

Agent for Windows User Guide

Arcserve® Unified Data Protection

Version 8.x

arcserve®

Legal Notice

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

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

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

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

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

The manufacturer of this Documentation is Arcserve.

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

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

Arcserve Product References

This document references the following Arcserve products:

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

Contents

Agent for Windows User Guide	1
Chapter 1: Understanding Arcserve UDP Agent (Windows)	13
Introduction	14
Arcserve UDP Agent (Windows) Documentation	15
About this Document	16
Features	17
Arcserve UDP Agent (Windows) Videos	26
How Arcserve UDP Agent (Windows) Works	27
How the Backup Process Works	28
How Block-Level Incremental Backups Work	30
How Infinite Incremental Backups Work	31
How File Level Restores Work	32
How Verify Backups Work	33
How Recovery Sets Work	34
How Bare Metal Recovery Works	37
How Arcserve UDP Agent (Windows) Update Works	38
Chapter 2: Installing/Uninstalling Arcserve UDP Agent (Windows)	43
How to Install Arcserve UDP Agent (Windows)	44
Review the Installation Considerations	46
Install Arcserve UDP Agent (Windows) Using the Installation Wizard	47
Install Arcserve UDP Agent (Windows) Silently	50
Verify that the Arcserve UDP Agent (Windows) Installation was Successful	54
How the Installation Process Affects Operating Systems	56
Arcserve UDP Agent (Windows) Installer Error Codes	70
How to Install Arcserve UDP Agent (Windows) Updates	75
Review the Considerations for Installing Updates	78
Specify Updates Preferences	83
Check for Updates and Download	88
Install the Arcserve UDP Agent (Windows) Updates	90
Verify that the Updates are Successfully Installed	92
(Optional) Install Arcserve UDP Agent (Windows) Updates Silently	93
Troubleshooting Update Issues	94
How to Uninstall Arcserve UDP Agent (Windows)	98

Review the Uninstallation Considerations	100
Uninstall Arcserve UDP Agent (Windows) Using Add or Remove Programs	101
Uninstall Arcserve UDP Agent (Windows) Using the Command Line	102
Remove Components Left Behind by the Uninstaller	103
Verify that the Arcserve UDP Agent (Windows) Uninstallation was Successful	104
Files Not Removed During Uninstallation	105
Troubleshooting Uninstall Issues	112
UDP Workstation Free	115
Chapter 3: Getting Started with Arcserve UDP Agent (Windows)	117
How to Navigate the Arcserve UDP Agent (Windows) User Interface	118
Accessing Arcserve UDP Agent (Windows)	121
Introducing the User Interface	122
Understanding the User Interface	124
Troubleshooting User Interface Issues	144
Chapter 4: Settings	147
Configure or Modify Backup Settings	148
Specify Protection Settings	149
Specify Schedule Settings	168
Specify Advanced Settings	193
Specify Pre/Post Backup Settings	197
Manage File Copy Settings	199
Specify the File Copy Source	200
Specify the File Copy Destination	207
Configure File Copy Settings to Optimize Performance	216
Specify the File Copy Schedule	219
Manage File Archive Settings	220
Specify the File Archive Source	221
Specify the File Archive Destination	228
Configure File Archive Settings to Optimize Performance	235
Specify the File Archive Schedule	238
Configure the Copy Recovery Point Settings	239
Copy Recovery Points - Example Scenarios	243
Specify Preferences	245
Specify General Preferences	246
Specify Email Preferences	248
Specify Updates Preferences	256

Manage Export/Import Settings	262
Export Settings	263
Import Settings	264
Chapter 5: Using Arcserve UDP Agent (Windows)	267
How to Perform a Backup	268
Review the Backup Prerequisites and Considerations	271
Configure or Modify Backup Settings	289
Perform a Backup	332
Verify that the Backup is Successful	336
How Arcserve UDP Agent (Windows) Works	337
Troubleshooting Backup Issues	346
Perform File Copy to Disk/Cloud	353
Perform a Restore	354
Restore Considerations	355
Restore Methods	357
How to Restore From a Recovery Point	360
How to Restore From a File Copy	384
How to Restore From a File Archive	400
How to Restore Files/Folders	415
How to Restore a Virtual Machine	443
How to Use Exchange Granular Restore (GRT)	466
How to Restore Microsoft Exchange Data	475
How to Restore a Microsoft Exchange Application	476
How to Restore a Microsoft SQL Server Application	488
How to Restore an Oracle Database	501
How to Restore an Active Directory	518
How to Perform an Authoritative Restore of an Active Directory after a BMR	525
How to Restore Microsoft Clustered Nodes and Shared Disks	533
Restore from Windows Explorer Using Arcserve UDP Recovery Point View	539
How to Copy a Recovery Point	542
Review the Prerequisites	543
Configure the Copy Recovery Point Settings	544
Copy a Recovery Point	550
Verify the Copied Recovery Point	558
Mount a Recovery Point	559
Create a VHD File from an Arcserve UDP Agent (Windows) Backup	565

View Logs	569
How to Download File/Folders without Restore	572
How to Create a Boot Kit	574
Launch the Create Boot Kit Utility	576
Determine the Method to Generate a BMR ISO Image	579
Create an Arcserve UDP Agent (Windows) BMR ISO Image for a CD/DVD	581
Create an Arcserve UDP Agent (Windows) BMR ISO Image for a USB Stick	585
Verify the Boot Kit is Created	590
How to Perform a Bare Metal Recovery Using a Backup	591
Review the BMR Prerequisites and Considerations	593
Define BMR Options	595
Verify that the BMR was Successful	612
BMR Reference Information	613
Troubleshooting BMR Issues	620
How to Perform a Bare Metal Recovery Using a Virtual Standby VM or Instant VM	627
Review the BMR Prerequisites and Considerations	628
Define BMR Options	630
Verify that the BMR was Successful	650
BMR Reference Information	651
Troubleshooting BMR Issues	658
Using the PowerShell Interface	665
How to use the PowerShell Interface	666
Add Arcserve UDP Agent (Windows) Licensing	686
Change Server Communication Protocol	688
Use Scripts to Backup and Restore MySQL Database	688
Restore MySQL Database	690
Modify Arcserve-MySQL-pre-post-snapshot-conf.bat	690
Use Scripts to Backup and Restore PostgreSQL Database	690
Prerequisites	690
Apply Scripts	691
Restore PostgreSQL Database	692
Chapter 6: Troubleshooting Arcserve UDP Agent (Windows)	695
Troubleshooting Overview	696
Arcserve UDP Agent Service could not be started because of port conflict	697
Reboot Not Required After Agent Deployment	700
Unable to Connect to Cloud	701

Unable to Change Destination to Removable Device	702
Unable to display Arcserve UDP Agent (Windows) UI in Firefox	704
Settings Disabled when Opening Agent UI	705
Unable to Open the SQL database in SQL Management Studio from Mounted Volume	706
Recovery of SQL Server Databases to Original Location fails	707
Login Link Does not Work at Arcserve UDP Agent Home	708
Troubleshooting Installation Issues	709
Unable to install/uninstall Arcserve UDP Agent (Windows) if a previous attempt was interrupted	710
Windows failed to start after Arcserve UDP Agent (Windows) is installed	712
Troubleshooting Update Issues	716
Unable to Access Arcserve UDP Agent (Windows) After Reboot	717
Unable to Connect to the Arcserve Download Server to Download Updates	718
Failed to Download Arcserve UDP Agent (Windows) Updates	719
Troubleshooting Uninstall Issues	720
Unable to install/uninstall Arcserve UDP Agent (Windows) if a previous attempt was interrupted	721
Troubleshooting User Interface Issues	723
Unable to display Arcserve UDP Agent (Windows) home page in IE	724
Job Monitor data speed displays a 0 or some other abnormal value	725
Troubleshooting Backup Issues	726
SQL Server backup failed due to "out of memory" error	727
Backup sessions do not include Microsoft SQL database information	728
Catalog Job fails Due to Less Space when Backing up Large Number of Files	729
Failed to create snapshot for selected volumes	730
Unable to change backup destination folder to Arcserve UDP Recovery Point View	731
Troubleshooting BMR Issues	733
Slow throughput performance during BMR	734
After BMR, dynamic volumes are not recognized by the operating system	735
Unable to Reboot Hyper-V VM After BMR	736
Unable to Reboot VMware VM After BMR	737
Unable to boot the server after performing a BMR	738
Failed to submit BMR job to Recovery Point Server	739
Troubleshooting Merge Issues	740
Merge Session is Skipped	741
Merge Job Failed when Configured to Retain Recovery Sets	742
Merge Job Fails After Being Paused by a Restore Job	743

Troubleshooting Exchange Issues	744
Fail to Restore Exchange Database in DAG Node to Original Location	745
Restore Job Fails During Dump Exchange Database	746
Unable to Connect across the Domain Live Mailbox from Exchange GRT utility	747
APPENDIX: Frequently Asked Questions (FAQ)	749
File Copy Related FAQ	750
Can I restore data if I lose the encryption password?	751
What is the maximum file size that can be backed up/restored?	752
What is not deleted during a File Copy – Delete Source job?	753
Does a File Copy job copy data directly from the local source disks?	754
What is the maximum file size that can be stored on Amazon S3 cloud?	755
For any file size less than 64K, will Arcserve UDP Agent (Windows) copy the entire file?	756
Can a File Copy job and a Backup run simultaneously?	757
During a File Copy job, will the stub files be copied again?	758
Does every File Copy job initiate a VSS snapshot like a regular Arcserve UDP Agent (Windows) Backup job?	759
Will a File Copy stored on an Amazon S3 cloud location be open source archive format?	760
If a File Copy – Delete Source job deletes files, will I be able to perform a BMR from the file copy destination?	761
For a File Copy job, is the Delete Source option enabled by default?	762
Encryption Related FAQ	763
If I change the encryption type or the encryption password and the maximum number of recovery points are then reached, what happens?	764
If I enter a new encryption password, will the old encryption password be asked for first?	765
What happens to data encrypted either using Windows or a third-party encryption system?	766
Exchange Granular Restore FAQ	767
Can Exchange search attachments in email?	768
Can I restore a mailbox without overwriting the existing data?	769
Service Related FAQ	770
How do I use a different account to start the Arcserve UDP Agent Service?	771
Updates Related FAQ	772
Can I use scripted information for specifying Updates proxy settings?	773
Can I use a workstation node as an Updates staging server?	774
Can I manage/operate Updates together or do I need to configure each node separately (one by one)?	775
Does an Updates staging server need a separate Arcserve UDP Agent (Windows) license if I am not using any Arcserve UDP Agent (Windows) functions on same staging server?	776

Can I continue to replicate my recovery points backed up in my local RPS server to the remote managed RPS server, after Upgrade?777

Can I continue to replicate backups from my production systems running Update 2 to a remotely managed RPS server running Update 1, after upgrade?778

APPENDIX: Using the RDX Cleaner Utilities779

What are the RDX Cleaner Utilities?780

How to Execute the RDX Cleaner Utility781

 Post Cleaning Verification (RDX Cleaner)783

How to Execute the RDX Force Cleaner Utility785

 Post Cleaning Verification (RDX Force Cleaner)787

APPENDIX: Arcserve UDP Terms and Definitions789

Agent-Based Backup790

Compression790

configuration791

Dashboard791

Destination791

Data Store791

Discovered Nodes791

Encryption791

Host-Based Agentless Backup792

HOTADD Transport Mode793

Job793

NBD Transport Mode793

NBDSSL Transport Mode793

Nodes793

Plan793

Protected Nodes793

Recent Event794

Recovery Point794

Recovery Point Server794

Replicate794

Resources794

SAN Transport Mode794

Systems794

Tasks794

Unprotected nodes795

Contact Arcserve Support

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

[Contact Support](#)

With Arcserve Support:

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

Chapter 1: Understanding Arcserve UDP Agent (Windows)

This section contains the following topics:

Introduction	14
Arcserve UDP Agent (Windows) Documentation	15
About this Document	16
Features	17
Arcserve UDP Agent (Windows) Videos	26
How Arcserve UDP Agent (Windows) Works	27

Introduction

Arcserve UDP Agent (Windows) is a disk-based backup product designed to provide a fast, simple, and reliable way to protect and recover critical business information. Arcserve UDP Agent (Windows) is a light-weight solution for tracking changes on a machine at the block level and then backing up only those changed blocks in an incremental fashion. As a result, Arcserve UDP Agent (Windows) lets you perform frequent backups (as frequently as every 15 minutes), reducing the size of each incremental backup (as well as the backup window) and providing a more up-to-date backup. Arcserve UDP Agent (Windows) also provides the capability to restore files/folders, volumes, and applications, and perform bare metal recovery from a single backup. In addition, Arcserve UDP Agent (Windows) also lets you copy and restore backed-up data from your specified cloud storage location.

Arcserve UDP Agent (Windows) Documentation

The latest Arcserve UDP Agent (Windows) documentation are:

- [Arcserve Unified Data Protection Agent for Windows User Guide](#)
- [Arcserve Unified Data Protection Release Notes](#)

The Arcserve UDP Release Notes contains information relating to system requirements, operating system support, application recovery support, and other information you may need to know before installation of this product. In addition, this document also contains a list of known issues that you should be aware of before using this product.

About this Document

This document provides the information to understand, install, use, and maintain Arcserve UDP Agent (Windows) in the most practical and efficient manner. This document is divided into multiple categories to help you easily identify and locate the specific information you are seeking.

The online help version of this document provides a link at the bottom of each topic to let you submit feedback to us about this document. We continually strive to make our documentation as complete, error free, and easy-to-read as possible. You can help by giving us feedback. Thank you in advance!

Understanding Arcserve UDP Agent (Windows)	This section contains an overview of the features of Arcserve UDP Agent (Windows), with process-flow descriptions of how some key features work. By understanding how the features work, it should be easier to understand and perform the related tasks.
Installing/Uninstalling Arcserve UDP Agent (Windows)	This section contains information about installing Arcserve UDP Agent (Windows), including any pre-installation considerations you should be familiar with, the installation procedure to be performed, and instructions if you want to perform a silent installation.
Getting Started with Arcserve UDP Agent (Windows)	This section contains an overview of the Arcserve UDP Agent (Windows) user interface, and detailed information about each individual area of this interface. Before you use Arcserve UDP Agent (Windows), it is important that you become familiar with the details of this interface.
Settings	This section contains information to understand and manage the various Arcserve UDP Agent (Windows) configuration settings.
Using Arcserve UDP Agent (Windows)	This section contains the step-by-step procedures for such tasks as performing ad-hoc backups, restoring from backups, copying recovery points, viewing logs, file copying, performing BMR, and installing any Arcserve UDP Agent (Windows) self-updates.
Troubleshooting Arcserve UDP Agent (Windows)	This section contains some fault-isolation information necessary to quickly identify and locate the source of a problem so that it can be remedied and allow Arcserve UDP Agent (Windows) to become fully operational again.
Arcserve UDP Agent (Windows) FAQ	This section provides answers to some of the more commonly asked questions.
Appendix	The appendix section at the end of this document contains a collection of useful and supplementary information which is not necessary for proper usage of Arcserve UDP Agent (Windows), but still may be of interest or use to you.

Features

The following features are provided with Arcserve UDP Agent (Windows):

BACKUP

The following backup features are provided with Arcserve UDP Agent (Windows):

- Lets you perform different types of backup jobs, such as full, incremental, or verify.
- Provides volume filtering capability to let you specify to back up only the selected volumes.
 - ◆ If the specified backup destination is on the local volume, a warning message displays notifying you this volume is not being backed up.
 - ◆ If system/boot volume is not selected for backup, a warning message displays notifying you the backup is unusable for Bare Metal Recovery (BMR).
 - ◆ If a data store is configured on the volume, then the volume cannot be selected as the backup source.
- Protects all specified volumes of your computer (except if the volume contains the backup destination).
- Lets you encrypt and protect (with encryption passwords) your sensitive data.
- Lets you set/change backup schedules (or immediately initiate a customized backup).
 - ◆ Lets you set advanced scheduling features. To use advanced scheduling, set your Backup Data Format to Advanced. Then, you can access the advanced schedule view, set the advanced schedule for the backup job, backup throttle, merge and daily/weekly/monthly retention.
- Provides a system tray monitor to display status/notification information and perform quick actions.
- The Arcserve UDP solution provides the capability to utilize a complimentary, limited version of Arcserve Backup to perform backups (agent-based and agentless-based) to tape.

Block Level Incremental Backups

- Only backs up the blocks on the source volumes that have changed after the last successful backup.
- Significantly reduces the amount of backup data.

If you have a large file and you only change a small portion of this file, Arcserve UDP Agent (Windows) backs up only the changed portion to the incremental backup. It does not back up the whole file.

- Consumes less disk space and less time.
- Lets you perform more frequent backups, making the backup images more up-to-date (as often as every 15 minutes) for recovery.

Infinite Incremental (I2) Snapshots

- Initially creates one full backup and then intelligently creates incremental snapshot backups forever (after the initial full backup).
- Uses less storage space, performs backups faster, and puts less load on your production servers.
- Can automatically collapse (merge) incremental changes optimizing the use of disk storage.

Application Consistent Backups

- Takes advantage of Windows Volume Shadow Copy Service (VSS) to ensure data consistency for any VSS-aware application.
- Provides recovery of both Microsoft SQL Server and Microsoft Exchange Server (without performing a full disaster recovery).

Ad-hoc Backups

An ad-hoc backup is one that is created when the situation makes it necessary, rather than being arranged in advance or being part of a plan.

- Provides you with the flexibility to perform "ad-hoc" backups outside of the scheduled backups.

For example, you have a repeat schedule for Full, Incremental, and Verify backups and you want to make major changes to your computer. You can perform an immediate backup without waiting for the next scheduled backup to occur.

- Provides you with the capability to add a customized (unscheduled) recovery point so that you can roll back to this previous point in time if necessary.

For example, you install a patch or service pack and then discover it adversely affects the performance of your computer. You can roll back to the ad-hoc backup session that does not include it.

Change Backup Destination

Provides you with the capability to change the backup destination, depending on the type of installation.

- **Arcserve Unified Data Protection - Full:**

Provides you with the capability to change the backup destination, when the destination is a data store on a Recovery Point Server (RPS). The next backup will be a full backup.

- **Arcserve Unified Data Protection - Agent:**

Provides you with the capability to change the backup destination, when the destination is a local disk or a remote shared folder. When the disk space on your destination volume reaches its maximum capacity, Arcserve UDP Agent (Windows) lets you change the destination and you can continue with full or incremental backups.

Note: You can also configure email alert notifications for Destination Threshold so that you can be informed if it reaches the specified threshold value.

Copy Recovery Points

Provides you with the capability to copy recovery point data and safely store it off-site for the purpose of restore in a catastrophe, or you can save your recovery points to multiple locations. In addition, if your destination is getting full you can consolidate your backups into a single recovery point which resembles the exact state at that point. When you select a recovery point to copy, you are capturing:

- Backup blocks that were created for that specified point-in-time.
- Whatever previous backup blocks are necessary to recreate a full and most recent backup image.

The Copy Recovery Points feature can be launched manually (ad-hoc) or automatically based upon your specified schedule.

Mount Recovery Points

Provides the ability to mount a recovery point to a drive letter (volume) or an NTFS folder, to view, browse, copy, or open the backup files directly in Windows Explorer.

Backup Speed Throttling

Provides you with the capability to specify the maximum speed (MB/min) at which your 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. Backup Speed Throttling only applies when the Backup Data Format is Standard. When the Backup Data Format is Advanced, a Backup Throttle Schedule is available.

Reserve Space on Destination

Provides you with the capability to specify a percentage of the calculated space that is necessary to perform a backup. This amount of continuous space is then reserved on the destination before the backup starts writing data and helps improve backup speed.

Note: Reserve Space on Destination only applies when the Backup Data Format is Standard. When the Backup Data Format is Advanced, this option does not exist.

Backup Status Monitoring

Arcserve UDP Agent (Windows) provides the capability to monitor:

- Last backup status
- Recovery Points
- Destination capacity
- Protection summary
- Most recent events
- License notifications

Job Status Monitoring

Arcserve UDP Agent (Windows) provides the capability to monitor:

- Details about the next scheduled job
- Details about the next scheduled job

RESTORE

Arcserve UDP Agent (Windows) provides the following restore features:

- Restore data from specific recovery points.
- Search/browse to a specific file/folder to restore.
- Restore from File Copy.
- Restore a virtual machine (VM) that you previously backed up.
- Set the restore destination to an alternate location or server.

-
- Restore encrypted backup data.
 - Granular-level restore of Exchange objects.
 - On Demand Catalog job for any catalog-less backup recovery point using the Find Files/Folders to Restore option.
 - Restore Active Directory to recover Active Directory objects and attributes.

Types of Restores

Arcserve UDP Agent (Windows) provides the following types of restores:

- **File-Level Restore**
Restores any backed up files/folders.
- **Recovery Point Restore**
Restores backed up data based from a specified point in time (recovery point).
- **File Copy Restore**
Restores File Copy data from a disk or cloud.
- **Exchange Granular-Level Restore**
Restores individual Exchange objects (mailboxes, mailbox folders, or mail).
- **Virtual Machine (VM) Recovery**
Restores a VM that you previously backed up.
- **Application Restore**
Restores backed up Microsoft SQL Server/Microsoft Exchange data at the database level.
- **Explorer Integration Restore**
Arcserve UDP Agent (Windows) provides the capability to browse directly and restore files/folder and Exchange objects (mailboxes, mail folders, mail) from Windows Explorer by using the "Change to Arcserve UDP Agent (Windows) View" option.
- **Change Restore Location**
Arcserve Unified Data Protection - Full:
Provides you with the ability to restore from a local disk, remote shared folder, and from a Recovery Point Server (RPS).
Arcserve Unified Data Protection - Agent:

Provides you with the ability to restore from a local disk or a remote shared folder.

- **Bare Metal Recovery (BMR)**
 - Recovers a computer system from "bare metal", and includes the operating system, applications, and data components necessary to rebuild or restore the entire backed-up system. BMR is used for disaster recovery or for migration from one server to another.
 - Restores to dissimilar hardware and resolves any hardware differences.
 - Expands and restores to bigger disks if necessary.
 - Provides the capability to perform the following types of V2P (Virtual to Physical) Bare Metal Recovery. This feature lets you perform V2P recovery from the latest state of a standby virtual machine and from any recovery point that has been previously converted from an Arcserve UDP Agent (Windows) backup session. This feature also helps you reduce the loss of your production computer.
 - ♦ BMR from a Hyper-V server
 - ♦ BMR from a VMware ESX or vCenter

Disk Resizing

- During a Bare Metal Recovery, you can restore the image to another disk and can resize the disk partitions if necessary (without losing any data that is stored on the drive).
- When restoring to another disk, the capacity of new disk must be the same size or larger than the original disk.

Note: Disk resizing if necessary is for basic disks only, and not for dynamic disks.

ALERT NOTIFICATIONS

Arcserve UDP Agent (Windows) provides the following email alert notifications:

- Missed jobs - Sends an alert notification for any scheduled job that did not run at the scheduled time.
- Backup, Catalog, File Copy, Restore, or Copy Recovery Point job failure/crash - Sends an alert notification for all unsuccessful job attempts. This category includes all failed, incomplete, and canceled jobs, and crashed attempts.

Note: These email alerts are sent with a high importance. The email alerts that have a high importance level setting display a visual indicator of an exclamation point in their Inbox.

- Backup, Catalog, File Copy, Restore, or Copy Recovery Point job success - Sends an alert notification for all successful job attempts.
- Merge job stopped, skipped, failed, or crashed - Sends an alert notification for all stopped, skipped, failed, or crashed merge jobs.
- Merge job success - Sends an alert notification for all successful merge jobs.
- Backup destination free space is less than - Sends an email notification when the amount of unused space at the backup destination is less than a specified value.
- New Updates Available - Sends an email notification when a new update for Arcserve UDP Agent (Windows) is available. Email notifications are also sent if a failure occurs during the check for updates or during the download.
- Resource threshold alerts - Sends an alert notification when any specified resource performance threshold is reached. The monitored resource levels are CPU Usage (percentage), Memory Usage (percentage), Disk Throughput (MB/second) and Network I/O (percentage of NIC bandwidth currently using).

ENCRYPTION/DECRYPTION SUPPORT

Arcserve UDP Agent (Windows) provides the capability to encrypt and protect (with encryption passwords) your sensitive data and also decrypt the encrypted data after recovery.

- Encryption support is provided for both uncompressed backup format and compressed backup format. (Uncompressed backup is no longer VHD format if encrypted).
- Windows built-in encryption libraries are used for data encryption and decryption.

For Windows 2003/Vista/2008: CAPI (CryptoAPI) is used for data encryption.

For Windows 7/2008 R2/Windows 2012: CNG (Cryptography API Next Generation) is used for data encryption.

Note: Data interoperability is supported both ways between CAPI and CNG, meaning that data that is encrypted on Windows 2003/Vista/2008 can be decrypted on Windows 7/2008 R2 (and vice versa). This data interoperability enables moving backups of any computer to a different computer, and to restore data from there.

- Encryption password management provides a memory feature so that you do not need to remember encryption passwords when attempting to restore encrypted data. For every encrypted backup, the encryption password is saved in a password list file.

As long as you can log in to Arcserve UDP Agent (Windows), there is no need to remember encryption passwords to restore data from current backups. (Current backups are defined as backups that were created from the same computer that you are logged in to). If you attempt to restore data from encrypted backups belonging to a different computer, you are always asked to provide the encryption password.

FILE COPY

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.

Arcserve UDP Agent (Windows) provides the following capabilities to copy or move files and help you reduce storage cost, meet compliance, and improve data protection.

Note: When you use the option File Copy - Delete Source, the data is moved from the source to the destination (deleted from source location). When you perform a file copy, the data is copied from the source to the destination (files remain intact on the original location).

- Copy files to disk or to cloud based upon your specified policies.
- Block-level file copying lets you save and store only the blocks of the source that have changed as of the last file copying. (Significantly reduces the amount of file copied data).
- Select the source to copy, which can be a specific volume, volumes, folder, or folders.
- Use filters to include or exclude files that are based upon your specific criteria or patterns.
- Specify a schedule for file copying that is based upon completion of a specified number of successful backups.
- File copy versions of the same source at the specified destination.
- Encrypt file copied data for security.
- Compress data before performing file copying process.

- Specify how long to retain file copy data.
- Specify how many versions of the data you can have on the destination.
Note: Arcserve UDP Agent (Windows) does not copy application files, files with system attributes, and files with temporary attributes. Only a current backed-up source is eligible for file copying.

Arcserve UDP Agent (Windows) UPDATES

Provides the following capabilities for downloading and installing self- updates to Arcserve UDP Agent (Windows):

- Check for new available updates to Arcserve UDP Agent (Windows) (manually initiated from the UI or system tray monitor or automatically as scheduled).
- Trigger automatic or manual downloading of updates.
- Specify a custom schedule to perform automatically periodic checks for updates.
- Trigger installation of updates either from the UI, the system tray monitor, or silently from the command line.
- Specify to send automatic email notifications when new updates become available (or when problems occur).
- Configure the client and or a staging server to connect to Arcserve Support (directly or by way of a proxy server) to download available updates. (A staging server is an Arcserve UDP Agent (Windows) installed computer which is used as a temporary storage location for downloading an update before it is installed into an Arcserve UDP client computer from that staging server).
- Use staging servers for clients that have limited access to the Internet.
- Configure multiple staging servers for downloading the updates. If the primary staging server is unavailable, the download function automatically transfers to the next specified staging server.
- Remote deploy from one computer to another and let you move all updates configuration and email settings from that first computer to the deployed computer.

Note: All updates that are released for Arcserve UDP Agent (Windows) are cumulative. As a result, each update also includes all previously released updates to ensure that your computer is always up-to-date.

Arcserve UDP Agent (Windows) Videos

For those of you who believe that "a picture is worth a thousand words" Arcserve UDP Agent (Windows) provides various how-to videos that are designed to simplify your understanding and performance of specific tasks. Watching step-by-step videos is a great way to help you learn how to use Arcserve UDP Agent (Windows) features to perform essential system protection procedures.

Note: These videos are meant to supplement (and not replace) the written procedures that they are related to. Refer to the actual procedures for all detailed information (precautions, notes, examples, and so on) associated with each task.

You can access these instructional videos from the Arcserve UDP Agent (Windows) user interface or from within the product documentation.

We provide a library of how-to videos that are designed to simplify your understanding and performance of specific tasks. You can access these instructional videos from either the arcserve.com website or from YouTube. The versions of the videos from arcserve.com and YouTube are identical, and only the viewing source is different:

[To view Arcserve UDP Agent \(Windows\) videos on YouTube](#)

The videos that are supplied are only a start, and we expect to have more created in the future. If you have any ideas for new videos, let us know. You can click the user interface link to Provide Feedback. You can even send Arcserve an email using the link at the bottom of all Online Help topics.

How Arcserve UDP Agent (Windows) Works

Arcserve UDP Agent (Windows) lets you perform frequent and periodic block level backups of your full machine. These backups can be stored on either an internal drive, an external drive, on a remote network share, or a data store on a Recovery Point Server (RPS), depending on the type of installation (Arcserve Unified Data Protection - Full or Arcserve Unified Data Protection - Agent). If the backup destination volume is also selected as the backup source volume, a never ending backup is not executed. During the backup, the backup destination volume is excluded and an entry is added to the Activity log. The Arcserve UDP Agent (Windows) provides the capability to perform Full, Incremental, or Verify type backups.

Arcserve Unified Data Protection - Full:

Available backup destinations include: internal drive, external drive, remote network share, or a data store on a Recovery Point Server (RPS). When you create a Plan from the Arcserve UDP server, you can select Data Store on Recovery Point Server as the destination and then deploy the plan to the agent node.

Arcserve Unified Data Protection - Agent:

Available backup destinations include: internal drive, external drive, or a remote network share.

Arcserve UDP Agent (Windows) also provides various methods to identify and locate the backed up data and allow you to restore it if necessary. Regardless of which restore method you select, Arcserve UDP Agent (Windows) lets you quickly identify the data you need and retrieve it from the appropriate backup location.

How the Backup Process Works

Arcserve UDP Agent (Windows) lets you perform frequent and periodic block level backups of your entire machine. These backups can be stored on either an internal drive, an external drive, on a remote network share, or a data store on a Recovery Point Server (RPS), depending on the type of installation (Arcserve Unified Data Protection - Full or Arcserve Unified Data Protection - Agent). The Arcserve UDP Agent (Windows) provides the capability to perform Full, Incremental, or Verify type backups.

The basic process for how Arcserve UDP Agent (Windows) performs a backup is simple. When you initiate a backup (either as scheduled or manually launched), Arcserve UDP Agent (Windows) captures a full VSS snapshot, and then backs up only those blocks that have been changed since the previous successful backup. (If it is a Full backup, all blocks are backed up). This block-level incremental backup process significantly reduces the amount of backup data. For example, if you have a large file and only change a small portion of this file, Arcserve UDP Agent (Windows) backs up only the changed portion to the incremental backup and not back up the entire file.

During this block-level incremental backup process, Arcserve UDP Agent (Windows) not only captures the data, but also creates a catalog containing all information related to the operating system, installed applications (Microsoft SQL and Microsoft Exchange only), configuration settings, necessary drivers, and so on. If necessary, you can then restore this backed-up image to recover your data or your entire machine. If the backup destination volume is also selected as the backup source volume, a never ending backup is not executed. During the backup, the backup destination volume is excluded and an entry is added to the Activity log.

Note: You can submit a faster backup job (catalog-less backup), since a catalog is not required after a backup job is complete. The backup settings option "Generate File System catalog for faster search after each backup" by default is unchecked, indicating it will perform a faster backup.

The details of what is being backed up, how it is being backed up, when it is being backed up, and so on, are controlled by the various backup configuration settings that you specify. These settings are applied to each backup job, regardless of how you initiate the backup (automatically or manually).


Arcserve Unified Data Protection - Full:

With this type of installation, configure your backup by creating a plan. Available backup destinations include: internal drive, external drive, remote network share, or a data store on a Recovery Point Server (RPS). When you create a Plan from the

Arcserve UDP server, you can select Data Store on Recovery Point Server as the destination and then deploy the plan to the agent node.

Arcserve Unified Data Protection - Agent:

With this type of installation, configure your backup by specifying the backup settings. Available backup destinations include: internal drive, external drive, or a remote network share.



How Block-Level Incremental Backups Work

When you start a backup, the specified volume is divided into a number of subordinate data blocks that are then backed up. The initial backup is considered the "parent backup" and will be a Full Backup of the entire volume to establish the baseline blocks to be monitored. Before performing the backup, a VSS snapshot is created, then an internal monitoring driver checks each block to detect any changes. As scheduled, Arcserve UDP Agent (Windows) will then incrementally back up only those blocks that have changed since the previous backup. You can schedule the subsequent block-level incremental backups ("child backups") as frequently as every 15 minutes to always provide accurate, up-to-date backup images.

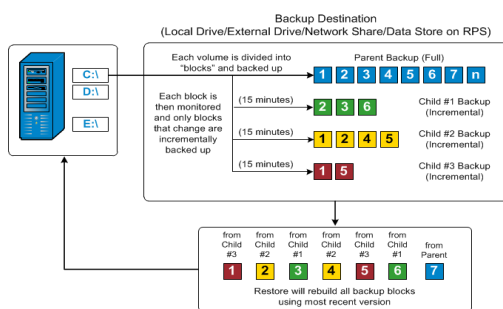
If you need to restore the volume information, the most recent backed up version of each block is located and the entire volume is rebuilt using these current blocks.

Arcserve Unified Data Protection - Full:

Available backup destinations include: internal drive, external drive, remote network share, or a data store on a Recovery Point Server (RPS). When you create a Plan from the Arcserve UDP server, you can select Data Store on Recovery Point Server as the destination and then deploy the plan to the agent node.

Arcserve Unified Data Protection - Agent:

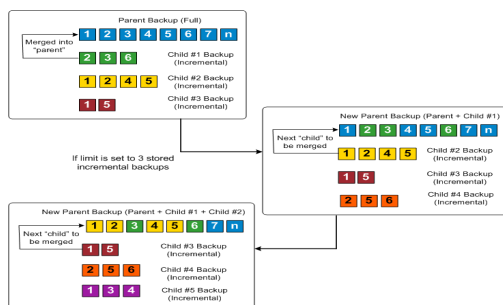
Available backup destinations include: internal drive, external drive, or a remote network share.



How Infinite Incremental Backups Work

If left alone, the incremental snapshots (backups) would continue, as often as 96 times each day (every 15 minutes). These periodic snapshots will accumulate a large chain of backed up blocks to be monitored each time a new backup is performed, and require added space to store these ever-growing backup images. To minimize this potential problem, Arcserve UDP Agent (Windows) utilizes the Infinite Incremental Backup process, which intelligently creates incremental snapshot backups forever (after the initial full backup) and uses less storage space, performs faster backups, and puts less load on your production servers. Infinite Incremental Backups allow you to set a limit for the number of incremental child backups to be stored. When the **Backup Data Format** is **Standard**, configure the **Recovery Points** option from the **Protection Settings** tab on the **Backup Settings** dialog. When the **Backup Data Format** is **Advanced** (default), configure the **Recovery Points** option from the **Schedule** tab on the **Backup Settings** dialog.

When the specified limit is exceeded, the earliest (oldest) incremental child backup is merged into the parent backup to create a new baseline image consisting of the "parent plus oldest child" blocks (unchanged blocks will remain the same). This cycle of merging the oldest child backup into the parent backup repeats for each subsequent backup, allowing you to perform Infinite Incremental (I2) snapshot backups while maintaining the same number of stored (and monitored) backup images.

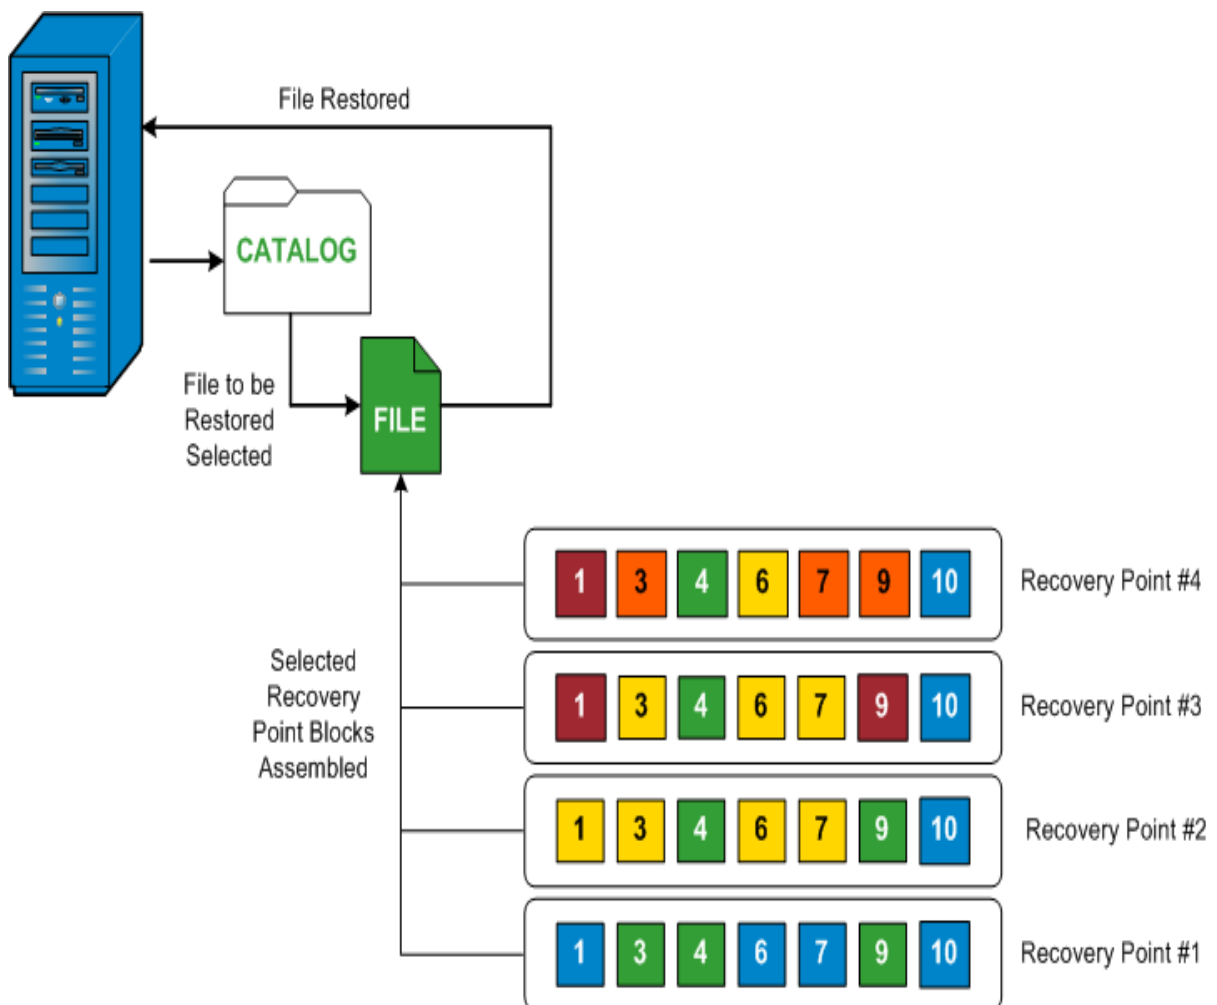


How File Level Restores Work

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:



How Verify Backups Work

Every so often (as scheduled or when manually initiated), Arcserve UDP Agent (Windows) can perform a Verify (resynchronization) type backup to provide a confidence check of the stored backup image and resynchronize that image if necessary. A Verify type 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 Agent (Windows) refreshes (resynchronizes) the backup of the block that does not match.

A Verify backup can also be used to get the same guarantee as a full backup without taking the space of full backup. The advantage of a Verify backup is that it is small when compared to full backup because only the changed blocks (blocks that do not match the last backup) are backed up. However, a Verify backup is also slower than an Incremental backup because Arcserve UDP Agent (Windows) has to compare all of source disk blocks with the blocks of the last backup.



How Recovery Sets Work

A Recovery Set is a storage setting where a group of recovery points (backup sessions) are backed-up over a specified period and then stored together as one compiled set. A Recovery Set includes a series of backups, always starting with a Full Backup, and then followed by a number of subsequent Incremental, Verify, or Full Backups. By using Recovery Sets (instead of Recovery Points), you disable infinite incremental backups and discontinue merging of backup sessions, thereby completely eliminating the time-consuming merge process.

Recovery Sets are generally used for large storage environments and helps you to manage your backup window time more efficiently when protecting large amounts of data. Recovery Sets are used when the backup time is more important than storage space constraints.

A Full Backup is required to start a Recovery Set. Therefore, the backup session that starts a Recovery 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 initial Full Backup is completed, all subsequent backups (regardless if which type of backup is performed) will be saved within the Recovery Set until the next new Recovery Set is launched (manually or automatically as scheduled).

You can configure the number of Recovery Sets to retain. When the number of Recovery Sets retained exceeds the specified retention count, the merge job deletes the oldest recovery set. A recovery set is considered complete only when the starting Full Backup for the next Recovery Set is completed. For example, if you specified to retain two Recovery Sets, Arcserve UDP Agent (Windows) deletes the first Recovery Set only after the Full Backup for the fourth Recovery Set is completed. This ensures that when the first backup is deleted, you already have two Recovery Sets (Recovery Set 2 and Recovery Set 3) retained on disk.

Notes:

- After reaching the retention count, the merge job gets triggered and the oldest Recovery Set gets deleted.
- If you want to delete a recovery set to save backup storage space, reduce the number of retained sets and Arcserve UDP Agent (Windows) automatically deletes the oldest recovery set. Do not attempt to delete the recovery set manually.

A flag in the status column on the Arcserve UDP Agent (Windows) home page **Most Recent Events** section indicates that a full backup is the starting backup of a recovery set. After the recovery set setting is changed (for example, changing the recov-

ery 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: Recovery sets are only available when using Arcserve UDP Agent (Windows) and you set the **Backup Data Format** to **Standard**. Recovery sets are not available if you set the **Backup Data Format** to **Advanced**. This is because merge jobs are very fast and efficient when using the **Advanced Backup Data Format**, therefore eliminating the need for recovery sets.

Default: 2

Minimum: 1

Maximum: 100

Example 1 - Retain 1 Recovery Set:

- Specify the number of recovery sets to retain as 1.

Arcserve UDP Agent (Windows) deletes the first recovery set when the third recovery set Full backup is completed.

Note: Even if you choose to retain only one recovery set, you need space for at least two full backups.

Example 2 - Retain 2 Recovery Sets:

- Specify the number of recovery sets to retain as 2.

Arcserve UDP Agent (Windows) deletes the first recovery set when the fourth recovery set full backup is completed. This ensures that when the first backup is deleted and the fourth recovery set Full backup is completed, you still have two recovery sets (recovery set 2 and recovery set 3) available on disk.

Example 3 - Retain 3 Recovery Sets:

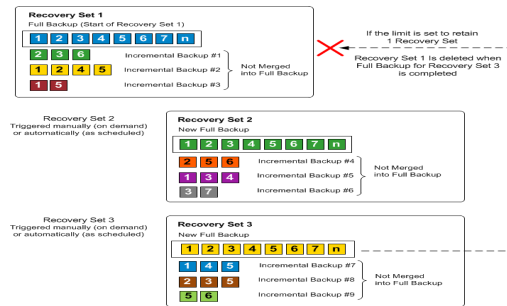
- Specify the number of recovery sets to retain as 3.

Arcserve UDP Agent (Windows) deletes the first recovery set when the fifth recovery set Full backup is completed.

- The backup start time is 6:00 AM, August 20, 2012.
- An incremental backup runs every 12 hours.
- A new recovery set starts at the last backup on Friday.
- You want to retain 3 recovery sets.

With the above configuration, an incremental backup runs 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

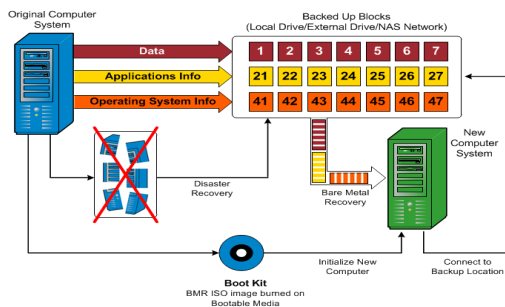
recovery set. When the backup scheduled at 6:00 PM on Friday runs, it will be converted to a full backup and marked as the starting backup of the recovery set.



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).


How Arcserve UDP Agent (Windows) Update Works

A product update lets Arcserve deliver product improvements to users. Updates are used to deliver bug fixes, new hardware support, and performance and security enhancements. Within Arcserve UDP Agent (Windows), the Updates function simplifies this process and provides a fast, easy, and reliable solution to keep your Arcserve UDP Agent (Windows) installation up-to-date with the latest available updates. The Updates function is your link between Arcserve and your Arcserve UDP Agent (Windows) installation.

Arcserve UDP Agent (Windows) Updates provide the following functions:

- Check for available updates (manual or scheduled)
- Download available updates from Arcserve (either directly to a client machine or to a staging server first and then to a client machine)
- Install successfully downloaded updates (to be initiated manually)
- Send email notifications when a new update is available

Note: When Arcserve UDP Agent (Windows) is managed by the Arcserve UDP Console, checking for updates is disabled from Arcserve UDP Agent (Windows). You need to check and deploy the update from the Arcserve UDP Console UI.



Check for Updates

When **Arcserve Server** is selected as download server, Arcserve UDP Agent (Windows) Updates provides the capability to connect to the Arcserve server either directly or using a proxy server (as configured manually) to check for new and available Arcserve UDP Agent (Windows) updates. The Arcserve UDP Agent (Windows) will directly connect to Arcserve server using the proxy settings configured by browser (only applicable for IE and Chrome). This check for updates function can be triggered either manually from the user interface/tray monitor or automatically as specified by the Scheduler. (The internal Scheduler is responsible for starting at a scheduled day and time and triggering an automatic check and download of available updates).

When triggered, the update manager contacts the Arcserve server to check the date /time stamp of a file that contains the available update information. If this available update information file has been modified since the last check, it will be downloaded from the server for comparison. The available update information is then compared to another file containing the already downloaded update information to determine if the available update is new and has not been previously downloaded. If the latest available update is not installed on your machine, Arcserve UDP Agent (Windows) displays an icon on the home page to inform you that a new update is available. In addition, an email notification can also be sent to inform you when a new Arcserve UDP Agent (Windows) update is available for downloading.

When **Staging Server** is selected as download server, Arcserve UDP Agent (Windows) downloads the available update information file from the staging server, and perform the same comparison check with the already available update information file. If the latest available update is not installed on your machine, Arcserve UDP Agent (Windows) will display an icon on the home page to inform you that a new update is available.

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

Download Updates

Arcserve UDP Agent (Windows) Updates provide the capability to download available Arcserve UDP Agent (Windows) updates either directly from the Arcserve server or from a staging server which in turn connects to the Arcserve server. This download process is triggered automatically when the check for updates process determines that a new update is available (unless this auto-download function is disabled). You can configure Arcserve UDP Agent (Windows) to download an update directly (or using a proxy server) to your client machine or to a staging server. A staging server can be used as a temporary storage location for downloading an update before it is downloaded and installed into an Arcserve UDP Agent (Windows) client machine. You may not want to expose your client machine to the internet to download updates from the Arcserve server. In this case, you can first download the update to a staging server and then allow other client machines to download the update from that staging server. The Arcserve UDP Agent (Windows) provides the capability to configure multiple staging servers for downloading the updates. If for some reason the primary staging server is unavailable, the download function will automatically transfer to the next specified staging server.

Note: If you are using a staging server for your Updates downloads, Arcserve UDP Agent (Windows) must be installed on that staging server, but does not need to be licensed unless you are using Arcserve UDP Agent (Windows) to protect that staging server.

When triggered, the Updates function contacts the Arcserve server and downloads the available update and places it in a holding directory (on either the staging server or the client machine) until directed to proceed with the subsequent installation process.

The default location for the download folder is: <Product Home>\Update Manager\EngineUpdates\7.0\

If for some reason, the download cannot be started, a popup message is displayed and Arcserve UDP Agent (Windows) waits a specified number of minutes and then attempt to download again. If after a specified number of retry attempts, the download still cannot continue, an error message will be displayed in the activity log indicating the most likely reason for the failure.

Install Updates

Arcserve UDP Agent (Windows) Updates provide the capability to install the available and successfully downloaded updates. This install process can only be triggered manually from the user interface/tray monitor (not automatically). When triggered, the update is installed from the holding directory to the applicable Arcserve UDP Agent (Windows) component directory of the client machine or the staging server. You cannot trigger the installation of the update directly from a staging server to a client machine. When you click install, the update is downloaded from the staging server to the client machine (if it has not been downloaded already), and then the installation process is triggered from the client machine.

Note: The installation only continues if no other active Arcserve UDP Agent (Windows) jobs are running. If another job is running, a message is displayed informing you of this condition and requesting that you try again at a later time.

If the installation is successful, the file containing the status information is updated for future use.

If the installation fails, an error message is displayed indicating the most likely reason for the failure.

Note: During the update installation Arcserve UDP Agent (Windows) will stop the Arcserve UDP Agent (Windows) Web service and will restart this web service after successful installation of update.

Email Notifications

Arcserve UDP Agent (Windows) Updates provide the capability to send automatic email notifications when a new update is available. The Arcserve UDP Agent (Windows) connects to an SMTP server (with appropriate credentials) to enable sending these email notifications over the Internet from Arcserve to your server. (The email recipients are specified from the **Preferences** dialog).

In addition, email notifications are also sent if a failure occurs during the check for updates or during the download.

Chapter 2: Installing/Uninstalling Arcserve UDP Agent (Windows)

This section contains the following topics:

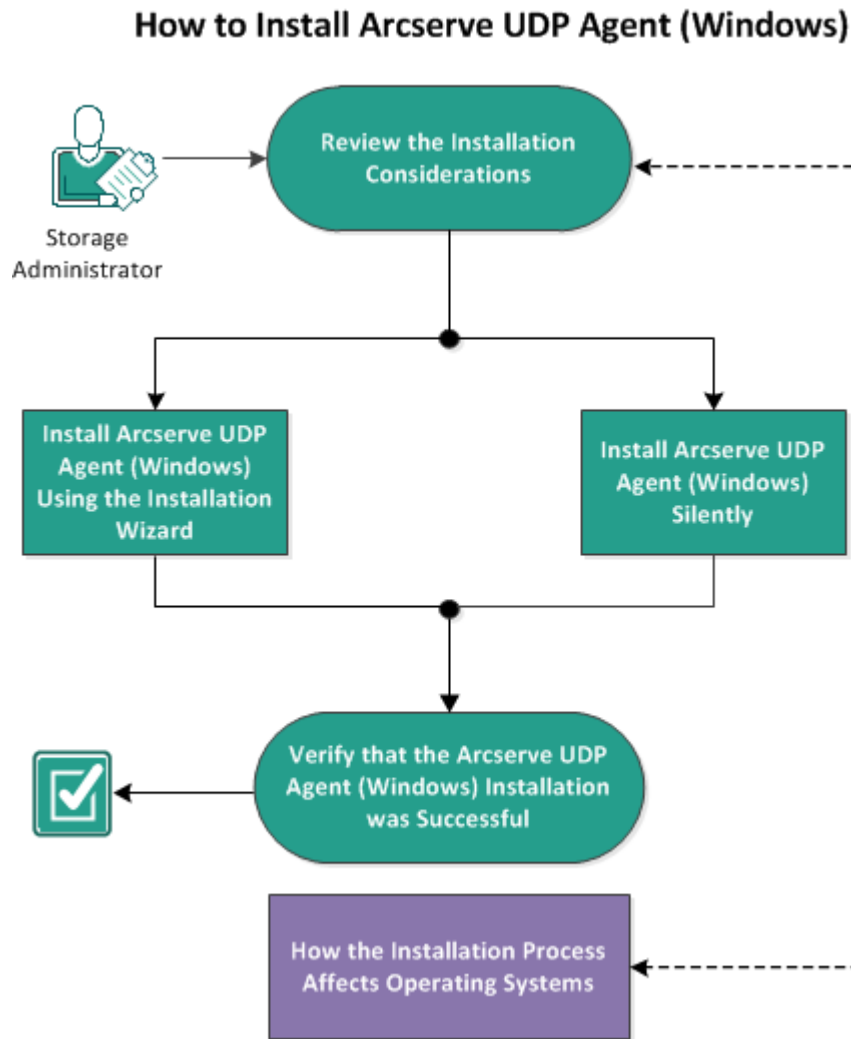
How to Install Arcserve UDP Agent (Windows)	44
How to Install Arcserve UDP Agent (Windows) Updates	75
How to Uninstall Arcserve UDP Agent (Windows)	98
UDP Workstation Free	115

How to Install Arcserve UDP Agent (Windows)

Installing Arcserve UDP Agent (Windows), a disk-based backup, allows you to protect and recover critical business information in a fast, simple, and reliable way. The Arcserve UDP Agent (Windows) is a light-weight solution for tracking changes on a machine at the block level and then backing up only those changed blocks in an incremental method. As a result, Arcserve UDP Agent (Windows) lets you perform frequent backups for every 15 minutes, reducing the size of each incremental backup as well as the backup window and provides a more up-to-date backup. The Arcserve UDP Agent (Windows) also provides the capability to restore files, folders, volumes, and applications, and perform bare metal recovery from a single backup. In addition, Arcserve UDP Agent (Windows) also lets you copy and restore backed-up data from your specified cloud storage location.

