point number>] [-Password < data store password >]
[-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>]

Options:

ChecksumUDPSession

Checks data integrity for the specified recovery point.

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

Regenerates the hash database by scanning index and reference files. Before you specify this option, manually stop the deduplication data store.

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

Examples:

```
ca_gddmgr.exe -Scan ChecksumUDPSession GDDDataStore1 -Node All -RecoveryPoint All  
-Password 123
```

```
ca_gddmgr.exe -Scan ChecksumUDPSession GDDDataStore1 -Node myComputer -RecoveryPoint  
1 -Password 123
```

```
ca_gddmgr.exe -Scan VerifyRefCount GDDDataStore1
```

```
ca_gddmgr.exe -Scan VerifyData GDDDataStore1 -Password 123
```

```
ca_gddmgr.exe -Scan VerifyAll GDDDataStore1
```

```
ca_gddmgr.exe -Scan RebuildHash GDDDataStore1
```

```
ca_gddmgr.exe -Scan RebuildHash GDDDataStore1 -NewHashPath C:\NewHashPath
```


Appendix D: 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

Related Topics:

[Data Deduplication](#) (see page 701)

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 UDP 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 UDP 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 UDP solution provides encryption feature for data.

Note: 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 UDP Agent.

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 UDP 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 UDP agent is backed up to data store.

- c. The Arcserve UDP solution has encryption password and session password.
 - The encryption password is required for data store.
 - The session password is required for node.
 - When you have backup, session password is mandatory for an agent node.

- The session password is required when you restore from a different machine.
- For Arcserve UDP, session password is mandatory.

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.

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 a Arcserve UDP action to back up, restore, create virtual standby, or replicate nodes.

jobs

A tab on the Arcserve UDP Console that lets you 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 Windows 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.

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 UDP. 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 UDP but a plan is not assigned. When a plan is not assigned, you cannot back up data and the node remains unprotected.