Note: 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 [Workstation Free Edition](#) with limited capabilities.

The following diagram illustrates the process to install Arcserve UDP Agent (Windows):



Perform the following tasks to install Arcserve UDP Agent (Windows):

1. [Review the Installation Considerations](#)
2. [Install Arcserve UDP Agent \(Windows\) Using the Installation Wizard](#)
3. [Install Arcserve UDP Agent \(Windows\) Silently](#)
4. [Verify that the Arcserve UDP Agent \(Windows\) Installation was Successful](#)
5. [\(Optional\) How the Installation Process Affects Operating Systems](#)

Review the Installation Considerations

Review the following installation considerations before installing Arcserve UDP Agent (Windows):

- The Arcserve UDP Agent (Windows) installation package is available through a web download and from the product installation CD.

Note: 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 [Workstation Free Edition](#) with limited capabilities.

- Verify that you have administrator privileges or the proper permissions to install software on the servers where you are installing Arcserve UDP Agent (Windows).
- If you uninstall and install a new Arcserve UDP Agent (Windows) build and specify the same backup destination as the previous build, the first backup after the installation runs as a Verify backup.

Note: After deploying the agent, you do not need to reboot to start backup. For details, refer to [Reboot Not Required After Agent Deployment](#).

- After installation, you can configure your antivirus software to exclude specific processes, folders, and files so that the antivirus software does not interfere with the proper operation of Arcserve UDP Agent (Windows). For a complete list of processes, folders, and files that must be excluded, see [Antivirus Configuration](#).
- If you install Arcserve UDP Agent (Windows) on a x64 Windows Core Operating System, you must also install Windows-on-Windows 64-bit (WOW64) on the Server Core for the Arcserve UDP Agent (Windows) setup to work.
- For a list of the possible error codes that the Arcserve UDP Agent (Windows) installer could return, see [Arcserve UDP Agent \(Windows\) Installer Error Codes](#).
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Install Arcserve UDP Agent (Windows) Using the Installation Wizard

This section describes installing Arcserve UDP Agent (Windows) on your local system using the installation wizard. The installation wizard is an interactive application that guides you through the installation process.

Follow these steps:

1. Access the Arcserve UDP Agent (Windows) installation package (arcserve_Unified_Data_Protection_Agent_Windows.exe) from either the Arcserve website or the product CD.

Notes:

If the installation is performed using the web downloaded installation package, the contents of the package are extracted to your local system.

If one of the supported non-English operating systems is detected, select the language for the product installation.

The **License Agreement** dialog opens.

2. Read and accept the terms of the Licensing Agreement on the **License Agreement** dialog and click **Next**.

The **Installation Type** dialog opens.

3. Select **Arcserve Unified Data Protection - Agent** from the available components to install.

The **Arcserve Unified Data Protection - Agent** installs only the Arcserve UDP Agent.

For instructions about how to install **Arcserve Unified Data Protection - Full**, see [Install Arcserve UDP Using the Setup Wizard](#) in the Solutions Guide.

The **Arcserve Unified Data Protection - Full** installs Arcserve UDP Console, Recovery Point Server, and Agent.

4. 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 verify/incremental backup.
- With this driver installed, you must have a valid Arcserve UDP Agent (Windows) license to perform a local backup.

- This driver is not required if this agent is used as Virtual Standby monitor or host-based VM backup proxy server.

Note: You can install this driver at any time after the installation is complete by running the InstallDriver.bat utility from the following location: <Arcserve Agent install folder>\Engine\BIN\DRIVER

The **Destination Folder** dialog opens.

5. Specify or browse to the location where you want to install Arcserve UDP Agent (Windows) and click **Next**.

Default location: C:\Program Files\Arcserve\Unified Data Protection\

Note: During Arcserve UDP Agent (Windows) installation, some files do not get installed at the default location. For a complete listing of these files, see [Installation of Files Outside the Default Location](#).

The **Configuration** dialog opens.

6. Enter the following information in the **Configuration** dialog:
 - a. Specify if you want to use HTTP or HTTPS for web communication.

Note: You can change the communication protocol at any time after installation. If you are concerned about the security of information that is communicated between these components including passwords, you can select this option to change the protocol being used to Hypertext Transfer Protocol Secure (HTTPS). To use the SSL protocol for a hostname that contains an underscore (_) character, you must manually run the following batch file before using Arcserve UDP Agent or Console:

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

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

- b. Specify the **Agent port number**. This port number is used to connect to the web-based UI.

Default Value: 8014.

Note: The available port numbers for Arcserve UDP Agent (Windows) installation are between 1024 and 65535. You must verify that the specified port number is free and available for use. Setup will not let you install Arcserve UDP Agent (Windows) for a port that is not available for use.

- c. Type the Windows Administrator Name and Password.
- d. Specify if you want to display the Arcserve UDP Agent monitor for all users or only the current user.

7. Click **Next**.

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

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

8. 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 summary** dialog is displayed and automatically performs the product configuration.

9. (Optional) Select the **Check for an update immediately** checkbox, to check if there are any product updates since the last release.

This option is checked by default.

10. (Optional) To install the **Arcserve UDP Agent for Linux**, click the provided link and follow the download instructions.

11. Click **Finish**.

When checking for updates, the **Check for Updates** dialog opens where you can download the updates from the Arcserve server or the staging server.

12. Click **Download and Install Updates**.

13. Click **Finish**.

An alert message is displayed, informing you that a system restart is required and asking if you want to reboot now or at a later time.

When the reboot is finished, Arcserve UDP Agent (Windows) is installed on your local system.

Note: You can access Arcserve UDP Agent (Windows) from either the Start menu or from the Arcserve UDP Agent (Windows) Monitor.

After the installation is complete, as a best practice, create a BMR ISO image using the Create Boot Kit utility. For more information about the BMR ISO image, see

[How to Create a Boot Kit](#).

Install Arcserve UDP Agent (Windows) Silently

You can install Arcserve UDP Agent (Windows) silently. The silent installation allows you to perform an unattended installation and does not prompt you for any input, which eliminates the need for user interaction. The silent installations are used when performing similar installations on multiple computers.

Use the Windows Command Line to silently install the application.

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 and use the following command to start the silent installation process:

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

Example:

```
"arcserve_Unified_Data_Protection_Agent_Windows.exe" -s -a -q -Products:Agent -User:administrator -Password:Password01
```

3. To configure the silent installation, use 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 must be -Password:"abc^*123".

-s

Specifies you to run the executable file package using the silent mode.

-a

Specifies any additional command line options.

-q

Specifies you to install the application in silent mode.

-Products:<ProductList>

Specifies the components to install silently. You can specify the following components:

Agent: Installs the Arcserve UDP Agent component.

Example:

Install Arcserve UDP Agent

-Products:Agent

-User:<UserName>

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

Note: The user name must be 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 the Arcserve UDP Agent.

Example:

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

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

-AgentPort:<Port Number>

(Optional) Specifies the communication port number for the 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 the 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 Arcserve Universal Agent service.

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

1: Stops the Arcserve Universal Agent service if it runs during the installation process.

Default: 0

Example:

-StopUA:1

Note: Use this option while you upgrade it to a new version. Verify that you set the value to 1 or stop the service before you start the upgrade process. This process ensures 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 that the Arcserve UDP Agent (Windows) Installation was Successful

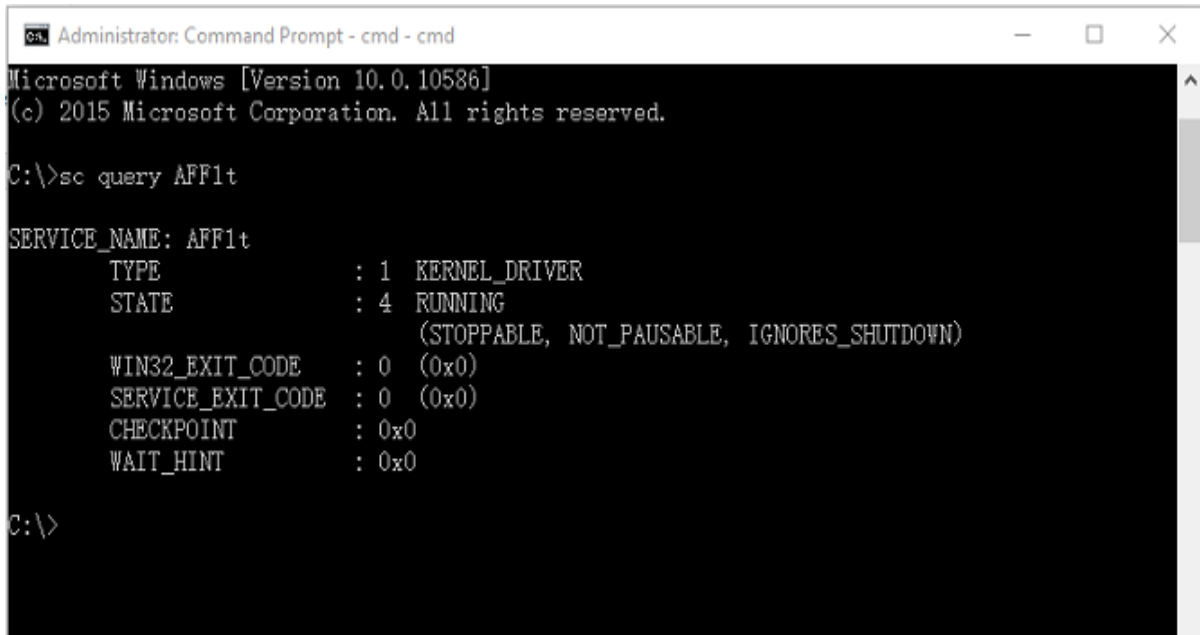
Follow these steps:

1. Verify that the Agent icon appears in the system tray.
2. From the command prompt tab, navigate to services.msc, and then click **OK**.
3. Verify that the Agent services is up and running from the Services Manager.



4. Open the command prompt window and type the following driver name to verify that the state is running:

sc query afflt



```
Administrator: Command Prompt - cmd - cmd
Microsoft Windows [Version 10.0.10586]
(c) 2015 Microsoft Corporation. All rights reserved.

C:\>sc query AFF1t

SERVICE_NAME: AFF1t
        TYPE               : 1  KERNEL_DRIVER
        STATE                : 4  RUNNING
                        (STOPPABLE, NOT_PAUSABLE, IGNORES_SHUTDOWN)
        WIN32_EXIT_CODE      : 0  (0x0)
        SERVICE_EXIT_CODE  : 0  (0x0)
        CHECKPOINT          : 0x0
        WAIT_HINT           : 0x0

C:\>
```

The Arcserve UDP Agent (Windows) is installed successfully.

How the Installation Process Affects Operating Systems

The Arcserve UDP Agent (Windows) installation process updates various Windows operating system components using an installation engine named the Microsoft Installer Package (MSI). The MSI components let Arcserve UDP Agent (Windows) perform custom actions such as upgrading, installing, or uninstalling Arcserve UDP Agent (Windows).

The following components describe the custom actions that you can perform:

Note: All Arcserve UDP Agent (Windows) MSI packages call the components listed in the following description when you install and uninstall Arcserve UDP Agent (Windows).

CallAllowInstall

Allows the installation process to check for conditions relating to the current Arcserve UDP Agent (Windows) installation.

CallPreInstall

Allows the installation process to read and write MSI properties. For example, read the Arcserve UDP Agent (Windows) installation path from the MSI.

CallPostInstall

Allows the installation process to perform various tasks relating to installation. For example, registering Arcserve UDP Agent (Windows) into the Windows Registry.

CallAllowUninstall

Lets the uninstallation process to check for conditions relating to the current Arcserve UDP Agent (Windows) installation.

CallPreUninstall

Lets the uninstallation process to perform various tasks relating to uninstallation. For example, un-registering Arcserve UDP Agent (Windows) from the Windows Registry.

CallPostUninstall

Lets the uninstallation process to perform various tasks after the installed files are uninstalled. For example, removing the remaining files.

ShowMsiLog

Displays the Windows Installer log file in Notepad if the end-user selects the Show the Windows Installer log check box in the SetupCompleteSuccess,

SetupCompleteError, or SetupInterrupted dialogs and then clicks **Finish**. This component works only with Windows Installer 4.0.

ISPrint

Prints the contents of a ScrollableText control on a dialog.

This component is a Windows Installer .dll custom action. The name of the .dll file is SetAllUsers.dll, and its entry point is PrintScrollableText.

CheckForProductUpdates

Uses FLEXnet Connect to check for product updates.

This custom action launches an executable file named Agent.exe, and it passes the following:

```
/au[ProductCode] /EndOfInstall
```

CheckForProductUpdatesOnReboot

Uses FLEXnet Connect to check for product updates on reboot.

This custom action launches an executable file named Agent.exe, and it passes the following:

```
/au[ProductCode] /EndOfInstall /Reboot
```

Directories Updated

The installation process installs and updates the Arcserve UDP Agent (Windows) files in the following directories.

Note: The default Operating Systems are x86 and x64.

C:\Program Files\Arcserve\Unified Data Protection\Engine

You can install Arcserve UDP Agent (Windows) into the default installation directory or into an alternate directory. The installation process copies various system files to the following directory:

C:\WINDOWS\SYSTEM32

Windows Registry Keys Updated

The installation process updates the following Windows registry keys:

- Default registry keys:
 - HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine
- The installation process creates new registry keys and modifies various other registry keys, based on the current configuration of your system.

Applications Installed

The installation process installs the following applications into your computer:

- Arcserve Licensing
- Microsoft Visual C++ 2013
- Java Runtime Environment (JRE) 1.8.0_65
- Tomcat 9

The following installation processes update various Windows Operating Systems:

Antivirus Configuration

Antivirus software can interfere with the functioning of Arcserve UDP Agent (Windows). For example, the Antivirus software can temporarily block access to files or quarantine or delete files that are incorrectly classified as suspicious or dangerous. You can configure the antivirus software to exclude particular processes, files, or folders so that you are not scanning data that does not need protection. For Arcserve UDP Agent (Windows), you must configure the antivirus software properly so that it does not interfere with backup and restore operations or any other processes such as merge and catalog generation.

Paths to be excluded for antivirus scanning:

- Backup Destination

Installation of Files Outside the Default Location

By default, Arcserve UDP Agent (Windows) is installed at the following location:

C:\Program Files\Arcserve\Unified Data Protection\Engine.

However, some Arcserve UDP Agent (Windows) files are installed outside this default folder.

File Path and Name	Reason
C:\Windows\Downloaded Installations\{D03BF724-4E4F-4DF4-A1BD-8497634F5589}\ASLicense.msi	Arcserve Licensing Component (shared)
C:\Windows\Downloaded Installations\{D03BF724-4E4F-4DF4-A1BD-8497634F5589}\1033.MST	Arcserve Licensing Component (shared)
C:\Windows\inf\oem9.inf (The digit number in file name may be changed in different computer)	Installed by mount driver in recommended location
C:\Windows\inf\oem9.PNF (The digit number in file name may be changed in different computer)	Installed by mount driver in recommended location
C:\Windows\inf\oem10.inf (The digit number in file name may be changed in different computer)	Installed by Interface driver in recommended location
C:\Windows\inf\oem10.PNF (The digit number in file name may be changed in different computer)	Installed by Interface driver in recommended location
C:\Windows\System32\drivers\AFStorHBA.sys	Installed by mount driver in recommended location
C:\Windows\System32\drivers\ARCFlashVolDrv.sys	Installed by Volume driver in recommended location
C:\Windows\System32\drivers\UMDF\AFStorHBATramp.dll	Installed by Interface driver

	in recommended location
C:\Windows\System32\DriverStore\FileRepository\afstorhba.inf_amd64_neutral_23f49884ad235baf\AFStorHBA.cat	Installed by mount driver in recommended location
C:\Windows\System32\DriverStore\FileRepository\afstorhba.inf_amd64_neutral_23f49884ad235baf\afstorhba.inf	Installed by mount driver in recommended location
C:\Windows\System32\DriverStore\FileRepository\afstorhba.inf_amd64_neutral_23f49884ad235baf\afstorhba.PNF	Installed by mount driver in recommended location
C:\Windows\System32\DriverStore\FileRepository\afstorhba.inf_amd64_neutral_23f49884ad235baf\AFStorHBA.sys	Installed by mount driver in recommended location
C:\Windows\System32\DriverStore\FileRepository\afstorhba.inf_amd64_neutral_23f49884ad235baf\WdfCoinstaller01009.dll	Installed by mount driver in recommended location
C:\Windows\System32\DriverStore\FileRepository\afstorhbatramp.inf_amd64_neutral_c8c319207a86e457\AFStorHBATramp.cat	Installed by Interface driver in recommended location
C:\Windows\System32\DriverStore\FileRepository\afstorhbatramp.inf_amd64_neutral_c8c319207a86e457\AFStorHBATramp.dll	Installed by Interface driver in recommended location
C:\Windows\System32\DriverStore\FileRepository\afstorhbatramp.inf_amd64_neutral_c8c319207a86e457\afstorhbatramp.inf	Installed by Interface driver in recommended location
C:\Windows\System32\DriverStore\FileRepository\afstorhbatramp.inf_amd64_neutral_c8c319207a86e457\afstorhbatramp.PNF	Installed by Interface driver in recommended location
C:\Windows\System32\DriverStore\FileRepository\afstorhbatramp.inf_	Installed by

amd64_neutral_c8c319207a86e457\WudfUpdate_01009.dll	Interface driver in recommended location
C:\Windows\System32\WdfCoinstaller01009.dll	Installed by BMR in recommended location
C:\Windows\System32\WudfUpdate_01009.dll	Installed by BMR in recommended location
C:\Windows\System32\atl100.dll	Microsoft Visual C++ component
C:\Windows\System32\mfc100.dll	Microsoft Visual C++ component
C:\Windows\System32\mfc100chs.dll	Microsoft Visual C++ component
C:\Windows\System32\mfc100cht.dll	Microsoft Visual C++ component
C:\Windows\System32\mfc100deu.dll	Microsoft Visual C++ component
C:\Windows\System32\mfc100enu.dll	Microsoft Visual C++ component
C:\Windows\System32\mfc100esn.dll	Microsoft Visual C++ component
C:\Windows\System32\mfc100fra.dll	Microsoft Visual C++ component
C:\Windows\System32\mfc100ita.dll	Microsoft Visual C++ component
C:\Windows\System32\mfc100jpn.dll	Microsoft Visual C++ component
C:\Windows\System32\mfc100kor.dll	Microsoft Visual C++ component
C:\Windows\System32\mfc100rus.dll	Microsoft Visual C++ component
C:\Windows\System32\mfc100u.dll	Microsoft Visual C++ component
C:\Windows\System32\mfcm100.dll	Microsoft Visual C++ component
C:\Windows\System32\mfcm100u.dll	Microsoft Visual C++ component
C:\Windows\System32\msvcp100.dll	Microsoft Visual C++ component

C:\Windows\System32\msvcr100.dll	Microsoft Visual C++ component
C:\Windows\System32\vcomp100.dll	Microsoft Visual C++ component

Installation of Unsigned Binary Files

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

Binary Name	Source
libbind9.dll	Arcserve Replication and High Availability
libdns.dll	Arcserve Replication and High Availability
libisc.dll	Arcserve Replication and High Availability
libiscfg.dll	Arcserve Replication and High Availability
liblwres.dll	Arcserve Replication and High Availability
win_nsupdate.exe	Arcserve Replication and High Availability
libeay32.dll	OpenSSL
msvcm90.dll	Microsoft
msvcp90.dll	Microsoft
msvcr90.dll	Microsoft
ssleay32.dll	OpenSSL
zlib10.dll	Zlib Compression Library
tcnative-1.dll	Tomcat
tomcat9.exe	Tomcat
UpdateData.exe	Arcserve License

Installation of Binary Files Containing Incorrect File Version Information

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

Binary Name	Source
libbind9.dll	Arcserve Replication and High Availability
libdns.dll	Arcserve Replication and High Availability
libisc.dll	Arcserve Replication and High Availability
libiscfg.dll	Arcserve Replication and High Availability
liblwres.dll	Arcserve Replication and High Availability
win_nsupdate.exe	Arcserve Replication and High Availability
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
gststreamer-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
gvmomi.dll	VMware
libcurl.dll	VMware
liblber.dll	VMware
libldap.dll	VMware
libldap_r.dll	VMware
libxml2.dll	VMware
zlib1.dll	Zlib Compression Library
zlib10.dll	Zlib Compression Library
UpdateData.exe	Arcserve License

Installation of Binary Files that Do Not Contain an Embedded Manifest

Arcserve UDP Agent (Windows) installs binary files that are developed by third parties, other Arcserve products, and Arcserve UDP Agent (Windows) that do not contain an embedded manifest and a text manifest. The following table describes these binary files:

Binary Name	Source
arcserve_Unified_Data_Protection_Agent_Windows.exe	Arcserve UDP Agent (Windows)
ARCFlashVolDrvINSTALL.exe	Arcserve UDP Agent (Windows)
BaseLicInst.exe	Arcserve License
UpdateData.exe	Arcserve License
vcredist_x64.exe	Microsoft
vcredist_x86.exe	Microsoft
tomcat9.exe	Tomcat

Installation of Binary Files That Require a Privilege Level of Administrator in Manifest

Arcserve UDP Agent (Windows) installs binary files developed by third parties, other Arcserve products, and Arcserve UDP Agent (Windows) that have a privilege level of Administrator or Highest Available. Log in using an administrative account or an account with the highest available permissions to run various Arcserve UDP Agent (Windows) services, components, and applications. The binaries corresponding to these services, components, and applications contain Arcserve UDP Agent (Windows) specific functionality that is not available to a basic user account. As a result, Windows prompts you to confirm an operation with your password or using an account with administrative privileges to complete the operation.

Administrative Privileges

Specifies that the administrative profile or an account with administrative privileges has read, write, and execute permissions to all Windows and system resources. If you do not have Administrative privileges, you are prompted to enter the user name / password of an administrator user to continue.

Highest Available Privileges

Specifies that an account with the highest-available privileges is a basic user account and a power user account with run-as administrative privileges.

The following table describes these binary files:

Binaries	Source
afbkw.exe	Arcserve UDP Agent (Windows)
AFBackend.exe	Arcserve UDP Agent (Windows)
Asremsvc.exe	Arcserve UDP Agent (Windows)
DeleteMe.exe	Arcserve UDP Agent (Windows)
MasterSetup.exe	Arcserve UDP Agent (Windows)
SetupFW.exe	Arcserve UDP Agent (Windows)
setup.exe	Arcserve UDP Agent (Windows)
silent.exe	Arcserve License
jbroker.exe	Java Runtime Environment
jucheck.exe	Java Runtime Environment

Installation of Driver API for Non-WDM Driver

Arcserve UDP Agent (Windows) implements "SetupInstallServicesFromInfSection" API to install the non-WDM driver.

User Mode Driver Framework

Arcserve UDP Agent (Windows) uses the "WUDFRd.sys" driver published by Microsoft and part of UMDF (User Mode Driver Framework). This driver is not WHQL (Windows Hardware Quality Labs) signed.

Arcserve UDP Agent (Windows) Installer Error Codes

The following table displays a list of the possible error codes that the Arcserve UDP Agent (Windows) installer could return:

Error Code	Description	Action
0	Install Successfully.	No action is required.
3010	A restart is required to complete the install.	Restart your system.
80000	Setup could not run. Another instance of Setup is running.	Quit and wait for the previous Setup to finish.
80009	The resource DLL is not found in current directory.	Download the package again and run Setup with the new downloaded package.
80015	Setup could not find setup.icf.	Quit and run Setup again.
80016	You must have administrative privileges to install this product.	Run Setup with an administrator account.
80018	Setup package cannot support the operating system on the target host.	Check the supported OS matrixes and use the full package to run Setup.
80031	Setup cannot continue. Setup detected a newer version of Arcserve UDP Agent (Windows) installed on this computer that you cannot upgrade.	Before you can install the current version of this product, you must uninstall the previous version from the target computer.
80032	Setup detected that the same version of Arcserve UDP Agent (Windows) is installed on this computer.	Quit Setup and use the installed product.
80043	The same version, a newer version, or a version of this product that is not supported is installed on the target computer.	Before you can install the current version of this product, you must uninstall the previous version from the target computer.
80044	Setup upgraded critical system files and a restart is required to continue.	Restart the computer and run Setup again.
80046	Internal error. Failed to launch process.	Reboot the machine and run Setup again.
80049	Setup requires Windows XP SP3 or later versions.	Check current OS and the supported OS matrixes.
80050	Arcserve UDP Agent (Windows) cannot be installed on Itanium-based system machines.	Check the supported OS matrixes and run Setup on another machine.
80051	Setup cannot continue. Setup detected jobs running on the target computer.	You must stop all jobs that are running on the target computer and then restart the installation.

80052	The specified installation path is invalid on target host.	<p>Please check the following and try again:</p> <ul style="list-style-type: none"> - Use the local disk directory and exclude the floppy disk, cd-rom, and mapped drive. - The path must not contain any special or non-English characters. - The path must not have a Read-Only attribute.
80053	There is insufficient amount of free space on the selected drive to complete this installation.	You must free space on the selected drive or specify a different installation path.
80057	An internal error has occurred. Failed to encrypt information.	Reboot the machine and run Setup again.
80058	An internal error has occurred. Failed to decrypt information.	Reboot the machine and run Setup again.
80060	A computer restart is required. You must restart this computer before you install the Arcserve UDP Agent (Windows).	Restart the computer and run Setup again.
80062	The installation failed. When the installation failed and Setup could not get the failed reason.	Check the Setup log file for detailed reason. Reboot the machine the run Setup again.
80063	Setup detected that the Arcserve Universal Agent service is running on the target computer.	To update Arcserve UDP Agent (Windows), stop the Arcserve Universal Agent service before you continue.
80064	Setup is unable to stop the Arcserve UDP Agent Mount Driver Service.	<p>You need to remove the Arcserve UDP Agent Mount Driver Service with the following steps:</p> <ol style="list-style-type: none"> 1. Find Arcserve UDP Agent (Windows) installation directory. 2. Enter BIN\\Driver directory. 3. Run "UninstallHBADriver.bat". <p>Upon completion of this operation, you must reboot the machine and then rerun setup.</p>
80065	Setup has removed the Arcserve UDP Agent Service and requires rebooting the system.	To continue with Setup, you must restart the system now and then rerun Setup.
80066	Failed to install Arcserve UDP Agent Service.	Reboot the machine and run Setup again.
80067	Failed to install Arcserve UDP Agent Volume	Reboot the machine and run

	Driver.	Setup again.
80068	Failed to install Arcserve UDP Agent Mount Driver.	Reboot the machine and run Setup again.
80069	Failed to install Arcserve UDP Agent Interface Driver.	Reboot the machine and run Setup again.
80070	The port number is invalid. It is being used by another program.	Please input a different value.
80071	The port number is invalid.	Specify a value between 1024 and 65535 for the port number.
80072	This port number is reserved for internal use.	Please input a different value for the port number.
80075	Arcserve UDP Agent (Windows) cannot be installed on a Windows XP (X86) machine unless service pack SP3 (or higher) has also been installed.	Install Windows XP SP3 and then run Setup again.
80076	Arcserve UDP Agent (Windows) cannot be installed on a Windows XP (X64) machine unless service pack SP1 (or higher) has also been installed.	Install Windows XP (x64) SP1 and then run Setup again.
80077	Setup has upgraded the Windows Driver Foundation files. To continue with Setup, you must restart the system and then rerun Setup.	Reboot the machine and run Setup again.
80078	Failed to upgrade the Windows Driver Foundation files.	Check the log file for detailed error: c:\windows\setupapi.log (before VISTA) c:\windows\inf\setupapi.app.log (Vista or later).
81002	Setup cannot continue because Arcserve Central Applications with a different version has been installed on the target host.	Remove Arcserve Central Applications and run Setup again. Or, run Setup with the full package.
81007	Failed to install Arcserve UDP RPS Port Sharing Service.	Reboot the machine and run Setup again.
90000	Failed to extract the setup package. The possible reasons include: 1. Not enough free disk space. 2. Input parameter is invalid. 3. The setup package is invalid.	Free up disk space, verify the input parameter is valid, or verify the setup package is valid.
0xE1010103	Cannot create the single instance event using WinAPI.	Reboot the system and try again.
0xE1010104	Cannot init the installer to write log file using	Verify that the system temp

	WinAPI.	folder exists (for example, C:\Windows\temp)
0xE1010105	Another installer is already running and two or more instances cannot run at the same time.	Wait for the other installer to complete and try again.
0xE1010107	Cannot find the resource file. The package is invalid.	Verify that your update package executable file is the same as the file on the Arcserve server.
0xE1010108	Cannot find the configuration inf file. The package is invalid.	Verify that your update package executable file is the same as the file on the Arcserve server.
0xE1010109	Cannot find the configuration XML file. The package is invalid.	Verify that your update package executable file is the same as the file on the Arcserve server.
0xE101010B	Cannot load the resource file. The package is invalid.	Verify that your update package executable file is the same as the file on the Arcserve server.
0xE101010C	The input parameter is invalid.	Verify that the input parameter is valid.
0xE101010D	The current user does not have administrator privileges. Setup cannot continue.	Verify the current user has administrative privileges.
0xE101020A	Cannot parse the configuration XML file. The package is invalid.	Verify the package is valid.
0xE1010501	Setup has detected that this machine does not meet the necessary requirements to install this update. A compatible version of Arcserve UDP Agent (Windows) is not detected. Note: Arcserve UDP Agent (Windows) is not installed on this machine.	Install a compatible version of Arcserve UDP Agent (Windows).
0xE1010503	Setup has detected that this machine does not meet the necessary requirements to install this update. The update does not apply to the installed version of Arcserve UDP Agent (Windows). Note: Arcserve UDP Agent (Windows) is installed on this machine, but the current update package does not match the installed version of Arcserve UDP Agent (Windows). For example, if the beta build is installed on this machine, and you are trying to apply the GM update, then the setup fails because the GM update build can only be applied to the GM build not for the beta	Verify the current update package is compatible with the installed version of Arcserve UDP Agent (Windows).

	build.	
0xE1010504	Failed to install update because Arcserve UDP Agent (Windows) has detected that the update is already installed on this machine.	No action is required.
0xE1010505	Setup has detected that this machine does not meet the necessary requirements to install this update. A newer version of the update has already been applied.	No action is required.
0xE1010506	Setup has detected that at least one active job is running on the machine. Setup cannot continue.	Stop all running jobs and try setup again.
0xE1010507	Setup has detected that this machine does not meet the necessary requirements to install this update. Not enough free disk space to install this update.	Free up disk space and try setup again.
0xE1010508	Setup has detected that this machine is currently deploying Arcserve UDP Agent (Windows) to another remote machine.	Finish the deployment and try again to launch this update.
0xE1010509	Setup has detected that this machine is currently creating a boot kit.	Finish the boot kit process and try to launch the update again.
0xE1010512	Setup detected that a reboot is required in a previous installation.	Reboot the system and try to launch the update again.
0xE101050A	Setup detected that Arcserve Universal Agent service is running on the target machine.	Stop the Arcserve Universal Agent service first and then try to launch the update again.
0xE101050B	Setup cannot stop the Arcserve Universal Agent service.	Wait for the active job to complete and then try to launch the update again.


How to Install Arcserve UDP Agent (Windows) Updates

The process of getting and installing Arcserve UDP Agent (Windows) updates is a three-part process: checking for a new update, downloading the update, and then installing the update.

Note: All updates that are released for Arcserve UDP Agent (Windows) 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 Agent (Windows) updates:

How to Install Arcserve UDP Agent (Windows) Updates



Perform the following tasks to install Arcserve UDP Agent (Windows) updates:

1. [Review the Considerations for Installing Updates](#)
2. [Specify Updates Preferences](#)
3. [Check for Updates and Download](#)
4. [Install the Arcserve UDP Agent \(Windows\) Updates](#)
5. [Verify that the Updates are Successfully Installed](#)

6. [\(Optional\) Install Arcserve UDP Agent \(Windows\) Updates Silently](#)
7. [\(Optional\) Troubleshooting Update Issues](#)

Review the Considerations for Installing Updates

Review the following considerations before installing Arcserve UDP Agent (Windows) updates:

- 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 Agent (Windows) updates.
- If you are not using Arcserve UDP Agent (Windows) for any function other than just as an Updates staging server, you do not need to have a separate Arcserve UDP Agent (Windows) license for the staging server.
- Verify that the Update preference settings are properly configured for each node.
- Updates can be installed either through the user interface or silently using the command line. For more information about installing Arcserve UDP Agent (Windows) updates silently, see [\(Optional\) Install Arcserve UDP Agent \(Windows\) Updates Silently](#).
- (Optional) Review the topic [How Arcserve UDP Agent \(Windows\) Updates Works](#).
- Review the [Compatibility Matrix](#) that provides the supported operating systems, databases, and browsers.

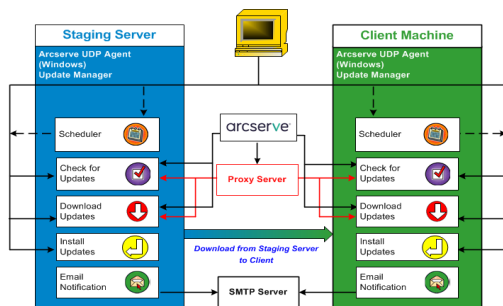
How Arcserve UDP Agent (Windows) Update Works

A product update lets Arcserve deliver product improvements to users. Updates are used to deliver bug fixes, new hardware support, and performance and security enhancements. Within Arcserve UDP Agent (Windows), the Updates function simplifies this process and provides a fast, easy, and reliable solution to keep your Arcserve UDP Agent (Windows) installation up-to-date with the latest available updates. The Updates function is your link between Arcserve and your Arcserve UDP Agent (Windows) installation.

Arcserve UDP Agent (Windows) Updates provide the following functions:

- Check for available updates (manual or scheduled)
- Download available updates from Arcserve (either directly to a client machine or to a staging server first and then to a client machine)
- Install successfully downloaded updates (to be initiated manually)
- Send email notifications when a new update is available

Note: When Arcserve UDP Agent (Windows) is managed by the Arcserve UDP Console, checking for updates is disabled from Arcserve UDP Agent (Windows). You need to check and deploy the update from the Arcserve UDP Console UI.



Check for Updates

When **Arcserve Server** is selected as download server, Arcserve UDP Agent (Windows) Updates provides the capability to connect to the Arcserve server either directly or using a proxy server (as configured manually) to check for new and available Arcserve UDP Agent (Windows) updates. The Arcserve UDP Agent (Windows) will directly connect to Arcserve server using the proxy settings configured by browser (only applicable for IE and Chrome). This check for updates function can be triggered either manually from the user interface/tray monitor or automatically as specified by the Scheduler. (The internal Scheduler is responsible for starting at a scheduled day and time and triggering an automatic check and download of available updates).

When triggered, the update manager contacts the Arcserve server to check the date /time stamp of a file that contains the available update information. If this available update information file has been modified since the last check, it will be downloaded from the server for comparison. The available update information is then compared to another file containing the already downloaded update information to determine if the available update is new and has not been previously downloaded. If the latest available update is not installed on your machine, Arcserve UDP Agent (Windows) displays an icon on the home page to inform you that a new update is available. In addition, an email notification can also be sent to inform you when a new Arcserve UDP Agent (Windows) update is available for downloading.

When **Staging Server** is selected as download server, Arcserve UDP Agent (Windows) downloads the available update information file from the staging server, and perform the same comparison check with the already available update information file. If the latest available update is not installed on your machine, Arcserve UDP Agent (Windows) will display an icon on the home page to inform you that a new update is available.

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

Download Updates

Arcserve UDP Agent (Windows) Updates provide the capability to download available Arcserve UDP Agent (Windows) updates either directly from the Arcserve server or from a staging server which in turn connects to the Arcserve server. This download process is triggered automatically when the check for updates process determines that a new update is available (unless this auto-download function is disabled). You can configure Arcserve UDP Agent (Windows) to download an update directly (or using a proxy server) to your client machine or to a staging server. A staging server can be used as a temporary storage location for downloading an update before it is downloaded and installed into an Arcserve UDP Agent (Windows) client machine. You may not want to expose your client machine to the internet to download updates from the Arcserve server. In this case, you can first download the update to a staging server and then allow other client machines to download the update from that staging server. The Arcserve UDP Agent (Windows) provides the capability to configure multiple staging servers for downloading the updates. If for some reason the primary staging server is unavailable, the download function will automatically transfer to the next specified staging server.

Note: If you are using a staging server for your Updates downloads, Arcserve UDP Agent (Windows) must be installed on that staging server, but does not need to be licensed unless you are using Arcserve UDP Agent (Windows) to protect that staging server.

When triggered, the Updates function contacts the Arcserve server and downloads the available update and places it in a holding directory (on either the staging server or the client machine) until directed to proceed with the subsequent installation process.

The default location for the download folder is: <Product Home>\Update Manager\EngineUpdates\7.0\

If for some reason, the download cannot be started, a popup message is displayed and Arcserve UDP Agent (Windows) waits a specified number of minutes and then attempt to download again. If after a specified number of retry attempts, the download still cannot continue, an error message will be displayed in the activity log indicating the most likely reason for the failure.

Install Updates

Arcserve UDP Agent (Windows) Updates provide the capability to install the available and successfully downloaded updates. This install process can only be triggered manually from the user interface/tray monitor (not automatically). When triggered, the update is installed from the holding directory to the applicable Arcserve UDP Agent (Windows) component directory of the client machine or the staging server. You cannot trigger the installation of the update directly from a staging server to a client machine. When you click install, the update is downloaded from the staging server to the client machine (if it has not been downloaded already), and then the installation process is triggered from the client machine.

Note: The installation only continues if no other active Arcserve UDP Agent (Windows) jobs are running. If another job is running, a message is displayed informing you of this condition and requesting that you try again at a later time.

If the installation is successful, the file containing the status information is updated for future use.

If the installation fails, an error message is displayed indicating the most likely reason for the failure.

Note: During the update installation Arcserve UDP Agent (Windows) will stop the Arcserve UDP Agent (Windows) Web service and will restart this web service after successful installation of update.

Email Notifications

Arcserve UDP Agent (Windows) Updates provide the capability to send automatic email notifications when a new update is available. The Arcserve UDP Agent (Windows) connects to an SMTP server (with appropriate credentials) to enable sending these email notifications over the Internet from Arcserve to your server. (The email recipients are specified from the **Preferences** dialog).

In addition, email notifications are also sent if a failure occurs during the check for updates or during the download.


Specify Updates Preferences

Arcserve UDP Agent (Windows) lets you specify the following Updates preferences:

Follow these steps:

1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Preferences** tab. When the **Preferences** dialog opens, select **Updates**.

The **Updates** preferences dialog opens.



2. Specify your **Updates** preference settings.

Download Server

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

▪ Arcserve Server

You can use this option to specify that Arcserve UDP Agent (Windows) updates are downloaded from the Arcserve server directly to your local server.

This is the default setting.

▪ Staging Server

You can use this option to specify the server that is used as a staging server.

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

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

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

Arcserve UDP Agent (Windows) updates are downloaded from the Arcserve server directly to the specified staging server location. After the updates are downloaded to this staging server, you can then further download the updates from the staging server to a client server. If you select the Staging Server location, you must also specify the host name or IP address for the staging server, along with the corresponding port number.

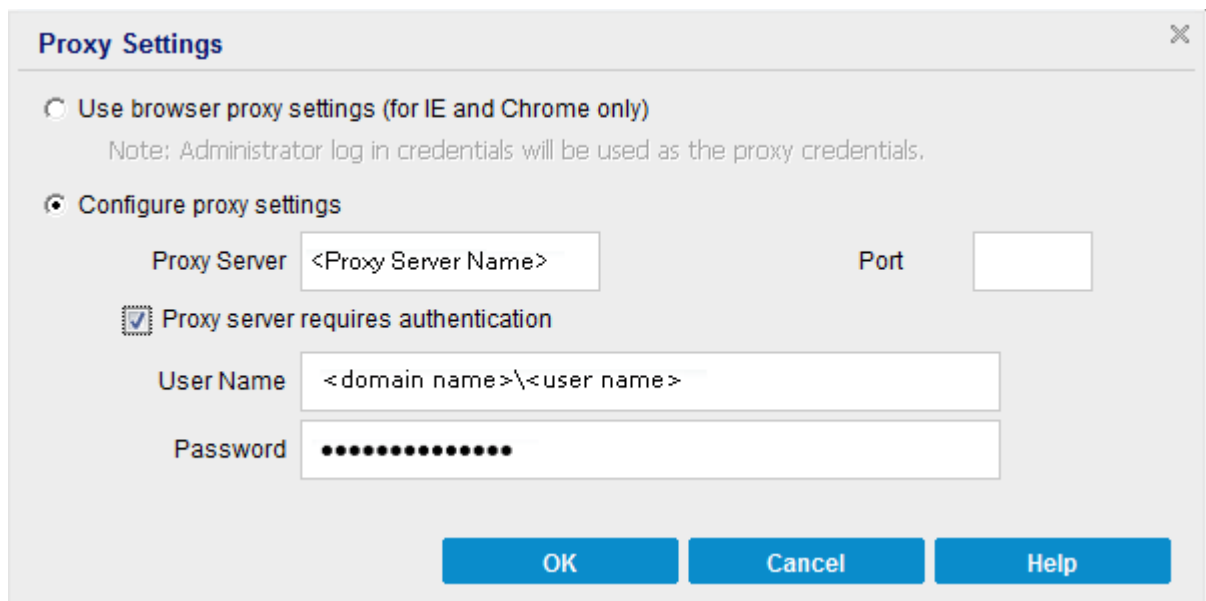
You cannot specify your same local client server as this staging server. This is an invalid configuration because the staging server cannot connect to itself to get and download the available updates from. If you attempt to use your local client server as the staging server, an error message is displayed.

▪ 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 Agent (Windows) 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 is the connection to the Arcserve server from which your download server gets the updates.

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



– **Use browser proxy settings**

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

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

– **Configure proxy settings**

When selected enables the specified proxy server to connect to the Arcserve server for Arcserve UDP Agent (Windows) 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 requires 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, a red icon is displayed on the status **Summary** section home page to provide a visual alert of this condition.

Note: The test connection is automatically performed when you launch the **Preferences Updates** dialog from the home page. When this auto test is performed it will check the latest connection status of the previously configured download server (either Arcserve Server or Staging Server(s), whichever is selected). If you previously configured more than one staging server, then this auto test is performed on all staging servers to get the latest connection status.

Update Schedule

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

- With this option selected, specifies to automatically check for new and available Arcserve UDP Agent (Windows) updates. When you select this option, you then have drop-down menu capabilities to specify when to perform this function (every day or weekly on a specified day) and the time of the day that it is performed.

Note: The default setting for the day or hour that these checks are automatically performed is randomly assigned by Arcserve UDP Agent (Windows) at the time of installation. After installation, you can use this **Update Schedule** setting to change the day and time for these checks.

By default, if this check determines that a new update is available, Arcserve UDP Agent (Windows) also automatically downloads the update.

- With this option not selected, specifies to disable all automatic check and download functions (and its status is displayed under status Summary section of the home page). With this option not selected, these update functions can only be triggered manually.

Notes:

If configured you get an email notification if the scheduled check for updates discovers that a new update is available. In addition, email notifications are sent if a failure occurs during the check for updates or during the download.

If the Arcserve UDP Agent (Windows) is managed by the Arcserve UDP Console, the **Automatically check for updates** option is disabled. Instead you can check updates from the Arcserve UDP Console and remote deploy updates to Arcserve UDP Agent (Windows).

3. Click **Save Settings**.

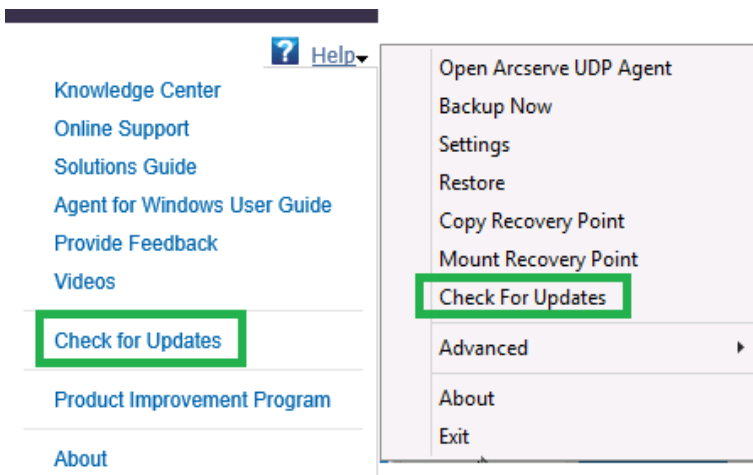
Your Updates preference settings are saved.

Check for Updates and Download

From the Arcserve UDP Agent (Windows) home page, you can select the **Check for Updates** option from the **Help** tab. Check for updates allows you to determine if any new updates are available.

Follow these steps:

1. Launch a check for updates to contact the Arcserve server or staging server. The check for updates can be launched automatically or manually from the Arcserve UDP Agent (Windows) **Help** menu or from the Arcserve UDP Agent (Windows) Monitor.



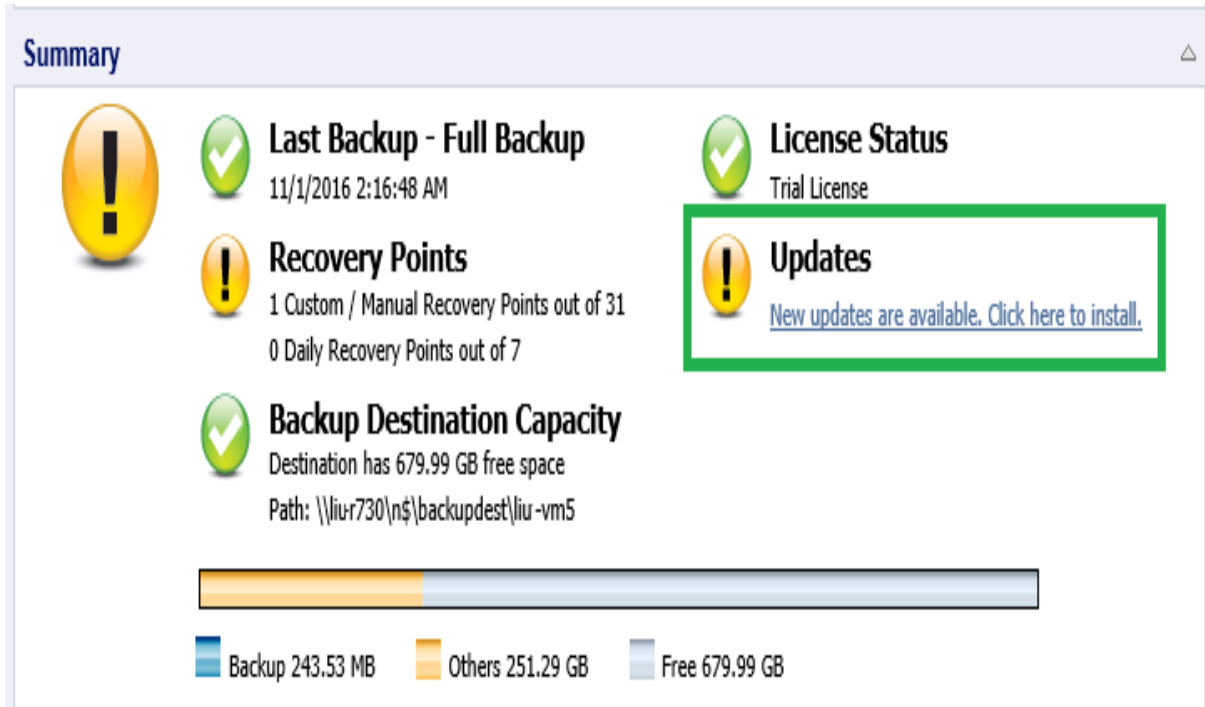
2. If a new update is available, it is automatically downloaded from Arcserve to the specified server staging or client.

A yellow **Updates** icon is displayed on the home page to provide a visual indication that a new update is ready to install.

Note: Update status balloon messages are also displayed from the Arcserve UDP

Agent (Windows) Monitor.

Summary



Last Backup - Full Backup
11/1/2016 2:16:48 AM

License Status
Trial License

Recovery Points
1 Custom / Manual Recovery Points out of 31
0 Daily Recovery Points out of 7

Updates
[New updates are available. Click here to install.](#)

Backup Destination Capacity
Destination has 679.99 GB free space
Path: \\liu-r730\n\$\backupdest\liu-vm5

Backup 243.53 MB Others 251.29 GB Free 679.99 GB

Install the Arcserve UDP Agent (Windows) Updates

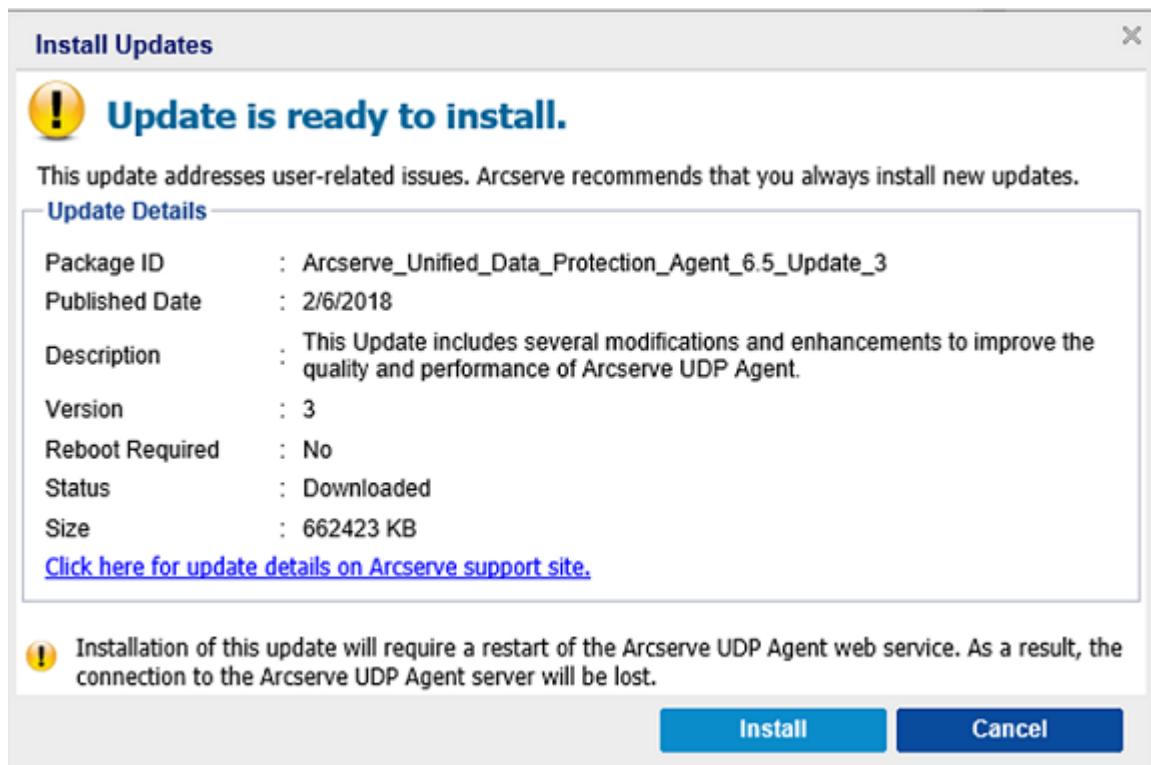
After you check for a new update and download the update, you can start installing the update.

Note: All updates that are released for Arcserve UDP Agent (Windows) 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.

Follow these steps:

1. Click the **Updates** icon.

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.



2. Review the update details, select **OK to reboot machine**, and click **Install** to trigger the installation of the Arcserve UDP Agent (Windows) update.


The new update is installed in your local computer. If the update requires a reboot and you selected the **OK to reboot machine** option, the computer reboots

automatically as part of the install process. Depending on the update status of each computer, you could have different reboot requirements for each computer.

Notes:

- ◆ If the update requires a computer reboot and you do not select the **OK to reboot machine** option, the **Install** button is disabled. You can then install the update at a more convenient time.
- ◆ During installation of the update, Arcserve UDP Agent (Windows) stops the Arcserve UDP Web Service and connection to UI is lost.

When the update is successfully installed, the **Updates** icon changes to a green status icon. The green status icon indicates that your computer has been updated and the updates function is enabled.




Updates can also be installed from the tray icon by clicking the New Update is available balloon message.

Installation of Arcserve UDP Agent (Windows) updates is complete.

Verify that the Updates are Successfully Installed

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

- From the Arcserve UDP Agent (Windows) home page UI, click **View Logs**, and then verify that the installed updates are listed in the **Activity Log**.



- From the Arcserve UDP Agent (Windows) home page, select **Help**, click **About Arcserve UDP Agent (Windows)**, and then verify that the about Arcserve UDP Agent (Windows) dialog displays the latest version updated.

(Optional) Install Arcserve UDP Agent (Windows) Updates Silently

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

The downloaded update installation file is under “<Product Home>\Update Manager\EngineUpdates\7.0”.

Follow these steps:

1. Launch the Arcserve UDP Agent (Windows) 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.

The update is configured and installed.

Troubleshooting Update 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.

- [Unable to Access Arcserve UDP Agent \(Windows\) After Reboot](#)
- [Unable to Connect to the Arcserve Download Server to Download Updates](#)
- [Failed to Download Arcserve UDP Agent \(Windows\) Updates](#)

Unable to Access Arcserve UDP Agent (Windows) After Reboot

If you are not able to access the Arcserve UDP Agent (Windows) UI, perform the following troubleshooting procedure:

1. From the **Add or Remove Programs** dialog, click the **Add/Remove Windows Components** option to access the **Windows Components Wizard** screen and remove the **Internet Explorer Enhanced Security Configuration** component.
2. Add the host name URL to the **Trusted Sites** in Internet Explorer.
3. Adjust the security level in Internet Explorer.

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 Connect to the Arcserve Download Server to Download Updates

If you are not able to connect to the Arcserve download server to download Arcserve UDP Agent (Windows) updates, follow these steps:

1. From the Arcserve UDP Agent (Windows) home page, click **View Logs**, and verify the error message.
2. Verify that you have a good network connection.
3. Open command line and ping the downloads.arcserve.com server.

Perform *one* of the following to establish connection with the download server:

- ◆ From the Arcserve UDP Agent (Windows) home page, select **Settings**, then **Preferences**, and click **Updates and Download Server**. Click on the proxy settings and verify that the default option **Use browser proxy settings** (for IE and Chrome only) is selected.
 - ◆ From the Arcserve UDP Agent (Windows) home page, select **Settings**, then **Preferences**, and click **Updates and Download Server**. Click on the proxy settings and select **Configure Proxy Settings** and enter the valid proxy server name, port number and credentials and click **OK**.
4. Click **Test Connection** to verify that the connection is established.

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 Download Arcserve UDP Agent (Windows) Updates

If you are not able to download Arcserve UDP Agent (Windows) updates, follow these steps:

1. From the Arcserve UDP Agent (Windows) home page, click **View Logs** and read the error message.
2. Verify that you have a good network connection.
3. Verify that there is enough disk space.
4. From the Arcserve UDP (Windows) installation home path, access the update Log file ("<Product Home>\Update Manager\logs\ARCUpdate.log").
5. Check the log entries for detailed error messages.

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.

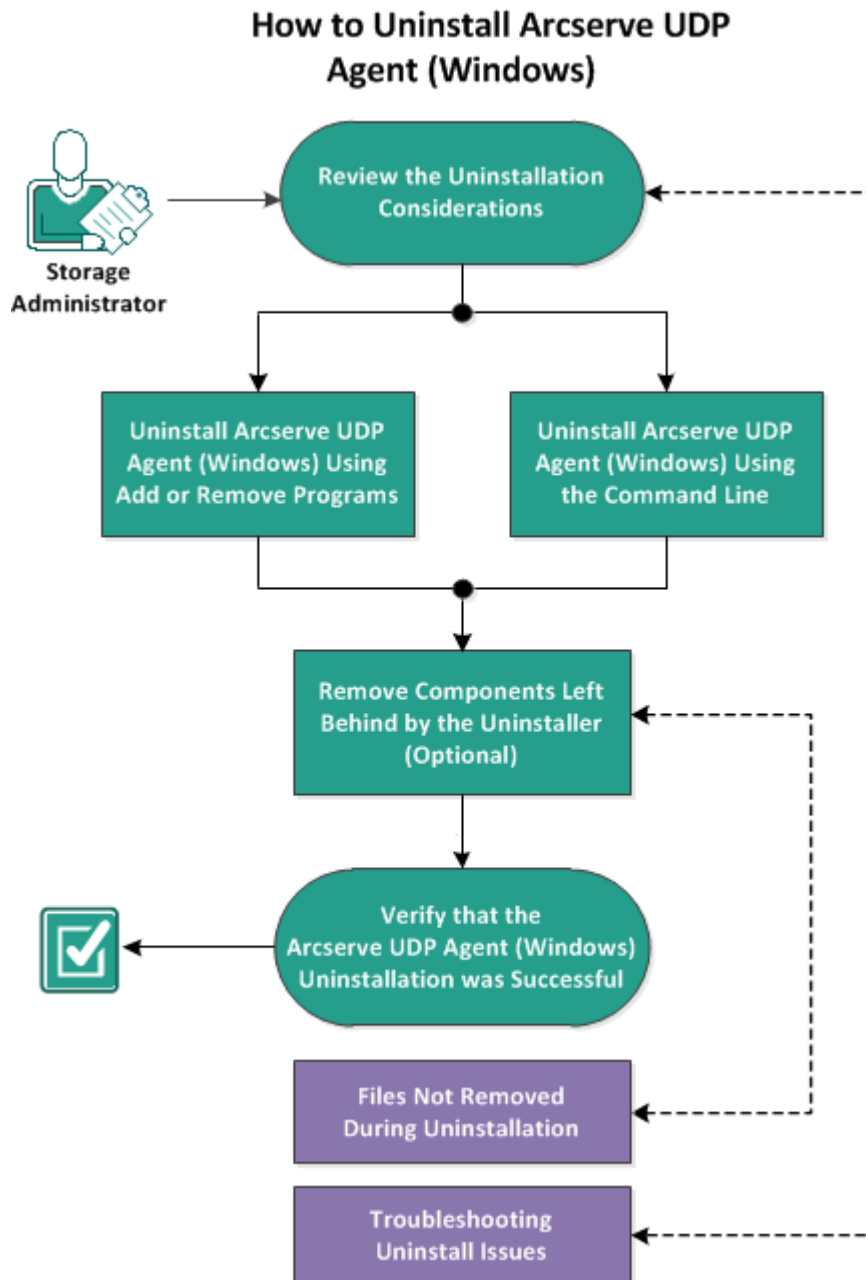
How to Uninstall Arcserve UDP Agent (Windows)

You can uninstall Arcserve UDP Agent (Windows) using the standard Add or Remove Programs application located in the Windows Control Panel and also by using the command line.

The uninstallation routine removes all Arcserve UDP Agent (Windows) directories and files from your computer, except for the following directories and all of their contents:

- Arcserve Licensing:
 - (x86 systems) C:\Program Files\Arcserve\SharedComponents\CA_LIC
 - (x64 systems) C:\Program Files(X86)\Arcserve\SharedComponents\CA_LIC

The following diagram illustrates the process to uninstall Arcserve UDP Agent (Windows):



Perform the following tasks to uninstall Arcserve UDP Agent (Windows):

1. [Review the Uninstallation Considerations](#)
2. [Uninstall Arcserve UDP Agent \(Windows\) Using Add or Remove Programs](#)
3. [Uninstall Arcserve UDP Agent \(Windows\) Using the Command Line](#)
4. [\(Optional\) Remove Components Left Behind by the Uninstaller](#)
5. [Verify that the Arcserve UDP Agent \(Windows\) Uninstallation was Successful](#)
6. [\(Optional\) Files Not Removed During Uninstallation](#)
7. [\(Optional\) Troubleshooting Uninstall Issues](#)

Review the Uninstallation Considerations

Review the following uninstallation considerations:

- When you upgrade Arcserve UDP Agent (Windows) to the next release, it is not necessary for you to uninstall Arcserve UDP Agent (Windows).
- A nodeID is kept after uninstallation to identify the server so that it does not change when you install Arcserve UDP agent (Windows) again. To get a different nodeID, you can delete the following file (if it exists) before the next installation:
`%windir%\Temp\Arcserve\Setup\UDP\Uninstall\Settings.ini`
- Review the [Compatibility Matrix](#) that provides the supported operating systems, databases, and browsers.

Uninstall Arcserve UDP Agent (Windows) Using Add or Remove Programs

You can uninstall Arcserve UDP Agent (Windows) using the standard **Add or Remove Programs** application located in the Windows **Control Panel**.

Follow these steps:

1. Click **Start, Settings, Control Panel, Add/Remove Programs**.

The **Add or Remove Programs** dialog opens. The list of installed programs is displayed.

2. Select **Arcserve Unified Data Protection**, and click **Remove**.

The **Components** dialog opens.

3. Select **Arcserve UDP Agent** and click **Remove**.

4. Click **Finish** to complete the uninstall process.

The application is uninstalled.

Uninstall Arcserve UDP Agent (Windows) Using the Command Line

A silent uninstallation eliminates the need for user interaction. The following steps describe how to uninstall the application using the Windows Command Line.

Follow these steps:

1. Log in to the computer where you want to uninstall Arcserve UDP components.
Note: You must log in to the computer using an administrative account.
2. Open the Windows Command Line.
3. Execute the syntax that corresponds with the architecture of the computer's operating system:

- ◆ x86 operating system:

```
"%ProgramFiles%\Arcserve\SharedComponents\Arcserve Unified Data Protection\Setup\uninstall.exe" /q /p {CAAD8AEA-A455-4A9F-9B48-C3838976646A}
```

- ◆ x64 operating system:

```
"%ProgramFiles(x86)%\Arcserve\SharedComponents\Arcserve Unified Data Protection\Setup\uninstall.exe" /q /p {CAAD1E08-FC33-462F-B5F8-DE9B765F2C1E}
```

Return Code:

0 = Uninstall was successful.

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

Other = Uninstall failed.

The uninstallation is complete. A reboot is required if the Arcserve UDP Agent (Windows) driver is installed.

Remove Components Left Behind by the Uninstaller

Uninstallation of Arcserve UDP Agent (Windows) may affect certain third party and proprietary components, some of which are installed and removed with the corresponding component, and others which may remain because they are shared components with other Arcserve products and numerous components. Note that if you remove the "shared" components, it may adversely affect your use and licensing of other Arcserve products, including but not limited to, losing the licensing for other Arcserve products installed on that machine. Further, if the "shared" components are removed, any programs that are installed after Arcserve UDP Agent (Windows) and depend on these components may not function properly.

Note: For a complete listing of all files (path and name) that are left behind by Arcserve UDP Agent (Windows) uninstallation, see [Files Not Removed During Uninstallation](#).

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

Remove Arcserve Licensing Component manually

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

Remove Microsoft Visual C++ manually

1. Access the standard **Add or Remove Programs** application located in the Windows **Control Panel** (Control Panel -> Programs and Features -> Remove Programs).
2. Select *Microsoft Visual C++ 2013 x86 Redistributable - 12.0.30501* and then click on **uninstall**.
3. Select *Microsoft Visual C++ 2013 x64 Redistributable - 12.0.30501* and then click on **uninstall**.

Verify that the Arcserve UDP Agent (Windows) Uninstallation was Successful

Follow these steps:

1. Verify that the Agent icon is deleted from the system tray.
2. Navigate to **services.msc** from the command prompt tab and click **OK**.
3. Verify that the Arcserve UDP Agent service is deleted from the **Services Manager**.
4. Open the **Control Panel** and verify if Arcserve UDP Agent (Windows) removed.
5. Go to **Start, Programs**, and verify if Arcserve UDP Agent (Windows) is deleted.

The Arcserve UDP Agent (Windows) was uninstalled successfully.

Files Not Removed During Uninstallation

Arcserve UDP Agent (Windows) is uninstalled using the standard **Add or Remove Programs** application located in the Windows **Control Panel** or using the Command Line. During the Arcserve UDP Agent (Windows) uninstallation process, some files may not be uninstalled or removed as expected.

Here is a list of the Arcserve UDP Agent (Windows) file names and corresponding paths that are not removed by the uninstaller after the uninstallation process is completed:

C:\Program Files (x86)\Arcserve\SharedComponents\CA_LIC\CA Licensing User Help.chm
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\CALicnse.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\CAMinfo.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\CAregit.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\countries.txt
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\countriesTrial.txt
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\ErrBox.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic_comp_codes.dat
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98.cap
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98.dat
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98.err
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98_64.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98_64_amd.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98_uninstaller.zip
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98FileSockLib.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98FileSockLib_amd64.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98FileSockLib_ia64.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98log.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\Lic98Msg.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98-port
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98Service.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98version.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\LicDebug.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\licinfo_win.zip
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\LicRCmd.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\licreg.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\licreg_64.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\licreg_64_amd.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\licregres.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\licregres_64.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\licregres_64_amd.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\LogWatNT.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\mergecalic.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\mergeolf.exe

C:\Program Files (x86)\CA\SharedComponents\CA_LIC\prod_codes.txt
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\silntreg.tmp
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\states.txt
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\statesTrial.txt
C:\Program Files (x86)\Arcserve\SharedComponents\CA_LIC\vendor.dat
C:\Program Files (x86)\Common Files\microsoft shared\VC\amd64\msdia80.dll
C:\Program Files (x86)\Common Files\microsoft shared\VC\msdia80.dll
C:\Users\Administrator.RIGONE\AppData\Local\IconCache.db
C:\User-
s\Ad-
min-
istrat-
or\Ap-
pData\LocalLow\Mi-
crosoft\CryptnetUrlCache\Content\696F3DE637E6DE85B458996D49D759AD
C:\User-
s\Ad-
min-
istrat-
or\Ap-
pData\LocalLow\Mi-
crosoft\CryptnetUrlCache\Content\B8CC409ACDBF2A2FE04C56F2875B1FD6
C:\User-
s\Ad-
min-
istrat-
or\Ap-
pData\LocalLow\Mi-
crosoft\CryptnetUrlCache\MetaData\696F3DE637E6DE85B458996D49D759AD
C:\User-
s\Ad-
min-
istrat-
or\Ap-
pData\LocalLow\Mi-
crosoft\CryptnetUrlCache\MetaData\B8CC409ACDBF2A2FE04C56F2875B1FD6
C:\Users\Administrator\arcserve Unified Data Protection Agent\TrayI-
con\ARCFashTrayIcon.log
C:\Users\Administrator\arcserve Unified Data Protection Agent\TrayIcon\
ARCFashTrayIcon_java.log
C:\Windows\Downloaded Installations\{D03BF724-4E4F-4DF4-A1BD-
8497634F5589}\1033.MST
C:\Windows\Downloaded Installations\{D03BF724-4E4F-4DF4-A1BD-
8497634F5589}\ASLicense.msi
C:\Windows\inf\WmiApRpl\0009\WmiApRpl.ini
C:\Windows\inf\WmiApRpl\WmiApRpl.h

C:\Windows\System32\config\COMPONENTS{016888b8-6c6f-11de-8d1d-001e0b-cde3ec}.TxR.0.regtrans-ms
C:\Windows\System32\config\COMPONENTS{016888b8-6c6f-11de-8d1d-001e0b-cde3ec}.TxR.1.regtrans-ms
C:\Windows\System32\config\COMPONENTS{016888b8-6c6f-11de-8d1d-001e0b-cde3ec}.TxR.2.regtrans-ms
C:\Windows\System32\config\COMPONENTS{016888b8-6c6f-11de-8d1d-001e0b-cde3ec}.TxR.blf
C:\Windows\System32\drivers\Msft_Kernel_AFStorHBA_01009.Wdf
C:\Windows\System32\drivers\Msft_Kernel_ARCFlashVolDrv_01009.Wdf
C:\Windows\System32\drivers\Msft_User_AFStorHBATramp_01_09_00.Wdf
C:\Windows\System32\LogFiles\WUDF\WUDFTrace.etl
C:\Windows\System32\winevt\Logs\Microsoft-Windows-DriverFrameworks-User-Mode%4Operational.evtx
C:\\$Mft
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\CALicnse.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\CALicnse.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\CAMinfo.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\CAMinfo.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\CAREgit.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\CAREgit.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\ErrBox.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\ErrBox.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98_64.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98_64.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98_64_amd.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98_64_amd.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98FileSockLib.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98FileSockLib.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98FileSockLib_amd64.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98FileSockLib_amd64.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98FileSockLib_ia64.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98FileSockLib_ia64.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98log.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98log.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\Lic98Msg.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\Lic98Msg.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98Service.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98Service.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98version.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\lic98version.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\LicDebug.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\LicDebug.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\LicRCmd.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\LicRCmd.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\licreg.dll

C:\Program Files (x86)\CA\SharedComponents\CA_LIC\licreg.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\licreg_64.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\licreg_64.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\licreg_64_amd.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\licreg_64_amd.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\licregres.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\licregres.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\licregres_64.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\licregres_64.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\licregres_64_amd.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\licregres_64_amd.dll
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\LogWatNT.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\LogWatNT.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\LogWatNT.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\mergecalic.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\mergecalic.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\mergeolf.exe
C:\Program Files (x86)\CA\SharedComponents\CA_LIC\mergeolf.exe
C:\Program Files (x86)\Common Files\microsoft shared\VC\msdia100.dll
C:\Users\Administrator.RIGONE\AppData\Local\Microsoft\Windows\UsrClass.dat
C:\User-
s\Administrator.RIGONE\AppData\Local\Microsoft\Windows\UsrClass.dat.LOG1
C:\Users\Administrator.RIGONE\NTUSER.DAT
C:\Users\Administrator.RIGONE\ntuser.dat.LOG1
C:\User-
s\Ad-
min-
istrat-
or\Ap-
pData\LocalLow\Mi-
crosoft\CryptnetUrlCache\Content\94308059B57B3142E455B38A6EB92015
C:\User-
s\Ad-
min-
istrat-
or\Ap-
pData\LocalLow\Mi-
crosoft\CryptnetUrlCache\MetaData\94308059B57B3142E455B38A6EB92015
C:\Users\Administrator\NTUSER.DAT
C:\Users\Administrator\ntuser.dat.LOG1
C:\Windows\AppCompat\Programs\RecentFileCache.bcf
C:\Windows\inf\setupapi.dev.log
C:\Win-
dows\Ser-
vicePro-
files\Net-

workSer-
 vice\AppData\Roaming\Microsoft\SoftwareProtectionPlatform\Cache\cache.dat
 C:\Windows\setupact.log
 C:\Windows\SoftwareDistribution\DataStore\DataStore.edb
 C:\Windows\SoftwareDistribution\DataStore\Logs\edb.chk
 C:\Windows\SoftwareDistribution\DataStore\Logs\edb.log
 C:\Windows\System32\7B296FB0-376B-497e-B012-9C450E1B7327-5P-
 0.C7483456-A289-439d-8115-601632D005A0
 C:\Windows\System32\7B296FB0-376B-497e-B012-9C450E1B7327-5P-
 1.C7483456-A289-439d-8115-601632D005A0
 C:\Windows\System32\catroot2\{127D0A1D-4EF2-11D1-8608-00C04FC295EE}\c-
 atdb
 C:\Windows\System32\catroot2\{F750E6C3-38EE-11D1-85E5-00C04FC295EE}\c-
 atdb
 C:\Windows\System32\catroot2\dberr.txt
 C:\Windows\System32\catroot2\edb.chk
 C:\Windows\System32\catroot2\edb.log
 C:\Windows\System32\config\COMPONENTS
 C:\Windows\System32\config\COMPONENTS.LOG1
 C:\Windows\System32\config\COMPONENTS\{016888b8-6c6f-11de-8d1d-001e0b-
 cde3ec}.TxR.0.regtrans-ms
 C:\Windows\System32\config\COMPONENTS\{016888b8-6c6f-11de-8d1d-001e0b-
 cde3ec}.TxR.blf
 C:\Windows\System32\config\COMPONENTS\{016888b9-6c6f-11de-8d1d-001e0b-
 cde3ec}.TMContainer00000000000000000001.regtrans-ms
 C:\Windows\System32\config\DEFAULT
 C:\Windows\System32\config\DEFAULT.LOG1
 C:\Windows\System32\config\SAM
 C:\Windows\System32\config\SAM.LOG1
 C:\Windows\System32\config\SOFTWARE
 C:\Windows\System32\config\SOFTWARE.LOG1
 C:\Windows\System32\config\SYSTEM
 C:\Windows\System32\config\SYSTEM.LOG1
 C:\Windows\System32\config\TxR\{016888cc-6c6f-11de-8d1d-001e0b-
 cde3ec}.TxR.0.regtrans-ms
 C:\Windows\System32\config\TxR\{016888cc-6c6f-11de-8d1d-001e0b-
 cde3ec}.TxR.blf
 C:\Windows\System32\config\TxR\{016888cd-6c6f-11de-8d1d-001e0b-
 cde3ec}.TMContainer00000000000000000001.regtrans-ms
 C:\Windows\System32\DriverStore\INF\CACHE.1
 C:\Windows\System32\DriverStore\infpub.dat
 C:\Windows\System32\DriverStore\infstor.dat
 C:\Windows\System32\DriverStore\infstrng.dat
 C:\Windows\System32\LogFiles\Scm\3cdb3c57-5945-4fa9-8e4d-f8bd141f0f8f
 C:\Windows\System32\LogFiles\Scm\63ee8552-a444-4ba2-8e1e-c8350d6d412a
 C:\Windows\System32\LogFiles\Scm\c7847981-48e6-476f-9581-4bbd8e73f7c5
 C:\Windows\System32\LogFiles\Scm\cd264f70-fd14-48ea-9d74-f52f1d1d3f89

C:\Windows\System32\perfc009.dat
C:\Windows\System32\perfh009.dat
C:\Windows\System32\PerfStringBackup.INI
C:\Windows\System32\SMI\Store\Machine\SCHEMA.DAT
C:\Windows\System32\SMI\Store\Machine\SCHEMA.DAT.LOG1
C:\Windows\System32\wbem\Performance\WmiApRpl.h
C:\Windows\System32\wbem\Performance\WmiApRpl.ini
C:\Windows\System32\wbem\Repository\INDEX.BTR
C:\Windows\System32\wbem\Repository\MAPPING1.MAP
C:\Windows\System32\wbem\Repository\OBJECTS.DATA
C:\Windows\System32\WdfCoinstaller01009.dll
C:\Windows\System32\winevt\Logs\Application.evtx
C:\Windows\System32\winevt\Logs\Microsoft-Windows-Bits-Client%4Operational.evtx
C:\Windows\System32\winevt\Logs\Microsoft-Windows-Diagnosis-DPS%4Operational.evtx
C:\Windows\System32\winevt\Logs\Microsoft-Windows-GroupPolicy%4Operational.evtx
C:\Windows\System32\winevt\Logs\Microsoft-Windows-Kernel-WHEA%4Operational.evtx
C:\Windows\System32\winevt\Logs\Microsoft-Windows-Known Folders API Service.evtx
C:\Windows\System32\winevt\Logs\Microsoft-Windows-Net-WorkProfile%4Operational.evtx
C:\Windows\System32\winevt\Logs\Microsoft-Windows-NlaSvc%4Operational.evtx
C:\Windows\System32\winevt\Logs\Microsoft-Windows-PrintService%4Admin.evtx
C:\Windows\System32\winevt\Logs\Microsoft-Windows-Resource-Exhaustion-Detector%4Operational.evtx
C:\Windows\System32\winevt\Logs\Microsoft-Windows-TaskScheduler%4Operational.evtx
C:\Windows\System32\winevt\Logs\Microsoft-Windows-TerminalServices-LocalSessionManager%4Operational.evtx
C:\Windows\System32\winevt\Logs\Microsoft-Windows-TerminalServices-RemoteConnectionManager%4Operational.evtx
C:\Windows\System32\winevt\Logs\Microsoft-Windows-User Profile Service%4Operational.evtx
C:\Windows\System32\winevt\Logs\Microsoft-Windows-Windows Firewall With Advanced Security%4Firewall.evtx
C:\Windows\System32\winevt\Logs\Microsoft-Windows-WindowsUpdateClient%4Operational.evtx
C:\Windows\System32\winevt\Logs\Microsoft-Windows-WinRM%4Operational.evtx
C:\Windows\System32\winevt\Logs\Security.evtx
C:\Windows\System32\winevt\Logs\Setup.evtx
C:\Windows\System32\winevt\Logs\System.evtx
C:\Windows\System32\winevt\Logs\Works with Tool.evtx
C:\Windows\System32\WudfUpdate_01009.dll
C:\Windows\WindowsUpdate.log
C:\Windows\System32\atl100.dll

Troubleshooting Uninstall 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.

Unable to install/uninstall Arcserve UDP Agent (Windows) if a previous attempt was interrupted

If during an attempt to install or uninstall Arcserve UDP Agent (Windows), the install/uninstall process was interrupted, you may not be able to successfully continue and complete the process.

For example, any of the following conditions could cause a partial install/uninstall state:

- Your computer is shut down in middle of install/uninstall process.
- You encounter a power outage during install/uninstall and there is no Uninterruptible Power Supply (UPS).

To resolve this problem, perform the following steps:

1. Enter "**regedit**" in the **Run** dialog and click **OK** to open **Registry Editor**.
2. Locate and delete the following entry:
"HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine"
3. Use the search option in the **Registry Editor** to locate and delete all occurrences of the following string:
 - ♦ [Arcserve UDP Agent (Windows) for x86]: {CAAD8AEA-A455-4A9F-9B48-C3838976646A}
 - ♦ [Arcserve UDP Agent (Windows) for x64]: {CAAD1E08-FC33-462F-B5F8-DE9B765F2C1E}
4. Use the search option in the **Registry Editor** to locate and delete all occurrences of the string "Arcserve UDP Agent " under the following key:
 - HKEY_CLASSES_ROOT\Installer\Products
 - HKEY_LOCAL_MACHINE\SOFTWARE\Classes\Installer\Products
 - HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Installer\UserData\S-1-5-18\Products
 - HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Uninstall
5. From the command line, delete the service by entering the following commands:
sc delete ShProvd
sc delete CASAD2DWebSvc
6. Run the command line to remove additional setup files.
 - ♦ x86 operating system:

```
"%ProgramFiles%\Arcserve\SharedComponents\Arcserve Unified Data Protection\Setup\uninstall.exe" /q
```

- ◆ x64 operating system:

```
"%ProgramFiles(x86)%\Arcserve\SharedComponents\Arcserve Unified Data Protection\Setup\uninstall.exe" /q
```

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.

UDP Workstation Free

Starting with Arcserve UDP Version 5.0 Update 2, at the end of the trial period, a free and fully functional Workstation Free Edition is provided to any user who has not yet obtained a proper license. This Workstation Free Edition 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) automatically reverts to the Workstation Free Edition.
- Your Workstation Free Edition 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 Workstation Free Edition, you cannot 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.
- Live Chat capabilities are unavailable, but you can use Online community based 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 Workstation Free Edition.

Q. What will happen if a Recovery Point Server (RPS) is selected as the destination for a Workstation Free Edition 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 Workstation Free Edition node and it would consume one of the available licenses (but no longer be a Workstation Free Edition node).

Q. Does Workstation Free Edition work for server-class operating systems such as Windows 2012?

A. No, Workstation Free Edition 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 Workstation Free Edition?

A. You can leverage support for Workstation Free Edition 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 Workstation Free Edition.

Chapter 3: Getting Started with Arcserve UDP Agent (Windows)

This section contains the following topics:

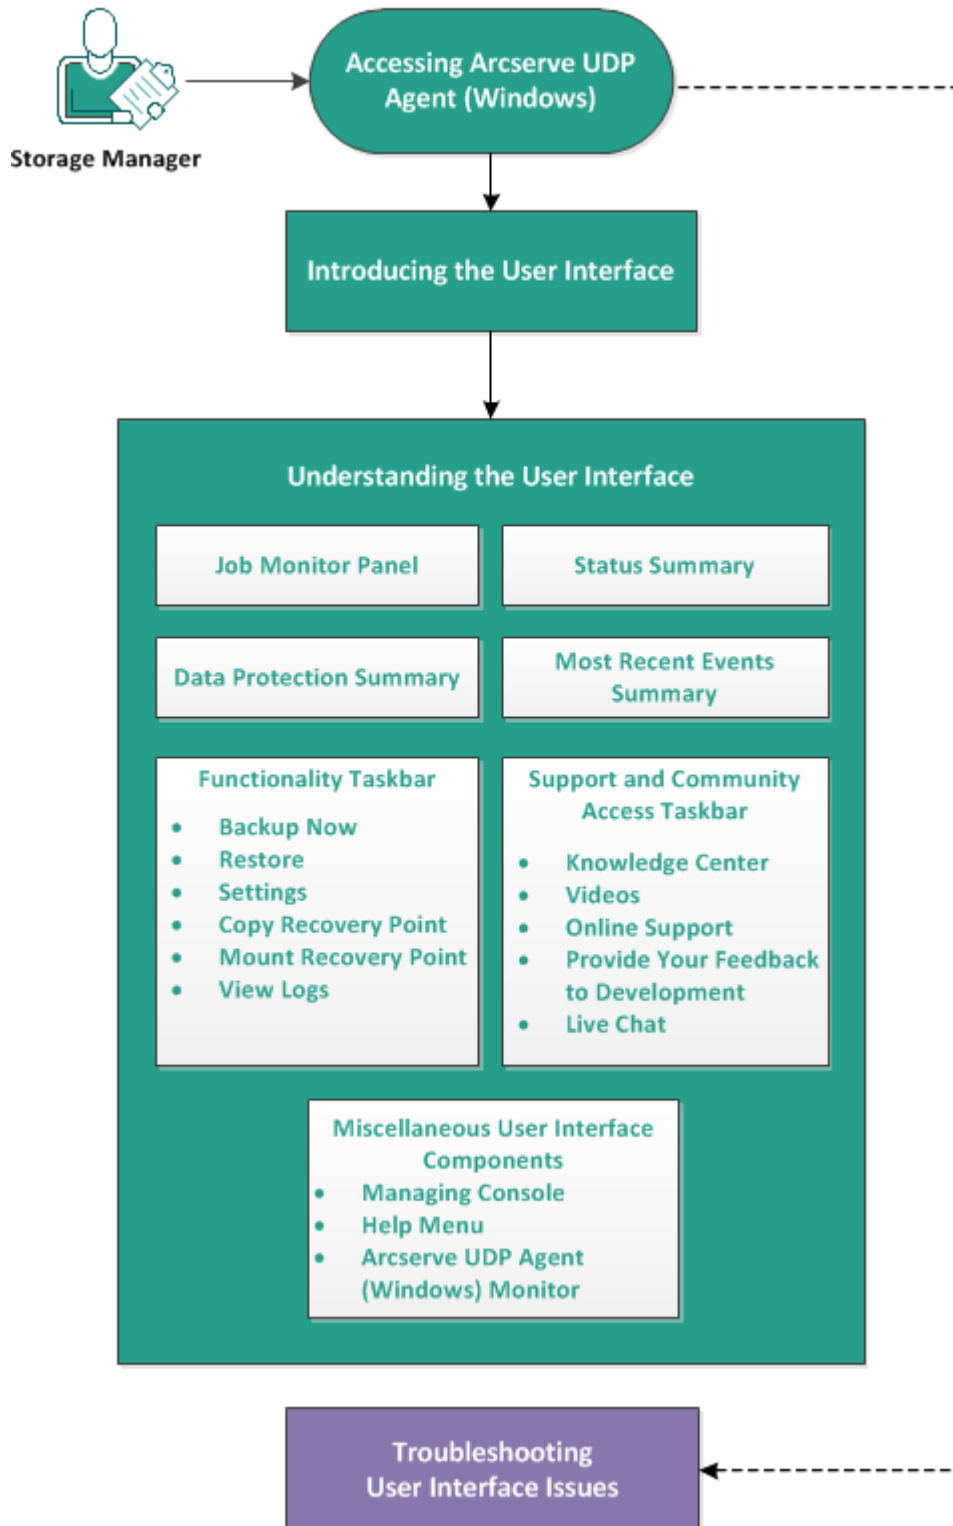
How to Navigate the Arcserve UDP Agent (Windows) User Interface	118
---	-----

How to Navigate the Arcserve UDP Agent (Windows) User Interface

Before you use Arcserve UDP Agent (Windows), you should learn how to navigate the related home page interface and become familiar with the tasks and monitoring functions that are performed from this single, easy-to-read home page. The home page interface can be accessed from either the Start menu or from the Arcserve UDP Agent (Windows) Monitor.

The following diagram illustrates the process for how to navigate the Arcserve UDP Agent (Windows) user interface:

How to Navigate the Arcserve UDP Agent (Windows) User Interface




Complete the following tasks to navigate the Arcserve UDP Agent (Windows) user interface:

1. [Accessing Arcserve UDP Agent \(Windows\)](#)
2. [Introducing the User Interface](#)
3. [Understanding the User Interface](#)
 - ♦ [Job Monitor Panel](#)
 - ♦ [Status Summary](#)
 - ♦ [Data Protection Summary](#)
 - ♦ [Most Recent Events Summary](#)
 - ♦ [Functionality Taskbar](#)
 - ♦ [Support and Community Access Taskbar](#)
 - ♦ [Miscellaneous User Interface Components](#)
4. [\(Optional\) Troubleshooting User Interface Issues](#)

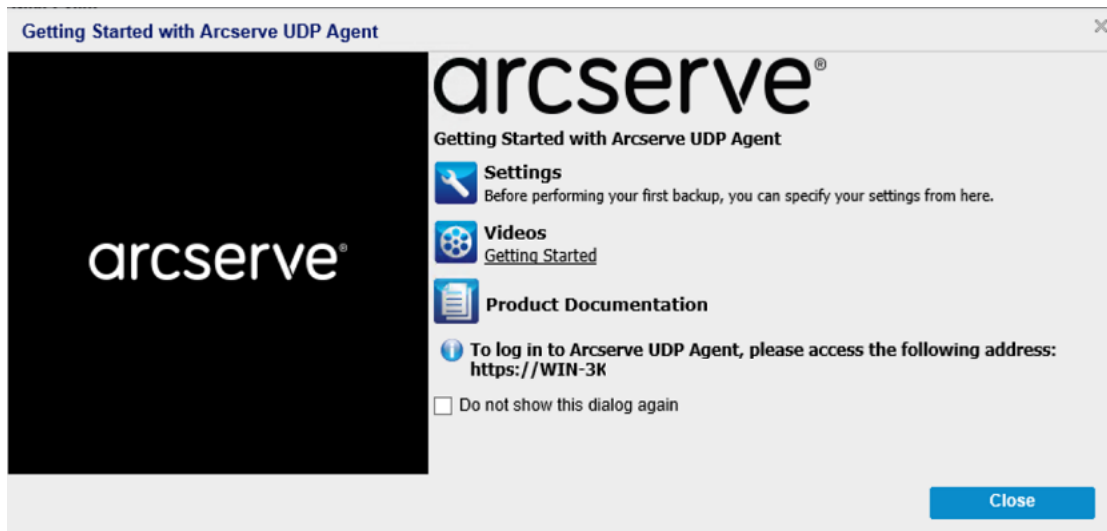
SUPPLEMENTAL VIDEO

This procedure contains a supplemental instructional video. Select either arcserve.com or YouTube as the source for viewing this video. The versions of the video are identical, and only the viewing source is different.

 Video	
YouTube:	Getting Started with Arcserve UDP Agent (Windows)

Accessing Arcserve UDP Agent (Windows)

When you first access Arcserve UDP Agent (Windows), the **Getting Started** dialog is displayed. From this dialog, you can access videos and the online help to learn more about Arcserve UDP Agent (Windows). In addition, you can also access the various dialogs to specify such configuration settings as your backup source and destination, parameters, schedule, alert notifications, file copy settings, copy recovery point settings, preferences, and other related tasks. You can also select the option to not show this **Getting Started** dialog again.



Introducing the User Interface

Before you use Arcserve UDP Agent (Windows), become familiar with the related home page interface. The Arcserve UDP Agent (Windows) interface lets you perform the following tasks all from a single, easy-to-read home page:

- Manage backup servers and workstations.
- Monitor job performance.
- Obtain backup statistics.
- Initiate data protection tasks.
- Socialize with the user community.
- Get help.

The Arcserve UDP Agent (Windows) home page displays various icon symbols to provide a quick visual indication of the recent status and indicates the urgency of any actions you must take.



Successful
(No action is necessary)



Caution
(Action may be necessary soon)



Warning
(Immediate action is necessary)

The Arcserve UDP Agent (Windows) home page consists of the following subsections:

- [Job Monitor Panel](#)
- [Status Summary](#)
- [Data Protection Summary](#)
- [Most Recent Events Summary](#)
- [Functionality Taskbar](#)
- [Support and Community Access Taskbar](#)
- [Help Menu Link](#)

arcserve UNIFIED DATA PROTECTION AGENT
Server: WIN-3K

Messages (1) | WIN-3K\1185.CQ\Administrator | Log Out | Help

Job Monitor Panel
Next Scheduled Backup: 4/9/2019 12:15:00 AM Incremental Backup

Status Summary

Summary

- Last Backup - Incremental Backup
4/8/2019 12:15:00 AM
- Recovery Points
2 Custom / Manual Recovery Points out of 2
1 Daily Recovery Points out of 1
- Backup Destination Capacity
Destination has 40.29 GB free space
Path: F:\Data\1185-3411\1185.D04
- License Status
This license
- Updates
The schedule for updates is enabled.

Protection Summary

Job Type	Count	Data Protected	Space Occupied	Last Successful Event	Next Event
Full Backup	1	11.76 GB	8.74 GB	4/7/2019 11:57:44 PM	
Incremental Backup	2	262.32 MB	75.96 MB	4/8/2019 12:15:01 AM	4/9/2019 12:15:00 AM
Verify Backup	0	0 Bytes	0 Bytes		
File Copy	0	0 Bytes	0 Bytes		
Copy Recovery Point	0	0 Bytes	0 Bytes		

Most Recent Events Summary

Status	Schedule Type	Backup Type	Date Time	Data Protected	Space Occupied	File Copy Status	Name
Success	Daily	Incremental Backup	4/8/2019 12:15:01 AM	123.03 MB	35.34 MB	N/A	
Success	Custom / Manual	Incremental Backup	4/8/2019 12:04:54 AM	138.47 MB	41.62 MB	N/A	Customized Incremental Backup

Managing Console

Help Link

Functionality Task Bar

- Backup Now
- Restore
- Settings
- Copy Recovery Point
- Mount Recovery Point
- View Logs

Support and Community Access Task Bar


- Knowledge Center
- Videos
- Online Support
- Live Chat

Understanding the User Interface

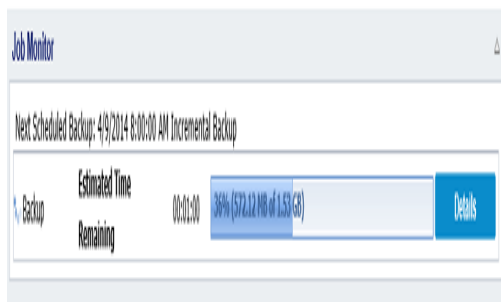
This section provides information on the following:

Job Monitor Panel

When no jobs are currently running, the **Job Monitor** panel displays the date and time of the next scheduled event, with the type of event to be performed.



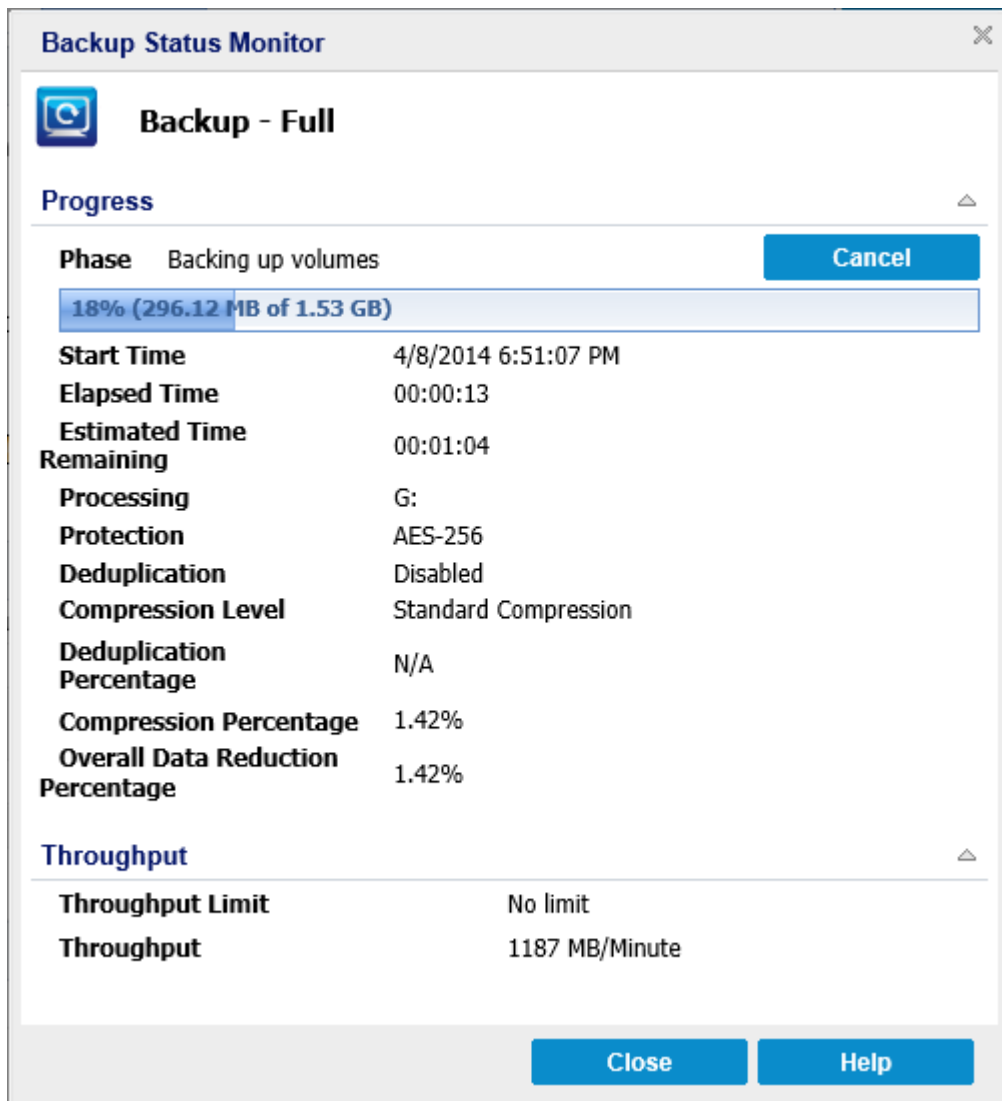
When a job is running, this panel expands to display information about the ongoing event such as the estimated time remaining to complete the job, the percentage and size of the job already completed, and the total size of the job when completed.



Note: When the Windows Performance Counter is disabled, the data speed of some Arcserve UDP Agent (Windows) jobs displayed in the job monitor may be 0 or some other abnormal value. If this occurs, see the troubleshooting section for more information.

When a job is running, you can click the **Detail** button to open the **Backup Status Monitor** and display more detailed information about the current job running. You can also click the **Cancel** button to stop the current job.

Note: If you want to stop the current job, click the **Detail** button first, to gain access to the **Cancel** button.




In addition to displaying more detailed information about the current job, the **Backup Status Monitor** also displays the throughput information for the job and the configured throttle limit.

- If the throughput is too fast, you could enable the **Throttle Backup** option to adjust and limit the job throughput. For more information about adjusting the **Throttle Backup** speed, see [Specify the Protection Settings](#) in the online help.

Note: Any changes made to the throttle settings take effect immediately when you save the settings.

- If the throughput is too slow, there could be various reasons for this reduced speed such as antivirus software may be scanning the machine, or some file is being copied, or the machine is being accessed by many users.

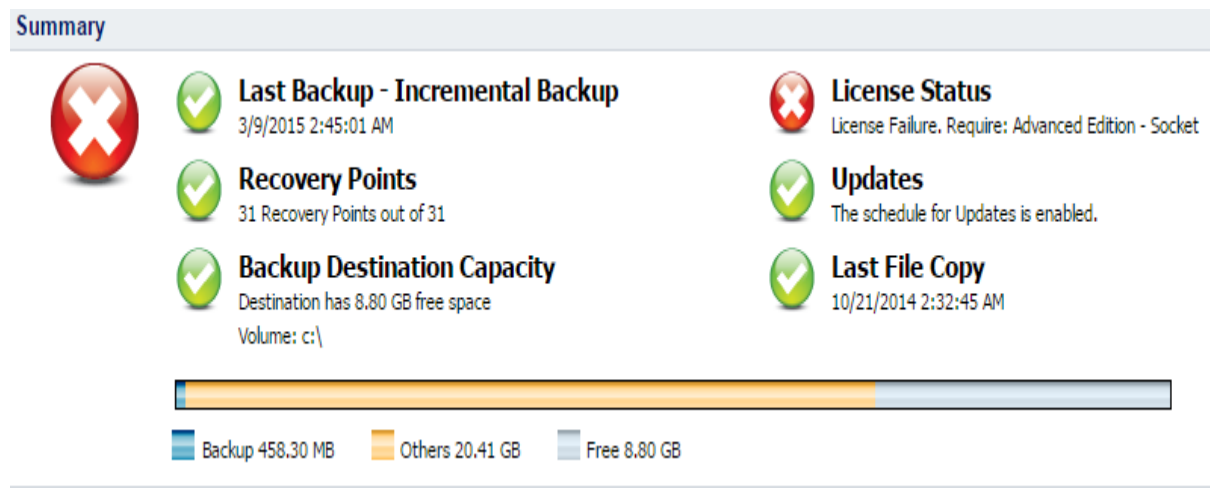
When a merge job is running, you can manually pause it from the Arcserve UDP Agent (Windows) home page **Job Monitor**.



If you manually pause a merge job, you must manually click **Resume** for the merge job to continue. For more information, see [Merge Job Guidelines](#) in the online help.

Status Summary

The **Status Summary** section of the home page provides a quick and easy, high-level status of your backup health.



Last Backup

Displays the date and time of the last backup, with the status of that backup.

- ♦ Green icon - Indicates that the last backup was successful and your computer is safely protected.
- ♦ Red icon - Indicates that the last backup was not successful, your most recent backup failed, and the computer cannot be restored with that recovery point.
- ♦ Yellow icon - Indicates that backups have not been performed for your computer and your computer is not protected.

Recovery Points/Recovery Sets

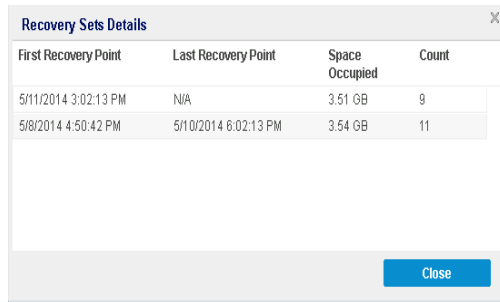
Displays the number of Recovery Points or Recovery Sets for your monitored server based on your specified retention settings.

- ♦ Green icon - Indicates that you have reached the specified number of recovery points or recovery sets.
- ♦ Red icon - Indicates that there are no saved recovery points or recovery sets and you have a potentially dangerous backup environment.
- ♦ Yellow icon - Indicates that you have at least one recovery point or one recovery set, but have not reached your specified number of recovery points or recovery sets.

If you specified your retention settings based on recovery sets, the status summary displays the number of recovery sets already retained and the number of recovery

sets in progress. In addition, click the link under **Recovery Sets** to display the **Recovery Sets Details** dialog. This dialog contains detailed information about the contents of the recovery set.

Note: The **Recovery Sets** option is available if you select **Standard** as your **Backup Data Format**. However, the **Recovery Sets** option is not available if you select **Advanced** as your **Backup Data Format**. For more information about recovery sets, see [Specify Retention Settings](#) in the online help.



First Recovery Point	Last Recovery Point	Space Occupied	Count
5/11/2014 3:02:13 PM	N/A	3.51 GB	9
5/8/2014 4:50:42 PM	5/10/2014 8:02:13 PM	3.54 GB	11

First Recovery Point

The date/time of the first backup in the recovery set.

Last Recovery Point

The date/time of the last backup in the recovery set. With the first/last recovery point time listed, you are able to determine the complete time range of the recovery set.

Space Occupied

The total size of the recovery set. This number can be used to calculate how much disk space is used by the recovery set.

Count

The number of recovery points that belong to the recovery set.

Destination Capacity

Displays the amount of free space available at your backup destination. The Destination Capacity display provides an additional indication of the amount of space that is used for the backups, the amount of space that others use, and the amount of available free space.

- ♦ Green icon - Indicates that the amount of available free space is above the safe level.
- ♦ Yellow icon - Indicates that the amount of available free space is reduced to 3% of your destination capacity. It can be configured from the Windows registry.
- ♦ Red icon - Indicates either of the following conditions:

- The specified destination is not accessible.
- The amount of available free space is reduced to 100 MB of your destination capacity. It can be configured from the Windows registry.
- Immediately increase the free space capacity of the backup destination or change the destination to another location which has adequate space.

Note: You can set up an email alert notification when the amount of unused space at the backup destination is less than a specified value. For more information about setting up this email alert notification, see [Specify Email Alert Preferences](#) in the online help.

License Failure

If a backup fails because of a license validation failure, the License Failure status is displayed indicating which license was the cause of the failure.

Updates

Displays the status of Arcserve UDP Agent (Windows) updates for your computer.

- ◆ Green icon - Indicates the Arcserve UDP Agent (Windows) Updates function is enabled. Your computer is able to connect to the download server, your **Update Schedule** is configured, and no new updates are available.
- ◆ Yellow icon - Indicates either of the following conditions:
 - The latest available update is not installed on your computer.
You can then click **Click here to install updates** to trigger the installation of the update.
 - Your **Update Schedule** has not been configured.
For more information about configuring the Update Schedule, see [Specify Updates Preferences](#) in the online help.

Note: All updates that are released for Arcserve UDP Agent (Windows) are cumulative. As a result, each update also includes all previously released updates to help ensure that your computer is always up-to-date.

- ◆ Red icon - Indicates Arcserve UDP Agent (Windows) is not able to connect to the download server. When this red icon is displayed, it means that you must provide valid download server details on the **Updates** tab of the **Preferences** dialog.

Last File Copy

Displays the date and time of the last File Copy job, with the status of the File Copy job that was performed.

- ◆ Green icon - Indicates that the last File Copy job was successful.
- ◆ Red icon - Indicates that the last File Copy job was not successful.
- ◆ Yellow icon - Indicates that the last File Copy job was incomplete or canceled.

In addition, the Last File Copy status indicator also displays the actual amount of space that is freed on the disk by the File Copy job. This space saved calculation is only displayed if you select to move the File Copy to a different location instead of copying the File Copy to as different location. This value is based upon the actual size of the backup that was moved from the computer to the specified destination. A File Copy job that does not move the copy to a different location, does not save any space.

Destination Space Usage Status Bar

- ◆ Backup - Total amount of space used for all backup sessions on the destination.

Note: For Windows Server 2012 NTFS, the Backup size displayed is the data deduplication un-optimized size. If the Arcserve UDP Agent (Windows) backup destination is configured with Windows NTFS data deduplication enabled, the Backup size may be larger than the actual data size on the disk.

- ◆ Others - Non-Arcserve UDP Agent (Windows) data size on the destination.
- ◆ Free - Amount of available space on the destination.

Note: If the backup destination is set to data store, this status bar does not display.

Data Protection Summary

The **Data Protection Summary** section of the home page displays status information for the available events (backups/file copy).

Protection Summary					
Job Type	Count	Data Protected	Space Occupied	Last Successful Event	Next Event
Full Backup	3	2.04 GB	N/A	9/23/2013:11:02 AM	9/24/2013 :11:00 PM
Incremental Backup	7	1.78 MB	N/A	9/23/2013 7:30:10 PM	9/24/2013 :11:00 PM
Verify Backup	0	0 Bytes	N/A		
File Copy	24	N/A	0 Bytes	9/23/2013 10:07:14 PM	
Copy Recovery Point	0	0 Bytes	0 Bytes		

For each type of backup job (Full, Incremental, and Verify) and each File Copy job, this summary displays the following:

Count

For each type of event, indicates the number of successful backup/File Copy jobs that were performed (scheduled or non-scheduled).

Data Protected

The amount of data protected from the source. This is the size of data that was backed up from the source volume during a backup job (without deduplication and compression).

Space Occupied

The amount of space occupied (saved) at the destination.

Note: For a deduplication data store, "N/A" will be displayed for the **Space Occupied** field. For a non-deduplication data store/share folder, the actual data size will be displayed.

Last Successful Event

For each type of event, indicates the date and time that the last successful event was performed.

Next Event

For each type of event, indicates the next scheduled event. If this column is blank, it indicates that you do not have a schedule for this type of event or the non-recurring schedule has been satisfied.

Most Recent Events Summary

The **Most Recent Events** section of the home page displays the most recent events (backup jobs), with the corresponding status, the type of event performed, the date and time of the event, the size of the data protected (backed up) from the source, the amount of space occupied (saved) at the destination, and the status of the corresponding File Copy job. It also includes the name of the event (if specified by the user). You can click on a specific date to display the corresponding events for that selected date.

A flag in the status column indicates that a full backup is the starting backup of a recovery set.

Most Recent Events		Status	Schedule Type	Backup Type	Date/Time	Data Protected	Space Occupied
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							
✓	Custom	Incremental Backup	4/8/2014 9:41:11 AM	2.81 MB	1.14 MB		
✓	Custom	Incremental Backup	4/8/2014 9:39:53 AM	2.81 MB	1.09 MB		
✓	Custom	Incremental Backup	4/8/2014 9:38:31 AM	2.94 MB	1.12 MB		
✓	Custom	Full Backup	4/8/2014 9:35:44 AM	253.56 MB	205.94 MB		

The calendar displays the dates of the most recent events, highlighted in corresponding status colors.

- Green - All backup attempts for that day were successful.
- Red - All backup attempts for that day were unsuccessful (failed or canceled).
- Yellow - Backup attempts for that day were not all successful or not all unsuccessful (mix of successful and unsuccessful backup attempts).


Note: A diagonal marker in the top left corner of a date indicates that the day

contains the start of a recovery set.



Functionality Taskbar

The Functionality taskbar section of the home page provides a means to initiate the various Arcserve UDP Agent (Windows) functions.



Backup Now

Lets you run a Full, Incremental, or Verify ad-hoc backup immediately, based on current backup settings. The backup settings are configured and defined through the **Backup Settings**. For more information, see [Perform Backup Manually \(Backup Now\)](#) in the online help.

Restore

Lets you perform a file level or application-level restores to the original location or to an alternate location. When you select this function, you then specify to which restore option you want to use to locate the backup image to be restored. For more information, see Restore Methods in the online help.

Settings

Lets you configure/modify the following settings:

- ♦ **Backup Settings** (backup destination, schedule, retention count, and so on). For more information, see [Configure or Modify Backup Settings](#) in the online help.
- ♦ **File Copy Settings** (source, destination, schedule, retention count, filters, and so on). For more information, see Manage File Copy Settings in the online help.

- ◆ **Copy-Recovery Points** (scheduled export of recovery points). For more information, see [Configure the Copy Recovery Point Settings](#) in the online help.
- ◆ **Preferences** (enable Email Alerts and Updates). For more information, see [Specify Preferences](#) in the online help.

Copy Recovery Point

Lets you view a list of available recovery points (successful backups) and select which recovery point you want to use to create a consolidated copy. This consolidated copy combines the blocks from the previous full and all incremental backups leading to the selected recovery point. In addition, the consolidated copy also removes any unused blocks (reducing the image size) to lets you gain more efficient use of your backup resources.

Each recovery point represents a point in time when a VSS snapshot image was captured and contains not only the data, but also all information relating to the operating system, installed applications, configuration settings, necessary drivers, and so on. For more information, see [How to Copy a Recovery Point](#) in the online help.

Mount Recovery Point


Provides the ability to mount a recovery point to a drive letter (volume) or an NTFS folder, to view, browse, copy, or open the backup files directly in Windows Explorer.

View Logs

Lets you view logs of activities that are performed during operations such as backup, restore, and copy. The **Activity Log** displays the status of the job, including such details as the throughput, compression size, elapsed time, encryption status, and so on. For more information, see [View Logs](#) in the online help.

Support and Community Access Taskbar

The **Support and Community Access** taskbar section of the home page provides a means to initiate the various support-related functions.



To avoid any delays in response and help ensure that your communication is routed to the proper destination, it is important to know and understand which of these sites are used for which functions.

For example:

- If you discover a bug in the product, you should select the **Online Support** link and post the problem there. By doing this, the Arcserve Support team can directly assist you in resolving your problem in a productive and efficient manner.
- If you have a suggestion or idea for improving the product for the next release, you should select the **Provide Feedback** link. By doing this, you can interact directly with the Arcserve team to work together to improve the product and help make it better for you.

Knowledge Center

Provides a complete "One Stop Knowledge Center" for all product related information. It can be accessed directly from the product and Arcserve Support.

Videos

Provides access to view various Arcserve UDP Agent (Windows)-related videos. These videos are for basic Arcserve UDP Agent (Windows) features and procedures.

Online Support

Provides access to "One Stop Support" from where you can resolve issues and get important product information.

Provide Feedback

Provides access to the Arcserve team from where you can view Frequently Asked Questions, ask your own questions, share ideas, and report any problems.

Note: Available in English only.

Live Chat

Provides real-time monitoring and live help/support. Lets you optimize intelligent conversation between you and the Arcserve UDP Agent (Windows) Technical Support team, allowing you to address your concerns and questions immediately, while still maintaining access to the product. Upon completion of the chat session, you can send a transcript of the chat to your email address.

Note: You may need to add the Live Chat link to your Trusted sites.

Miscellaneous User Interface Components


This section provides information about the following:

Managing Console

Managing Console refers to the **Managed By** field on the Arcserve UDP Agent (Windows) home page UI. If the agent is being centrally managed by an Arcserve UDP Console, this link allows you to open the Arcserve UDP Console. If it is not centrally managed, this field is not displayed.

Help Menu

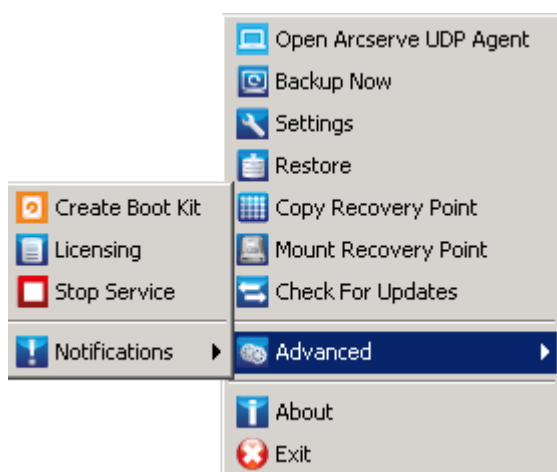
The **Help** menu on the home page provides a quick and easy access to the Arcserve UDP Knowledge Center, Arcserve UDP Agent (Windows) Help, User Guide, and the About Arcserve UDP Agent (Windows) page. In addition, you can launch a manual **Check for Updates** from this menu.



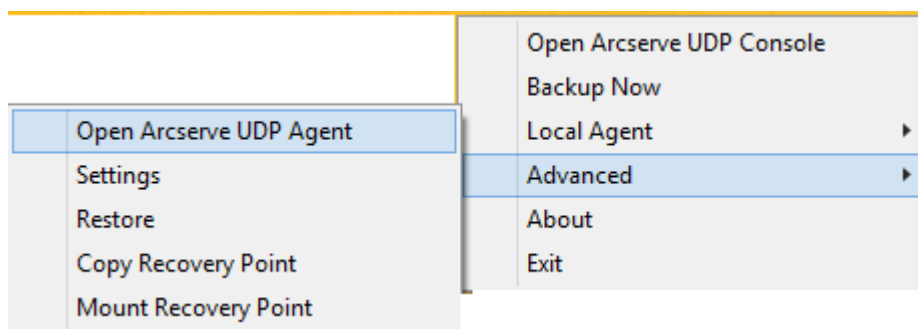
Arcserve UDP Agent (Windows) Monitor

From the Arcserve UDP Agent (Windows) Monitor, you can access many of the same task functions that you can access from the Arcserve UDP Agent (Windows) home page. You can launch the following tasks from the tray monitor: Open the Arcserve UDP Agent (Windows) home page, **Backup Now**, **Settings**, **Restore**, **Copy Recovery Point**, **Mount Recovery Point**, **Check for Updates**, and **Advanced**. From the **Advanced** option, you can access additional subordinate options such as **Create Boot Kit**, **Licensing**, **Start/Stop Service**, and configure alert Notifications (None, Errors and Warnings, or All).


When the Arcserve UDP Agent (Windows) is being managed by Console:



When the Arcserve UDP Agent (Windows) is managed by console and not protected in a plan:



In addition, when you hover your cursor over the Arcserve UDP Agent (Windows) Monitor icon, a backup status overview message is displayed. An animated icon indicates when any job is running and includes the progress completed. You can easily determine if a job (backup, restore, file copy, copy recovery point, catalog, or granular restore catalog) is running without logging in to Arcserve UDP Agent (Windows).



Troubleshooting User Interface 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. 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.

Unable to display Arcserve UDP Agent (Windows) home page in IE

If you are using an Internet Explorer (IE) web browser to access the Arcserve UDP Agent (Windows) home page and it does not display, it may be that the Arcserve UDP Agent (Windows) website is not included as a "Trusted Site" in your IE browser.

If this condition occurs, add this website as a Trusted Site in your IE browser. For more information about adding a website as a Trusted Site, see [Security zones: adding or removing websites](#).

If the problem persists, use [Live Chat](#) to contact Arcserve support. Live Chat lets you optimize intelligent conversation between you and the Technical Support team, allowing you to address your concerns and questions immediately, while still maintaining access to the product.

Job Monitor data speed displays a 0 or some other abnormal value

Symptom

Windows Performance Counters are disabled.

Solution

From the Registry Editor, delete or enable the following registry keys on all Windows versions:

- Perflib

Path: HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib

Name: "Disable Performance Counters"

Type: DWORD

Value: Set to 0 to enable performance counter.

- Performance

Path: HKLM\SYSTEM\CurrentControlSet\Services\PerfProc\Performance

Name: "Disable Performance Counters"

Type: DWORD

Value: Set to 0 to enable performance counter.

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.

Chapter 4: Settings

This section contains the following topics:

Configure or Modify Backup Settings	148
Manage File Copy Settings	199
Manage File Archive Settings	220
Configure the Copy Recovery Point Settings	239
Specify Preferences	245
Manage Export/Import Settings	262

Configure or Modify Backup Settings

Before you perform your first backup, you must configure the backup settings which are applied to each backup job. These settings can be retained for future backup or they can be modified at any time from the Arcserve UDP Agent (Windows) home page.

The settings let you specify behaviors such as:

- Backup source and destination.
- Schedule standard or advanced settings for each type of backup.
- Advanced settings for your backup jobs.
- Any pre or post backup operations.

Note: For more information about these Backup Settings, see [How to Perform a Backup](#).

To manage the backup settings, click the **Settings** link on the Arcserve UDP Agent (Windows) home page to display the **Backup Settings** dialogs and these subordinate tab options:

Specify Protection Settings

Protection settings for the information that is going to be backed up ensures that the backup data is reliably protected (copied and saved) against any form of data loss.

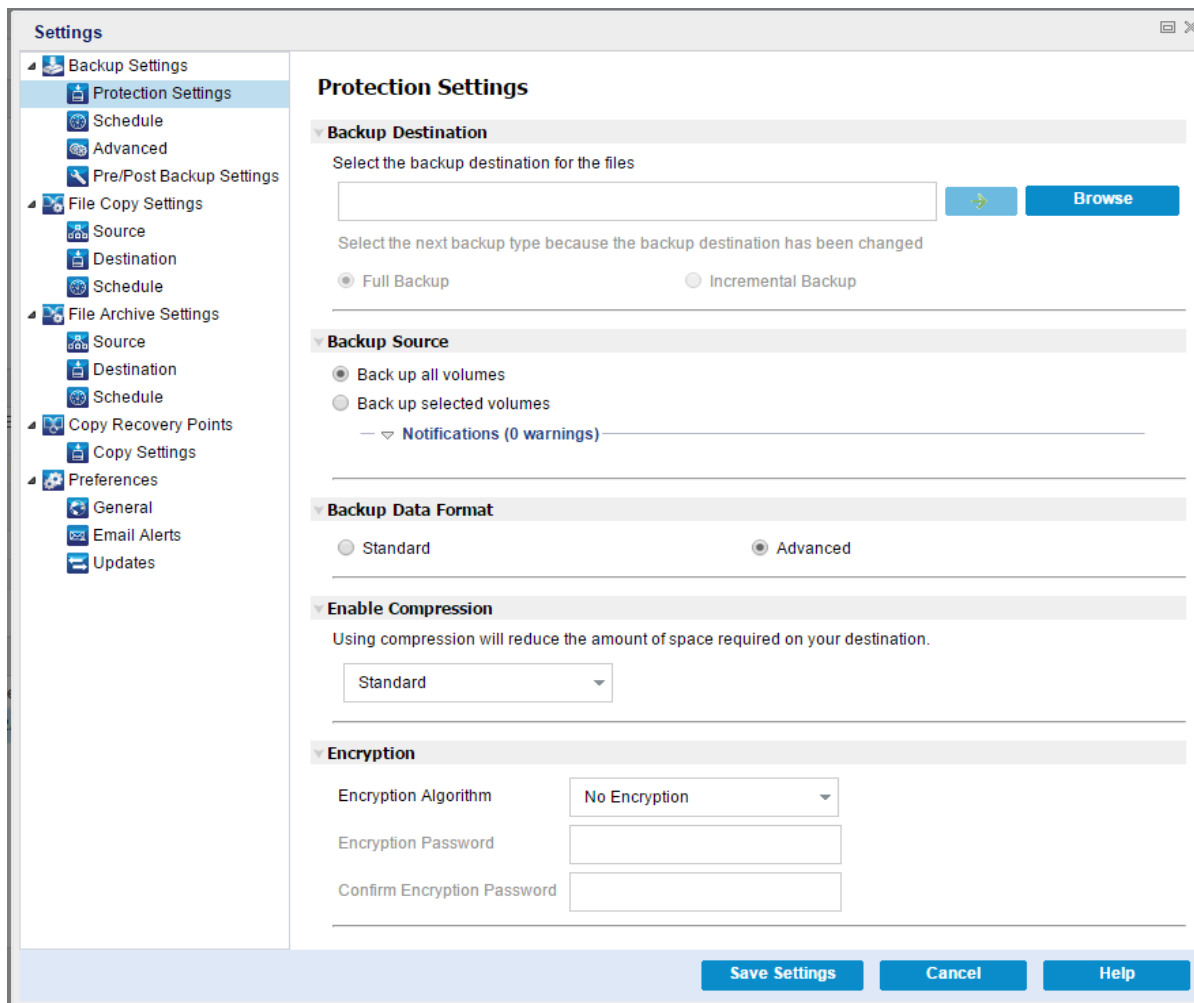
Specify the protection settings

1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Backup Settings** tab. When the **Backup Settings** dialog opens, select **Protection Settings**.

The **Protection Settings** dialog opens.

Notes:

- If the Arcserve UDP Agent (Windows) is being managed by Console, not all settings are available and will be displayed as read-only information.
- When agent is managed by console and not protected in a plan, all the settings are still available except the Preference > Updates panel.



2. Specify the **Backup Destination**.

♦ **Use local disk or shared folder**

You can specify a local path (volume or folder), or remote shared folder (or mapped drive) for the backup location, or you can browse to a backup location.

Click the green arrow icon button to verify the connection to the specified location.

- If you entered a local path for the destination, this button is disabled.
- If you enter a network path and click this button, you are prompted to provide the username and password.
- If you are already connected to this path successfully, and click the arrow button you can change the username and password you used to connect.

- If you do not click the arrow button, the destination path is verified. If needed, you are prompted for the username and password.
- a. If you want to back up to your local path (volume or folder), the specified backup destination cannot be the same location as your backup source. If you inadvertently include the source in your destination, the backup job ignores this portion of the source and it is not included in the backup.

Example: You attempt to back up your entire local machine consisting of Volumes C, D, and E and also specify Volume E as your destination. The Arcserve UDP Agent (Windows) only backs up Volumes C and D to Volume E. Data from Volume E is not included in the backup. If you want to back up all local volumes, specify a remote location for your destination.

Important! Verify that your specified destination volume does not contain system information. Or else it will not be protected (backed up) and your system will fail to recover after Bare Metal Recovery (BMR) if necessary.

Note: Dynamic disks are restored at disk-level only. If your data is backed up to a volume on a dynamic disk, you are not able to restore this dynamic disk during BMR.

- b. If you want to back up to a remote shared location, specify a location path or browse to the location. You also have to provide user credentials (Username and Password) to access the remote machine.
- c. If the backup destination has changed after the last backup was performed, select the backup type: Full Backup or Incremental Backup. These options are only enabled when you change your backup destination.

Default: Full Backup

Note: If the backup destination has changed and catalog jobs are pending, the catalog job first runs and completes on the old location before running on the new location.

Full Backup

The next backup that is performed is going to be a Full Backup. The new backup destination does not have any dependency on the old backup destination. If you continue with a full backup, the previous location is no longer needed for backups to continue. You can select to keep the old backup for any restores or if you do not want to perform

any restores from there you can delete it. The old backup will not affect future backups.

Incremental Backup

The next backup that is performed is going to be an Incremental Backup. The next incremental backup to the new destination is performed without copying all the backups from the previous destination. However, for this option, the new location is dependent on the previous location because the changes include only the incremental data (not the full backup data). Do not delete the data from the previous location. If you change the backup destination to another folder and attempt to perform an incremental backup, and the former backup destination does not exist, the backup fails.

Note: With the Full installation of Arcserve UDP, you can specify to use an Arcserve UDP Recovery Point Server as the backup location. If you do, the Protection Settings Backup Destination displays the Arcserve UDP Recovery Point Server Settings, including the Hostname, Username, Password, Port, Protocol, and the Plan Summary.

3. Specify the **Backup Source**.

You can back up the entire machine or selected volumes.

Back up the entire machine

Lets you back up the entire machine. All volumes on the machine are backed up.

Note: If you select the full machine backup option, Arcserve UDP Agent (Windows) automatically discovers all disks or volumes attached to the current machine and Arcserve UDP Agent (Windows) includes them in the backup.

Example: If a new disk is attached to the machine after the backup setting is configured, you do not need to change the backup settings and the data on the new disk will be protected automatically.

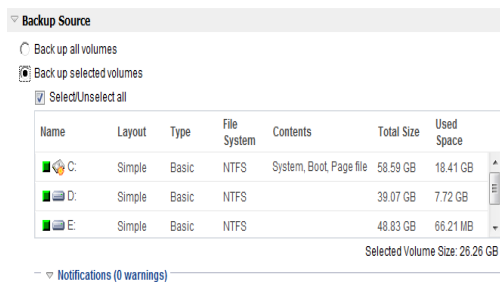
Select individual volumes to back up

This volume filtering capability lets you specify to back up only the selected volumes. You also have the option to Select or clear selection of all listed volumes.

Note: If some volumes are selected explicitly for backup, only the selected volumes are backed up. If a new disk or volume is attached to the machine, manually change the volume selection list to protect the data on the new disk or volume.

When you select this option, a listing of all available volumes display, with the corresponding volume information and notification messages.

Note: Computers that adhere to the Extensible Firmware Interface (EFI) use the EFI System Partition, which is a partition on a data storage device. The EFI System partition is critical for Bare Metal Recovery (BMR). Therefore, when you select boot volume "C" on a UEFI system, the EFI System Partition is selected automatically for the backup source for BMR and an information message is displayed.



Name

Specifies the name of the volume drive letter, mount point, volume GUID (Globally Unique Identifier) name.

Layout

Indicates the simple, spanned, mirror, striped, RAID5 (backup of a RAID 5 volume on Microsoft Dynamic Disks is not supported; but backup of hardware RAID is supported).

Type

Indicates the type, basic or dynamic.

File system

Lists the following file systems: NTFS, ReFS, FAT, FAT32 (backup of FAT, FAT32, and exFAT is not supported).

Contents

Indicates whether the application is (SQL/Exchange), System, Boot, Page file, Removable Device, VHD, 2-TB Disk.

Total size

Specifies the size or the capacity of the volume.

Used Space

Indicates the space, files or folders and volume data occupies.

The notification messages display for any of the following conditions:

- **Local volume related**

If the specified backup destination is on the local volume, a warning message displays notifying you that this volume is not backed up.

- **BMR related**

If system/boot volume is not selected for backup, a warning message displays notifying you that the backup is not usable for BMR.

If you select boot volume "C" on a UEFI system, the EFI system partition is selected automatically for the backup source for BMR and an information message is displayed.

- **Application related**

If the application data files are on a volume that is not selected for backup, the application name and database name display for reference.

4. Specify the **Backup Data Format**.

Standard

Standard Backup Data Format allows you to set the number of recovery points to retain or the number of recovery sets to retain and includes a basic repeat backup schedule. The Standard format is the legacy format used in releases of Arcserve D2D and Arcserve Central Applications.

Advanced

Advanced Backup Data Format allows you to set the number of recovery points to retain and includes advanced scheduling. The Advanced format is a new data storage format, dividing source disks into multiple logical segments. Compared to the Standard format, backup, restore, and merge job throughputs are greatly improved.

If the **Advanced Backup Data Format** is selected, advanced scheduling will be enabled. Advanced scheduling consists of the following:

- Week-based repeat backup schedule
- Week-based backup throttling schedule
- Week-based merge schedule
- Daily backup schedule
- Weekly backup schedule
- Monthly backup schedule

5. Specify the **Retention Setting** if you selected **Standard** as the **Backup Data Format**.

Note: If you selected **Advanced** as the **Backup Data Format**, the retention setting is specified on the **Advanced Schedule Settings** dialog.

You can set the retention setting based on the number of recovery points to retain (merges sessions) or based on the number of recovery sets to retain (deletes recovery sets and disables infinite incremental backups).

Default: Retain Recovery Points

Recovery Point

This is the recommended option. With this option selected, you can fully leverage the infinite incremental backup capabilities and save storage space.

Note: If you selected **Advanced** as the **Backup Data Format**, then you can only specify the number of recovery points to retain.

Recovery Set

This option is generally used for large storage environments. With this option selected, you can create and manage backup sets that help you manage your backup window time more efficiently when you are protecting a large amount of data. You can use this option when backup time is a priority over space constraints.

Note: Recovery sets are only available if you are backing up to a location that is not a data store. Recovery sets are not supported with RPS deduplication. They are also not available for Advanced format backup to non-RPS locations.

For more information about setting the Recovery Point and Recovery Set options, see [Specify Retention Settings](#).

6. Specify the type of **Compression**.

Specifies the type of compression that 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

No compression is performed. 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. Standard compression 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 or ZIP files), you can allocate additional storage space to handle such data. As a result, if you select any compression option and you have uncompressible data in your backup, it can 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 that is performed after this compression level change is automatically a Full Backup. After the Full Backup is performed, all future backups (Full, Incremental, or Verify) will be performed as scheduled.
- If your destination does not have sufficient free space, you can consider increasing the Compression setting of the backup.

7. Specify the **Encryption** settings.

- a. Select the type of encryption algorithm that is used for backups.

Data encryption is the translation of data into a form that is unintelligible without a deciphering mechanism. The Arcserve UDP Agent (Windows) data protection uses secure, AES (Advanced Encryption Standard) encryption algorithms to achieve maximum security and privacy of your specified data.

The available format options are No Encryption, AES-128, AES-192, and AES-256. (To disable encryption, select No Encryption).

- ◆ A full backup and all its related incremental and verify backups must use the same encryption algorithm.
- ◆ If the encryption algorithm for an incremental or verify backup is changed, a full backup must be performed. This means after changing encryption algorithm, the first backup will be full, despite the original backup type.

For example, if you change the algorithm format and you submit a customized incremental or verify backup manually, it is automatically converted to a full backup.

- 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 the 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 you submit a customized incremental or verify backup manually, it is automatically converted to a full backup.
- c. Arcserve UDP Agent (Windows) provides encryption password management so that you do not need to remember encryption passwords.
 - Password is also encrypted.
 - Password is remembered and not required (if you restore to the same machine).
 - Password is required if you restore to a different machine.
 - Password is not required if you are attempting to export a recovery point that contains encrypted data and the recovery point belongs to backups performed on the current machine.
 - Password is always required if you are attempting to recover encrypted data from an exported recovery point.
 - Password is not required to browse to an encrypted recovery point.
 - Password is required to perform a BMR.
- d. When encryption is enabled, the activity log is updated.
 - A message is recorded in the activity log to describe the selected encryption algorithm for every backup.
 - A message is recorded in the activity log to indicate why an incremental or verify backup was converted to a full backup (password change or algorithm change).

Note: Encryption settings do not have to remain the same for your backups. You can change these settings at any time, including after several backups of the same data.

8. Specify the **Throttle Backup**.

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. For a backup job, the Job Monitor on the Arcserve UDP Agent (Windows) home page displays the average Read and Write speed of the job in progress and the configured throttle speed limit.

Note: By default, the **Throttle Backup** option is not enabled and backup speed is not being controlled.

9. Calculate the **Estimated Backup Size**.

Displays the estimated usage of the destination volume.

Note: For more information about using these estimated backup calculations, see [Estimate Future Backup Space Requirements](#).

10. Click **Save Settings**.

Your backup protection settings are saved.

Specify Retention Settings

The retention setting for the **Standard Backup Data Format** can be based on the number of recovery points to retain (merges sessions) or based on the number of recovery sets to retain (deletes recovery sets and disables infinite incremental backups).

◆ Retain Recovery Points

Select this option to set your retention setting based on the number of recovery points to retain instead of on the number of recovery sets to retain.

Note: The recovery points to retain is set in the **Protection Backup Settings** if you select **Standard** as the **Backup Data Format**. The recovery points to retain is set in the **Advanced Schedule Settings** if you select **Advanced** as the **Backup Data Format**.

▼ **Backup Data Format**

Standard
 Advanced

▼ **Retention Setting**

Retain Recovery Points
 Retain Recovery Sets

Specify the number of recovery points to retain:

Run the merge job:

As soon as possible
 Each day during the following time range

From :
To :

Specify the number of recovery points to retain

Specifies the number of recovery points (full, incremental, and verify backup images) retained. When the number of recovery points present on the destination exceeds the specified limit, the earliest (oldest) incremental backups beyond the retention count are merged into the parent backup to generate a new baseline image consisting of the "parent plus oldest childs" blocks. If there are multiple sessions available for merge, the oldest child backups will be merged into the parent

backup in a single pass, if the backups are compressed. If the backups are not compressed, then only the oldest child backup will be merged into the parent backup and this cycle repeats for each subsequent child backup to be merged.

Specifying the number of recovery points to retain allows you to perform infinite incremental backups, while maintaining the same retention count. For more information, see [Merge Job Guidelines](#).

Note: If your destination does not have sufficient free space, you can consider reducing the number of saved recovery points.

Default Value: 31

Minimum Value: 1

Maximum Value: 1344

Note: The Arcserve UDP Agent (Windows) home page **Summary** section indicates how many recovery points are retained out of the number specified. For more information, see [Status Summary](#) in the online help.

Run the merge job:

As soon as possible

Select this option to run the merge job at any time.

Each day during the following time range

Select this option to run the merge job each day only within the specified time range. Setting a time range helps to avoid the merge job introducing too many I/O operations to the production server if the merge job runs for a long time.

Note: When setting the time range to run the merge job, ensure that you specify a time range that will allow the related backup jobs to complete prior to the start of the merge.

▪ **Retain Recovery Sets**

Select this option to set your retention setting based on the number of recovery sets to retain instead of on the number of recovery points to retain. With this setting you can disable infinite incremental backups, without merging any sessions. Using recovery sets helps reduce the amount of time it takes to complete merge jobs.

Note: The **Recovery Sets** option is available if you select **Standard** as your **Backup Data Format**. However, the **Recovery Sets** option is not

available if you select **Advanced** as your **Backup Data Format**.

▼ **Backup Data Format**

Standard
 Advanced

▼ **Retention Setting**

Retain Recovery Points
 Retain Recovery Sets

! When you specify a number of recovery sets to retain, ensure that you have enough free space available for the specified number plus two additional full backups.

! The retention setting has been changed. Use new backup destination to start backups with new retention setting.

Specify the number of recovery sets to retain.

Start a new recovery set on every:

Selected day of the week

Selected day of the month

Start a new recovery set with:

First backup on the selected day

 Last backup on the selected day

Specify the number of recovery sets to retain

Specifies the number of recovery sets retained. A recovery set is a series of backups, starting with a full backup, and then followed by a number of incremental, verify, or full backups.

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. A flag in the status column on the Arcserve UDP Agent (Windows) home page **Most Recent Events** section indicates that a full backup is the starting backup of a recovery set. 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.

When the specified limit is exceeded, the oldest recovery set is deleted (instead of merged).

Default Value: 2

Minimum Value: 1

Maximum Value: 100

Note: If you want to delete a recovery set to save backup storage space, reduce the number of retained sets and Arcserve UDP Agent (Windows) automatically deletes the oldest recovery set. Do not attempt to delete the recovery set manually.

Example 1 - Retain 1 Recovery Set:

- Specify the number of recovery sets to retain as 1.

Arcserve UDP Agent (Windows) always keeps two sets in order 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.

Arcserve UDP Agent (Windows) will delete 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 at the last backup on Friday.
- 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 PM on Friday is run, it will be converted to a full backup and marked as the starting backup of the recovery set.

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, since a given month may have 28, 29, 30, or 31 days, you can specify the last day of the month as the day to create the recovery set.

Start a new recovery set with:**First backup on the selected day**

Indicates you want to start a new recovery set with the first scheduled backup on the specified day.

Last backup on the selected day

Indicates you want to start a new recovery set with the last scheduled backup on the specified day. If the last backup is selected to start the set and for any reason the last backup did not run, then the next scheduled backup will start the set by converting it to a full backup. If the next backup is run ad-hoc (for example an emergency situation requires a quick incremental backup), you can decide if you want to run a full backup to start the recovery set or run an incremental backup so that the next backup starts the recovery set.

Note: The last backup may not be the last backup of the day if you run an ad-hoc backup.

The Arcserve UDP Agent (Windows) home page **Summary** section indicates how many recovery sets are retained (or in progress) out of the number specified. Click the link under **Recovery Sets** to display the **Recovery Sets Details** dialog. This dialog contains detailed information about the contents of the recovery set. For more information about this dialog, see [Status Summary](#) in the online help.

Estimate Future Backup Space Requirements

Arcserve UDP Agent (Windows) provides you with this tool to calculate the estimated amount of available free space that you will need for backups. The calculations are based on your estimate of future data change and on the space that is occupied from previous backups.

Estimated Backup Size

The graph below shows the estimated usage of the destination volume. You can change the Space Saved After Compression or the Change Rate to see their effect on the estimated backup size.

Estimated backup 0.72 GB Used 115.56 GB
Free 1362.28 GB



i Actual disk space used by current backups is: 1.70 GB.

Estimated Values

Space Saved After Compression	10%	▼
Change Rate	10%	▼
Space Saved After Windows Deduplication	0%	▼

Estimated Backup Size

Total Source Size	282.57 MB
Compressed Full Backup Size	254.31 MB
Compressed Incremental Backup Size	483.19 MB
Estimated Total Backup Size	737.50 MB

To use this estimating tool

1. Select the backup source. This can be your entire machine or selected volumes within your machine.

The actual size of the selected backup source is displayed in the **Total Source Size** field.

2. Estimate the anticipated **Change Rate** for future backups.

Base this estimate upon past performance of how much your total backup size has changed for each subsequent incremental backup.

With the Estimated Values defined, Arcserve UDP Agent (Windows) calculates and displays the estimated backup size required based on the configuration of the

backup destination and the recovery points. The pie chart also displays the amount of used space and free space.

3. Estimate the **Space Saved After Compression** percentage value.

Estimated Values

You can use estimated values to calculate the approximate overall backup size that is based on the number of recovery points. Base this estimate upon past performance of your backups with different Compression settings applied. As you change this value, you will see the corresponding size impact for your backup sizes.

Note: If necessary, you can perform some Full Backups, each with a different Compression setting (No Compression, Standard Compression, and Maximum Compression) to establish past performance values and help you to better calculate the percent of space saving that each setting produces for your backup

- ◆ **Space Saved After Compression**

This value indicates how much disk space is saved after compression.

Example: If the data size of a volume is 1000 MB and after backup the compressed data size is 800 MB, then the Space Saved After Compression is estimated to be 200 MB (20%).

- ◆ **Change Rate**

This value indicates the typical data size of an incremental backup.

Example: If an incremental backup data size is 100 MB and the full backup data size is 1000 MB, the change rate is estimated to be 10%.

- ◆ **Space Saved After Windows Deduplication**

This value indicates how much disk space is saved after Windows deduplication.

If the backup destination directory is located on a volume where Windows deduplication is enabled, the estimated backup size may exceed the total capacity of the volume. The reason is that with deduplication enabled, only one copy of multiple same size data blocks is kept. This value helps to estimate the size while taking deduplication into consideration.

Example: If the total size of the source backed up is 100 GB and it has 20 GB of data that is redundant, then the space saved after deduplication will be 20 GB.

Estimated Backup Size

Displays the estimated values for **Total Source Size**, **Compressed Full Backup Size**, **Compressed Incremental Backup Size**, and **Estimated Total Backup Size**.

- ◆ The **Compressed Full Backup Size** field displays a calculated value that is based upon:
 - Size of the backup source
 - Specified compression percentage.
 - ◆ The **Compressed Incremental Backup Size** field displays a calculated value that is based upon:
 - Estimated Change Rate
 - Number of recovery points to be saved
 - Specified compression percentage
 - ◆ The **Estimated Total Backup Size** field will display the anticipated space that you will need for future backups and is based upon:
 - Amount of space that is required for one Full Backup plus
 - Amount of space that is required for the number of Incremental Backups needed to satisfy the specified number of saved recovery points.
4. From this **Estimated Total Backup Size** value, you should be able to determine if your backup destination has sufficient space to fit your backup.

If your destination does not have sufficient free space, you can consider the following corrective actions:

- ◆ Reduce the number of saved recovery points.
- ◆ Increase the available free space at the backup destination.
- ◆ Change the backup destination to a larger capacity.
- ◆ Reduce the size of the backup source (maybe eliminate unnecessary volumes from the backup).
- ◆ Increase the Compression setting of the backup.

Specify Schedule Settings

Arcserve UDP Agent (Windows) lets you specify the schedule for your backups. If you set the **Protection Settings Backup Data Format** to **Standard**, the **Standard Schedule** dialog opens, where you can specify the standard schedule settings. If you set the **Protection Settings Backup Data Format** to **Advanced**, the **Advanced Backup Schedule** dialog opens, where you can specify the advanced schedule settings.

Specify Standard Schedule Settings

Arcserve UDP Agent (Windows) lets you specify the schedule for your backups. If you set the **Backup Data Format** option to **Standard** in **Protection Settings**, the **Standard Schedule** dialog opens, where you can specify the standard schedule settings.

Follow these steps:


1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Backup Settings** tab. When the **Backup Settings** dialog opens, select **Schedule**.

The **Backup Settings StandardSchedule** dialog opens.

Notes:

- If the Arcserve UDP Agent (Windows) is being managed by Console, not all settings are available and will be displayed as read-only information.

- When agent is managed by console and not protected in a plan, all the settings are still available except the Preference > Updates panel.



2. Specify your backup schedule options.

Set start date and time

The start date and start time for your scheduled backups.

Note: When setting the interval between repeat backup jobs, ensure that you leave enough time to allow the previous job and any related merge jobs to complete before the next backup job starts. This amount of time can be estimated based on your own specific backup environment and history.

Incremental Backup

Determines the backup schedule for Incremental Backups.

As scheduled, Arcserve UDP Agent (Windows) 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.

The available options are **Repeat** and **Never**. If you select the **Repeat** option, you must also specify the elapsed time period (in minutes, hours, or days) between backup attempts. The minimum setting for Incremental backups is every 15 minutes.

By default the schedule for Incremental backups is to repeat every 1 day.

Full Backup

Determines the backup schedule for Full Backups.

As scheduled, Arcserve UDP Agent (Windows) performs a Full backup of all used blocks from the source machine. The available options are **Repeat** and **Never**. If you select the **Repeat** option, you must also specify the elapsed time period (in minutes, hours, or days) between backup attempts. The minimum setting for Full backups is every 15 minutes.

By default the schedule for Full backups is **Never** (no scheduled repeat).

Verify Backup

Determines the backup schedule for Verify Backups.

As scheduled, Arcserve UDP Agent (Windows) 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 Agent (Windows) 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.

The available options are **Repeat** and **Never**. If you select the **Repeat** option, you must also specify the elapsed time period (in minutes, hours, or days)

between backup attempts. The minimum setting for Verify backups is every 15 minutes.

By default the schedule for **Verify** backups is **Never** (no scheduled repeat).

Catalogs

File System Catalog

When this option is selected, enables generation of the file system catalog. If your browse time is too slow (especially if the Arcserve UDP Agent (Windows) destination is over a WAN) or if your restore by search time is too slow, this option helps reduce your wait time. This catalog job will run for each scheduled backup job after this option is selected.

If this option is not selected, the restores can be performed immediately after backup without having to wait for the catalog job to finish. By default, this option is not enabled.

Note: When you generate a File System catalog for each backup job, it results in an increased amount of disk storage needed to store the metadata files and catalog files and an increase in CPU usage. In addition, if the backup source contains a large amount of files, the process of generating a catalog could be a time consuming task.

Note: If you selected an ReFS volume as the backup source, you will not be able to generate a catalog and a warning message will be displayed to inform you of this condition.

3. Click **Save Settings**.

Your settings are saved.

Note: If at a given time there are more than one type of backup scheduled to be performed simultaneously, the type of backup that will be performed is based upon the following priorities:

- ◆ Priority 1 - Full backup
- ◆ Priority 2 - Verify backup
- ◆ Priority 3 - Incremental backup

For example, if you schedule all three types of backups to be performed at the same time, Arcserve UDP Agent (Windows) will perform the Full Backup. If there is no Full Backup scheduled, but you scheduled a Verify Backup and Incremental Backup to be performed at the same time, Arcserve UDP Agent (Windows) will perform the Verify Backup. A scheduled Incremental Backup is performed only if there is no conflict with any other type of backup.

Specify Advanced Schedule Settings

Arcserve UDP Agent (Windows) lets you specify the schedule for your backups. If you set the **Backup Data Format** option in **Protection Settings** to **Advanced**, the **Advanced Backup Schedule** dialog opens, where you can view your Repeat Schedule and Daily/Weekly/Monthly Settings.

Advanced Scheduling allows you to set the Repeat Schedule and the Daily Weekly Monthly Schedule. Advanced scheduling consists of the following:

- Week-based repeat backup schedule
- Week-based backup throttling schedule
- Week-based merge schedule
- Daily backup schedule
- Weekly backup schedule
- Monthly backup schedule

Follow these steps:


1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Backup Settings** tab. When the **Backup Settings** dialog opens, select **Schedule**.

The **Backup Settings Advanced Schedule** dialog opens.


Notes:

- If the Arcserve UDP Agent (Windows) is being managed by Console, not all settings are available and will be displayed as read-only information.
- If the Arcserve UDP Agent (Windows) is managed by console and not protected in a plan, all the settings are still available except the Preference >

Updates panel.



- (Optional) Click **Add** to add a backup schedule, backup throttle schedule, or a merge schedule.



For more information, see the following topics:

- ◆ [Add Backup Job Schedule.](#)
- ◆ [Add Backup Throttle Schedule.](#)
- ◆ [Add Merge Schedule.](#)

- Specify the **Start Date and Time**.

The start date and start time for your scheduled backups.

Note: When setting the interval between repeat backup jobs, ensure that you leave enough time to allow the previous job and any related merge jobs to complete before the next backup job starts. This amount of time can be estimated based on your own specific backup environment and history.

4. Specify the **Number of recovery points to retain**.

The number of recovery points to retain can be set for Daily, Weekly, Monthly, and Custom/Manual.

Note: The total retention count (Daily + Weekly + Monthly + Custom/Manual), the maximum limitation is 1440.

5. Specify **File System Catalog** and **Exchange Catalog** generation.

File System Catalog

When this option is selected, enables generation of the file system catalog. If your browse time is too slow (especially if the Arcserve UDP Agent (Windows) destination is over a WAN) or if your restore by search time is too slow, this option helps reduce your wait time. This catalog job will run for each scheduled backup job after this option is selected.

If this option is not selected, the restores can be performed immediately after backup without having to wait for the catalog job to finish. By default, this option is not enabled.

Note: When you generate a File System catalog for each backup job, it results in an increased amount of disk storage needed to store the metadata files and catalog files and an increase in CPU usage. In addition, if the backup source contains a large amount of files, the process of generating a catalog could be a time consuming task.

Note: If you selected an ReFS volume as the backup source, you will not be able to generate a catalog and a warning message will be displayed to inform you of this condition.

6. Click **Save Settings**.

Your settings are saved.

Add Backup Job Schedule

Follow these steps:

1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Backup Settings** tab. When the **Backup Settings** dialog opens, select **Schedule**.

The **Backup Settings Advanced Schedule** dialog opens.

2. From the **Backup Settings Advanced Schedule** dialog, click **Add** and then click **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 with 'Custom' selected. Below this, the 'Backup Type' is set to 'Incremental'. The 'Start Time' is '8:00 AM'. A grid of days is shown with checkboxes for Sunday, Monday, Tuesday, Wednesday, Thursday, and Friday, all of which are checked. The 'Saturday' checkbox is unchecked. The 'Repeat' checkbox is checked. Below this, there is a section for frequency: 'Every 3 Hours' and 'Until 6:00 PM'. At the bottom, there are 'Help', 'Save', and 'Cancel' buttons.

3. From the drop-down list, select **Daily**, **Weekly**, **Monthly**, or **Custom**.
4. Enter the appropriate fields based on the schedule you selected:
 - ◆ To add a Daily Backup Schedule, see [Add Daily Backup Schedule](#).
 - ◆ To add a Weekly Backup Schedule, see [Add Weekly Backup Schedule](#).
 - ◆ To add a Monthly Backup Schedule, see [Add Monthly Backup Schedule](#).
 - ◆ To add a Custom/Manual Backup Schedule, see [Add Custom Backup Schedule](#).
5. Click **Save**.

Your settings are saved.

Notes:

- You can add up to 4 time windows for any week day.
- The time window cannot be set across multiple days. You can only configure the time window from 12:00 AM to 11:59 PM.
- For each time window, you can specify the time window and the repeat frequency.
- The default backup schedule is 1 daily backup at 10:00pm.

Add Custom Backup Schedule

Follow these steps:

1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Backup Settings** tab. When the **Backup Settings** dialog opens, select **Schedule**.

The **Backup Settings Advanced Schedule** dialog opens.

2. From the **Backup Settings Advanced Schedule** dialog, click **Add** and then click **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 with 'Custom' selected. Below this, the 'Backup Type' is set to 'Incremental'. The 'Start Time' is '8:00 AM'. A grid of days is shown with checkboxes for Sunday, Monday, Tuesday, Wednesday, Thursday, and Friday, all of which are checked. The 'Saturday' checkbox is unchecked. The 'Repeat' checkbox is checked. Below this, there is a section for frequency: 'Every 3 Hours' and 'Until 6:00 PM'. At the bottom, there are 'Help', 'Save', and 'Cancel' buttons.

3. From the dropdown, **Custom** is selected as the default.
4. Enter the following fields:

Backup Type

Select the type of scheduled backup (Full, Verify, or Incremental) from a dropdown menu.

Start Time

Specify the time of the day and which days to start applying the configured schedule settings.

Repeat Every

Specify the time interval (hours/minutes) for how often to repeat this backup schedule.

Until

Specify the time of the day to stop applying the configured schedule settings.

How many backups do you want to retain?

Specifies the number of recovery points (full, incremental, and verify backup images) retained. When the number of recovery points present on the destination exceeds the specified limit, the earliest (oldest) incremental backups beyond the retention count are merged into the parent backup to generate a new baseline image consisting of the "parent plus oldest child's" blocks. If there are multiple sessions available for merge, the oldest child backups will be merged into the parent backup in a single pass, if the backups are compressed. If the backups are not compressed, then only the oldest child backup will be merged into the parent backup and this cycle repeats for each subsequent child backup to be merged.

Specifying the number of recovery points to retain allows you to perform infinite incremental backups, while maintaining the same retention count. For more information, see [Merge Job Guidelines](#).

Note: If your destination does not have sufficient free space, you can consider reducing the number of saved recovery points.

Default Value: 31

Minimum Value: 1

Maximum Value: 1440

Note: The Arcserve UDP Agent (Windows) home page **Summary** section indicates how many recovery points are retained out of the number specified. For more information, see [Status Summary](#) in the online help.

5. Click Save.

Your settings are saved.

Notes:

- The time window cannot be set across multiple days. You can only configure the time window from 12:00 AM to 11:59 PM.
- For each time window, you can specify the time window and the repeat frequency.

Add Daily Backup Schedule

Follow these steps:

1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Backup Settings** tab. When the **Backup Settings** dialog opens, select **Schedule**.

The **Backup Settings Advanced Schedule** dialog opens.

2. From the **Backup Settings Advanced Schedule** dialog, click **Add** and then click **Add Backup Schedule**.

The **New Backup Schedule** dialog opens.

3. From the drop-down list, select **Daily**.

The screenshot shows the 'New Backup Schedule' dialog box. At the top, there is a title bar with the text 'New Backup Schedule' and a close button (X). Below the title bar is a dropdown menu with 'Daily' selected. Underneath, there are several fields: 'Backup Type' with a dropdown menu showing 'Incremental'; 'Start Time' with a text box containing '10:00 PM' and a calendar icon; a grid of days with checkboxes: Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, and Saturday, all of which are checked; and 'How many backups do you want to retain' with a text box containing '7'. At the bottom of the dialog, there are three buttons: 'Help', 'Save', and 'Cancel'.

4. Enter the following fields:

Backup Type

Select the type of scheduled backup (Full, Verify, or Incremental) from a drop-down menu.

Start Time

Specify the time of the day and which days to start applying the configured schedule settings.

How many backups do you want to retain?

Specifies the number of recovery points (full, incremental, and verify backup images) retained. When the number of recovery points present on the destination exceeds the specified limit, the earliest (oldest) incremental backups beyond the retention count are merged into the parent backup to generate a new baseline image consisting of the "parent plus oldest child's" blocks. If there are multiple sessions available for merge, the oldest child backups will be merged into the parent backup in a single pass, if the backups are compressed. If the backups are not compressed, then only the oldest child backup will be merged into the parent backup and this cycle repeats for each subsequent child backup to be merged.

Specifying the number of recovery points to retain allows you to perform infinite incremental backups, while maintaining the same retention count. For more information, see [Merge Job Guidelines](#).

Note: If your destination does not have sufficient free space, you can consider reducing the number of saved recovery points.

Default Value: 7

Minimum Value: 1

Maximum Value: 1440

Note: The Arcserve UDP Agent (Windows) home page **Summary** section indicates how many recovery points are retained out of the number specified. For more information, see [Status Summary](#) in the online help.

5. Click Save.

Your settings are saved.

Notes:

- The time window cannot be set across multiple days. You can only configure the time window from 12:00 AM to 11:59 PM.
- For each time window, you can specify the time window and the repeat frequency.
- The default backup schedule is 1 daily backup at 10:00pm.

Add Weekly Backup Schedule

Follow these steps:

1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Backup Settings** tab. When the **Backup Settings** dialog opens, select **Schedule**.

The **Backup Settings Advanced Schedule** dialog opens.

2. From the **Backup Settings Advanced Schedule** dialog, click **Add** and then click **Add Backup Schedule**.

The **New Backup Schedule** dialog opens.

3. From the drop-down list, select **Weekly**.

New Backup Schedule

Weekly

Backup Type: Incremental

Start Time: 10:00 PM, Friday

How many backups do you want to retain: 5

Help Save Cancel

4. Enter the following fields:

Backup Type

Select the type of scheduled backup (Full, Verify, or Incremental) from a drop-down menu.

Start Time

Specify the time of the day and which days to start applying the configured schedule settings.

How many backups do you want to retain?

Specifies the number of recovery points (full, incremental, and verify backup images) retained. When the number of recovery points present on the destination exceeds the specified limit, the earliest (oldest) incremental backups beyond the retention count are merged into the parent backup to generate a new baseline image consisting of the "parent plus oldest child's" blocks. If there are multiple sessions available for merge, the oldest child backups will be merged into the parent backup in a single pass, if the backups are compressed. If the backups are not compressed, then only the oldest child backup will be merged into the parent backup and this cycle repeats for each subsequent child backup to be merged.

Specifying the number of recovery points to retain allows you to perform infinite incremental backups, while maintaining the same retention count. For more information, see [Merge Job Guidelines](#).

Note: If your destination does not have sufficient free space, you can consider reducing the number of saved recovery points.

Default Value: 5

Minimum Value: 1

Maximum Value: 1440

Note: The Arcserve UDP Agent (Windows) home page **Summary** section indicates how many recovery points are retained out of the number specified. For more information, see [Status Summary](#) in the online help.

5. Click Save.

Your settings are saved.

Notes:

- The time window cannot be set across multiple days. You can only configure the time window from 12:00 AM to 11:59 PM.
- For each time window, you can specify the time window and the repeat frequency.

Add Monthly Backup Schedule

Follow these steps:

1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Backup Settings** tab. When the **Backup Settings** dialog opens, select **Schedule**.

The **Backup Settings Advanced Schedule** dialog opens.

2. From the **Backup Settings Advanced Schedule** dialog, click **Add** and then click **Add Backup Schedule**.

The **New Backup Schedule** dialog opens.

3. From the drop-down list, select **Monthly**.

The screenshot shows the 'New Backup Schedule' dialog box. The 'Monthly' option is selected in the top dropdown. The 'Backup Type' is set to 'Incremental'. The 'Start Time' is '10:00 PM'. The 'Day' radio button is selected, and the frequency is 'last day'. The 'Week' radio button is unselected, with a frequency of 'Last' and a day of 'Friday'. The number of backups to retain is set to '12'. The 'Save' button is highlighted in blue.

4. Enter the following fields:

Backup Type

Select the type of scheduled backup (Full, Verify, or Incremental) from a drop-down menu.

Start Time

Specify the time of the day and which days to start applying the configured schedule settings.

How many backups do you want to retain?

Specifies the number of recovery points (full, incremental, and verify backup images) retained. When the number of recovery points present on the destination exceeds the specified limit, the earliest (oldest) incremental backups beyond the retention count are merged into the parent backup to generate a new baseline image consisting of the "parent plus oldest child's" blocks. If there are multiple sessions available for merge, the oldest child backups will be merged into the parent backup in a single pass, if the backups are compressed. If the backups are not compressed, then only the oldest child backup will be merged into the parent backup and this cycle repeats for each subsequent child backup to be merged.

Specifying the number of recovery points to retain allows you to perform infinite incremental backups, while maintaining the same retention count. For more information, see [Merge Job Guidelines](#).

Note: If your destination does not have sufficient free space, you can consider reducing the number of saved recovery points.

Default Value: 12

Minimum Value: 1

Maximum Value: 1440

Note: The Arcserve UDP Agent (Windows) home page **Summary** section indicates how many recovery points are retained out of the number specified. For more information, see [Status Summary](#) in the online help.

5. Click Save.

Your settings are saved.

Notes:

- The time window cannot be set across multiple days. You can only configure the time window from 12:00 AM to 11:59 PM.
- For each time window, you can specify the time window and the repeat frequency.

Add Backup Throttle Schedule

Follow these steps:

1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Backup Settings** tab. When the **Backup Settings** dialog opens, select **Schedule**.

The Backup Settings Advanced Schedule dialog opens.

2. From the **Backup Settings Advanced Schedule** dialog, click **Add** and then click **Add Throttle Schedule**.

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

The screenshot shows the 'Add New Throttle Schedule' dialog box. It features a title bar with a close button (X). The main area contains the following elements:

- Throughput Limit:** A text input field is empty, followed by the unit 'MB/min'.
- Start Time:** A text input field contains '8:00 AM' and a small calendar icon to its right.
- Days of the Week:** A grid of seven days: Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, and Saturday. Each day has a checked checkbox to its left.
- Until:** A text input field contains '6:00 PM' and a small calendar icon to its right.
- Footer:** Three buttons: 'Help' (text), 'Save' (blue), and 'Cancel' (text).

3. Enter the following fields:

Throughput Limit

Specify the maximum speed (MB/min) at which the backup will be written.

You can throttle the backup speed to reduce CPU or network utilization.

However, by limiting the backup speed, it has an adverse effect on the backup window. As you lower the maximum backup speed it increases the amount of time to perform the backup. For a backup job, the Job Monitor on the home page will display the average Read and Write speed of the job in progress and the configured throttle speed limit.

Note: By default, the throttle backup speed option is not enabled and backup speed is not being controlled.

Start Time

Specify the time of the day to start applying the configured backup throttle settings.

Until

Specify the time of the day to stop applying the configured backup throttle settings.

4. Click Save

Your settings are saved.

Notes:

- You can add up to 4 time windows for any week day.
- The throttling value control the backup speed. For example, if you set 2 time windows, 1 from 8:00 AM to 6:00 PM, backup throughput limit is 1500 MB/minute, and 1 from 6:00 PM to 8:00 PM, backup throughput limit is 3000 MB/minute. If a backup job runs from 5:00 PM to 7:00 PM, the throughput will be 1500 MB/minute from 5:00 PM to 6:00 PM, and change to 3000 MB/minute from 6:00 PM to 7:00 PM.
- The time window cannot be set across multiple days. You can only configure the time window from 12:00 AM to 11:45 PM. If the throttle schedule ends at 11:45 PM, the schedule takes effect until the next day.
- Backup throttle schedule applies to repeat backup, as well as daily / weekly / monthly backups.

Add Merge Schedule

Follow these steps:

1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Backup Settings** tab. When the **Backup Settings** dialog opens, select **Schedule**.

The **Backup Settings Advanced Schedule** dialog opens.

2. From the **Backup Settings Advanced Schedule** dialog, click **Add** and then click **Add Merge Schedule**.

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

3. Enter the following fields:

Start Time

Specify the time of the day to start applying the configured backup throttle settings.

Until

Specify the time of the day to stop applying the configured backup throttle settings.

4. Click **Save**.

Your settings are saved.

Notes:

- You can add up to 2 time windows for any week day.
- If there is no merge schedule configured for any day, the merge job will be launched as soon as it is ready. If you configured any time window for the merge schedule, the merge job will only be launched within the time windows.

For example, if the merge schedule is 8:00 AM to 6:00 PM on Sunday, the merge job will only be launched during this time window.

- If the merge job launches within the configured time windows, and it runs to completion, despite the end time of the time windows. For example, if the merge time window is 8:00 AM to 6:00 PM on Sunday, a merge job started at 5:55 PM. It will continue to run after 6:00 PM to complete, even though the time is moving out of the defined time window.
- Merge schedule applies to repeat backup, as well as daily / weekly / monthly backups.
- When you configure a merge job schedule, the merge will only be triggered when the time is within the configured time window. If the merge is not within the configured time window, the merge will not run when you click the **Run a merge job manually now** link in the summary panel of the Arcserve UDP Agent (Windows) home page.

Schedule Considerations

The Arcserve UDP Agent (Windows) provides flexible settings to let you specify schedules for your backup. It consists of the following:

- Week-based repeat backup schedule
- Week-based backup throttling schedule
- Week-based merge schedule
- Daily backup schedule
- Weekly backup schedule
- Monthly backup schedule

However each backup, merge, or catalog job will consume system resources (CPU Usage, Memory Usage, IO Usage), occupy network bandwidth, and occupy disk space. Therefore, to help protect your system, consider the following:

What is the business processing time range of your server?

To avoid affecting your business processing, configure your system to run less jobs when the server is busy. For example, only configure to run backup jobs when the server is busy and leave merge jobs to run when the server is idle.

How about the data change frequency of your server?

Normally more frequent data change means more frequent backup is required. This is to reduce data lost to the minimum. When needed, you can recover the server to the last good known status.

How about your network bandwidth?

If your backup destination is configured to a network shared path, obviously the job occupies some of your network bandwidth when it is running. This might affect your business processing of this server. In case of this, specify a throttle schedule to limit the Arcserve UDP Agent (Windows) occupying network bandwidth.

How much disk storage is allocated for your backup destination?

More Full backups and more backups to retain means more disk storage is required. So when you configure how frequently to run a Full backup and how many backups to retain, consider the disk storage allocated for the backup destination.

How do you expect to use your backed up data?

Enable "File System Catalog" can shorten the browse time when you want to restore a file or a mailbox. But to generate catalogs, it also results in an

increased amount of disk storage needed to store the metadata files and catalog files and an increase in CPU usage. In addition, if the backup source contains a large amount of files, the process of generating a catalog could be a time consuming task. So whether to enable or disable catalogs is depending on how you would like to use the backed up data.

Based on the above considerations, the following is an example of using advanced scheduling to protect a build server, showing the situation and corresponding schedule settings:

- The build server is used to provide source code pre-compile service every working day. It's business process time slot is 9:00 AM – 7:00 PM of every work day (from Monday to Friday). During other times, it is idle.

Schedule Settings:

- Configure to run custom incremental backup from 9:00 AM to 7:00 PM, run merge job at night – 7:00 PM to 9:00 AM of next day.
- The pre-compile service is launched every 2 hours, and there are lots of data changes at that time.

Schedule Settings:

- Configure to run custom incremental backup every 2 hours.
- Every time to run pre-compile, the build server need to fetch source code from a remote source code repository server.

Schedule Settings:

- Limit backup throttle to 500 MB/Minute during 9:00 AM to 7:00 PM and no limitation during other time slots.
- Due to the poor disk storage, there is no requirement to retain a lot of recovery points. Only need to keep recovery points in one release cycle; 6 months is enough. But there is a requirement to keep the recovery point in the last 24 hours, so that once needed you can recover to the last good known status.

Schedule Settings:

- Specify to retain last 12 manual backups (the backups of the last 24 hours).
- Configure to run Daily Incremental backup at 9:00 PM of every day. And keep the last 7 Daily backups.
- Configure to run Weekly Full Backup at 11:00 PM of every Friday. And keep the last 4 Weekly backups.

- Configure to run Monthly Full Backup at 12:00 PM on last Saturday of month. And keep the last 6 monthly backups.

Finally, there are 6 monthly backups, 4 weekly backups, 7 daily backups and 12 most recent backups. There are enough choices to recover the build server to a good known status.

- For the build server, there is no requirement to quickly browse and restore files. Once needed, perform a BMR to restore the build server to the last good known status. That is enough.

Schedule Settings:

- Disable options to generate "File System Catalog".

Specify Advanced Settings

Arcserve UDP Agent (Windows) lets you specify the **Advanced Settings** for your backups.

Follow these steps:

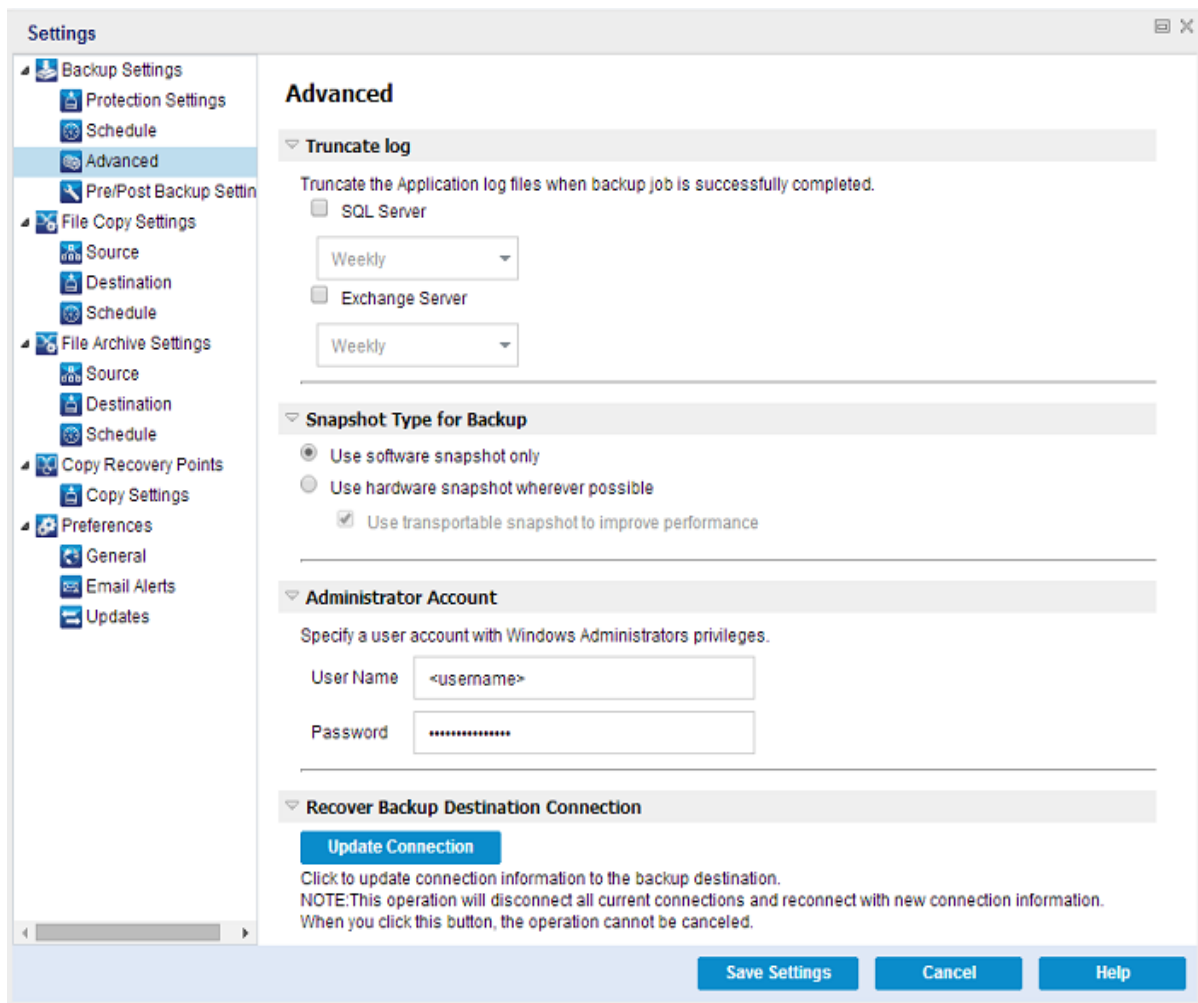
1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Backup Settings** tab. When the **Backup Settings** dialog opens, select **Advanced**.

The **Advanced** dialog opens.

Notes:

- If the Arcserve UDP Agent (Windows) is being managed by Console, not all settings are available and will be displayed as read-only information.
- If the Arcserve UDP Agent (Windows) is managed by console and not protected in a plan, all the settings are still available except the Preference >

Updates panel.



2. Specify your advanced backup settings options.

Truncate Log

Truncates the accumulated transaction log files for the selected applications after the next successful backup.

Arcserve UDP Agent (Windows) backups consist of a snapshot image and the transaction log files that were created for it. At some point in time, the older (committed) transaction log files are no longer needed and can be purged to make space for new log files. The process of purging these log files is truncating the log. This option enables truncating of committed transaction log files, which conserves disk space.

The available options are SQL Server and Exchange Server. You can select either, both, or none of these applications. If you select any of these applications, you can also specify a scheduled time period (Daily, Weekly, or Monthly) for automatic log truncation:

Note: The transaction log files cannot be truncated without performing a successful backup.

- **Daily** - Specifies that each day after the backup completes successfully, the committed transaction logs will be purged immediately.
- **Weekly** - Specifies that after seven days, the committed transaction logs will be purged immediately after the backup completes successfully.
- **Monthly** - Specifies that after 30 days, the committed transaction logs will be purged immediately after the backup completes successfully.

If a backup job is already running at the same time the purging is scheduled to be performed, the purging operation is moved to the next scheduled job.

Example:

You scheduled an Incremental Backup to run automatically every day at 5:00 pm, and then started a Full Backup manually at 4:55 pm. You assume that the backup successfully finishes at 5:10 pm.

In this case, the Incremental Backup that is scheduled for 5:00 pm is skipped because the ad-hoc Full Backup is still in progress. Now the committed transaction log files are purged after the next successful backup job and be performed on the next day after the scheduled Incremental Backup completes successfully at 5:00 pm.

Snapshot Type for Backup

You can select the required option from software snapshot or hardware snapshot.

Use software snapshot only

Specifies that the backup type uses only the software snapshot. Arcserve UDP will not check for hardware snapshot. The software snapshot utilizes less resources on the virtual machines. You can use this option if the server has lower configurations and processing speed.

Use hardware snapshot wherever possible

Specifies that the backup type first checks for a hardware snapshot. If all the criteria are met, the backup type uses hardware snapshot.

Note: For more information on the hardware snapshot criteria, see the prerequisite.

Administrator Account

Specifies the User Name and Password with access rights to perform the backup. The Arcserve UDP Agent (Windows) verifies that the name and password are valid and the user belongs to an administrator group.

Important! If the Administrator Account credential information for the Arcserve UDP Agent (Windows) server is changed (User Name/Password), you must also reconfigure/update the Administrator Account information in this dialog.

Note: To specify a domain account, the format for the user name is a fully qualified domain user name in the form of "<domain name>\<user name>".

Recover Backup Destination Connection

Lets you update (resynchronize) the connection information to your backup destination.

You can use this option if you are performing periodic backups to a remote share computer and then you can change the access credentials (user name/password) for that remote computer. In this case, typically your next backup would fail because the access credentials configured at your local computer do not match the new credentials at the remote computer.

Note: When you click the **Update Connection** button and the resynchronize process begins, you cannot cancel it.

Before you click this **Update** button, perform the following tasks:

- a. Log in to the remote destination computer and use the following net session command to disconnect the connection between the local Arcserve UDP Agent (Windows) computer and the remote computer:

```
net session \\<computer name or IP address> /d
```

- b. Return to the Arcserve UDP Agent (Windows) computer, and click the **Update Connection** button.
- c. Enter new password for destination.

Arcserve UDP Agent (Windows) updates your configured credentials to match the new credential information at the remote share destination. A pop-up confirmation screen appears informing you that the credentials have been updated.

3. Click **Save Settings**.

Your advanced backup settings are saved.

Specify Pre/Post Backup Settings

Arcserve UDP Agent (Windows) lets you specify the **Pre/Post Backup Settings**.

Specify the Pre/Post Backup Settings

1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Backup Settings** tab. When the **Backup Settings** dialog opens, select **Pre/Post Backup**.

The **Pre/Post Backup Settings** dialog opens.

Notes:

- If the Arcserve UDP Agent (Windows) is being managed by Console, not all settings are available and will be displayed as read-only information.
- When agent is managed by console and not protected in a plan, all the settings are still available except the Preference > Updates panel.

The screenshot shows the 'Settings' dialog box with the 'Pre/Post Backup Settings' tab selected. The left-hand navigation pane lists various settings categories, with 'Pre/Post Backup Settings' highlighted. The main content area is titled 'Pre/Post Backup Settings' and features an 'Actions' section. This section includes a text box for 'User Name', a 'Password' field, and several checkboxes: 'Run a command before backup is started', 'On exit code' (with a value of 0), 'Run Job' (selected), 'Fail Job', 'Run a command after snapshot is taken', 'Run a command after backup is over', and 'Run the command even when the job fails'. At the bottom of the dialog, there are three buttons: 'Save Settings', 'Cancel', and 'Help'.

2. Specify your pre/post backup setting options.

Actions

Runs script commands for actions to take before the start of the backup, after the snapshot image is captured, and/or upon the completion of the backup. You can also trigger the script command based upon specific exit codes and select the action to be taken (run job or fail job) when that exit code is returned.

- A "run job" action directs Arcserve UDP Agent (Windows) to continue to run the job if the specified exit code is returned.
- A "fail job" action directs Arcserve UDP Agent (Windows) to cancel the job if the specified exit code is returned.

3. Click **Save Settings**.

Your pre/post backup settings are saved.

Manage File Copy Settings

Arcserve UDP Agent (Windows) lets you copy selected source files from a backup session to a destination (disk or cloud) based upon your specified file copy and retention criteria. File Copy can be used for copying critical data to secondary locations.

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 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.

Before you perform your first File Copy job, specify the File Copy settings and plans. These configurations allow you to specify behaviors such as the source of your file copy data, destination for your copied files, the schedule for each file copy job, and the settings and filters applied to your file copy jobs. These settings can be modified at any time from the Arcserve UDP Agent (Windows) home page.

Note: To improve the performance (upload speed and server load), File Copy can upload the data to the specified destination in parallel chunks. To configure the number of chunks that are simultaneously sent to the destination, see [Configure File Copy Chunk Value](#).

To manage the File Copy settings, click the Settings link on the Arcserve UDP Agent (Windows) home page and select the File Copy Settings tab. The File Copy Settings dialogs consist of the following subordinate tab options:

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

Specify the File Copy Source

Arcserve UDP Agent (Windows) lets you specify the source settings for your information to be file copied.

Note: For more information about the File Copy Settings, see [Manage File Copy Settings](#).

Follow these steps:

1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **File Copy Settings** tab. When the **File Copy Settings** dialog opens, select **Source**.

The **File Copy Sources** dialog opens.

Notes:

- If the Arcserve UDP Agent (Windows) is being managed by Console, not all settings are available and will be displayed as read-only information.
 - When agent is managed by console and not protected in a plan, all the settings are still available except the Preference > Updates panel.
2. Specify your file copy source settings.

Enable File Copy

Enables the scheduled copying of files after the specified number of backups. If this option is not selected (File Copy disabled), no scheduled file copying is performed and any changes to the File Copy settings are not validated and saved.

Note: ReFS and deduplicated NTFS volumes will not be listed as selectable backup sources to be file copied. As a result, if all the volumes for a specified backup source are only ReFS or deduplicated NTFS volumes, the file copy option will be disabled.

Recovery Points to Copy From

Specifies the recovery point that you want to copy. You have two options to specify the recovery points. You can copy the recovery points from a specific backup number. For example, you can specify that the recovery points must be copied from every fifth backup. Another option is, you can specify to copy recovery points from a daily, weekly, or monthly backup.

File Copy Sources

Displays a selectable listing of all sources, with the corresponding policy (filters) and type of file copy (copy and retain or copy and move) to be performed

after each successful Arcserve UDP Agent (Windows) backup. These File Copy Sources can be added, removed, or modified.

Note: Arcserve UDP Agent (Windows) 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. From the File Copy Restore user interface, the actual path will be displayed instead of the symbolic link.

- **Add**

When clicked, the Plan type dialog opens to let you initially select the type of file copy job to perform (copy and retain or copy and move). After you select the plan type, the corresponding File Copy Plan dialog opens to let you add a source to copy and specify the corresponding plans for that source. For more information, see [Specify File Copy Plans](#).

Note: The File Copy source can only be selected from a volume that is currently selected in Backup Settings. If the source contains ReFS or deduplicated NTFS volumes, these volumes will not be available for selection.

- **Remove**

When clicked removes the selected source from this displayed list.

- **Modify**

When clicked, the File Copy Plans dialog opens to let you change the plan settings for the selected source. For more information, see [Specify File Copy Plans](#).

3. Click **Save Settings**.

Your File Copy settings are saved.

Specify File Copy Plans

When you click the Add Source option for File Copy, the Plan type dialog opens to let you initially select the type of File Copy job to be performed.

The available type is File Copy. In File copy plan, data is copied from the source to the destination (remains on source location) and provides multiple stored versions on the destination.

If you want to Add a new File Copy source or you want to Modify an existing File Copy source, the File Copy Plans dialog lets you specify the details.

Depending upon the plan type that is selected, a different File Copy Plans dialog opens; however, the selections are similar.

File Copy Selected:

File copy plans [Close]

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.

[Text Input] [Browse]

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] [Text Input]

Type	Variable	Value
------	----------	-------

[Add] [Remove]

You can use wildcard characters '*' and '?' in File/Folder Patterns

[OK] [Cancel] [Help]

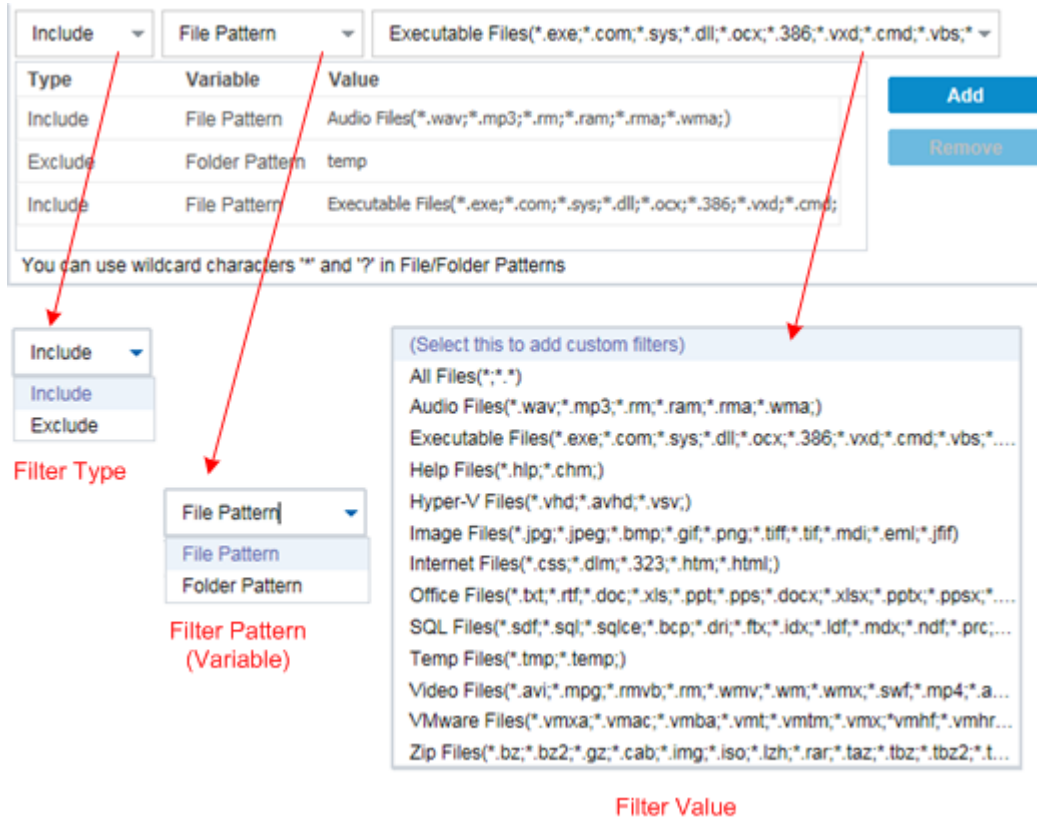
File Copy - Source Selection

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 pre-defined filters for many commonly used files is available (MS-Office files, Image files, Executable files, Temp files, etc.). After choosing any of the pre-defined filters, you can still append or modify the corresponding values.

How File Copy Filters Work

The File Copy source filters for files and folders work as follows:

- File with "d2darc" and "ASBUARC" extensions are always skipped.
- Files with system and temporary attributes are always skipped.
- Windows, Program Files, and Arcserve UDP Agent (Windows) installation folders (for both File Copy and File Copy - delete source policies) are always skipped.
- The following precedence order is used for filtering (with the highest precedence listed first):
 - Exclude directory filters
 - Exclude file filters
 - Include directory filters
 - Include file filters
 - Include Criteria
 - Exclude system and application files (Exchange and SQL only) present in any location. (This filter is only applicable for File Copy - Delete Source policies).
- A file is copied only if the Include Folder or the Include File filter matches, it does not need to satisfy both filter requirements.
- The File filter works only on the file name, and does not depend on the path.

For example, if you have the three files "Test.txt", "Hellotest.txt", and "TestHello.txt" these filters produce the following results:

 - Test*.txt filter matches only Test.txt and TestHello.txt
 - Test* filter matches Test.txt and TestHello.txt
 - Test filter does not match anything
 - *.txt filter matches all of them
 - *test filter does not match anything
- A Folder filter works on the policy source level.

For example, if you have the following directory structure:

```
C:  
->Z99  
-> ->A00  
-> -> ->B01
```

-> -> ->C01

-> -> ->D01

- If you configure your File Copy Source as "C:\Z99\A00" and you apply an Include Folder **b*** filter, then all files under c:\Z99\A00\B01 are copied.

In this example, the Source includes the parent folder and the asterisk is located after the "b". As a result all files in any folder subordinate to "A00" that starts with "b" are copied.

- If you configure your File Copy Source as "C:\Z99" and you apply an Include Folder **b*** filter, then this filter does not match any folder and no files are copied.

In this example, the Source does includes the "Z99" grandparent folder, but not the "A00" parent folder. As a result, there are no "b" folders directly subordinate to "Z99" and no files are copied.

- However if you specify a ***b*** filter, it now matches any subordinate folder that starts with "b" and all files within these "b" folders are then copied.

In this example, the asterisk is now located before the "b". As a result, all files in any folder subordinate to "C:\Z99" (regardless of the root level) that starts with "b" are copied.

- If you configure your File Copy Source as "C:\Z99" and you apply an Include Folder ***01** filter, then all subordinate folders containing "01" (B01, C01, and D01) are copied.

In this example, the asterisk is located before the "01". As a result, all files in any subordinate folders (regardless of the root level) that contains "01" are copied.

Note: Folder filters are always relative to the source folder path specified in the policy.

Specify the File Copy Destination

Arcserve UDP Agent (Windows) lets you specify the destination settings for your information to be file copied.

Note: For more information about the File Copy Settings, see [Manage File Copy Settings](#).

Specify the file copy destination

1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **File Copy Settings** tab. When the **File Copy Settings** dialog opens, select **Destination**.

The **File Copy Settings Destination** dialog opens.

Note: If the Arcserve UDP Agent (Windows) is being managed by Console, not all settings are available and will be displayed as read-only information.

2. Specify your File Copy destination settings.

Destination

Specifies the destination location for the file copy job. You can only select one destination.

Arcserve UDP Agent (Windows) 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. The Arcserve UDP Agent (Windows) 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](#).

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. You should 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 Agent (Windows) 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 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. The Arcserve UDP Agent (Windows) 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.

Files Retention

Retains the files in the file copy destination if the specified criteria is met.

File created within the last

Specifies the amount of time (years, months, days) that the stored data is retained at the destination location. At the end of the specified retention time period, the stored data is purged from the destination.

Important! At the end of the specified retention time when the data is purged from the destination, all of this purged data is no longer stored or saved.

Note: The Retention Time purge process is only triggered if the File Copy Schedule option is enabled.

File version less than

Specifies the number of copies retained and stored at the destination location. After this number is exceeded, the earliest (oldest) version will be discarded. This cycle of discarding the oldest stored version repeats as newer versions are added to the destination, allowing you to always maintain the specified number of stored versions.

For example, if your specified File Versions retention count is set to 5 and you perform five file copies at times t1, t2, t3, t4, and t5, these file copies become the five file copy versions retained and available to recover. After the sixth file copy is performed (new version is saved), Arcserve UDP Agent (Windows) will remove the t1 copy and the five available versions to recover are now t2, t3, t4, t5, and t6.

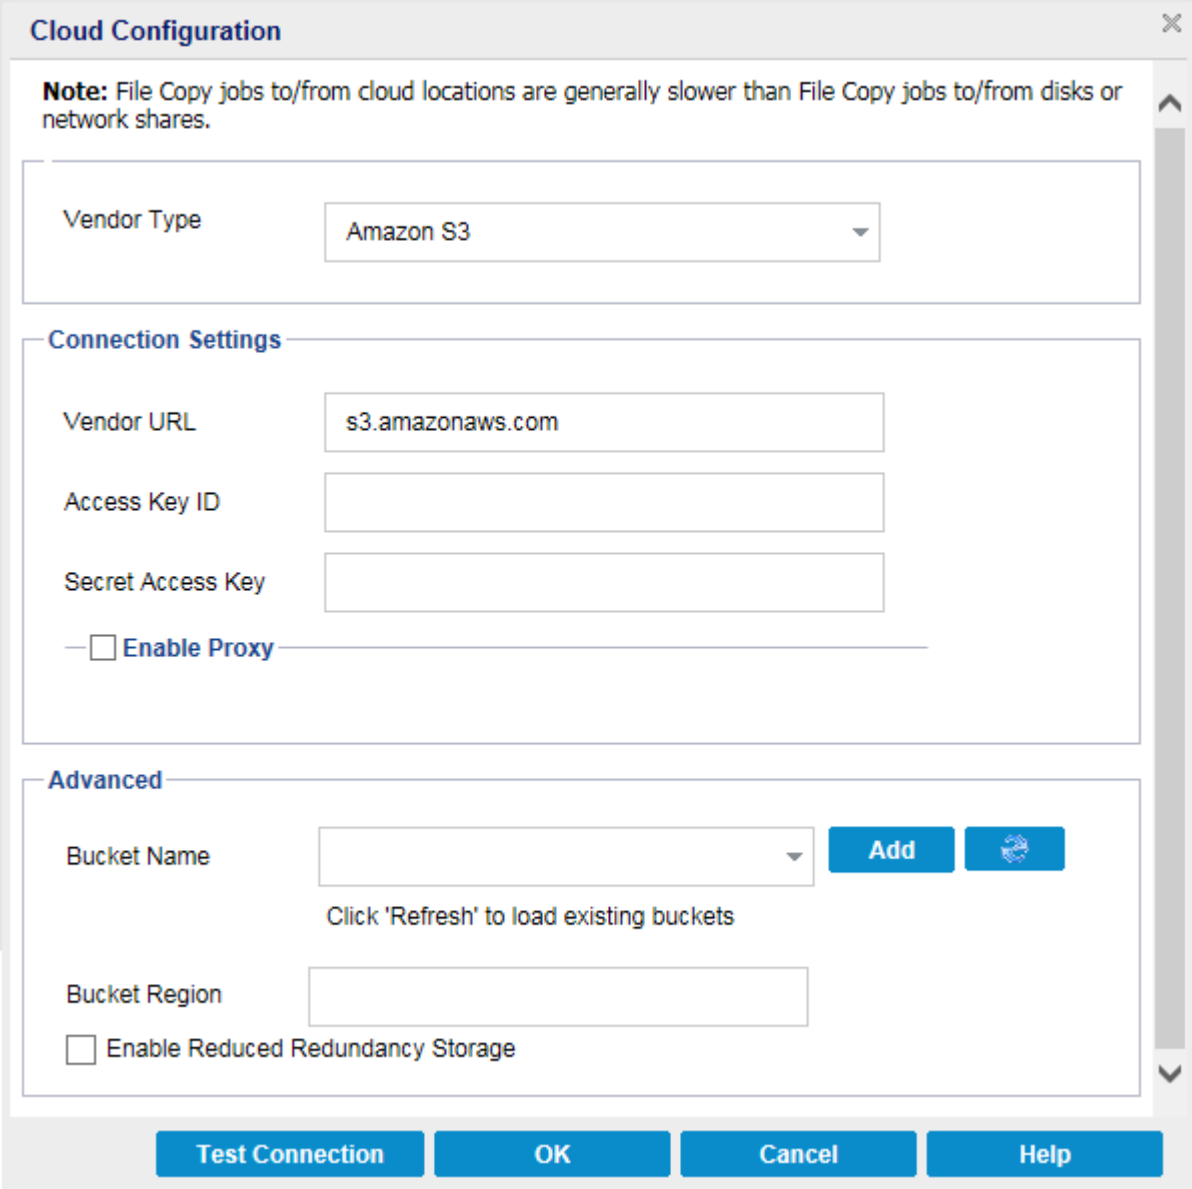
By default, the number of copies retained at the destination location before discarding is 15.

3. Click Save Settings.

Your File Copy settings are saved.

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 screenshot shows the 'Cloud Configuration' dialog box. At the top, there is a note: 'Note: File Copy jobs to/from cloud locations are generally slower than File Copy jobs to/from disks or network shares.' Below the note, the 'Vendor Type' is set to 'Amazon S3'. The 'Connection Settings' section includes fields for 'Vendor URL' (s3.amazonaws.com), 'Access Key ID', and 'Secret Access Key', along with an unchecked 'Enable Proxy' checkbox. The 'Advanced' section features a 'Bucket Name' dropdown menu with an 'Add' button and a 'Refresh' icon, a 'Bucket Region' text box, and an unchecked 'Enable Reduced Redundancy Storage' checkbox. At the bottom, there are four buttons: 'Test Connection', 'OK', 'Cancel', and 'Help'.

From this dialog you can use the drop-down menu to select a cloud vendor type that you want to use for the 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 cannot 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.

Follow these steps:

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, you must enter manually the Vendor URL 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.

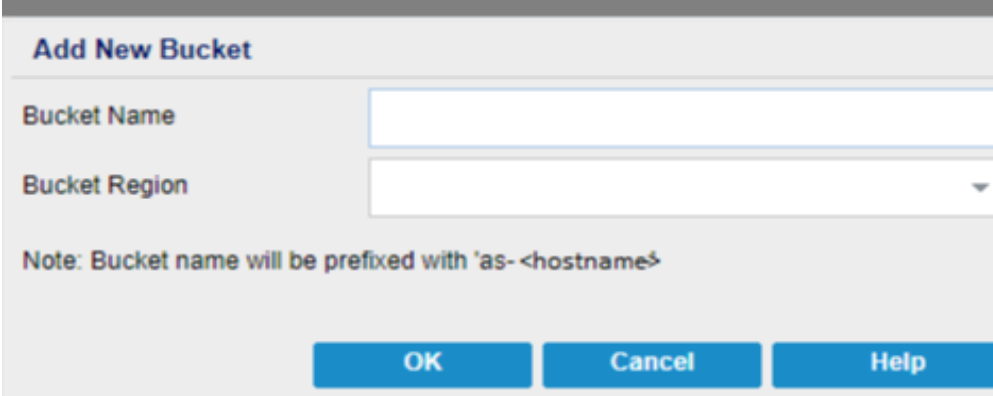
(For this field, Amazon S3 and Eucalyptus-Walrus use Bucket Name. Windows Azure and Fujitsu Cloud (Windows Azure) use Container).

Note: For the remainder of this step, you can also apply all references of Buckets to Containers unless specified.

You can either select a bucket name from the drop-down list, or you can add a new bucket name. If necessary, you can click the refresh button to update the list of available buckets.

To add a new bucket name:

- a. Click the **Add** button placed next to the Bucket Name field to display the Add New Bucket dialog.



The screenshot shows a dialog box titled "Add New Bucket". It contains two input fields: "Bucket Name" and "Bucket Region". Below the "Bucket Name" field, there is a note: "Note: Bucket name will be prefixed with 'as- <hostname>". At the bottom of the dialog, there are three buttons: "OK", "Cancel", and "Help".

- b. Enter a unique Bucket Name.

The new Bucket Name is automatically prefixed with name *as- <hostname>*. This format is applicable to the Bucket Name you create and is used as your File Copy destination.

Note: When creating a new bucket, Arcserve UDP Agent (Windows) only uses the *as- <hostname>* prefix and Arcserve UDP Agent (Windows) supports restoring from previous file copy destinations having *arcserve- <hostname>*, *d2dfilecopy- <hostname>*, or *d2d-filecopy- <hostname>* prefixes.

A bucket name should be unique, easily identifiable, and compliant with internet domain naming rules. No two buckets can have the same name. We recommend to have proper understanding about valid syntax for bucket names.

For more information about bucket naming requirements of Amazon S3 and Eucalyptus-Walrus, refer to the Amazon S3 documentation.

For more information about container naming requirements of Windows Azure and Fujitsu Cloud (Windows Azure), refer to the Microsoft documentation.

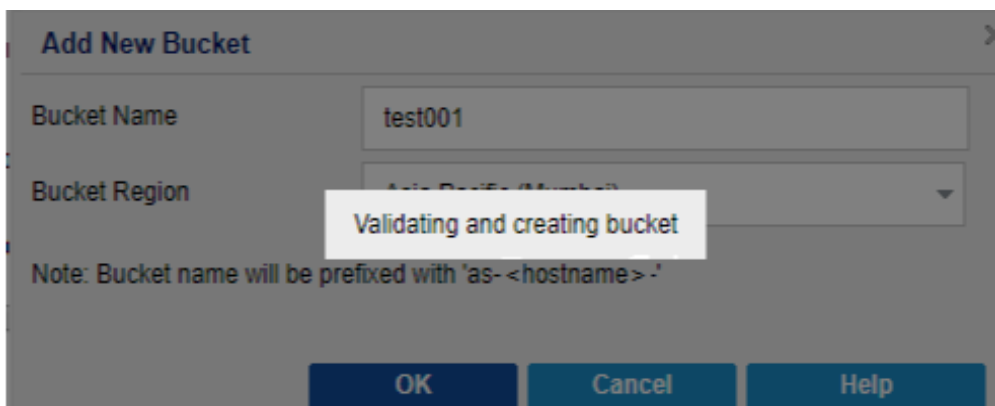
- c. For Amazon S3 only, select an available region from the drop-down menu. By default, all available regions are included in the drop-down menu and you can select the region where you want the new bucket to be created.

Regions allow you to select the geographical region where Amazon S3 stores the buckets that you create. You should select a Region that provides you with fast access to your data and allows you to optimize latency, minimize costs, or address regulatory requirements.

(For Windows Azure, Fujitsu Cloud (Windows Azure), and Eucalyptus-Walrus, the region is not selectable).

- d. After you have specified your values, click **OK**.

The Bucket name is validated and created at the cloud.



After you successfully create the new bucket, the main Cloud Configuration dialog is displayed again, with the new bucket information (name and region) included in the Advanced Settings fields.

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.

Configure File Copy Settings to Optimize Performance

To improve the performance (upload speed and server load), File Copy can upload data to the specified destination in parallel chunks and or parallel threads.

Chunk Value

You can set the number of 1-MB chunks that will be simultaneously sent to the destination. By increasing the number of parallel chunks, you will decrease the time to complete the job, but will also have an adverse affect on server performance. Configure this value as necessary to obtain optimal performance.

For example, if you are performing File Copy for a 10-MB file and set the number of 1-MB chunks to 2, then File Copy will write 10 chunks, two at a time. If you see that this is taking too long to complete the job, change this value to 4. The time to complete the job will then decrease because File Copy will now be writing 10 chunks, four at a time, but the load on your server will increase.

Threads for Archive Value

File Copy lets you copy more than one file at a time. By default, File Copy transfers 8 files in parallel when the destination is configured to File Systems and transfers 32 files in parallel when the destination is configured to Cloud. If you see that File Copy is taking too long to transfer the data, increase the number of threads up to 32, to optimize performance. However, if you experience a problem on a machine with less memory, decrease the number of threads.

Chunk Value and Threads for Archive Value can be used together to control the speed of the File Copy. If you increase Chunk Value and Threads for Archive Value, you see File Copy is performed faster.

For example, if you are transferring 8 files with 10 MB each and set the number of 1-MB chunks to 2, then File Copy will write 16 at a time (8 files X 2-MB chunks), but the load on your server will increase. When you see the load on the server has increased to a point where it becomes a problem, decrease the number of threads. If the destination is a Cloud location, it is recommended that you configure these settings in such a way that produces at least 20 writes, to optimize performance.

Threads for Restore Value

Restore from a File Copy lets you download more than one file at a time. By default, restores from file copies downloads 8 files when the File Copy location is configured to File Systems and downloads 32 files in parallel when the file copy location is configured to Cloud. If you see that the restore from a File Copy is taking too long to transfer the data, increase the number of threads up to 32.

Note: Chunk Value does not apply to restore jobs.

Threads for Catalog Synchronization Value

Catalog Synchronization jobs lets you use multiple threads to optimize performance.

If you see that the catalog synchronization job is taking too long to transfer the data, increase the number of threads up to 10. You will see that the job is performed faster and the load on the server increases. When you see the load on the server has increased to a point where it becomes a problem, decrease the number of threads.

To configure the file copy settings to optimize performance, set the corresponding DWORD values as follows:

1. Start edit registry.
2. Locate key:

"HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AfArchiveDll"

Note: The same registry key is used when your File Copy destination is either File Systems or a Cloud location.

3. To modify the value for the number of 1-MB chunks that will be simultaneously sent to the destination, follow these steps:
 - a. Manually create a DWORD value for "ArchMultChunkIO".
 - b. Assign a DWORD value:

The available range for the number of chunks is 1 - 4 chunks.

Default: 4 chunks

Maximum: 4 chunks
4. To modify the value for the number of threads (files) that will be transferred in parallel to the copy destination, follow these steps:
 - a. Manually create a DWORD value for "ThreadsForArchive".
 - b. Assign a DWORD value:

The available range for the number of files is 1 - 32 files.

Default: 8 files when the destination is configured to File Systems and 32 files when the destination is configured to a Cloud location.

Maximum: 32
5. To modify the value for the number of file copies that can be downloaded in parallel from the copy destination, follow these steps:

- a. Manually create a DWORD value for "ThreadsForRestore".
- b. Assign a DWORD value:

The available range for the number of files is 1 - 32 files.

Default: 8 files when the copy destination is File Systems and 32 files when the copy destination is a Cloud location.

Maximum: 32

6. To modify the value for the number of threads (streams) that can be used in parallel to perform catalog synchronization, follow these steps:

- a. Manually create a DWORD value for "ThreadForCatalogSync".
- b. Assign a DWORD value:

The available range for the number of files is 1 - 10 threads.

Default: 8 threads

Maximum: 10

Specify the File Copy Schedule

Arcserve UDP Agent (Windows) lets you specify the schedule settings for your information to be file copied.

Note: For more information about the File Copy Settings, see [Manage File Copy Settings](#).

Follow these steps:

1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **File Copy Settings** tab. When the **File Copy Settings** dialog opens, select **Schedule**.

The **File Copy Settings Schedule** dialog opens.

Note: If the Arcserve UDP Agent (Windows) is being managed by Console, not all settings are available and will be displayed as read-only information.

2. Specify your File Copy schedule settings.

Schedule

Enables the file copying of data after the specified number of backups.

The file copy process will be launched automatically after the specified number of successful backups (Full, Incremental, and Verify) and will be based on your selected File Copy policies.

You can use this setting to control how many times a File Copy job is triggered each day. For example, if you specify to run a backup job every 15 minutes, then if you specify to run a File Copy job after every 4 backups, there will be 24 File Copy jobs performed each day (1 each hour).

The number of backups that can be specified before the File Copy job runs must be in the range 1 - 700. By default, the schedule for file copying is after every five successful backups are completed.

3. Click Save Settings.

Your File Copy settings are saved.

Manage File Archive Settings

Arcserve UDP Agent (Windows) lets you copy selected source files from a backup session to a destination (disk or cloud) based upon your specified file copy and retention criteria. File Copy can be used for copying critical data to secondary locations.

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 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.

Before you perform your first File Copy job, specify the File Copy settings and plans. These configurations allow you to specify behaviors such as the source of your file copy data, destination for your copied files, the schedule for each file copy job, and the settings and filters applied to your file copy jobs. These settings can be modified at any time from the Arcserve UDP Agent (Windows) home page.

Note: To improve the performance (upload speed and server load), File Copy can upload the data to the specified destination in parallel chunks. To configure the number of chunks that are simultaneously sent to the destination, see [Configure File Copy Chunk Value](#).

To manage the File Copy settings, click the Settings link on the Arcserve UDP Agent (Windows) home page and select the File Copy Settings tab. The File Copy Settings dialogs consist of the following subordinate tab options:

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

Specify the File Archive Source

Arcserve UDP Agent (Windows) lets you specify the source settings for your information to be file copied.

Note: For more information about the File Copy Settings, see [Manage File Copy Settings](#).

Follow these steps:

1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **File Copy Settings** tab. When the **File Copy Settings** dialog opens, select **Source**.

The **File Copy Sources** dialog opens.

Note: If the Arcserve UDP Agent (Windows) is being managed by Console, not all settings are available and will be displayed as read-only information.

2. Specify your file copy source settings.

Enable File Copy

Enables the scheduled copying of files after the specified number of backups. If this option is not selected (File Copy disabled), no scheduled file copying is performed and any changes to the File Copy settings are not validated and saved.

Note: ReFS and deduplicated NTFS volumes will not be listed as selectable backup sources to be file copied. As a result, if all the volumes for a specified backup source are only ReFS or deduplicated NTFS volumes, the file copy option will be disabled.

Recovery Points to Copy From

Specifies the recovery point that you want to copy. You have two options to specify the recovery points. You can copy the recovery points from a specific backup number. For example, you can specify that the recovery points must be copied from every fifth backup. Another option is, you can specify to copy recovery points from a daily, weekly, or monthly backup.

File Copy Sources

Displays a selectable listing of all sources, with the corresponding policy (filters) and type of file copy (copy and retain or copy and move) to be performed after each successful Arcserve UDP Agent (Windows) backup. These File Copy Sources can be added, removed, or modified.

Note: Arcserve UDP Agent (Windows) 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. From the File Copy Restore user interface, the actual path will be displayed instead of the symbolic link.

■ Add

When clicked, the Plan type dialog opens to let you initially select the type of file copy job to perform (copy and retain or copy and move). After you select the plan type, the corresponding File Copy Plan dialog opens to let you add a source to copy and specify the corresponding plans for that source. For more information, see [Specify File Copy Plans](#).

Note: The File Copy source can only be selected from a volume that is currently selected in Backup Settings. If the source contains ReFS or deduplicated NTFS volumes, these volumes will not be available for selection.

■ Remove

When clicked removes the selected source from this displayed list.

■ Modify

When clicked, the File Copy Plans dialog opens to let you change the plan settings for the selected source. For more information, see [Specify File Copy Plans](#).

3. Click **Save Settings**.

Your File Copy settings are saved.

Specify File Archive Plans

When you click the Add Source option for File Copy, the Plan type dialog opens to let you initially select the type of File Copy job to be performed.

The available type is File Copy. In File copy plan, data is copied from the source to the destination (remains on source location) and provides multiple stored versions on the destination.

If you want to Add a new File Copy source or you want to Modify an existing File Copy source, the File Copy Plans dialog lets you specify the details.

Depending upon the plan type that is selected, a different File Copy Plans dialog opens; however, the selections are similar.

File Copy Selected:

File copy plans

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.

Browse

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

Add
Remove

You can use wildcard characters '*' and '?' in File/Folder Patterns

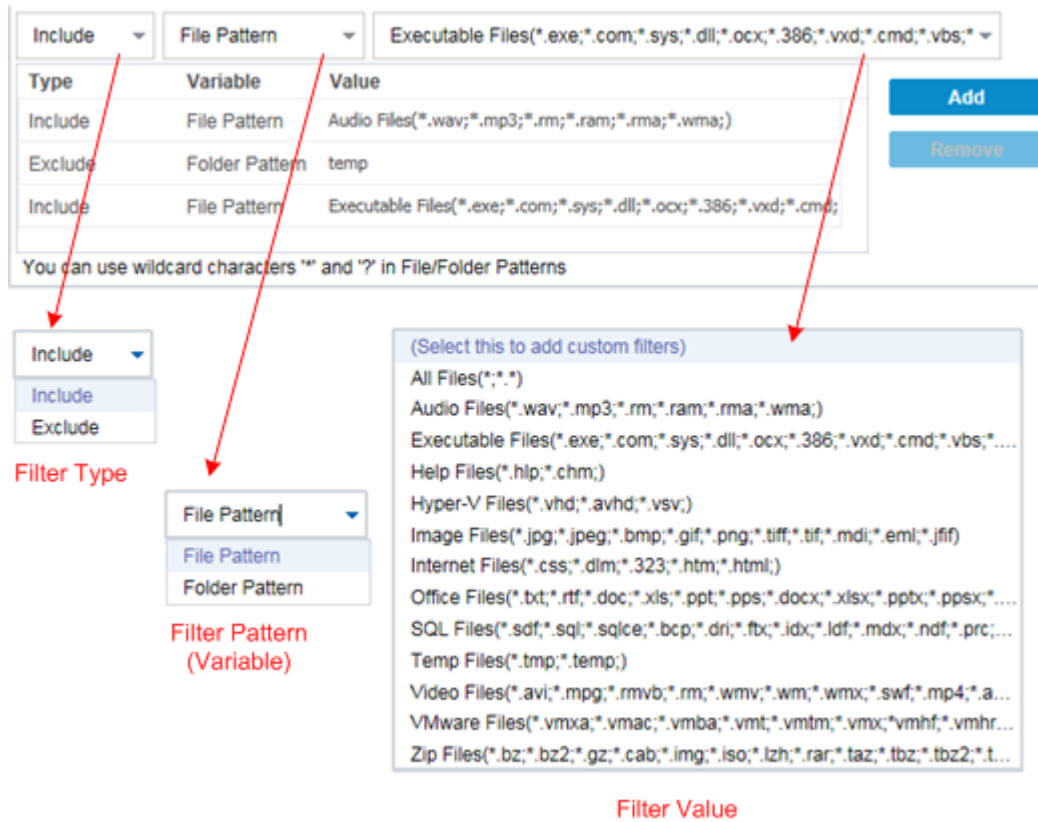
OK **Cancel** **Help**

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 pre-defined filters for many commonly used files is available (MS-Office files, Image files, Executable files, Temp files, etc.). After choosing any of the pre-defined filters, you can still append or modify the corresponding values.

How File Archive Filters Work

The File Copy source filters for files and folders work as follows:

- File with "d2darc" and "ASBUARC" extensions are always skipped.
- Files with system and temporary attributes are always skipped.
- Windows, Program Files, and Arcserve UDP Agent (Windows) installation folders (for both File Copy and File Copy - delete source policies) are always skipped.
- The following precedence order is used for filtering (with the highest precedence listed first):
 - Exclude directory filters
 - Exclude file filters
 - Include directory filters
 - Include file filters
 - Include Criteria
 - Exclude system and application files (Exchange and SQL only) present in any location. (This filter is only applicable for File Copy - Delete Source policies).
- A file is copied only if the Include Folder or the Include File filter matches, it does not need to satisfy both filter requirements.
- The File filter works only on the file name, and does not depend on the path.
For example, if you have the three files "Test.txt", "Hellotest.txt", and "TestHello.txt" these filters produce the following results:
 - Test*.txt filter matches only Test.txt and TestHello.txt
 - Test* filter matches Test.txt and TestHello.txt
 - Test filter does not match anything
 - *.txt filter matches all of them
 - *test filter does not match anything
- A Folder filter works on the policy source level.

For example, if you have the following directory structure:

```
C:  
->Z99  
-> ->A00  
-> -> ->B01
```

-> -> ->C01

-> -> ->D01

- If you configure your File Copy Source as "C:\Z99\A00" and you apply an Include Folder **b*** filter, then all files under c:\Z99\A00\B01 are copied.

In this example, the Source includes the parent folder and the asterisk is located after the "b". As a result all files in any folder subordinate to "A00" that starts with "b" are copied.

- If you configure your File Copy Source as "C:\Z99" and you apply an Include Folder **b*** filter, then this filter does not match any folder and no files are copied.

In this example, the Source does includes the "Z99" grandparent folder, but not the "A00" parent folder. As a result, there are no "b" folders directly subordinate to "Z99" and no files are copied.

- However if you specify a ***b*** filter, it now matches any subordinate folder that starts with "b" and all files within these "b" folders are then copied.

In this example, the asterisk is now located before the "b". As a result, all files in any folder subordinate to "C:\Z99" (regardless of the root level) that starts with "b" are copied.

- If you configure your File Copy Source as "C:\Z99" and you apply an Include Folder ***01** filter, then all subordinate folders containing "01" (B01, C01, and D01) are copied.

In this example, the asterisk is located before the "01". As a result, all files in any subordinate folders (regardless of the root level) that contains "01" are copied.

Note: Folder filters are always relative to the source folder path specified in the policy.

Specify the File Archive Destination

Arcserve UDP Agent (Windows) lets you specify the destination settings for your information to be file copied.

Note: For more information about the File Copy Settings, see [Manage File Copy Settings](#).

Specify the file copy destination

1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **File Copy Settings** tab. When the **File Copy Settings** dialog opens, select **Destination**.

The **File Copy Settings Destination** dialog opens.

Notes:

- If the Arcserve UDP Agent (Windows) is being managed by Console, not all settings are available and will be displayed as read-only information.
 - When agent is managed by console and not protected in a plan, all the settings are still available except the Preference > Updates panel.
2. Specify your File Copy destination settings.

Destination

Specifies the destination location for the file copy job. You can only select one destination.

Arcserve UDP Agent (Windows) 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. The Arcserve UDP Agent (Windows) 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](#).

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. You should 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 Agent (Windows) 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 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. The Arcserve UDP Agent (Windows) 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.

Files Retention

Retains the files in the file copy destination if the specified criteria is met.

File created within the last

Specifies the amount of time (years, months, days) that the stored data is retained at the destination location. At the end of the specified retention time period, the stored data is purged from the destination.

Important! At the end of the specified retention time when the data is purged from the destination, all of this purged data is no longer stored or saved.

Note: The Retention Time purge process is only triggered if the File Copy Schedule option is enabled.

3. Click **Save Settings**.

Your File Copy settings are saved.

Specify Cloud Configuration for File Archive

From the **File Copy Settings Destination** dialog, you can 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: [Empty]

Secret Access Key: [Empty]

Enable Proxy

Advanced

Bucket Name: [Empty] Add Refresh

Click 'Refresh' to load existing buckets

Bucket Region: [Empty]

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 cannot 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.

Follow these steps:

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.

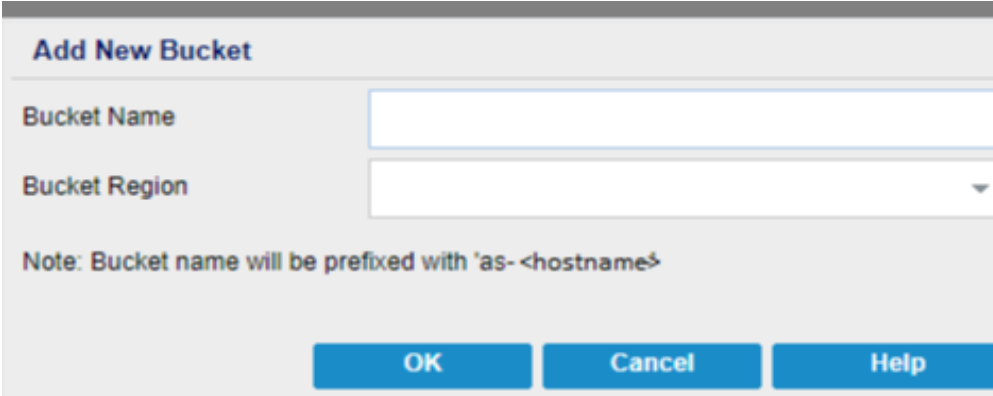
(For this field, Amazon S3 and Eucalyptus-Walrus use Bucket Name. Windows Azure and Fujitsu Cloud (Windows Azure) use Container).

Note: For the remainder of this step, all references to Buckets can also be applied to Containers unless specified.

You can either select a bucket name from the drop-down list, or you can add a new bucket name. If necessary, you can click the refresh button to update the list of available buckets.

To add a new bucket name:

- a. Click the **Add** button placed next to the Bucket Name field to display the Add New Bucket dialog.



- b. Enter a unique Bucket Name.

The new Bucket Name is automatically prefixed with name *as-`<hostname>`*. This format is applicable to the Bucket Name you create and is used as your File Copy destination.

Note: When creating a new bucket, Arcserve UDP Agent (Windows) only uses the *as-`<hostname>`* prefix and Arcserve UDP Agent (Windows) supports restoring from previous file copy destinations having *arcserve-`<hostname>`*-*d2dfilecopy-`<hostname>`*- or *d2d-filecopy-`<hostname>`*- prefixes.

A bucket name should be unique, easily identifiable, and compliant with internet domain naming rules. No two buckets can have the same name. We recommend to understand valid syntax for bucket names.

For more information about bucket naming requirements of Amazon S3 and Eucalyptus-Walrus, refer to the Amazon S3 documentation.

For more information about container naming requirements of Windows Azure and Fujitsu Cloud (Windows Azure), refer to the Microsoft documentation.

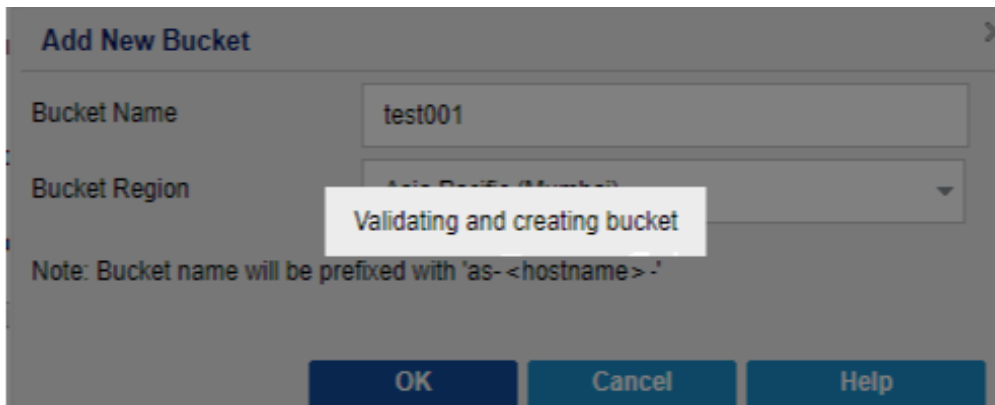
- c. For Amazon S3 only, select an available region from the drop-down menu. By default, all available regions are included in the drop-down menu and you can select the region where you want the new bucket to be created.

Regions allow you to select the geographical region where Amazon S3 stores the buckets that you create. You should select a Region that provides you with fast access to your data and allows you to optimize latency, minimize costs, or address regulatory requirements.

(For Windows Azure, Fujitsu Cloud (Windows Azure), and Eucalyptus-Walrus, the region is not selectable).

- d. After you have specified your values, click **OK**.

The Bucket name is validated and created at the cloud.



After you successfully create the new bucket, the main Cloud Configuration dialog is displayed again, with the new bucket information (name and region) included in the Advanced Settings fields.

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.

Configure File Archive Settings to Optimize Performance

To improve the performance (upload speed and server load), File Copy can upload data to the specified destination in parallel chunks and or parallel threads.

Chunk Value

You can set the number of 1-MB chunks that will be simultaneously sent to the destination. By increasing the number of parallel chunks, you will decrease the time to complete the job, but will also have an adverse affect on server performance. Configure this value as necessary to obtain optimal performance.

For example, if you are performing File Copy for a 10-MB file and set the number of 1-MB chunks to 2, then File Copy will write 10 chunks, two at a time. If you see that this is taking too long to complete the job, change this value to 4. The time to complete the job will then decrease because File Copy will now be writing 10 chunks, four at a time, but the load on your server will increase.

Threads for Archive Value

File Copy lets you copy more than one file at a time. By default, File Copy transfers 8 files in parallel when the destination is configured to File Systems and transfers 32 files in parallel when the destination is configured to Cloud. If you see that File Copy is taking too long to transfer the data, increase the number of threads up to 32, to optimize performance. However, if you experience a problem on a machine with less memory, decrease the number of threads.

Chunk Value and Threads for Archive Value can be used together to control the speed of the File Copy. If you increase Chunk Value and Threads for Archive Value, you see File Copy is performed faster.

For example, if you are transferring 8 files with 10 MB each and set the number of 1-MB chunks to 2, then File Copy will write 16 at a time (8 files X 2-MB chunks), but the load on your server will increase. When you see the load on the server has increased to a point where it becomes a problem, decrease the number of threads. If the destination is a Cloud location, it is recommended that you configure these settings in such a way that produces at least 20 writes, to optimize performance.

Threads for Restore Value

Restore from a File Copy lets you download more than one file at a time. By default, restores from file copies downloads 8 files when the File Copy location is configured to File Systems and downloads 32 files in parallel when the file

copy location is configured to Cloud. If you see that the restore from a File Copy is taking too long to transfer the data, increase the number of threads up to 32.

Note: Chunk Value does not apply to restore jobs.

Threads for Catalog Synchronization Value

Catalog Synchronization jobs lets you use multiple threads to optimize performance.

If you see that the catalog synchronization job is taking too long to transfer the data, increase the number of threads up to 10. You will see that the job is performed faster and the load on the server increases. When you see the load on the server has increased to a point where it becomes a problem, decrease the number of threads.

To configure the file copy settings to optimize performance, set the corresponding DWORD values as follows:

1. Start edit registry.

2. Locate key:

"HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine\AfArchiveDll"

Note: The same registry key is used when your File Copy destination is either File Systems or a Cloud location.

3. To modify the value for the number of 1-MB chunks that will be simultaneously sent to the destination, follow these steps:

- a. Manually create a DWORD value for "ArchMultChunkIO".
- b. Assign a DWORD value:

The available range for the number of chunks is 1 - 4 chunks.

Default: 4 chunks

Maximum: 4 chunks

4. To modify the value for the number of threads (files) that will be transferred in parallel to the copy destination, follow these steps:

- a. Manually create a DWORD value for "ThreadsForArchive".
- b. Assign a DWORD value:

The available range for the number of files is 1 - 32 files.

Default: 8 files when the destination is configured to File Systems and 32 files when the destination is configured to a Cloud location.

Maximum: 32

5. To modify the value for the number of file copies that can be downloaded in parallel from the copy destination, follow these steps:
 - a. Manually create a DWORD value for "ThreadsForRestore".
 - b. Assign a DWORD value:

The available range for the number of files is 1 - 32 files.

Default: 8 files when the copy destination is File Systems and 32 files when the copy destination is a Cloud location.

Maximum: 32

6. To modify the value for the number of threads (streams) that can be used in parallel to perform catalog synchronization, follow these steps:
 - a. Manually create a DWORD value for "ThreadForCatalogSync".
 - b. Assign a DWORD value:

The available range for the number of files is 1 - 10 threads.

Default: 8 threads

Maximum: 10

Specify the File Archive Schedule

Arcserve UDP Agent (Windows) lets you specify the schedule settings for your information to be file copied.

Note: For more information about the File Copy Settings, see [Manage File Copy Settings](#).

Follow these steps:

1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **File Copy Settings** tab. When the **File Copy Settings** dialog opens, select **Schedule**.

The **File Copy Settings Schedule** dialog opens.

Notes:

- If the Arcserve UDP Agent (Windows) is being managed by Console, not all settings are available and will be displayed as read-only information.
 - When agent is managed by console and not protected in a plan, all the settings are still available except the Preference > Updates panel.
2. Specify your File Copy schedule settings.

Schedule

Enables the file copying of data after the specified number of backups.

The file copy process will be launched automatically after the specified number of successful backups (Full, Incremental, and Verify) and will be based on your selected File Copy policies.

You can use this setting to control how many times a File Copy job is triggered each day. For example, if you specify to run a backup job every 15 minutes, then if you specify to run a File Copy job after every 4 backups, there will be 24 File Copy jobs performed each day (1 each hour).

The number of backups that can be specified before the File Copy job runs must be in the range 1 - 700. By default, the schedule for file copying is after every five successful backups are completed.

3. Click Save Settings.

Your File Copy settings are saved.

Configure the Copy Recovery Point Settings

Arcserve UDP Agent (Windows) lets you specify the recovery point copy settings. Before you copy a recovery point, configure the copy recovery point settings. For a better understanding about how the options on this dialog can be used to configure your recovery point copy schedule, see [Copy Recovery Points - Example Scenarios](#)

Note: The recovery point copy process is a copy and paste operation only and not a cut and paste operation. As a result, whenever a scheduled copy recovery point job is performed Arcserve UDP Agent (Windows) creates an additional copy of the recovery point to the specified copy destination, while still retaining the original copy of the recovery point at the backup destination that was specified in Backup Settings.

Follow these steps:


1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Copy Recovery Points** tab. When the **Copy Recovery Points** dialog opens, select **Copy Settings**.

The **Copy Recovery Points** dialog opens.

Notes:

- If the Arcserve UDP Agent (Windows) is being managed by Console, not all settings are available and will be displayed as read-only information.

- When agent is managed by console and not protected in a plan, all the settings are still available except the Preference > Updates panel.



2. Select **Enable Copy Recovery Points**.

When selected, enables the copying of recovery points.

Note: If you do not select this option, no scheduled copying of recovery points is performed.

3. Specify the following recovery point copy schedule settings:

Destination

Specifies the storage location for the copy of the selected recovery point. (Optional) You can click the green arrow button to verify the connection to the specified location.

Note: The maximum length for the specified destination path is 158 characters.

Copy Recovery Points job will run after the specified number of backups are performed

Specifies when the scheduled recovery point copy process is automatically launched. This process is launched based on your selected copy policies and specified number of successful backups (Full, Incremental, and Verify).

Note: Number of successful backup is counted for any custom, daily, weekly monthly backups that are configured.

You can use this setting to control how many times a recovery point copy process is triggered each day. For example, if you schedule to run a backup job every 15 minutes, and copy job after every 4 backups, then it performs 24 recovery point copy jobs each day (1 each hour).

Default: 8

Minimum: 1

Maximum: 1440

Important! If you schedule backup and copy jobs to run at regular intervals and if the copy job is currently running (in active state) when the scheduled time for the backup job time arrives, the backup job fails. (The next backup job will run as scheduled and should be successful if it does not conflict with another copy job). Because the copy operation takes almost same amount of time as performing a full backup, the best practice is not to set a frequent schedule for your recovery point copy jobs.

Specify the number of recovery points to keep

Specifies the number of recovery points that are retained and stored at the specified copy destination. Discards the oldest recovery point, when this number is exceeded.

Note: If you do not have sufficient free space at the target destination, reduce the number of saved recovery points.

Default: 1

Maximum: 1440

4. Select the **Compression** level.

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.

5. If you also want the copied recovery point to be encrypted, specify the following information:

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

Lets you specify and confirm the encryption password being used to encrypt the destination session.

6. Click **Save Settings**.

Your recovery point copy settings are saved.

The copy recovery point settings are successfully configured.

Copy Recovery Points - Example Scenarios

The following example scenarios are provided to give you a better understanding of how the various options can affect your scheduled copying of recovery points.

For this example, assume that you configured your Arcserve UDP Agent (Windows) backup schedule as follows:

- Full Backup - Every 7 days
- Incremental Backup - Every 1 hour
- Verify Backup - Every 3 days

and assume:

- First backup is on Day #1 at 5:00PM (by default, the first backup is always a Full Backup)
- First Incremental Backup will be on Day #1 at 6:00PM (and every hour after)
- Recovery Points retention count is set to 31 (default number)
- Location "D" is configured as the copy destination.

Scenario #1

For this scenario, the Copy Recovery Point settings are as follows:

- Copy after 4 backups
- Retain 1 recovery point

Result:

- At 8:00PM (after the 4th backup), the scheduled copy job will run and consolidate all 4 recovery points into a single recovery point and store it at destination D.
- At 12:00 midnight (after the 8th backup), the next scheduled copy job will run and consolidate all 8 recovery points into a single recovery point and store it at destination D.

The previous recovery point is removed from destination D because the setting is to retain only 1 recovery point at the destination.

Scenario #2

For this scenario, the Copy Recovery Point settings are as follows:

- Copy after 4 backups
- Retain 4 recovery points

Result:

- At 8:00PM (after the 4th backup), the scheduled copy job will run and consolidate all 4 recovery points into a single recovery point (Recovery Point #1) and store it at destination D.
- At 12:00 midnight (after the 8th backup), the next scheduled copy job will run to create Recovery Point #2 and store it at destination D.
- At 4:00AM on Day #2 (after the 12th backup), the next scheduled copy job will run to create Recovery Point #3 and store it at destination D.
- At 8:00AM on Day #2 (after the 16th backup), the next scheduled copy job will run to create Recovery Point #4 and store it at destination D.
- At 12:00 noon on Day #2 (after the 20th backup), the next scheduled copy job will run. A new recovery point will be created and the first recovery point (created after the 8:00PM backup on previous day) is removed from destination D, because the setting is to retain only 4 recovery points at the destination.

Scenario #3

For this scenario, the Copy Recovery Point settings are as follows:

- Copy after 1 backup
- Retain 4 recovery points

Result:

- At 5:00PM (after the 1st backup), the scheduled copy job will run to create a single recovery point (Recovery Point #1) and store it at destination D.
- At 6:00PM (after the 2nd backup), the next scheduled copy job will run to create Recovery Point #2 and store it at destination D.
- At 7:00PM (after the 3rd backup), the next scheduled copy job will run to create Recovery Point #3 and store it at destination D.
- At 8:00PM (after the 4th backup), the next scheduled copy job will run to create Recovery Point #4 and store it at destination D.
- At 9:00PM (after the 5th backup), the next scheduled copy job will run. A new recovery point will be created and the first recovery point (created after the 5:00PM backup) is removed from destination D, because the setting is to retain only 4 recovery points at the destination.

Specify Preferences

The **Preferences** dialog page provides a quick and easy way to specify various options for the behavior of your Arcserve UDP Agent (Windows). When clicked, the Preferences dialog opens with the following subordinate tabs:

Specify General Preferences

Arcserve UDP Agent (Windows) lets you specify your General preferences:

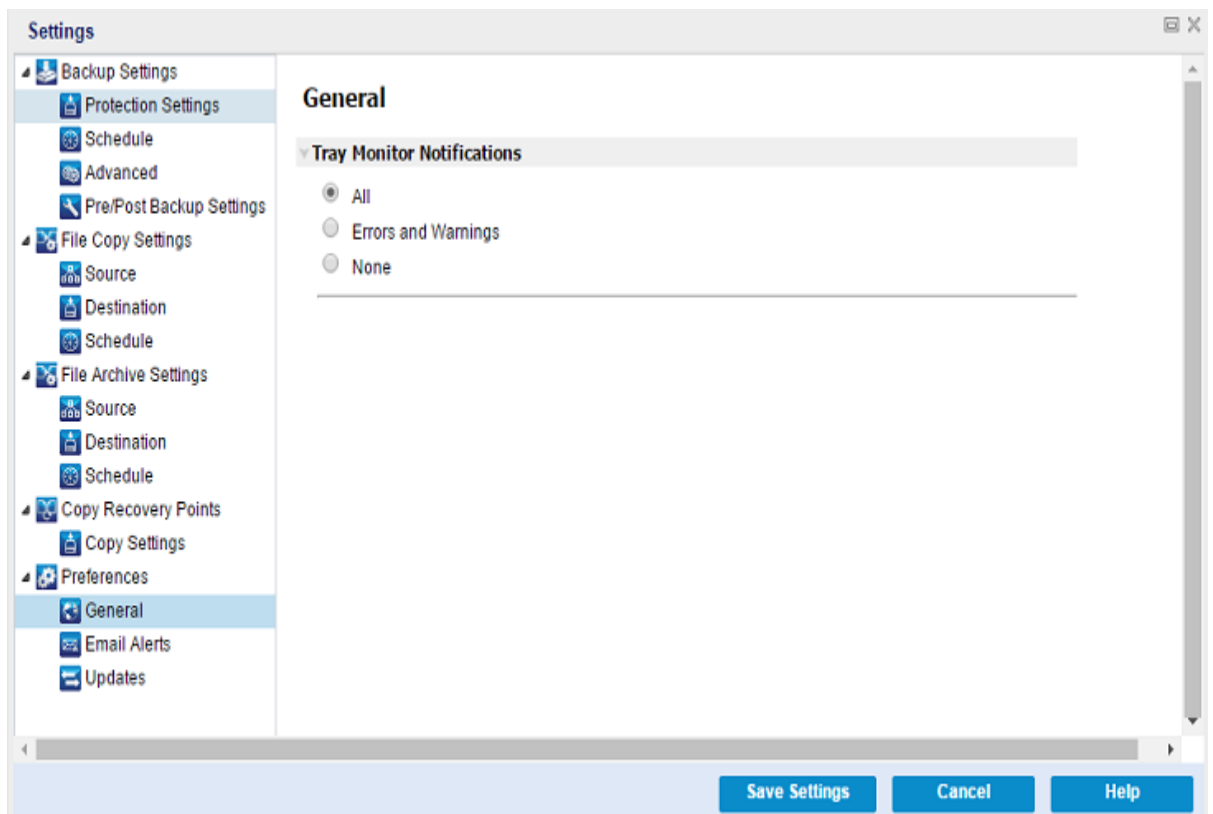
Specify the General Preferences

1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Preferences** tab. When the **Preferences** dialog opens, select **General**.

The **General** preferences dialog opens.

Notes:

- If the Arcserve UDP Agent (Windows) is being managed by Console, not all settings are available and will be displayed as read-only information.
- When agent is managed by console and not protected in a plan, all the settings are still available except the Preference > Updates panel.



2. Specify your General preference settings.

Tray Monitor Notifications

Select the type of Alert notifications you want to display. The available options are **All**, **Errors and Warnings**, and **None**.

3. Click Save Settings.

Your General preference settings are saved.

Specify Email Preferences

Arcserve UDP Agent (Windows) lets you specify the following Email Alert preferences:

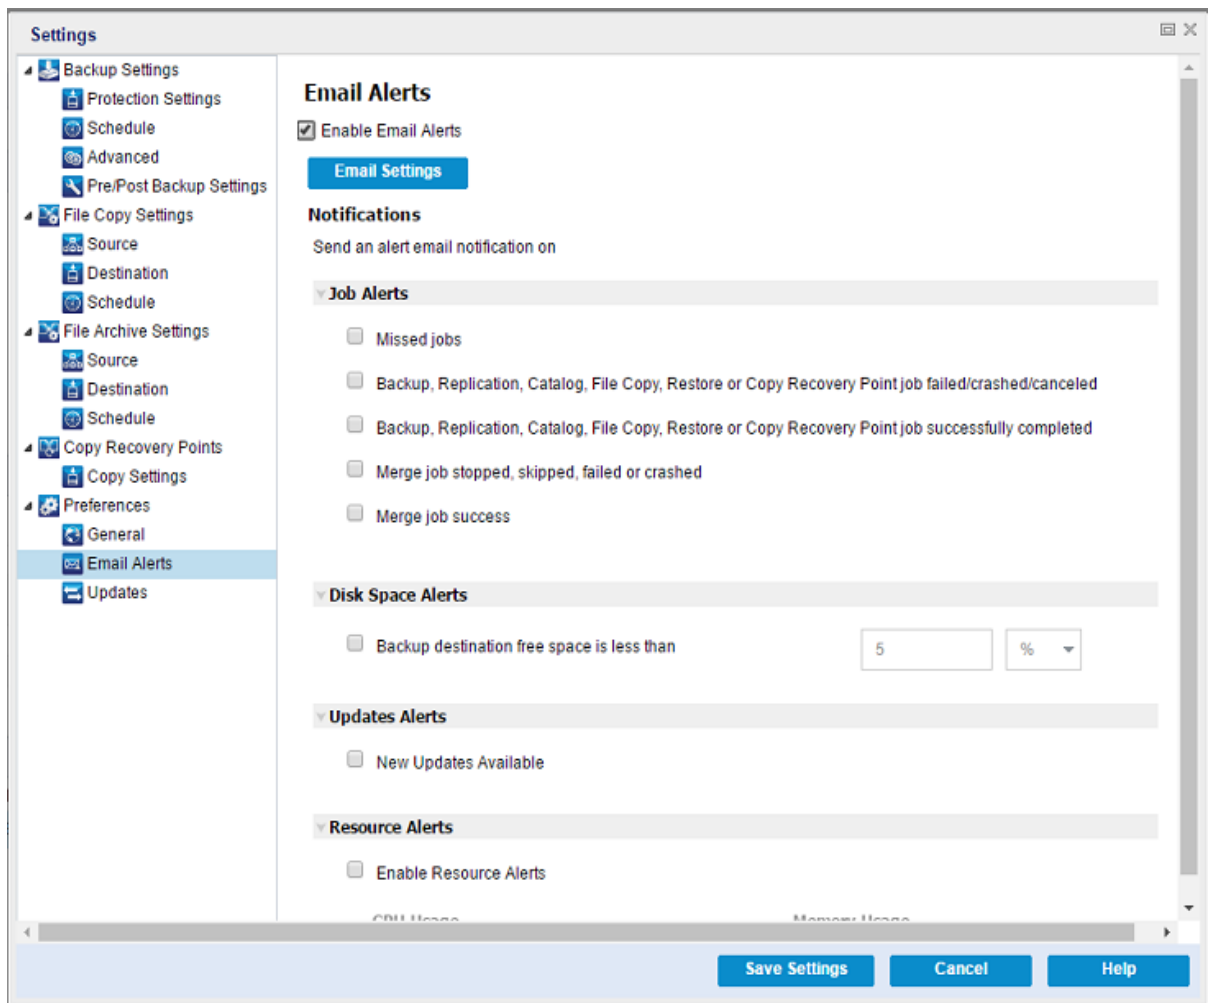
Specify the Email Alerts Preferences

1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Preferences** tab. When the **Preferences** dialog opens, select **Email Alerts**.

The **Email Alerts** preferences dialog opens.

Notes:

- If the Arcserve UDP Agent (Windows) is being managed by Console, not all settings are available and will be displayed as read-only information.
- When agent is managed by console and not protected in a plan, all the settings are still available except the Preference > Updates panel.



-
2. Select the **Enable Email Alerts** checkbox to send an email notification for **Job Alerts, Disk Space Alerts, Updates Alerts, and Resource Alerts**.
 3. Specify your Email Alerts notification settings.

Automatic email alert notifications are sent upon the completion of selected events. You can select any or all of the available options.

Note: If you do not need specific notifications for successful jobs, you can configure Arcserve UDP Agent (Windows) to only send email alerts for failed and missed jobs. This configuration could help you reduce the amount of email notifications and also monitor any failures.

The available options are to send an alert notification for the following events:

Missed jobs

Sends an email alert notification for all missed jobs. A missed job is any scheduled job that did not run at the scheduled time. A missed job could happen when some other job of the same type is running or previous job that started earlier did not finish yet.

Arcserve UDP Agent (Windows) allows different types of jobs to be run in parallel; however, only one job of each type can be run at the same time. For example, if a copy job is still running at the scheduled time for another copy job then the scheduled copy job is missed, but another backup job can still run.

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

Sends an alert notification for all unsuccessful backup, catalog, file copy, restore, or copy recovery point job attempts. This category includes all failed, incomplete, and canceled jobs, and crashed attempts.

Note: These email alerts are sent with a high importance. The email alerts that have a high importance level setting display a visual indicator of an exclamation point in their Inbox.

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

Sends an alert notification for all successful backup, catalog, file copy, restore, or copy recovery point job attempts.

Merge job stopped, skipped, failed, or crashed

Sends an alert notification for all stopped, skipped, failed, or crashed merge jobs. If you enable this alert, you are informed once a merge job is unsuccessful.

A merge failure can occur for the following reasons:

- The session is mounted.
To solve the problem, you can dismount the session.
- The session is locked by a catalog job.
The next backup job automatically merges this session.
- The session is locked due to other reasons.

If you disable this alert, you only know when a merge was unsuccessful from the balloon message in the tray monitor or the Recovery Points Summary on the Arcserve UDP Agent (Windows) Home Page.

Merge job success

Sends an alert notification for all successful merge jobs.

Backup Destination free space is less than

Sends an email alert notification when the amount of unused space at the backup destination is less than a specified value. For this option, you can further select either a percentage of the total capacity or a specific value (in MB) for the threshold level of when the alert notification is sent.

New Updates Available

Sends an email notification when a new update for Arcserve UDP Agent (Windows) is available. Email notifications are also sent if a failure occurs during the check for updates or during the download.

Enable Resource Alerts

Sends an email notification when any specified resource threshold level is reached. To ensure that your server is efficient and reliable, continually be aware of the performance to identify possible problems and quickly address bottleneck situations.

Defining threshold levels for these resource indicators is strictly up to you and your knowledge of your server. You cannot specify right or wrong settings and could base these alert notifications upon "normal" and acceptable performance. For example, if your system typically runs at an 80 percent CPU load, then setting a CPU Usage threshold at 75 percent would not be useful or efficient.

Each of these resource parameters can be separately configured to send an alert notification when the corresponding threshold level is reached. The maximum number that each resource alert email is sent is 5 per day.

- **CPU Usage**

The specified CPU Usage alert threshold indicates the percentage of CPU usage for your Arcserve UDP Agent (Windows) protected server. You can use this alert notification to ensure that your server does not become overloaded too often.

If your CPU usage is too high, your server response time can become slow or unresponsive. Therefore, consider spreading out (balancing) your load.

– **Disk Throughput**

The specified Disk Throughput alert threshold indicates the disk throughput (MB/second) for your Arcserve UDP Agent (Windows) protected server. You can use this alert notification to ensure that you are maximizing the capability of your disk.

If your disk throughput is close to the maximum value your disk can handle, consider upgrading to a disk that better matches your needs. Generally a faster disk leads to better performance.

Memory Usage

The specified Memory Usage alert threshold indicates the percentage of memory in use on your Arcserve UDP Agent (Windows) protected server. Utilization is how much of your memory capacity you are using. The higher the percentage the worse your server performance is going to be.

If your memory use continually becomes too high, determine the process causing this high usage. You can use this indicator setting to alert you of when an application or server upgrade can be necessary.

Network I/O

The specified Network I/O alert threshold indicates the percentage of NIC bandwidth you are currently using on your Arcserve UDP Agent (Windows) protected server. Utilization is how much of your network interface card (or NIC) capacity you are using. The higher the percentage the worse your network performance is going to be.

If your network use continually becomes too high, determine the process causing this high usage and remedy the problem. In addition, if based on your specific network capacity the percentage of your network use is too high during backup time, you can upgrade your NIC card to handle the higher throughput requirements.

4. Click Save Settings.

Your Email Alerts preference settings are saved.

5. After you select to send an email notification, you can then click Email Settings to display the related dialog.

Specify Email Settings

From the **Email Settings** dialog, you can specify email settings such as the:

- Mail server
- Subject title
- Who the email is coming from
- Email recipients

You can also enable and define your proxy settings. These settings apply to all email alert notifications and can be modified at any time.

After you establish your email settings, you can test the settings using the **Test Email** button. After a job runs successfully or fails, it is too late to receive an email alert for having the wrong settings. Therefore, testing the email information that is provided validates the settings and tries to send an email using the specified settings. If the email settings are valid, you receive an email indicating so. If the email

settings are invalid, you receive a failure message.

Email Settings

Email Settings

Service: Other

Mail Server: [] Port: 25

Requires Authentication

Account Name: []

Password: []

Subject: Arcserve UDP Agent Alert

From: []

Recipients: []

Use SSL Send STARTTLS Use HTML format

Enable Proxy Settings

Proxy Server: [] Port: 1080

Requires Authentication

Proxy User Name: []

Proxy Password: []

Test Email OK Cancel Help

Service

The email provider service to use for sending the alert notifications. The available options are Google Mail, Yahoo Mail, Live Mail, and Other.

- ◆ If you select Other, identify the mail server and corresponding port number used.
- ◆ If you select Google Mail, Yahoo Mail, or Live Mail the mail server and port number fields are automatically populated.

Default: Other

Mail Server

The host name of the SMTP mail server that Arcserve UDP Agent (Windows) can use to send the email alerts.

Port

The output port number for the mail server.

Requires Authentication

Specifies if this mail server requires authentication when attempting to send an email by way of the Internet. When this option is selected, provide the corresponding user Account Name and Password.

Subject

Subject description for the email alert notifications that Arcserve UDP Agent (Windows) sends.

Default: "Arcserve UDP Agent Alert"

From

The email address that Arcserve UDP Agent (Windows) uses to send the email alert notifications.

Recipients

Email address for the recipients of the email alert notifications sent.

Note: To enter multiple email addresses, separate each address with a semi-colon character.

Use SSL

Email server requires an SSL (Secure Sockets Layer) connection to transmit data securely by way of the Internet.

Send STARTTLS

Email server requires a STARTTLS (Start TLS extension) command that is issued to initiate a secure SMTP connection between servers.

Use HTML format

Email alert notifications are sent as HTML. If this option is not selected, the alerts are sent as plain text. By default, this option is selected.

Enable Proxy Settings

Specifies if you want to connect to a proxy server for sending your email alert notifications. When this option is selected, provide the corresponding name of the proxy server and port number.


Specify Updates Preferences

Arcserve UDP Agent (Windows) lets you specify the following Updates preferences:

Follow these steps:

1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Preferences** tab. When the **Preferences** dialog opens, select **Updates**.

The **Updates** preferences dialog opens.



2. Specify your **Updates** preference settings.

Download Server

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

▪ Arcserve Server

You can use this option to specify that Arcserve UDP Agent (Windows) updates are downloaded from the Arcserve server directly to your local server.

This is the default setting.

▪ Staging Server

You can use this option to specify the server that is used as a staging server.

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

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

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

Arcserve UDP Agent (Windows) updates are downloaded from the Arcserve server directly to the specified staging server location. After the updates are downloaded to this staging server, you can then further download the updates from the staging server to a client server. If you select the Staging Server location, you must also specify the host name or IP address for the staging server, along with the corresponding port number.


You cannot specify your same local client server as this staging server. This is an invalid configuration because the staging server cannot connect to itself to get and download the available updates from. If you attempt to use your local client server as the staging server, an error message is displayed.

▪ 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 Agent (Windows) 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 is the connection to the Arcserve server from which your download server gets the updates.

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



– **Use browser proxy settings**

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

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

– **Configure proxy settings**

When selected enables the specified proxy server to connect to the Arcserve server for Arcserve UDP Agent (Windows) 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 requires 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, a red icon is displayed on the status **Summary** section home page to provide a visual alert of this condition.

Note: The test connection is automatically performed when you launch the **Preferences Updates** dialog from the home page. When this auto test is performed it will check the latest connection status of the previously configured download server (either Arcserve Server or Staging Server(s), whichever is selected). If you previously configured more than one staging server, then this auto test is performed on all staging servers to get the latest connection status.

Update Schedule

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

- With this option selected, specifies to automatically check for new and available Arcserve UDP Agent (Windows) updates. When you select this option, you then have drop-down menu capabilities to specify when to perform this function (every day or weekly on a specified day) and the time of the day that it is performed.

Note: The default setting for the day or hour that these checks are automatically performed is randomly assigned by Arcserve UDP Agent (Windows) at the time of installation. After installation, you can use this **Update Schedule** setting to change the day and time for these checks.

By default, if this check determines that a new update is available, Arcserve UDP Agent (Windows) also automatically downloads the update.

- With this option not selected, specifies to disable all automatic check and download functions (and its status is displayed under status Summary section of the home page). With this option not selected, these update functions can only be triggered manually.

Notes:

If configured you get an email notification if the scheduled check for updates discovers that a new update is available. In addition, email notifications are sent if a failure occurs during the check for updates or during the download.

If the Arcserve UDP Agent (Windows) is managed by the Arcserve UDP Console, the **Automatically check for updates** option is disabled. Instead you can check updates from the Arcserve UDP Console and remote deploy updates to Arcserve UDP Agent (Windows).

3. Click **Save Settings**.

Your Updates preference settings are saved.

How to create a Staging Server

Staging server is a node on which the Arcserve UDP Agent or Console is installed. Once this node finishes downloading updates from arcserve download server, it can work as a staging server to provide updates for others.

Adding a Staging Server:

You can add a staging server manually, consider the following notes:

- For other nodes, to download updates from staging server you need to specify the server name. By default, console is 8015 and agent is 8014.
- To work as a staging server, the node can use 'http' or 'https' protocol.
- Arcserve UDP console can only download updates from console staging server.
- Arcserve UDP agent can download updates from the console or agent staging server.

Manage Export/Import Settings

Arcserve UDP Agent (Windows) lets you export and import settings using the JSON file. This section provides information about how to perform the export settings from your agent and import settings to the same or different Windows agents.

This section contains the following topics:

Export Settings

This section explains how to export the Windows Agent settings as a JSON file.

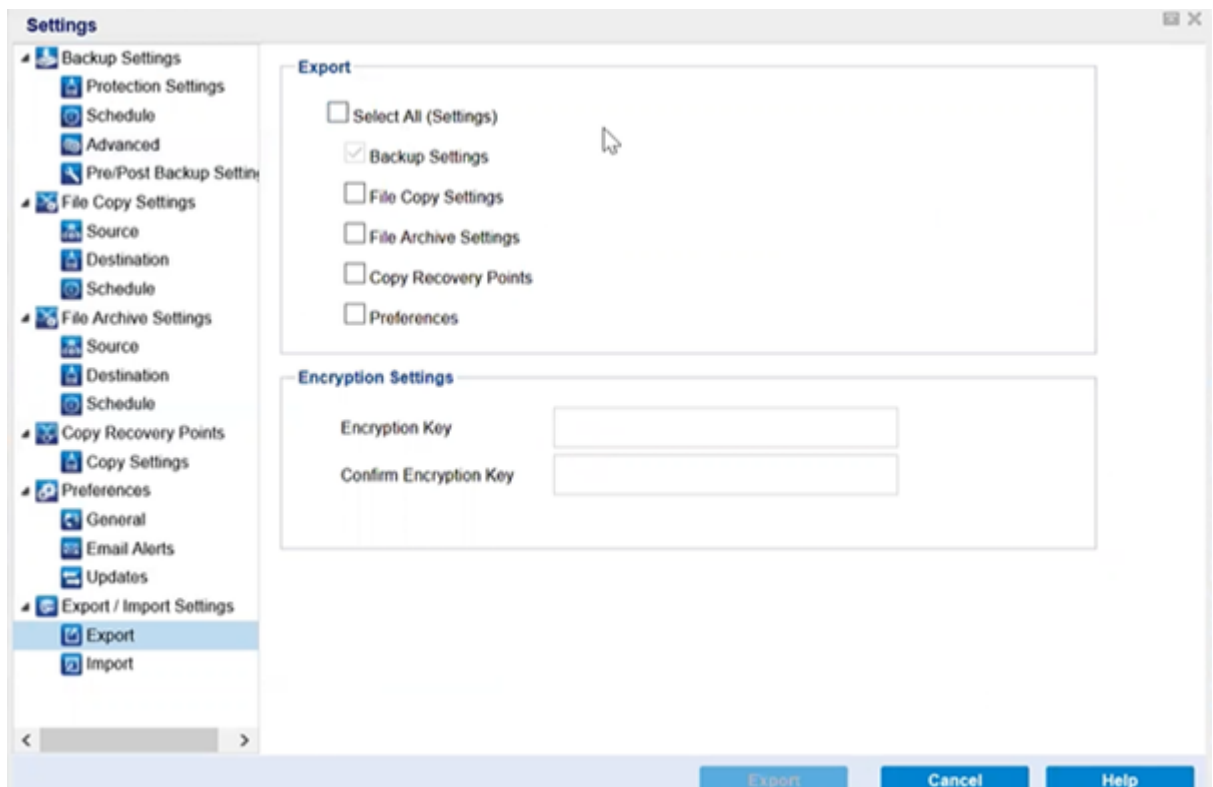
Follow these steps:

1. Navigate to the Arcserve UDP Agent (Windows) home page or Arcserve UDP Agent (Windows) Monitor, and then select **Settings** from the taskbar.

The Settings window opens.

Note: If the Arcserve UDP Agent (Windows) is being managed by console, not all settings are available and will be displayed as read-only information.

2. Under the Export/Import Settings side menu, select **Export**.



3. For Export, select all or individual settings as needed.

Note: The Backup Settings option is selected by default because it is a prerequisite for all other settings.

4. For Encryption Settings, type an Encryption Key, and then retype to confirm.
5. Click **Export**.

The Agent settings are successfully exported as a JSON file.

Import Settings

This section explains how to import the JSON file that contains the Windows Agent settings to the same or different Windows Agents.

Note: Before importing, you may modify the configuration values in the JSON file as required.

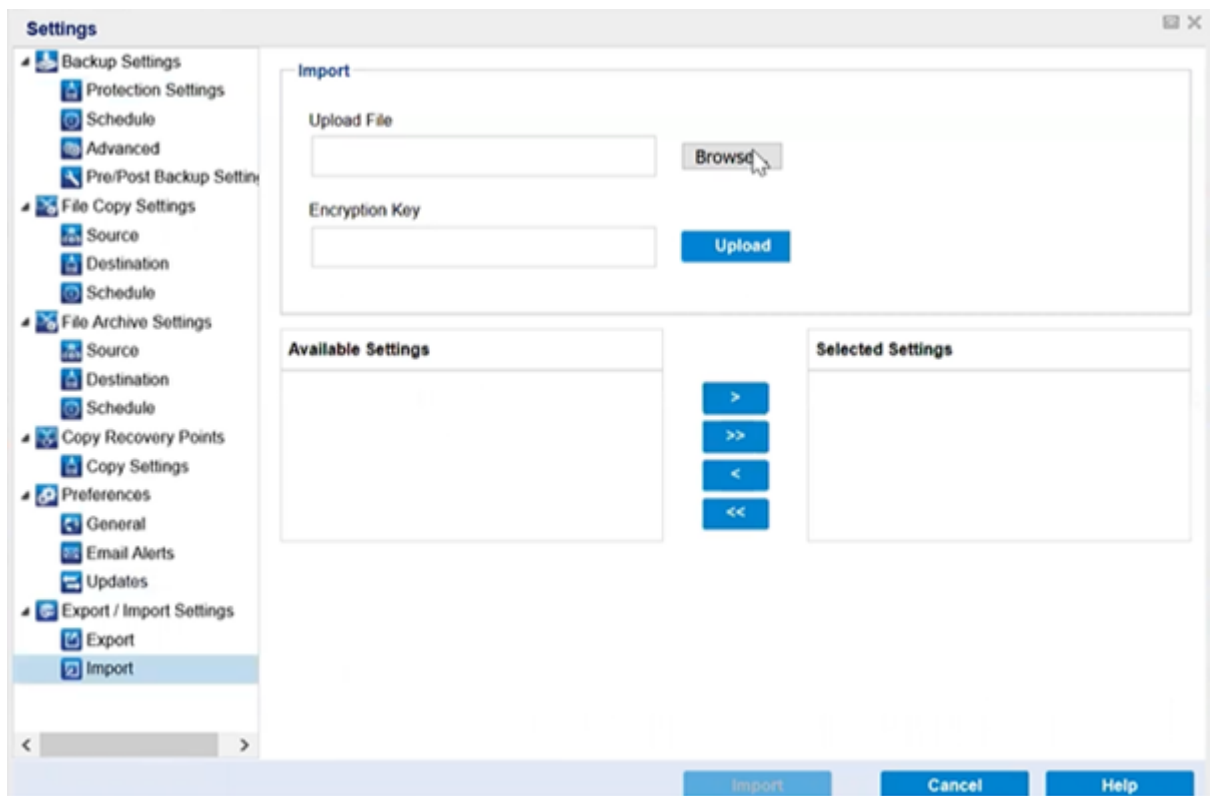
Follow these steps:

1. Navigate to the Arcserve UDP Agent (Windows) home page or Arcserve UDP Agent (Windows) Monitor, and then select **Settings** from the taskbar.

The Settings window opens.

Note: When agent is managed by console and not protected in a plan, all the settings are still available.

2. Under the Export/Import Settings side menu, select **Import**.



3. For Import, do the following, and then click **Upload**:

Upload File: Click the **Browse** button to locate and select the JSON file.

Encryption Key: Type the encryption password you provided while exporting the settings.

The settings display in the Available Settings box.

4. To import the required settings, from the Available Settings box, select the preferred settings, and then click the right-arrow to move the settings to the Selected Settings box.
5. Click **Import**.

The Agent settings are successfully imported.

Chapter 5: Using Arcserve UDP Agent (Windows)

This section contains the following topics:

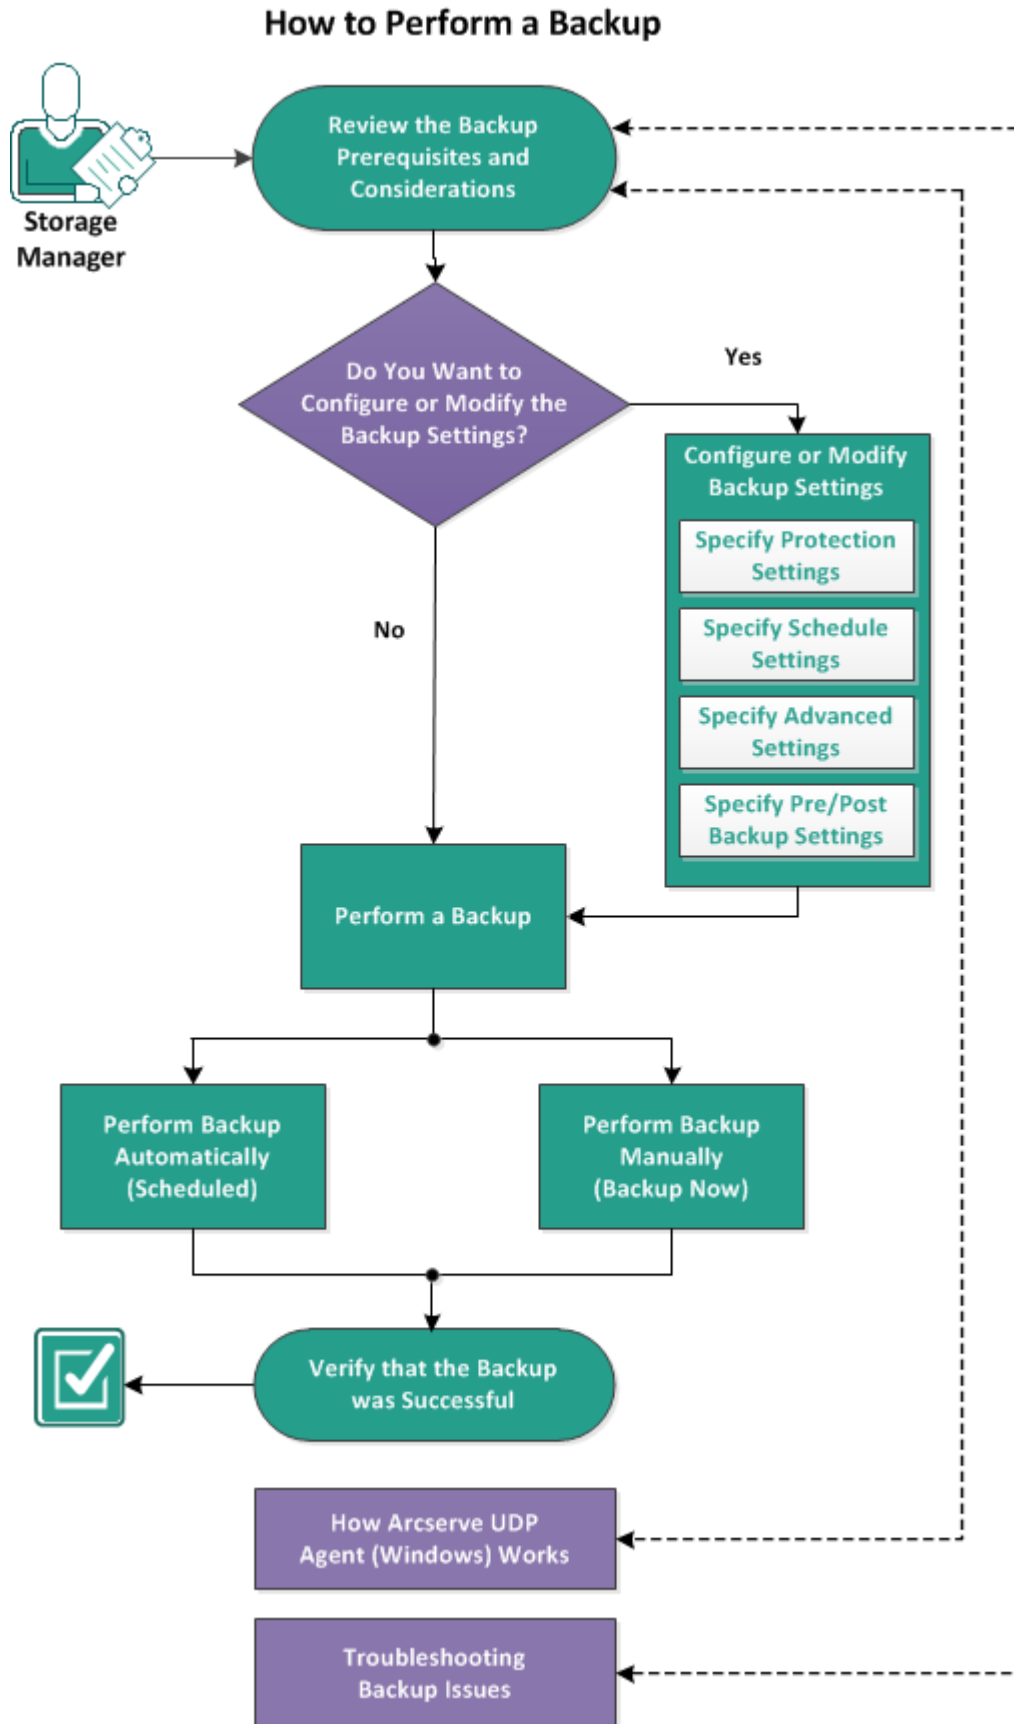
How to Perform a Backup	268
Perform File Copy to Disk/Cloud	353
Perform a Restore	354
How to Copy a Recovery Point	542
Mount a Recovery Point	559
Create a VHD File from an Arcserve UDP Agent (Windows) Backup	565
View Logs	569
How to Download File/Folders without Restore	572
How to Create a Boot Kit	574
How to Perform a Bare Metal Recovery Using a Backup	591
How to Perform a Bare Metal Recovery Using a Virtual Standby VM or Instant VM	627
Using the PowerShell Interface	665
Add Arcserve UDP Agent (Windows) Licensing	686
Change Server Communication Protocol	688
Use Scripts to Backup and Restore MySQL Database	688
Use Scripts to Backup and Restore PostgreSQL Database	690

How to Perform a Backup

Arcserve UDP Agent (Windows) lets you perform frequent backups (as frequently as every 15 minutes), reducing the size of each incremental backup (as well as the backup window) and providing a more up-to-date backup.

Before you perform your first backup, review the backup prerequisites and considerations and then configure or modify the backup settings which are applied to each backup job. A backup job can be initiated either automatically based upon your schedule settings (Scheduled), or manually as an immediate or ad-hoc backup (Backup Now).

The following diagram illustrates the process for how to perform a backup:



Complete the following tasks to perform a backup:

Review the Backup Prerequisites and Considerations

Review the following backup considerations before you perform an Arcserve UDP Agent (Windows) backup:

- **Destination available free space**

If your destination does not have sufficient free space, you can consider the following corrective actions:

- Reduce the number of saved recovery points.
- Increase the available free space at the backup destination.
- Change the backup destination to a larger capacity.
- Reduce the size of the backup source (maybe eliminate unnecessary volumes from the backup).
- Increase the Compression setting of the backup.

- **Verify that you have proper licensing**

When using Arcserve UDP Agent (Windows) to perform backups (especially for Microsoft SQL Server and Microsoft Exchange Server), it is important to verify that you have the proper licenses.

Arcserve UDP Agent (Windows) uses all the VSS writers during backup to ensure consistent backups. The only exceptions are Microsoft SQL Server, Microsoft Exchange, and Hyper-V writers which are only included when they are properly licensed.

- **Backup disk size**

A volume is skipped from a backup if it is on a disk larger than 2 TB and the compression option is disabled. However, there is no size limitation if compression is enabled (which is the default setting). As a result, if you want to back up source volumes larger than 2 TB, you must keep the compression option enabled.

Note: The 2 TB limitation only remains with the VHD format backup.

The minimum size for a block-level incremental (BLI) backup is 64K. For any file size less than 64K, Arcserve UDP Agent (Windows) will copy the entire file.

- **Verify that you are using a supported disk**

Different types of disks are supported as Arcserve UDP Agent (Windows) backup source disks and destination disks.

For more information, see [Disks Supported by Arcserve UDP Agent \(Windows\)](#).

- **Selecting your backup frequency and retention count**

If your scheduled backups are occurring more frequently than the amount of time it takes to generate a file system catalog for previous backup jobs, your recovery point retention count may be exceeded before the file system catalog being generated for the oldest session is completed. If this happens there may be a delay in the catalog generation for all pending recovery points. As a result, the retained recovery points could accumulate (beyond the specified maximum retention number), and you could see a lack of disk space at the destination, on the Status Summary. To avoid this problem, you can increase the schedule interval time for your incremental backups.

- **Backup destination folder manual operations**

Manual operations (such as copy, cut, paste, or drag-and-drop) for the backup destination folder are not successful if a job is active or a user is browsing recovery points using the Arcserve UDP Recovery Point View. Verify that no active jobs are running or browsing of recovery points (using Arcserve UDP Recovery Point View) is being performed before attempting any of these manual operations.

- **Proper drivers installed**

Verify that you have latest drivers or firmware installed for all the devices.

- **Verify that your machine is properly shutdown**

Even when backup jobs are not running, Arcserve UDP Agent (Windows) is constantly monitoring changes that are related to the operating system and data. Any detected changes are then compiled and saved in a list to include as an Incremental Backup after the next machine start-up. If your machine was not properly shut down and all of the changed information was not saved, Arcserve UDP Agent (Windows) may perform a more lengthy Verify Backup for the next backup, even if a Verify Backup was not scheduled.

- **Arcserve UDP Agent (Windows) in a Microsoft Hyper-V Environment**

Arcserve UDP Agent (Windows) provides both host-level and virtual machine (VM) level protection for Microsoft Hyper-V environments. For more information about situations that you can encounter and offer protection solutions using Arcserve UDP Agent (Windows), see [Arcserve UDP Agent \(Windows\) in a Microsoft Hyper-V Environment](#).

- **How running backup jobs on a Hyper-V server affects the tasks that can be performed**

When the Arcserve UDP Agent (Windows) backup job runs on a Hyper-V server, the status of the VMs is "Backing up" and the following tasks cannot be performed:

- Power on
- Power off
- Save
- Pause
- Reset
- Snapshot
- Move
- Rename
- Enable replication

▪ **How changing the machine host name can affect saving your settings**

When you enter a backup path, Arcserve UDP Agent (Windows) appends the host name to that path to use as the destination and this host name is also displayed in the settings dialog. When the name of the machine is changed, you must also change the destination path (backup, file copy, copy recovery point) by removing the old host name from the path before you attempt to save the settings.

For example: If your host name is "Host_A" and your backup destination is X:\ and you change your host name to "Host_B", any changes that are made to your backup settings are not saved unless you first change the backup destination from x:\Host_A to x:\ again.

If you do not change the backup destination host name and attempt to save the settings, Arcserve UDP Agent (Windows) thinks the backup destination "x:\Host_A" is already in use by Host_A and that Host_A is another machine and does not accept any changes to the settings.

▪ **How changing your backup destination can affect saved recovery points**

When you continue performing Incremental Backups to the changed destination and the specified number of saved recovery points is reached, Arcserve UDP Agent (Windows) merges the earliest backup sessions to the first destination to maintain the specified number of recovery points. As this merging process repeats, the number of recovery points that are saved to the first destination decreases, and at the same time the number of recovery points for the changed destination increases. Eventually, there are no recovery points for the first destination and all of the sessions are merged to the changed destination.

▪ **How changing backup destinations can affect continued backups**

If you configure and perform a Full Backup (and maybe some Incremental Backups) to a destination and then you decide to change your backups to a different destination, you can reconfigure your backup settings and continue performing Incremental Backups to the new destination without any problems.

If you later decide to change your backup destination again, you can simply reconfigure your backup settings again and continue performing Incremental Backups to the new destination without any problems.

For example:

- If you have a machine and it is configured to back up to Folder A on your local or remote volume, and after you perform a Full Backup and some Incremental backups, your destination is getting full and you want to change to a different destination (Folder B). You can reconfigure the backup settings to the Folder B destination and Arcserve UDP Agent (Windows) continues performing Incremental backups to that new destination. As a result, you have your Full Backup and some Incremental Backups on the original Folder A destination, and you have some Incremental Backups on the new Folder B destination.
- If after performing some Incremental Backups to Folder B, you decide to reconfigure to another new destination (Folder C), Arcserve UDP Agent (Windows) will continue performing Incremental Backups to the Folder C destination because the link to the original Full Backup location (Folder A) has been maintained.

If you configure and perform a Full Backup (and maybe some Incremental Backups) to a destination and then you decide to change your backups to a different destination, you can copy or move the contents from the original destination to the new destination, and then reconfigure your backup settings and continue performing Incremental Backups to the new destination without any problems.

However, if you have Full Backups in one location and Incremental Backups in a second location and then move the contents from the second location to a third location and attempt to continue performing Incremental Backups, then these backups fail because the link to the first location has been lost.

For example:

- If you have a machine and it is configured to back up to Folder A on your local or remote volume, and after you perform a Full Backup and some Incremental Backups, your destination is getting full and you want to change to a different destination (Folder B). You can move the contents of Folder A to Folder B and reconfigure the backup settings to the new Folder B destination. The Arcserve UDP Agent (Windows) continues performing Incremental Backups to that new

Folder B destination. As a result, you have your Full Backup and Incremental Backups all on the new Folder B destination.

- However, if your first destination is in Folder A (which now contains a Full Backup and some Incremental Backups) and you change the destination to Folder B using Arcserve UDP Agent (Windows) backup settings and continue to perform Incremental Backups, in this scenario, Folder B now only contains Incremental Backups. Then, if you move the contents from Folder B to another new destination in Folder C (moving only the Incremental Backups from Folder B without a Full Backup included), in this scenario, if you continue to perform Incremental Backups to Folder C, these Incremental Backups fails because the link to the original Full Backup location (Folder A) has been lost.

- **How your Retention Settings can affect your merge performance**

If you configure the backup format to Advanced, the merge performance is improved significantly.

- **How volume defragmentation can affect continued backups**

Volume defragmentation by Windows native tool affects the size of the block-level backups because Arcserve UDP Agent (Windows) will continue to incrementally back up all changed blocks. This means that blocks that shifted during defragmentation will also be included in the backup, even if no data has changed in the files. As a result, the backup size may increase. This is expected behavior. If you do not want the increased backup size and the added backup time is a problem, you can exclude volumes from defragmentation or stop any schedules for defragmentation.

- **How to configure backups of replicated volumes**

If you are backing up volumes that were replicated using Arcserve Replication and High Availability, you should verify that the spool has been created on a separate volume and configure your backup settings to exclude the spool volume. This helps to avoid the backing up of unnecessary temp spool data.

- **Restrictions for a Microsoft SQL Server Backup**

Due to Microsoft SQL Server VSS writer restrictions, some Microsoft SQL Server databases with a special status are automatically skipped and not backed up.

The Microsoft SQL server database includes:

- Database with 'Restoring' status. This status indicates that the database may be the log shipping secondary database, the mirror database, or the database waiting for more backed-up data to be restored.
- Database with 'Offline' status. This status indicates that the database is not available for general use.

- If your database is configured in one volume and the logs are configured in another volume and you select only one volume to back up, the Microsoft SQL application backup is skipped for that particular database.
- If you install Microsoft SQL Server after Arcserve UDP Agent (Windows) has been installed, and no backup has been performed yet, the Microsoft SQL Server may not be detected. As a result if you unselect a volume that has that application installed you may not get a warning notification that you are missing that application from the backup. This condition will be automatically remedied after you stop and start the Arcserve UDP Agent Services or perform the next backup.

▪ **Restrictions for a Microsoft Exchange Server Backup**

- If your database is configured in one volume and the logs are configured in another volume and you select only one volume to back up, the Microsoft Exchange application backup is skipped for that particular database.
- Any database in a dismounted state is skipped from Microsoft Exchange application backup.
- If you install Microsoft Exchange after Arcserve UDP Agent (Windows) has been installed, and no backup has been performed yet, the Microsoft Exchange server may not be detected. As a result if you unselect a volume that has that application installed you may not get a warning notification that you are missing that application from the backup. This condition will be automatically remedied after you stop and start the D2D services or perform the next backup.

▪ **Restrictions for VSS Writers**

Arcserve UDP Agent (Windows) uses all the VSS writers during backup to ensure consistent backups. The only exceptions are Microsoft SQL Server, Microsoft Exchange, and Hyper-V writers which are only included when they are properly licensed.

▪ **VHD Restrictions for Compression and Encryption**

If both compression and encryption are disabled, then Arcserve UDP Agent (Windows) can only back up the files in .VHD format. The Arcserve UDP Agent (Windows) cannot back up the files to .VHDX format.

▪ **Active Directory Backup Prerequisites**

An Active Directory restore requires an agent-based backup.

▪ **Oracle Backup Prerequisites**

For more information, see the following topic:

[Review the prerequisites to back up an Oracle database.](#)

▪ **Microsoft Clustered Nodes and Shared Disks Backup Prerequisites**

For more information, see the following topic:

[Review the Prerequisites to Back Up Microsoft Clustered Nodes and Shared Disks.](#)

- **How Arcserve UDP Agent (Windows) and the Backup Process Works**

(Optional) Understand how the restore process works. For more information, see the following topics:

- [How Arcserve UDP Agent \(Windows\) Works](#)
- [How the Backup Process Works](#)
- [How Block-Level Incremental Backups Work](#)
- [How Infinite Incremental Backups Work](#)
- [How Verify Backups Work](#)

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

Disks Supported by Arcserve UDP Agent (Windows)

Different types of disks are supported for Arcserve UDP Agent (Windows) backup source and destination disks. The following matrix lists the types of disks that are supported for each function.

Disk (Volume) Type	As Backup or File Copy Source	As Backup Destination	BMR Support	
			Data Volume	System and Boot Volume
System Reserved Partition	Yes*2	Not applicable	Not applicable	Yes
Mounted Volume (No drive letter/NTFS formatted)	Yes	Yes	Yes	Yes
RAW Volume (No drive letter/Not formatted)	No	No	No	No
VHD and VHDX Mounted Volume	No	Yes*5	No	No
GPT Disk:				
GPT (GUID Partition Table) Data Disk	Yes	Yes	Yes	Not applicable
GPT (GUID Partition Table) Boot Disk	Yes – R16 Update 5 or higher required	Yes – Not recommended to put Arcserve UDP Agent (Windows) backups on boot disk	Yes	Yes
Dynamic Disk:				
• No RAID	Yes	Yes	Yes*6	Yes*3, *4
• Software RAID (RAID-0 (Stripe))	Yes	Yes	Yes*6	Not applicable
• Software RAID (RAID-1 (Mirrored))	Yes	Yes	Yes*6	No
• Software RAID-5	No	Yes	No	Not applicable
Hardware RAID (include Embedded RAID)	Yes	Yes	Yes	Yes*4

File System:				
• FAT/FAT32	No	Yes*1	No	No
• NTFS	Yes	Yes	Yes	Yes
• REFS	Backup: Yes File Copy: No	Yes	Yes	Not applicable
• Deduplication-enabled NTFS	Backup: Yes File Copy: No	Yes	Yes	Not applicable
Shared Volume:				
Windows Shared Volume	No	Yes	No	No
Linux Shared Volume (samba shared)	No	Yes	No	No
Device Type:				
• Removable Disk (Ex. Memory Stick, RDX)	No	Yes	No	No

Notes:

- Any non-removable disk that appears as a local drive to a server protected by Arcserve UDP Agent (Windows) will also be protected. This includes any Fibre Channel (FC) attached Storage Area Network (SAN) disks, or iSCSI disks. For iSCSI disks, Arcserve UDP Agent (Windows) can protect system and data disks; however, iSCSI system disks are not supported for BMR. Therefore, you should only use iSCSI for data disks.
- Supported file copy destinations include Amazon S3, Fujitsu Cloud (Windows Azure), Eucalyptus-Walrus, and NTFS local or network volumes (could be on iSCSI).
- MSCS Shared Volume and CSV are not supported.

*1 FAT/FAT32 cannot hold a single file larger than 4 GB. If after compression the Arcserve UDP Agent (Windows) file is larger than 4 GB (because the source is very large), the backup will fail.

* 2 For Agent-based backup job, Arcserve UDP supports backing up UEFI system boot partition (FAT32 volume), but does not support backing up other FAT32 data volume.

*3 Arcserve UDP Agent (Windows) supports BIOS and UEFI systems.

*4 A spanned volume cannot be used as a boot volume.

*5 The VHD mounted volume used as the backup destination should not reside on a volume which is selected as the backup source.

*6 When your system has multiple dynamic disks, BMR can fail. But, as long as the system volume is on a basic disk, the system should be able to boot. You can perform a restore of dynamic disks after the system has booted, by using the How to Restore Files/Folders procedure.

Arcserve UDP Agent (Windows) in a Microsoft Hyper-V Environment

Arcserve UDP Agent (Windows) provides both host-level and virtual machine (VM) level protection for Microsoft Hyper-V environments. The following scenarios describe situations that you can encounter and offer protection solutions using Arcserve UDP Agent (Windows).

Note: Verify that you apply the appropriate license for each of these scenarios.

Scenario 1 - You want to protect your Hyper-V host server:

1. Install Arcserve UDP Agent (Windows) on the Hyper-V Host server.
2. Verify that you back up the system and boot volume (if you are choosing volume filtering).
3. If the Hyper-V host server goes down, follow the standard Bare Metal Recovery procedure in Arcserve UDP Agent (Windows) to recover your Hyper-V host server.
4. If you want to restore selected files, use the standard Arcserve UDP Agent (Windows) restore procedure.

Scenario 2 - You want to protect your Hyper-V host server and the Virtual Machines which are hosted on that server:

1. Install Arcserve UDP Agent (Windows) on the Hyper-V host server.
2. Verify that you back up the entire machine to provide complete protection of both the host server and VMs.
3. Submit the backup jobs.
4. To restore a VM from an Arcserve UDP Agent (Windows) backup, there are two possible solutions:
 - a. **Restore the VM to original location:**
 - From the restore window in Arcserve UDP Agent (Windows), select the VM files (.vhd, .avhd, configuration files, and so on).
 - Select **Restore to original location** as the destination and select the **Overwrite existing files** option for resolving conflicts.
 - Submit the restore job.

Note: The recommendation is to turn off the VM before submitting the restore job because if the older file is active, it will be overwritten by Arcserve UDP Agent (Windows) only after you reboot the Hyper-V host server to complete the restore process.

- When the restore job has finished, open the Hyper-V Manager and start the VM.
 - If the VM is not yet registered in the Hyper-V Manager, create a VM. During the VM creation process, point the path of VM's configuration and .vhd file to the same path of destination location where the restore was performed.
- b. **Restore the VM to alternate location on the same Hyper-V Host server:**
- From the restore window in Arcserve UDP Agent (Windows), select the VM files (.vhd, .avhd, configuration files, and so on).
 - Select **Restore to alternate location** as the destination and provide a destination path.
 - Submit the restore job.
 - Open the Hyper-V Manager and create a VM when the restore job has finished. During VM creation process, point the path of VM's configuration and vhd file to the same path of destination location where the restore was performed.
 - Start the VM, when the VM is created.

Note: For more detailed information about restoring Hyper-V Virtual Machines, see the Microsoft Hyper-V documentation.

Scenario 3 - You want to protect your Hyper-V Virtual Machines:

To protect your Hyper-V Virtual Machines (VM) using Arcserve UDP Agent (Windows), there are two possible solutions:

- a. **Install Arcserve UDP Agent (Windows) on the Hyper-V host server**
- Using the Arcserve UDP Agent (Windows) Backup Settings, select the volume where the VM files (.vhd, .avhd, configuration files, and so on) are located.
 - Submit a backup job.
 - To restore a Hyper-V Virtual Machine from an Arcserve UDP Agent (Windows) backup, follow the steps for either of the restore solutions provided in Scenario 2.

- b. **Install Arcserve UDP Agent (Windows) inside the Windows Virtual Machine**

Follow the standard backup and restore procedure to protect the VM, the same as a physical machine.

Note: For scenarios 2 and 3a, if you attached/mounted an iSCSI LUN directly inside the VM, the data inside the LUN is not backed up using Arcserve UDP Agent

(Windows) Hyper-V host level backups. You can overcome this limitation by using the same approach as the Install Arcserve UDP Agent (Windows) inside the Windows Virtual Machine solution in scenario 3b.

Merge Job Guidelines

Review the following merge job guidelines:

- A merge job has the lowest priority. When a merge job is running, if any other job comes in, the merge job will be stopped. After that job completes, the merge will be resumed or restarted.
- You can manually stop or pause the merge job when it is running. If the merge job is manually stopped/paused, you must manually start or resume it from the Arcserve UDP Agent (Windows) home page. It will not be resumed/restarted automatically. As a result, the launching of all scheduled merge jobs will be suspended until you manually resume them.
- If the merge job is automatically stopped it will be automatically started when no other job is running.
- When a merge job is resumed, Arcserve UDP Agent (Windows) will know exactly where to start the process. If the merge job crashed or the machine was abruptly shut down, the job is resumed from the previous merge state.

Example 1: When a merge job is started and crashed at 20%, the next time when the job is restarted, it will start to merge sessions again from 0%.

Example 2: When a merge job is started and paused at 10%, the next time when the job is restarted, it will start to merge sessions from the 10% point. If it crashes at 20%, then the merge job will be restarted from the 10% point.

- When a merge job is resumed or restarted, if the list of sessions being merged is not changed since the time it was paused, the merge is resumed. This means it resumes and continues the merge from the point where it was paused.
- When a merge job is resumed or restarted, if the list of sessions being merged is changed since the time it was paused, the original merge is resumed without any added or modified sessions. This means the original merge resumes and continues the merge from the point where it was paused. When the original merge is completed, a new merge of the added or modified sessions will then be performed.

Example: The original merge job contains 4 backup sessions and is paused when it completes 90% of the merge. When the merge is resumed, Arcserve UDP Agent (Windows) will complete the remaining 10% of the original merge and then a new merge will be performed for the added or modified sessions.

- The Arcserve UDP Agent (Windows) home page Job Monitor indicates the merge job status. It shows the percentage complete and displays more details if needed. For more information, see [Job Monitor Panel](#) in the online help.

- The merge process must be able to keep the recovery point in a consistent state. You can restore a file from any visible session even if the session is partially merged. If there is a session merge not completed, the merge job will run in the background to merge the session.

Review the Prerequisites for Oracle Database

To back up an Oracle database with consistent data, ensure that the ARCHIVELOG mode is enabled to archive the Redo logs.

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

- a. Log into the Oracle server as an Oracle user with SYSDBA privileges.

- b. Enter the following command at the SQL*Plus prompt:

```
ARCHIVE LOG LIST;
```

Archive log settings for the current instance is displayed.

- c. Configure the following settings:

Database log mode: Archive Mode

Automatic archival: Enabled

- d. Start the ARCHIVELOG mode.

Note: If the ARCHIVELOG mode is not enabled, 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 is 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: Archive Mode

Automatic archival: Enabled

Archive destination: E:\app\oracle\oradata\\arch

Oldest online log sequence: 21

Current log sequence: 23

- d. 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.

- e. Arcserve UDP Agent (Windows) is installed and a plan is scheduled.

Ensure that you have selected the volumes that include all the Oracle data files, server parameter file, control files, archived redo logs, and online redo logs for the backup.

- f. 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](#) that provides the supported operating systems, databases, and browsers.

Note: The shared disks will be backed up along with the agent which owns the shares disks. If the shared disk is moved from Node A to Node B during a failover, then for the next backup job on Node B, the disk will be backed up as a full disk even though the job itself appears as an incremental. After another failover if the shared disk moves back to Node A, even then the disk will be backed up as a full disk even though the job itself appears as an incremental.

Configure or Modify Backup Settings

Before you perform your first backup, you must configure the backup settings which are applied to each backup job. These settings can be retained for future backup or they can be modified at any time from the Arcserve UDP Agent (Windows) home page.

The settings let you specify behaviors such as:

- Backup source and destination.
- Schedule standard or advanced settings for each type of backup.
- Advanced settings for your backup jobs.
- Any pre or post backup operations.

Note: For more information about these Backup Settings, see [How to Perform a Backup](#).

To manage the backup settings, click the **Settings** link on the Arcserve UDP Agent (Windows) home page to display the **Backup Settings** dialogs and these subordinate tab options:

Specify Protection Settings

Protection settings for the information that is going to be backed up ensures that the backup data is reliably protected (copied and saved) against any form of data loss.

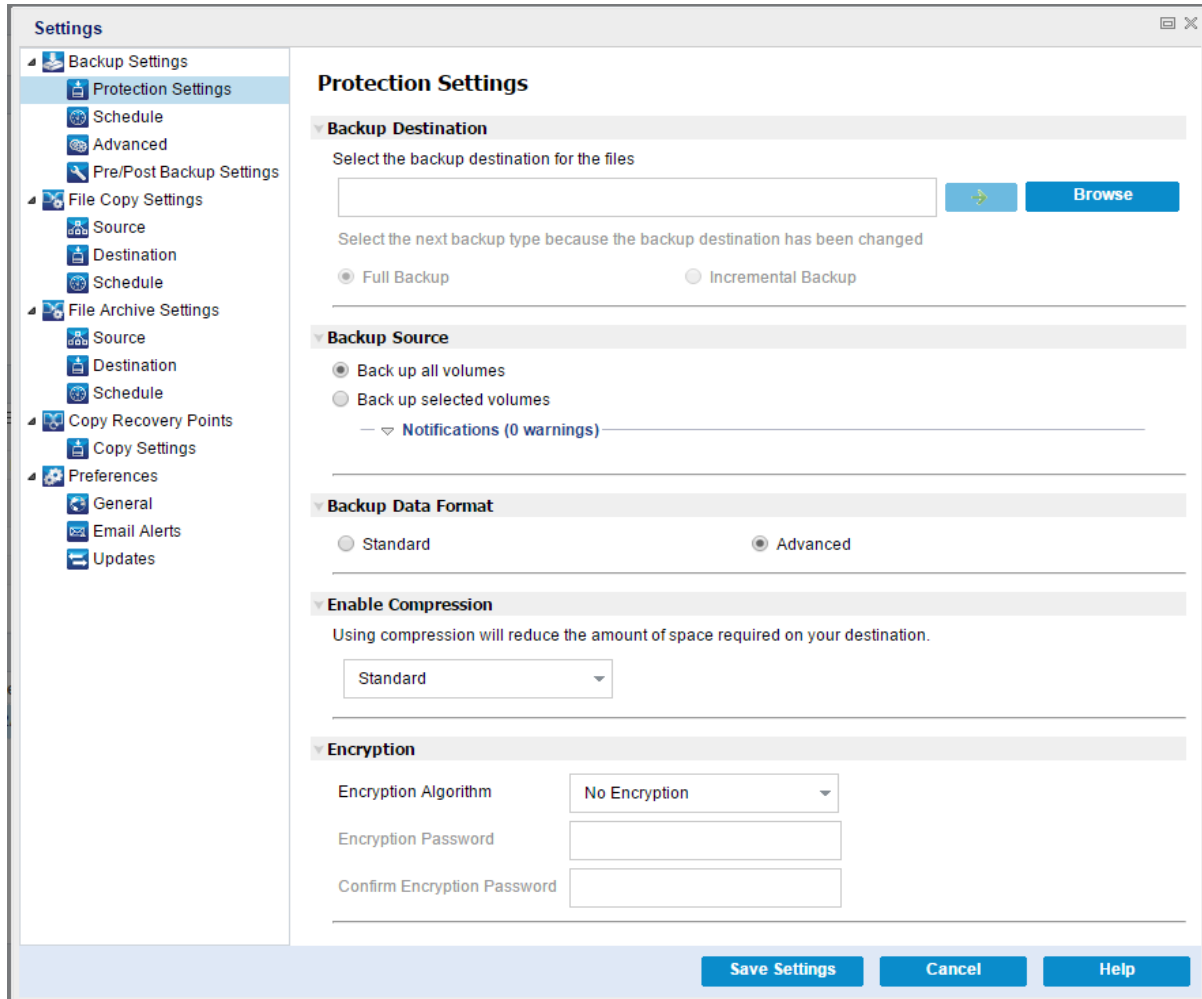
Specify the protection settings

1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Backup Settings** tab. When the **Backup Settings** dialog opens, select **Protection Settings**.

The **Protection Settings** dialog opens.

Notes:

- If the Arcserve UDP Agent (Windows) is being managed by Console, not all settings are available and will be displayed as read-only information.
- When agent is managed by console and not protected in a plan, all the settings are still available except the Preference > Updates panel.



2. Specify the **Backup Destination**.

♦ Use local disk or shared folder

You can specify a local path (volume or folder), or remote shared folder (or mapped drive) for the backup location, or you can browse to a backup location.

Click the green arrow icon button to verify the connection to the specified location.

- If you entered a local path for the destination, this button is disabled.
- If you enter a network path and click this button, you are prompted to provide the username and password.
- If you are already connected to this path successfully, and click the arrow button you can change the username and password you used to connect.

- If you do not click the arrow button, the destination path is verified. If needed, you are prompted for the username and password.
- a. If you want to back up to your local path (volume or folder), the specified backup destination cannot be the same location as your backup source. If you inadvertently include the source in your destination, the backup job ignores this portion of the source and it is not included in the backup.

Example: You attempt to back up your entire local machine consisting of Volumes C, D, and E and also specify Volume E as your destination. The Arcserve UDP Agent (Windows) only backs up Volumes C and D to Volume E. Data from Volume E is not included in the backup. If you want to back up all local volumes, specify a remote location for your destination.

Important! Verify that your specified destination volume does not contain system information. Or else it will not be protected (backed up) and your system will fail to recover after Bare Metal Recovery (BMR) if necessary.

Note: Dynamic disks are restored at disk-level only. If your data is backed up to a volume on a dynamic disk, you are not able to restore this dynamic disk during BMR.

- b. If you want to back up to a remote shared location, specify a location path or browse to the location. You also have to provide user credentials (Username and Password) to access the remote machine.
- c. If the backup destination has changed after the last backup was performed, select the backup type: Full Backup or Incremental Backup. These options are only enabled when you change your backup destination.

Default: Full Backup

Note: If the backup destination has changed and catalog jobs are pending, the catalog job first runs and completes on the old location before running on the new location.

Full Backup

The next backup that is performed is going to be a Full Backup. The new backup destination does not have any dependency on the old backup destination. If you continue with a full backup, the previous location is no longer needed for backups to continue. You can select to keep the old backup for any restores or if you do not want to perform

any restores from there you can delete it. The old backup will not affect future backups.

Incremental Backup

The next backup that is performed is going to be an Incremental Backup. The next incremental backup to the new destination is performed without copying all the backups from the previous destination. However, for this option, the new location is dependent on the previous location because the changes include only the incremental data (not the full backup data). Do not delete the data from the previous location. If you change the backup destination to another folder and attempt to perform an incremental backup, and the former backup destination does not exist, the backup fails.

Note: With the Full installation of Arcserve UDP, you can specify to use an Arcserve UDP Recovery Point Server as the backup location. If you do, the Protection Settings Backup Destination displays the Arcserve UDP Recovery Point Server Settings, including the Hostname, Username, Password, Port, Protocol, and the Plan Summary.

3. Specify the **Backup Source**.

You can back up the entire machine or selected volumes.

Back up the entire machine

Lets you back up the entire machine. All volumes on the machine are backed up.

Note: If you select the full machine backup option, Arcserve UDP Agent (Windows) automatically discovers all disks or volumes attached to the current machine and Arcserve UDP Agent (Windows) includes them in the backup.

Example: If a new disk is attached to the machine after the backup setting is configured, you do not need to change the backup settings and the data on the new disk will be protected automatically.

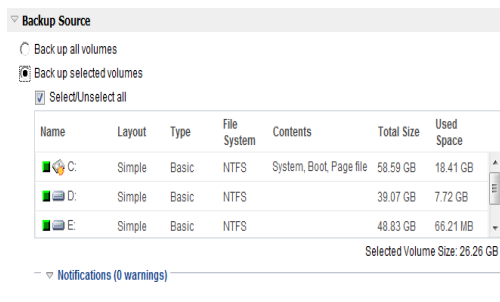
Select individual volumes to back up

This volume filtering capability lets you specify to back up only the selected volumes. You also have the option to Select or clear selection of all listed volumes.

Note: If some volumes are selected explicitly for backup, only the selected volumes are backed up. If a new disk or volume is attached to the machine, manually change the volume selection list to protect the data on the new disk or volume.

When you select this option, a listing of all available volumes display, with the corresponding volume information and notification messages.

Note: Computers that adhere to the Extensible Firmware Interface (EFI) use the EFI System Partition, which is a partition on a data storage device. The EFI System partition is critical for Bare Metal Recovery (BMR). Therefore, when you select boot volume "C" on a UEFI system, the EFI System Partition is selected automatically for the backup source for BMR and an information message is displayed.



Name

Specifies the name of the volume drive letter, mount point, volume GUID (Globally Unique Identifier) name.

Layout

Indicates the simple, spanned, mirror, striped, RAID5 (backup of a RAID 5 volume on Microsoft Dynamic Disks is not supported; but backup of hardware RAID is supported).

Type

Indicates the type, basic or dynamic.

File system

Lists the following file systems: NTFS, ReFS, FAT, FAT32 (backup of FAT, FAT32, and exFAT is not supported).

Contents

Indicates whether the application is (SQL/Exchange), System, Boot, Page file, Removable Device, VHD, 2-TB Disk.

Total size

Specifies the size or the capacity of the volume.

Used Space

Indicates the space, files or folders and volume data occupies.

The notification messages display for any of the following conditions:

- **Local volume related**

If the specified backup destination is on the local volume, a warning message displays notifying you that this volume is not backed up.

- **BMR related**

If system/boot volume is not selected for backup, a warning message displays notifying you that the backup is not usable for BMR.

If you select boot volume "C" on a UEFI system, the EFI system partition is selected automatically for the backup source for BMR and an information message is displayed.

- **Application related**

If the application data files are on a volume that is not selected for backup, the application name and database name display for reference.

4. Specify the **Backup Data Format**.

Standard

Standard Backup Data Format allows you to set the number of recovery points to retain or the number of recovery sets to retain and includes a basic repeat backup schedule. The Standard format is the legacy format used in releases of Arcserve D2D and Arcserve Central Applications.

Advanced

Advanced Backup Data Format allows you to set the number of recovery points to retain and includes advanced scheduling. The Advanced format is a new data storage format, dividing source disks into multiple logical segments. Compared to the Standard format, backup, restore, and merge job throughputs are greatly improved.

If the **Advanced Backup Data Format** is selected, advanced scheduling will be enabled. Advanced scheduling consists of the following:

- Week-based repeat backup schedule
- Week-based backup throttling schedule
- Week-based merge schedule
- Daily backup schedule
- Weekly backup schedule
- Monthly backup schedule

5. Specify the **Retention Setting** if you selected **Standard** as the **Backup Data Format**.

Note: If you selected **Advanced** as the **Backup Data Format**, the retention setting is specified on the **Advanced Schedule Settings** dialog.

You can set the retention setting based on the number of recovery points to retain (merges sessions) or based on the number of recovery sets to retain (deletes recovery sets and disables infinite incremental backups).

Default: Retain Recovery Points

Recovery Point

This is the recommended option. With this option selected, you can fully leverage the infinite incremental backup capabilities and save storage space.

Note: If you selected **Advanced** as the **Backup Data Format**, then you can only specify the number of recovery points to retain.

Recovery Set

This option is generally used for large storage environments. With this option selected, you can create and manage backup sets that help you manage your backup window time more efficiently when you are protecting a large amount of data. You can use this option when backup time is a priority over space constraints.

Note: Recovery sets are only available if you are backing up to a location that is not a data store. Recovery sets are not supported with RPS deduplication. They are also not available for Advanced format backup to non-RPS locations.

For more information about setting the Recovery Point and Recovery Set options, see [Specify Retention Settings](#).

6. Specify the type of **Compression**.

Specifies the type of compression that 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

No compression is performed. 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. Standard compression 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 or ZIP files), you can allocate additional storage space to handle such data. As a result, if you select any compression option and you have uncompressible data in your backup, it can 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 that is performed after this compression level change is automatically a Full Backup. After the Full Backup is performed, all future backups (Full, Incremental, or Verify) will be performed as scheduled.
- If your destination does not have sufficient free space, you can consider increasing the Compression setting of the backup.

7. Specify the **Encryption** settings.

- a. Select the type of encryption algorithm that is used for backups.

Data encryption is the translation of data into a form that is unintelligible without a deciphering mechanism. The Arcserve UDP Agent (Windows) data protection uses secure, AES (Advanced Encryption Standard) encryption algorithms to achieve maximum security and privacy of your specified data.

The available format options are No Encryption, AES-128, AES-192, and AES-256. (To disable encryption, select No Encryption).

- ◆ A full backup and all its related incremental and verify backups must use the same encryption algorithm.
- ◆ If the encryption algorithm for an incremental or verify backup is changed, a full backup must be performed. This means after changing encryption algorithm, the first backup will be full, despite the original backup type.

For example, if you change the algorithm format and you submit a customized incremental or verify backup manually, it is automatically converted to a full backup.

- 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 the 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 you submit a customized incremental or verify backup manually, it is automatically converted to a full backup.
- c. Arcserve UDP Agent (Windows) provides encryption password management so that you do not need to remember encryption passwords.
 - Password is also encrypted.
 - Password is remembered and not required (if you restore to the same machine).
 - Password is required if you restore to a different machine.
 - Password is not required if you are attempting to export a recovery point that contains encrypted data and the recovery point belongs to backups performed on the current machine.
 - Password is always required if you are attempting to recover encrypted data from an exported recovery point.
 - Password is not required to browse to an encrypted recovery point.
 - Password is required to perform a BMR.
- d. When encryption is enabled, the activity log is updated.
 - A message is recorded in the activity log to describe the selected encryption algorithm for every backup.
 - A message is recorded in the activity log to indicate why an incremental or verify backup was converted to a full backup (password change or algorithm change).

Note: Encryption settings do not have to remain the same for your backups. You can change these settings at any time, including after several backups of the same data.

8. Specify the **Throttle Backup**.

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. For a backup job, the Job Monitor on the Arcserve UDP Agent (Windows) home page displays the average Read and Write speed of the job in progress and the configured throttle speed limit.

Note: By default, the **Throttle Backup** option is not enabled and backup speed is not being controlled.

9. Calculate the **Estimated Backup Size**.

Displays the estimated usage of the destination volume.

Note: For more information about using these estimated backup calculations, see [Estimate Future Backup Space Requirements](#).

10. Click **Save Settings**.

Your backup protection settings are saved.

Specify Retention Settings

The retention setting for the **Standard Backup Data Format** can be based on the number of recovery points to retain (merges sessions) or based on the number of recovery sets to retain (deletes recovery sets and disables infinite incremental backups).

◆ Retain Recovery Points

Select this option to set your retention setting based on the number of recovery points to retain instead of on the number of recovery sets to retain.

Note: The recovery points to retain is set in the **Protection Backup Settings** if you select **Standard** as the **Backup Data Format**. The recovery points to retain is set in the **Advanced Schedule Settings** if you select **Advanced** as the **Backup Data Format**.

▼ **Backup Data Format**

Standard Advanced

▼ **Retention Setting**

Retain Recovery Points Retain Recovery Sets

Specify the number of recovery points to retain:

Run the merge job:

As soon as possible

Each day during the following time range

From : To :

Specify the number of recovery points to retain

Specifies the number of recovery points (full, incremental, and verify backup images) retained. When the number of recovery points present on the destination exceeds the specified limit, the earliest (oldest) incremental backups beyond the retention count are merged into the parent backup to generate a new baseline image consisting of the "parent plus oldest childs" blocks. If there are multiple sessions available for merge, the oldest child backups will be merged into the parent

backup in a single pass, if the backups are compressed. If the backups are not compressed, then only the oldest child backup will be merged into the parent backup and this cycle repeats for each subsequent child backup to be merged.

Specifying the number of recovery points to retain allows you to perform infinite incremental backups, while maintaining the same retention count. For more information, see [Merge Job Guidelines](#).

Note: If your destination does not have sufficient free space, you can consider reducing the number of saved recovery points.

Default Value: 31

Minimum Value: 1

Maximum Value: 1344

Note: The Arcserve UDP Agent (Windows) home page **Summary** section indicates how many recovery points are retained out of the number specified. For more information, see [Status Summary](#) in the online help.

Run the merge job:

As soon as possible

Select this option to run the merge job at any time.

Each day during the following time range

Select this option to run the merge job each day only within the specified time range. Setting a time range helps to avoid the merge job introducing too many I/O operations to the production server if the merge job runs for a long time.

Note: When setting the time range to run the merge job, ensure that you specify a time range that will allow the related backup jobs to complete prior to the start of the merge.

▪ **Retain Recovery Sets**

Select this option to set your retention setting based on the number of recovery sets to retain instead of on the number of recovery points to retain. With this setting you can disable infinite incremental backups, without merging any sessions. Using recovery sets helps reduce the amount of time it takes to complete merge jobs.

Note: The **Recovery Sets** option is available if you select **Standard** as your **Backup Data Format**. However, the **Recovery Sets** option is not

available if you select **Advanced** as your **Backup Data Format**.

The screenshot shows a configuration interface with two main sections:

- Backup Data Format:** Contains two radio buttons: Standard and Advanced.
- Retention Setting:** Contains two radio buttons: Retain Recovery Points and Retain Recovery Sets. Below these are two warning icons with text:
 - ! When you specify a number of recovery sets to retain, ensure that you have enough free space available for the specified number plus two additional full backups.
 - ! The retention setting has been changed. Use new backup destination to start backups with new retention setting.

A large text box contains the following configuration options:

- Specify the number of recovery sets to retain:
- Start a new recovery set on every:
 - Selected day of the week:
 - Selected day of the month:
- Start a new recovery set with:
 - First backup on the selected day
 - Last backup on the selected day

Specify the number of recovery sets to retain

Specifies the number of recovery sets retained. A recovery set is a series of backups, starting with a full backup, and then followed by a number of incremental, verify, or full backups.

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. A flag in the status column on the Arcserve UDP Agent (Windows) home page **Most Recent Events** section indicates that a full backup is the starting backup of a recovery set. 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.

When the specified limit is exceeded, the oldest recovery set is deleted (instead of merged).

Default Value: 2

Minimum Value: 1

Maximum Value: 100

Note: If you want to delete a recovery set to save backup storage space, reduce the number of retained sets and Arcserve UDP Agent (Windows) automatically deletes the oldest recovery set. Do not attempt to delete the recovery set manually.

Example 1 - Retain 1 Recovery Set:

- Specify the number of recovery sets to retain as 1.

Arcserve UDP Agent (Windows) always keeps two sets in order 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.

Arcserve UDP Agent (Windows) will delete 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 at the last backup on Friday.
- 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 PM on Friday is run, it will be converted to a full backup and marked as the starting backup of the recovery set.

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, since a given month may have 28, 29, 30, or 31 days, you can specify the last day of the month as the day to create the recovery set.

Start a new recovery set with:

First backup on the selected day

Indicates you want to start a new recovery set with the first scheduled backup on the specified day.

Last backup on the selected day

Indicates you want to start a new recovery set with the last scheduled backup on the specified day. If the last backup is selected to start the set and for any reason the last backup did not run, then the next scheduled backup will start the set by converting it to a full backup. If the next backup is run ad-hoc (for example an emergency situation requires a quick incremental backup), you can decide if you want to run a full backup to start the recovery set or run an incremental backup so that the next backup starts the recovery set.

Note: The last backup may not be the last backup of the day if you run an ad-hoc backup.

The Arcserve UDP Agent (Windows) home page **Summary** section indicates how many recovery sets are retained (or in progress) out of the number specified. Click the link under **Recovery Sets** to display the **Recovery Sets Details** dialog. This dialog contains detailed information about the contents of the recovery set. For more information about this dialog, see [Status Summary](#) in the online help.

Estimate Future Backup Space Requirements

Arcserve UDP Agent (Windows) provides you with this tool to calculate the estimated amount of available free space that you will need for backups. The calculations are based on your estimate of future data change and on the space that is occupied from previous backups.

Estimated Backup Size

The graph below shows the estimated usage of the destination volume. You can change the Space Saved After Compression or the Change Rate to see their effect on the estimated backup size.

■ Estimated backup 0.72 GB ■ Used 115.56 GB
■ Free 1362.28 GB



i Actual disk space used by current backups is: 1.70 GB.

Estimated Values

Space Saved After Compression	<input type="text" value="10%"/> ▼
Change Rate	<input type="text" value="10%"/> ▼
Space Saved After Windows Deduplication	<input type="text" value="0%"/> ▼

Estimated Backup Size

Total Source Size	282.57 MB
Compressed Full Backup Size	254.31 MB
Compressed Incremental Backup Size	483.19 MB
Estimated Total Backup Size	737.50 MB

To use this estimating tool

1. Select the backup source. This can be your entire machine or selected volumes within your machine.

The actual size of the selected backup source is displayed in the **Total Source Size** field.

2. Estimate the anticipated **Change Rate** for future backups.

Base this estimate upon past performance of how much your total backup size has changed for each subsequent incremental backup.

With the Estimated Values defined, Arcserve UDP Agent (Windows) calculates and displays the estimated backup size required based on the configuration of the

backup destination and the recovery points. The pie chart also displays the amount of used space and free space.

3. Estimate the **Space Saved After Compression** percentage value.

Estimated Values

You can use estimated values to calculate the approximate overall backup size that is based on the number of recovery points. Base this estimate upon past performance of your backups with different Compression settings applied. As you change this value, you will see the corresponding size impact for your backup sizes.

Note: If necessary, you can perform some Full Backups, each with a different Compression setting (No Compression, Standard Compression, and Maximum Compression) to establish past performance values and help you to better calculate the percent of space saving that each setting produces for your backup

- ♦ **Space Saved After Compression**

This value indicates how much disk space is saved after compression.

Example: If the data size of a volume is 1000 MB and after backup the compressed data size is 800 MB, then the Space Saved After Compression is estimated to be 200 MB (20%).

- ♦ **Change Rate**

This value indicates the typical data size of an incremental backup.

Example: If an incremental backup data size is 100 MB and the full backup data size is 1000 MB, the change rate is estimated to be 10%.

- ♦ **Space Saved After Windows Deduplication**

This value indicates how much disk space is saved after Windows deduplication.

If the backup destination directory is located on a volume where Windows deduplication is enabled, the estimated backup size may exceed the total capacity of the volume. The reason is that with deduplication enabled, only one copy of multiple same size data blocks is kept. This value helps to estimate the size while taking deduplication into consideration.

Example: If the total size of the source backed up is 100 GB and it has 20 GB of data that is redundant, then the space saved after deduplication will be 20 GB.

Estimated Backup Size

Displays the estimated values for **Total Source Size**, **Compressed Full Backup Size**, **Compressed Incremental Backup Size**, and **Estimated Total Backup Size**.

- ◆ The **Compressed Full Backup Size** field displays a calculated value that is based upon:
 - Size of the backup source
 - Specified compression percentage.
 - ◆ The **Compressed Incremental Backup Size** field displays a calculated value that is based upon:
 - Estimated Change Rate
 - Number of recovery points to be saved
 - Specified compression percentage
 - ◆ The **Estimated Total Backup Size** field will display the anticipated space that you will need for future backups and is based upon:
 - Amount of space that is required for one Full Backup plus
 - Amount of space that is required for the number of Incremental Backups needed to satisfy the specified number of saved recovery points.
4. From this **Estimated Total Backup Size** value, you should be able to determine if your backup destination has sufficient space to fit your backup.

If your destination does not have sufficient free space, you can consider the following corrective actions:

- ◆ Reduce the number of saved recovery points.
- ◆ Increase the available free space at the backup destination.
- ◆ Change the backup destination to a larger capacity.
- ◆ Reduce the size of the backup source (maybe eliminate unnecessary volumes from the backup).
- ◆ Increase the Compression setting of the backup.

Specify Schedule Settings

Arcserve UDP Agent (Windows) lets you specify the schedule for your backups. If you set the **Protection Settings Backup Data Format** to **Standard**, the **Standard Schedule** dialog opens, where you can specify the standard schedule settings. If you set the **Protection Settings Backup Data Format** to **Advanced**, the **Advanced Backup Schedule** dialog opens, where you can specify the advanced schedule settings.

Specify Standard Schedule Settings

Arcserve UDP Agent (Windows) lets you specify the schedule for your backups. If you set the **Backup Data Format** option to **Standard** in **Protection Settings**, the **Standard Schedule** dialog opens, where you can specify the standard schedule settings.

Follow these steps:


1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Backup Settings** tab. When the **Backup Settings** dialog opens, select **Schedule**.

The **Backup Settings StandardSchedule** dialog opens.

Notes:

- If the Arcserve UDP Agent (Windows) is being managed by Console, not all settings are available and will be displayed as read-only information.

- When agent is managed by console and not protected in a plan, all the settings are still available except the Preference > Updates panel.



2. Specify your backup schedule options.

Set start date and time

The start date and start time for your scheduled backups.

Note: When setting the interval between repeat backup jobs, ensure that you leave enough time to allow the previous job and any related merge jobs to complete before the next backup job starts. This amount of time can be estimated based on your own specific backup environment and history.

Incremental Backup

Determines the backup schedule for Incremental Backups.

As scheduled, Arcserve UDP Agent (Windows) 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.

The available options are **Repeat** and **Never**. If you select the **Repeat** option, you must also specify the elapsed time period (in minutes, hours, or days) between backup attempts. The minimum setting for Incremental backups is every 15 minutes.

By default the schedule for Incremental backups is to repeat every 1 day.

Full Backup

Determines the backup schedule for Full Backups.

As scheduled, Arcserve UDP Agent (Windows) performs a Full backup of all used blocks from the source machine. The available options are **Repeat** and **Never**. If you select the **Repeat** option, you must also specify the elapsed time period (in minutes, hours, or days) between backup attempts. The minimum setting for Full backups is every 15 minutes.

By default the schedule for Full backups is **Never** (no scheduled repeat).

Verify Backup

Determines the backup schedule for Verify Backups.

As scheduled, Arcserve UDP Agent (Windows) 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 Agent (Windows) 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.

The available options are **Repeat** and **Never**. If you select the **Repeat** option, you must also specify the elapsed time period (in minutes, hours, or days)

between backup attempts. The minimum setting for Verify backups is every 15 minutes.

By default the schedule for **Verify** backups is **Never** (no scheduled repeat).

Catalogs

File System Catalog

When this option is selected, enables generation of the file system catalog. If your browse time is too slow (especially if the Arcserve UDP Agent (Windows) destination is over a WAN) or if your restore by search time is too slow, this option helps reduce your wait time. This catalog job will run for each scheduled backup job after this option is selected.

If this option is not selected, the restores can be performed immediately after backup without having to wait for the catalog job to finish. By default, this option is not enabled.

Note: When you generate a File System catalog for each backup job, it results in an increased amount of disk storage needed to store the metadata files and catalog files and an increase in CPU usage. In addition, if the backup source contains a large amount of files, the process of generating a catalog could be a time consuming task.

Note: If you selected an ReFS volume as the backup source, you will not be able to generate a catalog and a warning message will be displayed to inform you of this condition.

3. Click **Save Settings**.

Your settings are saved.

Note: If at a given time there are more than one type of backup scheduled to be performed simultaneously, the type of backup that will be performed is based upon the following priorities:

- ◆ Priority 1 - Full backup
- ◆ Priority 2 - Verify backup
- ◆ Priority 3 - Incremental backup

For example, if you schedule all three types of backups to be performed at the same time, Arcserve UDP Agent (Windows) will perform the Full Backup. If there is no Full Backup scheduled, but you scheduled a Verify Backup and Incremental Backup to be performed at the same time, Arcserve UDP Agent (Windows) will perform the Verify Backup. A scheduled Incremental Backup is performed only if there is no conflict with any other type of backup.

Specify Advanced Schedule Settings

Arcserve UDP Agent (Windows) lets you specify the schedule for your backups. If you set the **Backup Data Format** option in **Protection Settings** to **Advanced**, the **Advanced Backup Schedule** dialog opens, where you can view your Repeat Schedule and Daily/Weekly/Monthly Settings.

Advanced Scheduling allows you to set the Repeat Schedule and the Daily Weekly Monthly Schedule. Advanced scheduling consists of the following:

- Week-based repeat backup schedule
- Week-based backup throttling schedule
- Week-based merge schedule
- Daily backup schedule
- Weekly backup schedule
- Monthly backup schedule

Follow these steps:


1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Backup Settings** tab. When the **Backup Settings** dialog opens, select **Schedule**.

The **Backup Settings Advanced Schedule** dialog opens.


Notes:

- If the Arcserve UDP Agent (Windows) is being managed by Console, not all settings are available and will be displayed as read-only information.
- If the Arcserve UDP Agent (Windows) is managed by console and not protected in a plan, all the settings are still available except the Preference >

Updates panel.



- (Optional) Click **Add** to add a backup schedule, backup throttle schedule, or a merge schedule.



For more information, see the following topics:

- ◆ [Add Backup Job Schedule.](#)
- ◆ [Add Backup Throttle Schedule.](#)
- ◆ [Add Merge Schedule.](#)

- Specify the **Start Date and Time.**

The start date and start time for your scheduled backups.

Note: When setting the interval between repeat backup jobs, ensure that you leave enough time to allow the previous job and any related merge jobs to complete before the next backup job starts. This amount of time can be estimated based on your own specific backup environment and history.

4. Specify the **Number of recovery points to retain**.

The number of recovery points to retain can be set for Daily, Weekly, Monthly, and Custom/Manual.

Note: The total retention count (Daily + Weekly + Monthly + Custom/Manual), the maximum limitation is 1440.

5. Specify **File System Catalog** and **Exchange Catalog** generation.

File System Catalog

When this option is selected, enables generation of the file system catalog. If your browse time is too slow (especially if the Arcserve UDP Agent (Windows) destination is over a WAN) or if your restore by search time is too slow, this option helps reduce your wait time. This catalog job will run for each scheduled backup job after this option is selected.

If this option is not selected, the restores can be performed immediately after backup without having to wait for the catalog job to finish. By default, this option is not enabled.

Note: When you generate a File System catalog for each backup job, it results in an increased amount of disk storage needed to store the metadata files and catalog files and an increase in CPU usage. In addition, if the backup source contains a large amount of files, the process of generating a catalog could be a time consuming task.

Note: If you selected an ReFS volume as the backup source, you will not be able to generate a catalog and a warning message will be displayed to inform you of this condition.

6. Click **Save Settings**.

Your settings are saved.

Add Backup Job Schedule

Follow these steps:

1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Backup Settings** tab. When the **Backup Settings** dialog opens, select **Schedule**.

The **Backup Settings Advanced Schedule** dialog opens.

2. From the **Backup Settings Advanced Schedule** dialog, click **Add** and then click **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 currently set to 'Custom'. Below this, the 'Backup Type' is set to 'Incremental'. The 'Start Time' is set to '8:00 AM'. A grid of checkboxes for the days of the week (Sunday through Saturday) is shown, with all checkboxes checked. Below the days, there is a 'Repeat' checkbox which is also checked. Underneath, there is a section for frequency: 'Every 3 Hours' and 'Until 6:00 PM'. At the bottom of the dialog, there are three buttons: 'Help', 'Save', and 'Cancel'.

3. From the drop-down list, select **Daily**, **Weekly**, **Monthly**, or **Custom**.
4. Enter the appropriate fields based on the schedule you selected:
 - ◆ To add a Daily Backup Schedule, see [Add Daily Backup Schedule](#).
 - ◆ To add a Weekly Backup Schedule, see [Add Weekly Backup Schedule](#).
 - ◆ To add a Monthly Backup Schedule, see [Add Monthly Backup Schedule](#).
 - ◆ To add a Custom/Manual Backup Schedule, see [Add Custom Backup Schedule](#).
5. Click **Save**.

Your settings are saved.

Notes:

- You can add up to 4 time windows for any week day.
- The time window cannot be set across multiple days. You can only configure the time window from 12:00 AM to 11:59 PM.
- For each time window, you can specify the time window and the repeat frequency.
- The default backup schedule is 1 daily backup at 10:00pm.

Add Backup Throttle Schedule

Follow these steps:

1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Backup Settings** tab. When the **Backup Settings** dialog opens, select **Schedule**.

The Backup Settings Advanced Schedule dialog opens.

2. From the **Backup Settings Advanced Schedule** dialog, click **Add** and then click **Add Throttle Schedule**.

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

3. Enter the following fields:

Throughput Limit

Specify the maximum speed (MB/min) at which the backup will be written.

You can throttle the backup speed to reduce CPU or network utilization.

However, by limiting the backup speed, it has an adverse effect on the backup window. As you lower the maximum backup speed it increases the amount of time to perform the backup. For a backup job, the Job Monitor on the home page will display the average Read and Write speed of the job in progress and the configured throttle speed limit.

Note: By default, the throttle backup speed option is not enabled and backup speed is not being controlled.

Start Time

Specify the time of the day to start applying the configured backup throttle settings.

Until

Specify the time of the day to stop applying the configured backup throttle settings.

4. Click **Save**

Your settings are saved.

Notes:

- You can add up to 4 time windows for any week day.
- The throttling value control the backup speed. For example, if you set 2 time windows, 1 from 8:00 AM to 6:00 PM, backup throughput limit is 1500 MB/minute, and 1 from 6:00 PM to 8:00 PM, backup throughput limit is 3000 MB/minute. If a backup job runs from 5:00 PM to 7:00 PM, the throughput will be 1500 MB/minute from 5:00 PM to 6:00 PM, and change to 3000 MB/minute from 6:00 PM to 7:00 PM.
- The time window cannot be set across multiple days. You can only configure the time window from 12:00 AM to 11:45 PM. If the throttle schedule ends at 11:45 PM, the schedule takes effect until the next day.
- Backup throttle schedule applies to repeat backup, as well as daily / weekly / monthly backups.

Add Merge Schedule

Follow these steps:

1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Backup Settings** tab. When the **Backup Settings** dialog opens, select **Schedule**.

The **Backup Settings Advanced Schedule** dialog opens.

2. From the **Backup Settings Advanced Schedule** dialog, click **Add** and then click **Add Merge Schedule**.

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

3. Enter the following fields:

Start Time

Specify the time of the day to start applying the configured backup throttle settings.

Until

Specify the time of the day to stop applying the configured backup throttle settings.

4. Click **Save**.

Your settings are saved.

Notes:

- You can add up to 2 time windows for any week day.
- If there is no merge schedule configured for any day, the merge job will be launched as soon as it is ready. If you configured any time window for the merge schedule, the merge job will only be launched within the time windows.

For example, if the merge schedule is 8:00 AM to 6:00 PM on Sunday, the merge job will only be launched during this time window.

- If the merge job launches within the configured time windows, and it runs to completion, despite the end time of the time windows. For example, if the merge time window is 8:00 AM to 6:00 PM on Sunday, a merge job started at 5:55 PM. It will continue to run after 6:00 PM to complete, even though the time is moving out of the defined time window.
- Merge schedule applies to repeat backup, as well as daily / weekly / monthly backups.
- When you configure a merge job schedule, the merge will only be triggered when the time is within the configured time window. If the merge is not within the configured time window, the merge will not run when you click the **Run a merge job manually now** link in the summary panel of the Arcserve UDP Agent (Windows) home page.

Schedule Considerations

The Arcserve UDP Agent (Windows) provides flexible settings to let you specify schedules for your backup. It consists of the following:

- Week-based repeat backup schedule
- Week-based backup throttling schedule
- Week-based merge schedule
- Daily backup schedule
- Weekly backup schedule
- Monthly backup schedule

However each backup, merge, or catalog job will consume system resources (CPU Usage, Memory Usage, IO Usage), occupy network bandwidth, and occupy disk space. Therefore, to help protect your system, consider the following:

What is the business processing time range of your server?

To avoid affecting your business processing, configure your system to run less jobs when the server is busy. For example, only configure to run backup jobs when the server is busy and leave merge jobs to run when the server is idle.

How about the data change frequency of your server?

Normally more frequent data change means more frequent backup is required. This is to reduce data lost to the minimum. When needed, you can recover the server to the last good known status.

How about your network bandwidth?

If your backup destination is configured to a network shared path, obviously the job occupies some of your network bandwidth when it is running. This might affect your business processing of this server. In case of this, specify a throttle schedule to limit the Arcserve UDP Agent (Windows) occupying network bandwidth.

How much disk storage is allocated for your backup destination?

More Full backups and more backups to retain means more disk storage is required. So when you configure how frequently to run a Full backup and how many backups to retain, consider the disk storage allocated for the backup destination.

How do you expect to use your backed up data?

Enable “File System Catalog” can shorten the browse time when you want to restore a file or a mailbox. But to generate catalogs, it also results in an

increased amount of disk storage needed to store the metadata files and catalog files and an increase in CPU usage. In addition, if the backup source contains a large amount of files, the process of generating a catalog could be a time consuming task. So whether to enable or disable catalogs is depending on how you would like to use the backed up data.

Based on the above considerations, the following is an example of using advanced scheduling to protect a build server, showing the situation and corresponding schedule settings:

- The build server is used to provide source code pre-compile service every working day. It's business process time slot is 9:00 AM – 7:00 PM of every work day (from Monday to Friday). During other times, it is idle.

Schedule Settings:

- Configure to run custom incremental backup from 9:00 AM to 7:00 PM, run merge job at night – 7:00 PM to 9:00 AM of next day.
- The pre-compile service is launched every 2 hours, and there are lots of data changes at that time.

Schedule Settings:

- Configure to run custom incremental backup every 2 hours.
- Every time to run pre-compile, the build server need to fetch source code from a remote source code repository server.

Schedule Settings:

- Limit backup throttle to 500 MB/Minute during 9:00 AM to 7:00 PM and no limitation during other time slots.
- Due to the poor disk storage, there is no requirement to retain a lot of recovery points. Only need to keep recovery points in one release cycle; 6 months is enough. But there is a requirement to keep the recovery point in the last 24 hours, so that once needed you can recover to the last good known status.

Schedule Settings:

- Specify to retain last 12 manual backups (the backups of the last 24 hours).
- Configure to run Daily Incremental backup at 9:00 PM of every day. And keep the last 7 Daily backups.
- Configure to run Weekly Full Backup at 11:00 PM of every Friday. And keep the last 4 Weekly backups.

- Configure to run Monthly Full Backup at 12:00 PM on last Saturday of month. And keep the last 6 monthly backups.

Finally, there are 6 monthly backups, 4 weekly backups, 7 daily backups and 12 most recent backups. There are enough choices to recover the build server to a good known status.

- For the build server, there is no requirement to quickly browse and restore files. Once needed, perform a BMR to restore the build server to the last good known status. That is enough.

Schedule Settings:

- Disable options to generate "File System Catalog".

Specify Advanced Settings

Arcserve UDP Agent (Windows) lets you specify the **Advanced Settings** for your backups.

Follow these steps:

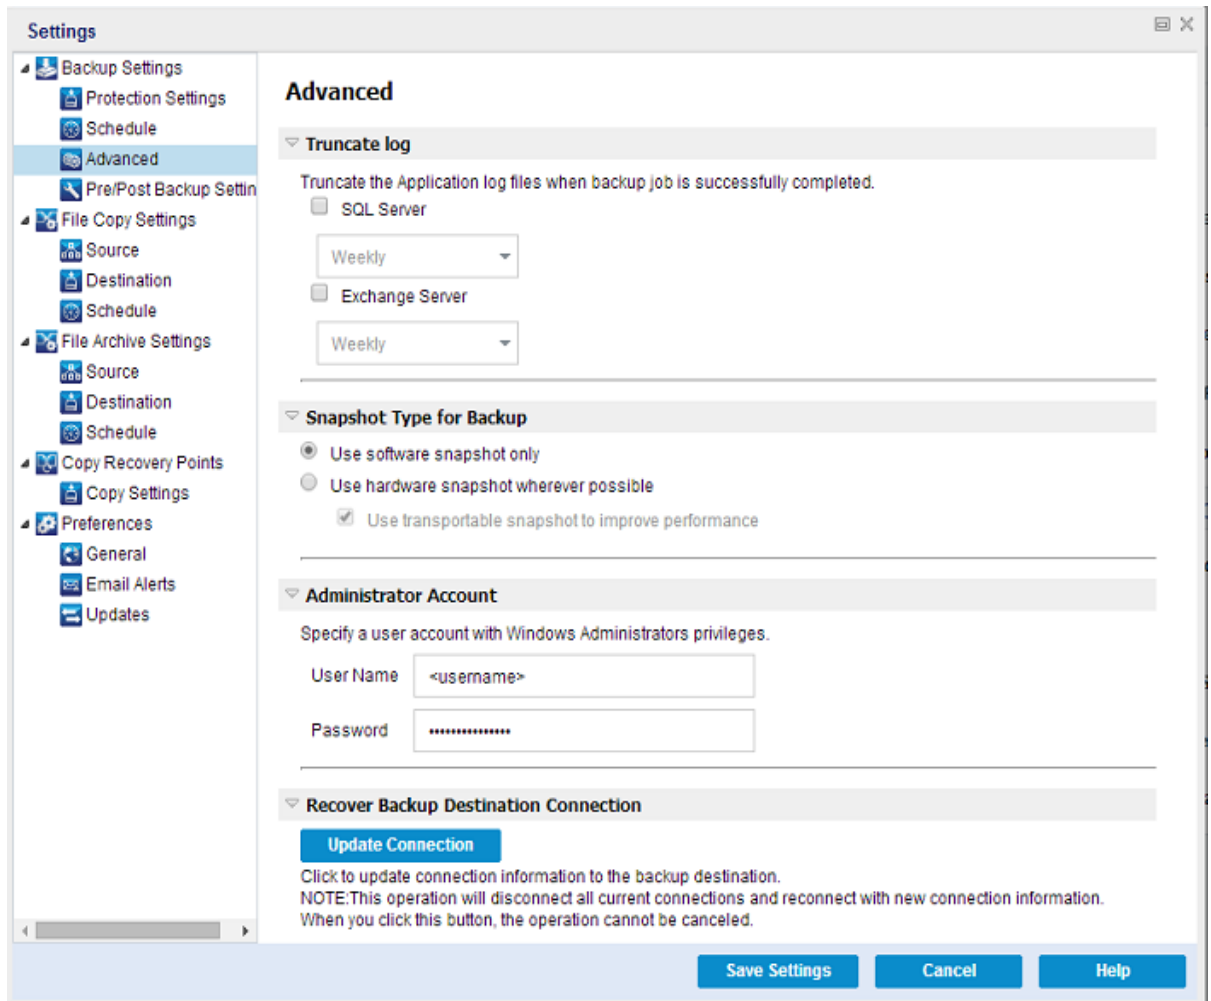
1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Backup Settings** tab. When the **Backup Settings** dialog opens, select **Advanced**.

The **Advanced** dialog opens.

Notes:

- If the Arcserve UDP Agent (Windows) is being managed by Console, not all settings are available and will be displayed as read-only information.
- If the Arcserve UDP Agent (Windows) is managed by console and not protected in a plan, all the settings are still available except the Preference >

Updates panel.



- Specify your advanced backup settings options.

Truncate Log

Truncates the accumulated transaction log files for the selected applications after the next successful backup.

Arcserve UDP Agent (Windows) backups consist of a snapshot image and the transaction log files that were created for it. At some point in time, the older (committed) transaction log files are no longer needed and can be purged to make space for new log files. The process of purging these log files is truncating the log. This option enables truncating of committed transaction log files, which conserves disk space.

The available options are SQL Server and Exchange Server. You can select either, both, or none of these applications. If you select any of these applications, you can also specify a scheduled time period (Daily, Weekly, or Monthly) for automatic log truncation:

Note: The transaction log files cannot be truncated without performing a successful backup.

- **Daily** - Specifies that each day after the backup completes successfully, the committed transaction logs will be purged immediately.
- **Weekly** - Specifies that after seven days, the committed transaction logs will be purged immediately after the backup completes successfully.
- **Monthly** - Specifies that after 30 days, the committed transaction logs will be purged immediately after the backup completes successfully.

If a backup job is already running at the same time the purging is scheduled to be performed, the purging operation is moved to the next scheduled job.

Example:

You scheduled an Incremental Backup to run automatically every day at 5:00 pm, and then started a Full Backup manually at 4:55 pm. You assume that the backup successfully finishes at 5:10 pm.

In this case, the Incremental Backup that is scheduled for 5:00 pm is skipped because the ad-hoc Full Backup is still in progress. Now the committed transaction log files are purged after the next successful backup job and be performed on the next day after the scheduled Incremental Backup completes successfully at 5:00 pm.

Snapshot Type for Backup

You can select the required option from software snapshot or hardware snapshot.

Use software snapshot only

Specifies that the backup type uses only the software snapshot. Arcserve UDP will not check for hardware snapshot. The software snapshot utilizes less resources on the virtual machines. You can use this option if the server has lower configurations and processing speed.

Use hardware snapshot wherever possible

Specifies that the backup type first checks for a hardware snapshot. If all the criteria are met, the backup type uses hardware snapshot.

Note: For more information on the hardware snapshot criteria, see the prerequisite.

Administrator Account

Specifies the User Name and Password with access rights to perform the backup. The Arcserve UDP Agent (Windows) verifies that the name and password are valid and the user belongs to an administrator group.

Important! If the Administrator Account credential information for the Arcserve UDP Agent (Windows) server is changed (User Name/Password), you must also reconfigure/update the Administrator Account information in this dialog.

Note: To specify a domain account, the format for the user name is a fully qualified domain user name in the form of "<domain name>\<user name>".

Recover Backup Destination Connection

Lets you update (resynchronize) the connection information to your backup destination.

You can use this option if you are performing periodic backups to a remote share computer and then you can change the access credentials (user name/password) for that remote computer. In this case, typically your next backup would fail because the access credentials configured at your local computer do not match the new credentials at the remote computer.

Note: When you click the **Update Connection** button and the resynchronize process begins, you cannot cancel it.

Before you click this **Update** button, perform the following tasks:

- a. Log in to the remote destination computer and use the following net session command to disconnect the connection between the local Arcserve UDP Agent (Windows) computer and the remote computer:

```
net session \\<computer name or IP address> /d
```

- b. Return to the Arcserve UDP Agent (Windows) computer, and click the **Update Connection** button.
- c. Enter new password for destination.

Arcserve UDP Agent (Windows) updates your configured credentials to match the new credential information at the remote share destination. A pop-up confirmation screen appears informing you that the credentials have been updated.

3. Click **Save Settings**.

Your advanced backup settings are saved.

Specify Pre/Post Backup Settings

Arcserve UDP Agent (Windows) lets you specify the **Pre/Post Backup Settings**.


Specify the Pre/Post Backup Settings

1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Backup Settings** tab. When the **Backup Settings** dialog opens, select **Pre/Post Backup**.

The **Pre/Post Backup Settings** dialog opens.

Notes:

- If the Arcserve UDP Agent (Windows) is being managed by Console, not all settings are available and will be displayed as read-only information.
- When agent is managed by console and not protected in a plan, all the settings are still available except the Preference > Updates panel.



2. Specify your pre/post backup setting options.

Actions

Runs script commands for actions to take before the start of the backup, after the snapshot image is captured, and/or upon the completion of the backup. You can also trigger the script command based upon specific exit codes and select the action to be taken (run job or fail job) when that exit code is returned.

- A "run job" action directs Arcserve UDP Agent (Windows) to continue to run the job if the specified exit code is returned.
- A "fail job" action directs Arcserve UDP Agent (Windows) to cancel the job if the specified exit code is returned.

3. Click **Save Settings**.

Your pre/post backup settings are saved.

Perform a Backup

Before you perform your first backup, specify the backup settings to be applied to and control all subsequent backup jobs. These settings are applied to each backup job, regardless of how you initiate the backup. For more information, see [Configure or Modify Backup Settings](#).

A backup job can be initiated either automatically (based upon your schedule settings) or manually (immediate ad-hoc backup).

Perform Backup Automatically (Scheduled)

Automatic backup jobs are the same as manual backup jobs, except they are triggered at pre-configured days and times. You can configure automatic backup jobs using the **Backup Schedule** dialog. For more information, see [Specify Schedule Settings](#).

The process for scheduling an automatic backup is as follows:

1. Based upon the configured time settings, Arcserve UDP Agent (Windows) triggers the launching of each type of scheduled backup job (Full, Incremental, and Verify).
2. Configuration settings specified in the **Backup Settings** dialogs are applied to the job.
3. If configured, an email notification is sent to the recipients informing them when the backup job is completed (or if a problem occurred that prevented the scheduled backup job from being completed).

Perform Backup Manually (Backup Now)

Backups are performed automatically and are controlled by the schedule settings. However, there can be times when you have to perform an ad-hoc backup (Full, Incremental, or Verify) immediately.

An ad-hoc backup is need-based, rather than being scheduled in advance as part of a backup plan. 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 ad-hoc backup without waiting for the next scheduled backup to occur.

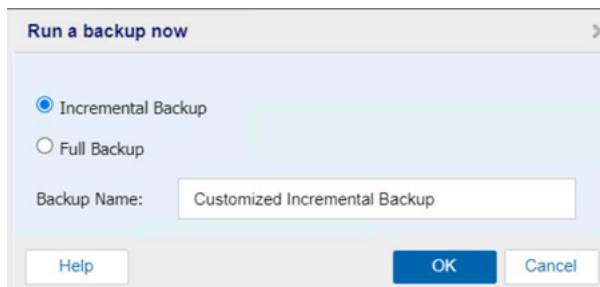
An ad-hoc backup also allows you to add a customized (unscheduled) recovery point so that you can roll back to this previous point in time if necessary. For example, if you install a patch or service pack and then discover that it adversely affects the performance of your machine, you may want to roll back to the ad-hoc backup session that does not include the patch or service pack.

Follow these steps:

1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), click **Backup Now**.

The Run a backup now dialog opens.

2. On the Run a backup now dialog, select the type of backup you want to perform, and then specify the backup name if required.



The available options are:

Incremental Backup

Initiates an Incremental Backup of your machine. An Incremental Backup backs up only those blocks that have changed after the previous backup.

The advantages of Incremental Backups are that it is a fast backup and produces a small backup image. This is the most optimal way to perform backups.

Full Backup

Initiates a Full Backup of your entire machine or the selected volumes.

Notes:

- If you add a new volume to the backup source, the newly added volume is fully backed up regardless of the overall backup method selected.
 - If no backup name is specified, by default, it is automatically named as Customized Full/Incremental Backup.
3. If necessary, specify a backup name and click **OK**. If no backup name is specified, it is automatically named Customized Full/Incremental Backup by default.
 4. Click **OK**.

The Progress Information dialog appears.

5. Wait until the backup job completes, and then click **OK**.

All configuration settings specified in the Backup Settings dialogs are applied to the job.

Notes:

- Only one job can be run at a time. If you manually attempt to launch a backup job when another job is currently running, an alert message informs you that another job is running and requests that you try again at a later time.
- If a custom (ad-hoc) backup job fails, no makeup job is created. A makeup job is only created for a failed scheduled job.

The manual backup is successfully performed.

Verify that the Backup is Successful

To verify that the backup process of the data to the specified destination is successful, perform one of the following procedures:

Follow these steps:

1. Navigate to the Arcserve UDP Agent (Windows) backup destination you specified.

A list of folders appears.

2. Verify that the size of the folder matches the size as displayed in the protection **Summary** list.

Note: The size of the folder should be equal to the sum of the Full backup, Incremental backups, and any Verify backups.

The Arcserve UDP Agent (Windows) backup process is successful.

Follow these steps:

1. From the Arcserve UDP Agent (Windows) home page UI, click **Restore** wizard.

The Restore dialog opens.

2. Click the **Browse Recovery Points** and verify the data that you have backed up is listed correctly.

The Arcserve UDP Agent (Windows) backup process is successful.

Follow these steps:

1. From the Arcserve UDP Agent (Windows) home page UI, click **Mount Recovery Point** wizard.

The Mount Recovery Point dialog opens.

2. Verify that the data in the mount backup volume is mounted correctly.

The Arcserve UDP Agent (Windows) backup process is successful.

How Arcserve UDP Agent (Windows) Works

Arcserve UDP Agent (Windows) lets you perform frequent and periodic block level backups of your full machine. These backups can be stored on either an internal drive, an external drive, on a remote network share, or a data store on a Recovery Point Server (RPS), depending on the type of installation (Arcserve Unified Data Protection - Full or Arcserve Unified Data Protection - Agent). If the backup destination volume is also selected as the backup source volume, a never ending backup is not executed. During the backup, the backup destination volume is excluded and an entry is added to the Activity log. The Arcserve UDP Agent (Windows) provides the capability to perform Full, Incremental, or Verify type backups.

Arcserve Unified Data Protection - Full:

Available backup destinations include: internal drive, external drive, remote network share, or a data store on a Recovery Point Server (RPS). When you create a Plan from the Arcserve UDP server, you can select Data Store on Recovery Point Server as the destination and then deploy the plan to the agent node.

Arcserve Unified Data Protection - Agent:

Available backup destinations include: internal drive, external drive, or a remote network share.

Arcserve UDP Agent (Windows) also provides various methods to identify and locate the backed up data and allow you to restore it if necessary. Regardless of which restore method you select, Arcserve UDP Agent (Windows) lets you quickly identify the data you need and retrieve it from the appropriate backup location.

How the Backup Process Works

Arcserve UDP Agent (Windows) lets you perform frequent and periodic block level backups of your entire machine. These backups can be stored on either an internal drive, an external drive, on a remote network share, or a data store on a Recovery Point Server (RPS), depending on the type of installation (Arcserve Unified Data Protection - Full or Arcserve Unified Data Protection - Agent). The Arcserve UDP Agent (Windows) provides the capability to perform Full, Incremental, or Verify type backups.

The basic process for how Arcserve UDP Agent (Windows) performs a backup is simple. When you initiate a backup (either as scheduled or manually launched), Arcserve UDP Agent (Windows) captures a full VSS snapshot, and then backs up only those blocks that have been changed since the previous successful backup. (If it is a Full backup, all blocks are backed up). This block-level incremental backup process significantly reduces the amount of backup data. For example, if you have a large file and only change a small portion of this file, Arcserve UDP Agent (Windows) backs up only the changed portion to the incremental backup and not back up the entire file.

During this block-level incremental backup process, Arcserve UDP Agent (Windows) not only captures the data, but also creates a catalog containing all information related to the operating system, installed applications (Microsoft SQL and Microsoft Exchange only), configuration settings, necessary drivers, and so on. If necessary, you can then restore this backed-up image to recover your data or your entire machine. If the backup destination volume is also selected as the backup source volume, a never ending backup is not executed. During the backup, the backup destination volume is excluded and an entry is added to the Activity log.

Note: You can submit a faster backup job (catalog-less backup), since a catalog is not required after a backup job is complete. The backup settings option "Generate File System catalog for faster search after each backup" by default is unchecked, indicating it will perform a faster backup.

The details of what is being backed up, how it is being backed up, when it is being backed up, and so on, are controlled by the various backup configuration settings that you specify. These settings are applied to each backup job, regardless of how you initiate the backup (automatically or manually).


Arcserve Unified Data Protection - Full:

With this type of installation, configure your backup by creating a plan. Available backup destinations include: internal drive, external drive, remote network share, or a data store on a Recovery Point Server (RPS). When you create a Plan from the

Arcserve UDP server, you can select Data Store on Recovery Point Server as the destination and then deploy the plan to the agent node.

Arcserve Unified Data Protection - Agent:

With this type of installation, configure your backup by specifying the backup settings. Available backup destinations include: internal drive, external drive, or a remote network share.



How Block-Level Incremental Backups Work

When you start a backup, the specified volume is divided into a number of subordinate data blocks that are then backed up. The initial backup is considered the "parent backup" and will be a Full Backup of the entire volume to establish the baseline blocks to be monitored. Before performing the backup, a VSS snapshot is created, then an internal monitoring driver checks each block to detect any changes. As scheduled, Arcserve UDP Agent (Windows) will then incrementally back up only those blocks that have changed since the previous backup. You can schedule the subsequent block-level incremental backups ("child backups") as frequently as every 15 minutes to always provide accurate, up-to-date backup images.

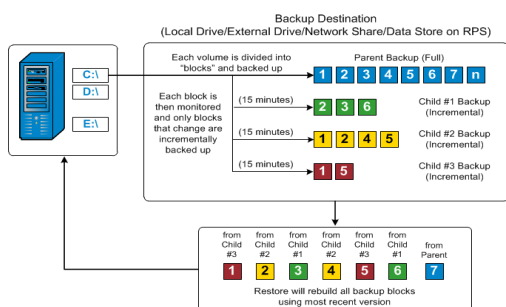
If you need to restore the volume information, the most recent backed up version of each block is located and the entire volume is rebuilt using these current blocks.

Arcserve Unified Data Protection - Full:

Available backup destinations include: internal drive, external drive, remote network share, or a data store on a Recovery Point Server (RPS). When you create a Plan from the Arcserve UDP server, you can select Data Store on Recovery Point Server as the destination and then deploy the plan to the agent node.

Arcserve Unified Data Protection - Agent:

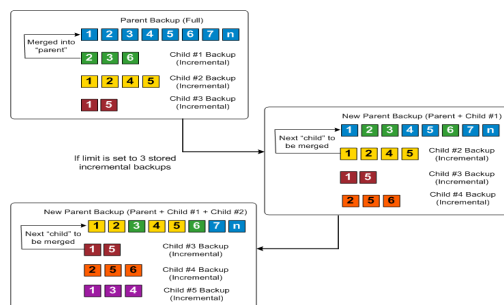
Available backup destinations include: internal drive, external drive, or a remote network share.



How Infinite Incremental Backups Work

If left alone, the incremental snapshots (backups) would continue, as often as 96 times each day (every 15 minutes). These periodic snapshots will accumulate a large chain of backed up blocks to be monitored each time a new backup is performed, and require added space to store these ever-growing backup images. To minimize this potential problem, Arcserve UDP Agent (Windows) utilizes the Infinite Incremental Backup process, which intelligently creates incremental snapshot backups forever (after the initial full backup) and uses less storage space, performs faster backups, and puts less load on your production servers. Infinite Incremental Backups allow you to set a limit for the number of incremental child backups to be stored. When the **Backup Data Format** is **Standard**, configure the **Recovery Points** option from the **Protection Settings** tab on the **Backup Settings** dialog. When the **Backup Data Format** is **Advanced** (default), configure the **Recovery Points** option from the **Schedule** tab on the **Backup Settings** dialog.


When the specified limit is exceeded, the earliest (oldest) incremental child backup is merged into the parent backup to create a new baseline image consisting of the "parent plus oldest child" blocks (unchanged blocks will remain the same). This cycle of merging the oldest child backup into the parent backup repeats for each subsequent backup, allowing you to perform Infinite Incremental (I2) snapshot backups while maintaining the same number of stored (and monitored) backup images.



How Verify Backups Work

Every so often (as scheduled or when manually initiated), Arcserve UDP Agent (Windows) can perform a Verify (resynchronization) type backup to provide a confidence check of the stored backup image and resynchronize that image if necessary. A Verify type 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 Agent (Windows) refreshes (resynchronizes) the backup of the block that does not match.

A Verify backup can also be used to get the same guarantee as a full backup without taking the space of full backup. The advantage of a Verify backup is that it is small when compared to full backup because only the changed blocks (blocks that do not match the last backup) are backed up. However, a Verify backup is also slower than an Incremental backup because Arcserve UDP Agent (Windows) has to compare all of source disk blocks with the blocks of the last backup.



How Recovery Sets Work

A Recovery Set is a storage setting where a group of recovery points (backup sessions) are backed-up over a specified period and then stored together as one compiled set. A Recovery Set includes a series of backups, always starting with a Full Backup, and then followed by a number of subsequent Incremental, Verify, or Full Backups. By using Recovery Sets (instead of Recovery Points), you disable infinite incremental backups and discontinue merging of backup sessions, thereby completely eliminating the time-consuming merge process.

Recovery Sets are generally used for large storage environments and helps you to manage your backup window time more efficiently when protecting large amounts of data. Recovery Sets are used when the backup time is more important than storage space constraints.

A Full Backup is required to start a Recovery Set. Therefore, the backup session that starts a Recovery 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 initial Full Backup is completed, all subsequent backups (regardless if which type of backup is performed) will be saved within the Recovery Set until the next new Recovery Set is launched (manually or automatically as scheduled).

You can configure the number of Recovery Sets to retain. When the number of Recovery Sets retained exceeds the specified retention count, the merge job deletes the oldest recovery set. A recovery set is considered complete only when the starting Full Backup for the next Recovery Set is completed. For example, if you specified to retain two Recovery Sets, Arcserve UDP Agent (Windows) deletes the first Recovery Set only after the Full Backup for the fourth Recovery Set is completed. This ensures that when the first backup is deleted, you already have two Recovery Sets (Recovery Set 2 and Recovery Set 3) retained on disk.

Notes:

- After reaching the retention count, the merge job gets triggered and the oldest Recovery Set gets deleted.
- If you want to delete a recovery set to save backup storage space, reduce the number of retained sets and Arcserve UDP Agent (Windows) automatically deletes the oldest recovery set. Do not attempt to delete the recovery set manually.

A flag in the status column on the Arcserve UDP Agent (Windows) home page **Most Recent Events** section indicates that a full backup is the starting backup of a recovery set. After the recovery set setting is changed (for example, changing the recov-

ery 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: Recovery sets are only available when using Arcserve UDP Agent (Windows) and you set the **Backup Data Format** to **Standard**. Recovery sets are not available if you set the **Backup Data Format** to **Advanced**. This is because merge jobs are very fast and efficient when using the **Advanced Backup Data Format**, therefore eliminating the need for recovery sets.

Default: 2

Minimum: 1

Maximum: 100

Example 1 - Retain 1 Recovery Set:

- Specify the number of recovery sets to retain as 1.

Arcserve UDP Agent (Windows) deletes the first recovery set when the third recovery set Full backup is completed.

Note: Even if you choose to retain only one recovery set, you need space for at least two full backups.

Example 2 - Retain 2 Recovery Sets:

- Specify the number of recovery sets to retain as 2.

Arcserve UDP Agent (Windows) deletes the first recovery set when the fourth recovery set full backup is completed. This ensures that when the first backup is deleted and the fourth recovery set Full backup is completed, you still have two recovery sets (recovery set 2 and recovery set 3) available on disk.

Example 3 - Retain 3 Recovery Sets:


- Specify the number of recovery sets to retain as 3.

Arcserve UDP Agent (Windows) deletes the first recovery set when the fifth recovery set Full backup is completed.

- The backup start time is 6:00 AM, August 20, 2012.
- An incremental backup runs every 12 hours.
- A new recovery set starts at the last backup on Friday.
- You want to retain 3 recovery sets.

With the above configuration, an incremental backup runs 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 PM on Friday runs, it will be converted to a full backup and marked as the starting backup of the recovery set.



Troubleshooting Backup 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.

Note: If you convert a basic disk to a dynamic disk, and then you restart the server. When you perform an Incremental backup, the backup will be as large as a full backup for that disk. The reason is, when you change the disk from basic to dynamic, Arcserve UDP considers the dynamic disks as a new disk and performs a full backup for the first time. From the next backup, the backup job will be an Incremental backup.

- [SQL Server backup failed due to Out of memory error](#)
- [Backup sessions for Arcserve UDP Agent \(Windows\) do not include any Microsoft SQL database information](#)
- [Catalog Job fails when backing up a large number of files because of less space](#)
- [Catalog Job fails when backing up a large number of files on Windows 2003 x86 machine](#)
- [Failed to create snapshot for selected volumes](#)
- [Unable to change backup destination folder to Arcserve UDP Recovery Point View](#)

SQL Server backup failed due to "out of memory" error

This is caused by a Microsoft known issue: Volume Shadow Copy Service (VSS) cannot create a volume snapshot even when VSS has sufficient memory space.

To resolve this problem, apply the Microsoft [patch](#).

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 sessions do not include Microsoft SQL database information

After upgrading from a previous release, the backup sessions of Arcserve UDP Agent (Windows) do not include any Microsoft SQL database information. This may be caused by the SQL server not starting automatically in a virtual environment. If this occurs, verify that the SQL database is in a good state and retry the backup.

If the problem persists, you can change the startup type of the SQL server to "Automatic (Delayed Start)".

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.

Catalog Job fails Due to Less Space when Backing up Large Number of Files

If you are attempting to back up a large number of files and the catalog generation job fails because there is not enough available space in the Arcserve UDP Agent (Windows) home folder, perform the following task to create a new temp location:

Important! Verify that this new location contains enough free space to hold all of your catalog temporary data.

1. Within the Arcserve UDP Agent (Windows) home folder, access the **Configuration** folder. (The Arcserve UDP Agent (Windows) home folder is located on the Arcserve UDP Agent (Windows) install path).

Program Files\Arcserve\Unified Data Protection\Engine\Configuration

2. Within the **Configuration** folder, create a **switch.ini** file. (File name is case sensitive).
3. Within the new **switch.ini** file, add the following content:

```
[CatalogMgrDll.DLL]
```

```
Common.TmpPath4Catalog="I:\catalogtemp"
```

4. Run the backup job again.

The catalog generation part of the job will now go to the newly created temp folder.

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 create snapshot for selected volumes

If a volume does not have enough disk space, the backup job can fail with the error message "Failed to create snapshot for selected volumes". If the backup job fails, you can perform either task:

- Free up some disk space on the volumes being backed up.
- Reconfigure the **Volume Shadow Copy** settings to save shadow copy to a volume with sufficient free disk space.

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 change backup destination folder to Arcserve UDP Recovery Point View

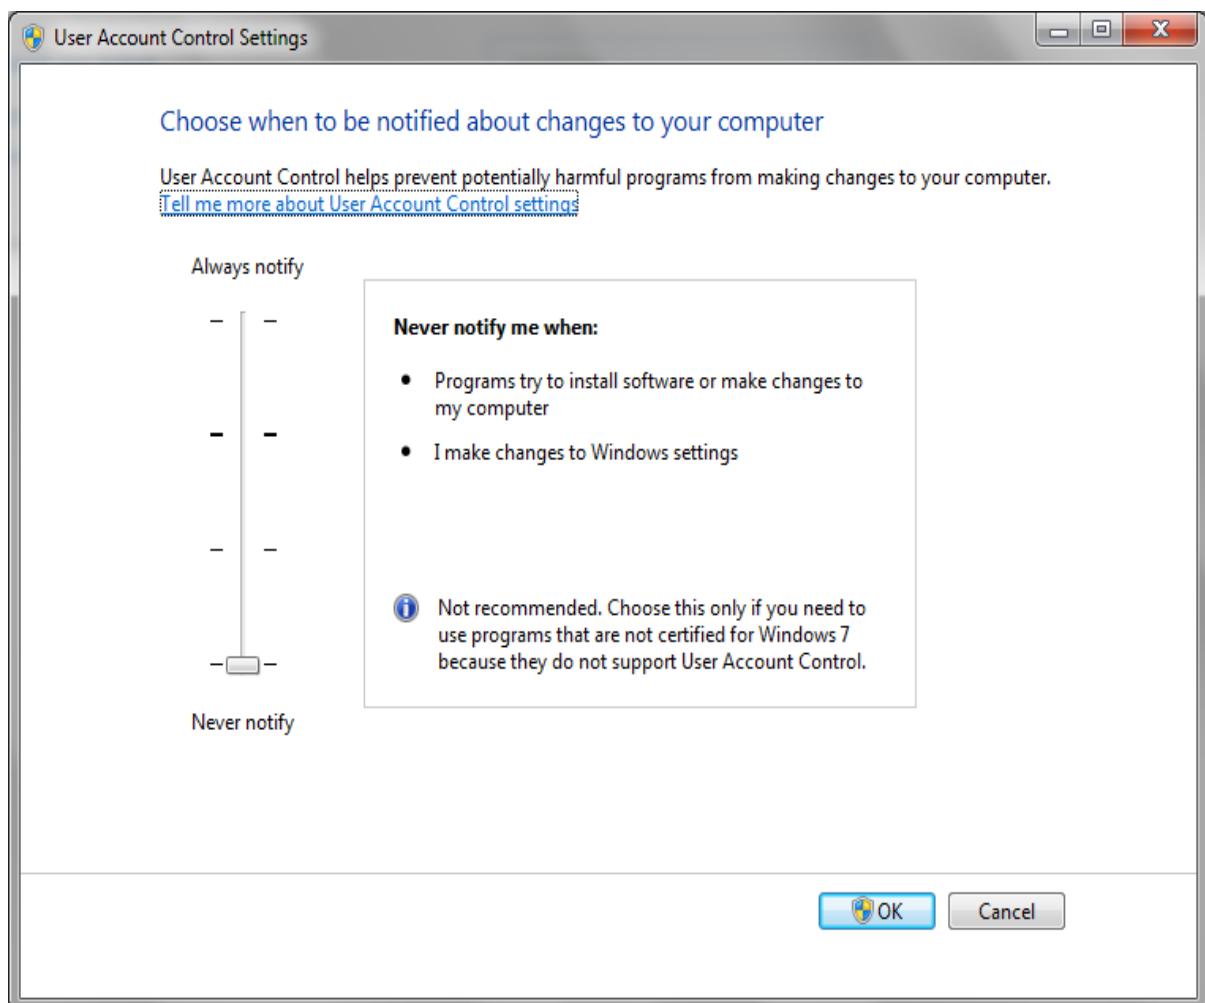
On Windows Vista and later operating systems, if you create an account which belongs to a local administrator group, and from this new account you attempt to change the Arcserve UDP Agent (Windows) backup destination folder to Arcserve UDP Recovery Point View, the folder view cannot be changed and no error message is displayed. This can happen when the **User Account Control** is enabled.

If this condition occurs, you can either disable the **User Account Control** or you can grant Modify privileges to the created Windows account.

To disable the User Account Control, perform the following task:

1. From the Windows **Control Panel**, select **User Accounts**, **User Accounts**, and then **Change User Account Control Settings**.

The **User Account Control Settings** dialog displays.




2. For the **Choose when to be notified about changes to your computer** option, drag the slide bar to the bottom (Never notify).
3. When you disable the **User Account Control**, reboot your computer.

To grant Modify privileges to the created Windows account, perform the following task:

1. From the **Windows Explorer** view, navigate to the specified backup destination.
2. Right-click on the backup destination folder, select **Properties**, and click the **Security** tab.
3. Click **Edit** and Add a user for this destination folder.

The **Permissions** dialog is displayed.



4. For this user, check the **Modify** permissions option to allow control specifically to this user and add it to the folder security list.

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.

Perform File Copy to Disk/Cloud

Arcserve UDP Agent (Windows) provides the capability to copy/move files to and from the cloud or disk, which is based on your specified File Copy and retention policies. 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 process of file copying Arcserve UDP Agent (Windows) backup information lets you specify the file copy source, destination, and corresponding policies for the type of File Copy job performed. The available types are File Copy and File Copy - Delete Source. The two processes are similar, with the exception that when you perform a File Copy - Delete Source job, 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 File Copy job, the data is copied from the source to the destination (remains on source destination) and provides multiple stored versions. For more information about configuring the File Copy settings and policies, see [Manage File Copy Settings](#).

The File Copy process can only be launched automatically as scheduled. For more information about configuring the schedule settings, see [Specify the File Copy Schedule](#).

Note: For a compressed File Copy job, the activity log displays only the uncompressed size.

Perform a Restore

Arcserve UDP provides you with various tools and options that you can use to restore data. The aim of running a successful restore job is to quickly identify the data you need and to retrieve it from the appropriate backup media. Each restore job requires a source and destination.

Restore Considerations

Before you perform an Arcserve UDP Agent (Windows) restore, review the following restore considerations:

- **Restore Considerations for a remote destination**

If all the drive letters (A - Z) are occupied, the restore to a remote path will not succeed because Arcserve UDP Agent (Windows) needs to use one drive letter to mount the remote destination path.

- **Restore Considerations for Hyper-V servers**

On a Hyper-V server (even if you have the proper VM license), you must manually restore VHD files of VM and then re-register them with Hyper-V Manager.

Note: After the VHDs are restored they are not directly registered with Hyper-V Manager. You can either attach them to existing VM or create a new VM and attach those to them.

- **Restore Considerations for a Microsoft SQL Server 2008 database with FILESTREAM data**

Both the database and its related FILESTREAM BLOB data can be automatically backed-up by Arcserve UDP Agent (Windows), but the FILESTREAM BLOB data cannot be restored automatically with the database. This is because FILESTREAM feature is not supported by the latest SQL Server Writer. As a result, when one database with FILESTREAM BLOB data is restored, just restoring the database is no longer enough, and the folder of FILESTREAM BLOB data also needs to be restored.

Note: FILESTREAM is a feature introduced by Microsoft SQL Server 2008, which provides the capability of storing binary large object (BLOB) data (MP3, Word, Excel, PDF, etc.) in the NTFS file system, rather than in a database file.

- **Restore Considerations for session dismount time**

When you browse the volume of one recovery point which does not have a catalog, the volume will be mounted. After the volume is mounted, the volume status is queried every 10 minutes to check if it is used. If it is not used, it will be dismounted.

To change the default session dismount time of 10 minutes, modify the Registry Key, using the following information:

- **Registry key path:** Arcserve UDP Agent (Windows) Installation path
- **Registry key name:** SessionDismountTime

- **Registry key type:** String
- **Registry key value unit:** second

For example: If you set the registry value to 60, the mounted volume status is queried every 60 seconds and if it is not used for the last 60 seconds, it will be dismounted.

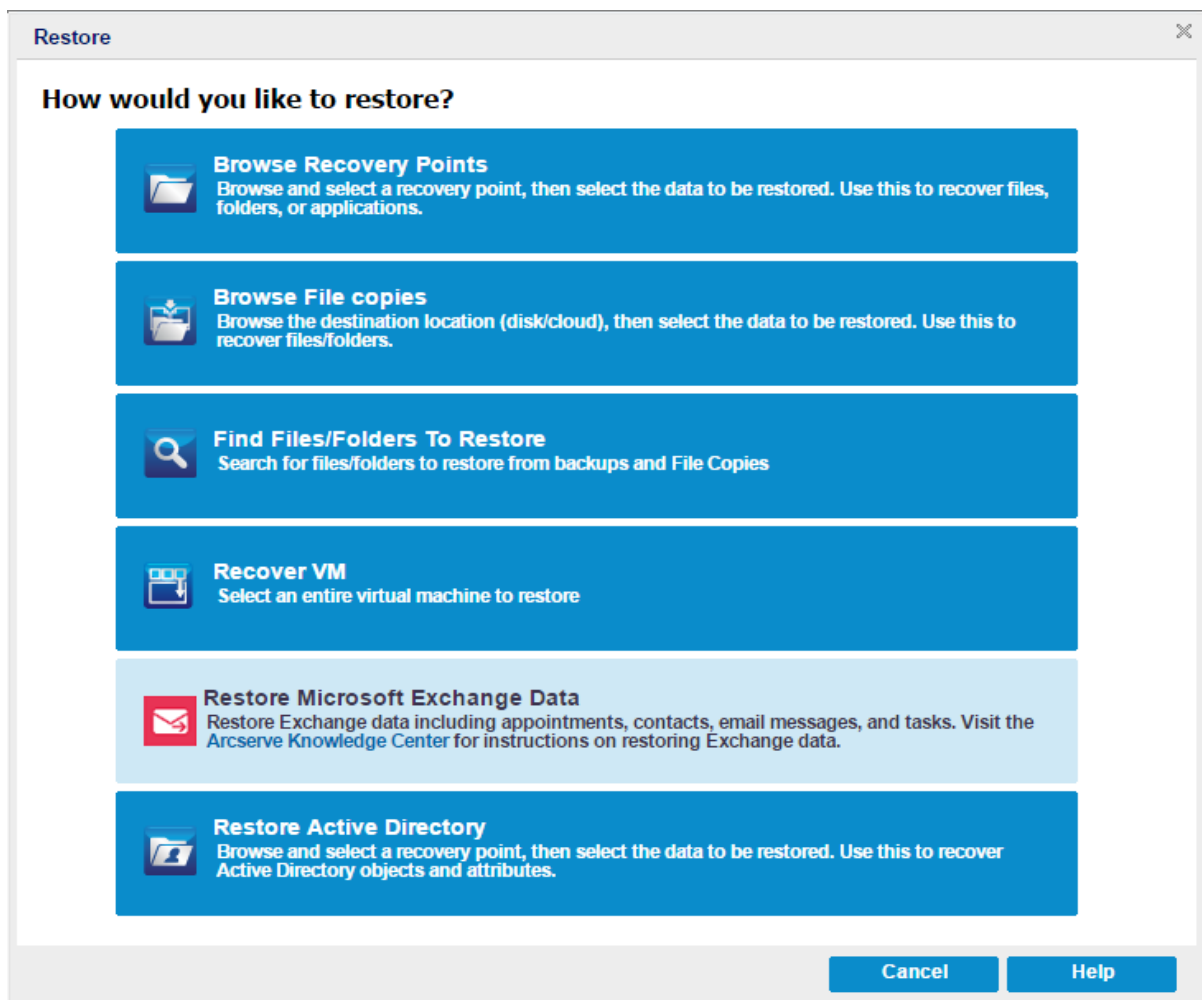
Restore Methods

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. The Arcserve UDP Agent (Windows) provides various methods to identify and locate the backed up data and allow you to restore it. Regardless of the restore method you select, Arcserve UDP Agent (Windows) uses visual indicators (restore markers) of the objects that are or are not selected for restore. For more information, see [Restore Markers](#).

Restore data

1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select Restore.

The restore methods selection dialog opens.



2. Select the type of restore you want to perform.

The available options are Browse Recovery Points, Browse File Copies, Find Files/Folders to Restore, Recover VM, and Restore Exchange Mails.

Note: Use the Browse Recovery Points if you want to restore any applications.

Browse Recovery Points

Lets you browse the available recovery points (successful backups) from a calendar view. The calendar dates that contain valid recovery points are highlighted in green. When you select a recovery date, all the associated recovery points for that date are displayed. You can then browse and select the backup content (including applications) to be restored.

Browse File Copies

Lets you browse the available File Copy locations (local disk/network drive or cloud) to locate the specific file or folder that is going to be restored.

Note: This option is not available for VM backup proxy.

Find Files/Folders to Restore

Lets you search for a file name pattern in a specific location and or the file version that is going to be restored.

Recover VM

Lets you browse the available virtual machine recovery points from a calendar view. The calendar dates that contain valid recovery points are highlighted in green. When you select a recovery date, all the associated virtual machine recovery points for that date are displayed. You can then browse and select the virtual machine content that is going to be restored.

Restore Microsoft Exchange Data

Lets you restore exchange data that includes appointments, contacts, email messages, and so on.

Restore Active Directory

Lets you recover Active Directory objects and attributes.

Restore Markers

Regardless of which restore method you select, when you navigate to a specific volume, folder, or file to restore, each object displayed in the restore window has a green or gray box to its left called a marker. These markers are visual indicators of the objects that are or are not selected for restore.




Green marker

Lets you control the extent of the restore for an object directly. Click a marker to exclude an object from a restore or to indicate that you want the restore for the object to be full or partial. As you click the marker, you fill or empty the marker of color, indicating the extent of the restore.

Gray marker

These markers are associated with objects that are not real and that you cannot restore. Typically, these items serve as placeholders under which other objects are grouped and displayed. As you click the green markers under a gray marker item, the fill proportion of the gray marker changes automatically from empty to partial to full depending on the proportion of files you have chosen to restore.

The following table describes the different marker configurations and corresponding restore levels:

Marker	Configuration	Description
	Completely filled center.	Full restore.
	Partially filled center.	Partial restore.
	Empty center.	Do not restore.

Note: Gray marker configurations follow the same pattern as green marker configurations, but reflect the proportion of files under them that are selected for restore.

The fill proportion of a marker at a higher level of the directory tree depends on the fill proportions of the markers of the objects at the lower levels.

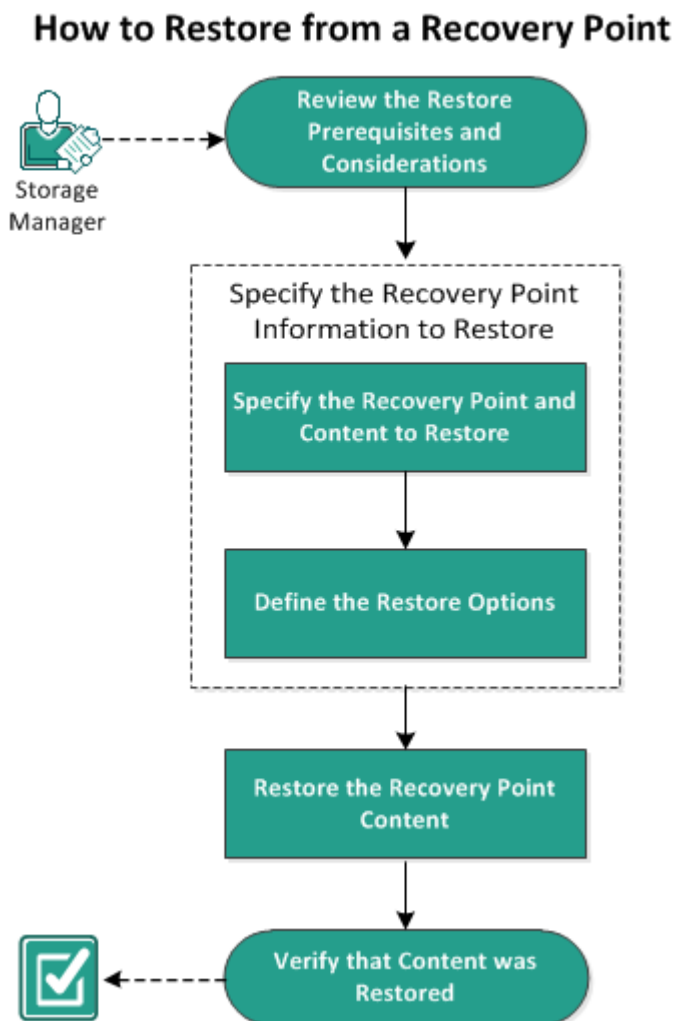
- If all the markers at the lower levels are completely filled, then the marker at the higher level is also automatically completely filled.
- If the markers at the lower levels are a mix of completely filled and partially filled, then the marker at the higher level is automatically partially filled.

If you click a marker at a higher level so that it is completely filled, then all the markers at the lower levels are automatically filled completely.

How to Restore From a Recovery Point

Each time Arcserve UDP performs a successful backup, a point-in-time snapshot image of your backup is created (recovery point). This collection of recovery points allows you to locate and specify exactly which backup image you want to restore. If at some later time, you suspect any of the backed up information is missing, corrupted, or not reliable, you can then locate and restore from a previous known good version.

The following diagram illustrates the process to restore from a recovery point:



Perform the following tasks to restore from a recovery point:

1. [Review the Restore Prerequisites and Considerations](#)
2. [Specify the Recovery Point Information to Restore](#)
 - a. [Specify the Recovery Point and Content to Restore](#)
 - b. [Define the Restore Options](#)

3. [Restore the Recovery Point Content](#)
4. [Verify that Content is Restored](#)

Review the Restore Prerequisites and Considerations

Verify that the following prerequisites exist before performing a restore:

- You have at least one recovery point available to restore.
- You have a valid and accessible recovery point destination to restore the recovery point content from.
- You have a valid and accessible target location to restore the recovery point content to.
- Review the [Compatibility Matrix](#) that provides the supported operating systems, databases, and browsers.

Review the following restore considerations:

- If the restore is to a remote destination and if all the drive letters (A - Z) are occupied, the restore to a remote path will not succeed. Arcserve UDP Agent (Windows) needs to use one drive letter to mount the remote destination path.
- (Optional) Understand how the restore process works. For more information, see [How File Level Restores Work](#).
- (Optional) Review the files skipped during restore. For more information, see [Files Skipped During Restore](#).
- When you attempt to restore an optimized backup session to a non-empty volume (unoptimized restore), the restore job may take more time than the estimated time displayed in the job monitor. The amount of data that is processed and the elapsed time may increase based on the data that is optimized on the volume.

Example:

The backup volume size is 100 GB and after optimization the volume size is reduced to 50 GB.

When you perform an unoptimized restore of this volume the restore job monitor displays 100% after restoring 50 GB, but it will take more time to restore the entire 100 GB.

- The following Activity log message will be displayed when restoring the system files:


"System files were skipped. If necessary, you can use the Bare Metal Recovery (BMR) option to restore them."

How File Level Restores Work

During a block-level backup, each backed up file is made up of a collection of blocks that define that particular file. A catalog file is created containing a list of the backed up files, along with the individual blocks that were used for each file and the available recovery points for these files. When you need to restore a particular file, you can search your backup and select the file you want to restore and the recovery point you want to restore from. Then Arcserve UDP collects the version of the blocks that were used for the recovery point of the specified file, and reassembles and restores the file.

Note: You can also perform a restore without a catalog file from a catalog-less backup recovery point.

The following flow diagram shows the process of how Arcserve UDP restores a specific file:



Files Skipped During Restore

During restore by Arcserve D2D some files may be skipped intentionally.

The files and folders in the tables below 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 are skipped when being an important component for Windows or Arcserve D2D.

OS	Folder or Location	File or Folder Name	Remark
All	Root folder of each volume	CAVolTrc.dat	Used by the 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 who 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 contains 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 when you restore to the original or alternate location:

OS	Folder or Location	File or Folder Name	Remark
All	Folder specified in value record under: HKLM\Software\Microsoft\Windows NT\CurrentVersion\WinLogon\SfcDllCache	All files/folders (recursively)	Folder contains a cached dll file

			which is used for System File Checker (SFC) and contents of the system dll cache directory are rebuilt by using SFC.
--	--	--	--

<p>%SystemRoot%\SYSTEM32\dllCache</p>		
<p>Root folder of quorum_device</p>	<p>MSCS*</p>	<p>Used for Microsoft Cluster Server.</p>
<p>%SystemRoot%\SYSTEM32\</p>	<p>perf?00?.dat</p>	<p>Performance data used by the Windows</p>
	<p>perf?00?.bak</p>	<p>performance counter.</p>
	<p>CATROOT*</p>	<p>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.</p>
<p>%SystemRoot%\inetsrv\</p>	<p>metabase.bin</p>	<p>Metabase binary file of earlier IIS versions before 6.0.</p>
<p>File or folder specified in value except "SIS Common Store" under HKLM\SYSTEM\CurrentControlSet\Control\BackupRestore\FilesNotToBackup</p>	<p>All files/folders (recursively)</p>	<p>Do not back up and restore Files and folders . For more inform-</p>

			ation, see the link .
XP W2003	System volume	NTLDR	The main boot loader.
		BOOT.INI	Contains boot configuration (if missing, NTLDR will default to \Windows on the first partition of the first hard drive).
		NTDETECT.COM	Required for booting an NT-based OS. Detects basic hardware information needed for a successful boot.
Vista and later	Root folder of system volume	boot*	Boot folder for Windows.
		bootmgr	Windows boot manager file.
		EFI\Microsoft\Boot*	Used for EFI boot.
	%SystemRoot%\SYSTEM32\	LogFiles\WMI\RTBackup*	Stores ETW trace files (extension .etl) for real time event trace ses-

			sions.
		config\RegBack*	Backup of current registry table.
Win-8 and later	System volume	swapfile.sys	System controller file, normally around 256 MB. It is used by Metro style applications that do not fit the traditional paging characteristics (such as usage pattern, growth, space reservation) of pagefile.sys.
		BOOTNXT	Used to boot from OS, other than Windows 8. Created when enabling the startup options, and updated by Windows.

The Activity log provides the following information:

- Date Time Information: jobxxxx System Files skipped. You can use Bare-Metal Recovery Option (BMR) to restore them.
- Date Time Information: jobxxxx Files or Directories skipped. Which files or directories were skipped can be found in: C:\Program Files\Arcserve\Unified Data Protection\Engine\Logs\Restore-<YYYYMMDD>-<hhmmss>-<Process ID>-<Job ID>.log.

Specify the Recovery Point Information to Restore

Arcserve UDP provides you with an option to restore data from a recovery point. The goal of running a successful restore job is to identify quickly the data you need and to retrieve it from the appropriate backup location. Each restore job requires a source and destination.

The process involved in restoring from a recovery point is as follows:

1. [Specify the Recovery Point and Content to Restore](#)
2. [Define the Restore Options](#)

Specify the Recovery Point and Content to Restore

Each time you perform a backup, a recovery point is created. Specify the recovery point information in the Restore dialog so that you can recover the exact data that you want. You can restore specific files or all files depending on your requirement.

Follow these steps:

1. Access the Restore dialog in one of the following ways:

From Arcserve UDP:


- a. Log into Arcserve UDP console.
- b. Navigate to **resources > Node > All Nodes**.
All the added nodes are displayed in the center pane.
- c. Right-click the node, and then click **Restore**.

You are automatically logged into the agent node console, and the Restore dialog is opened.

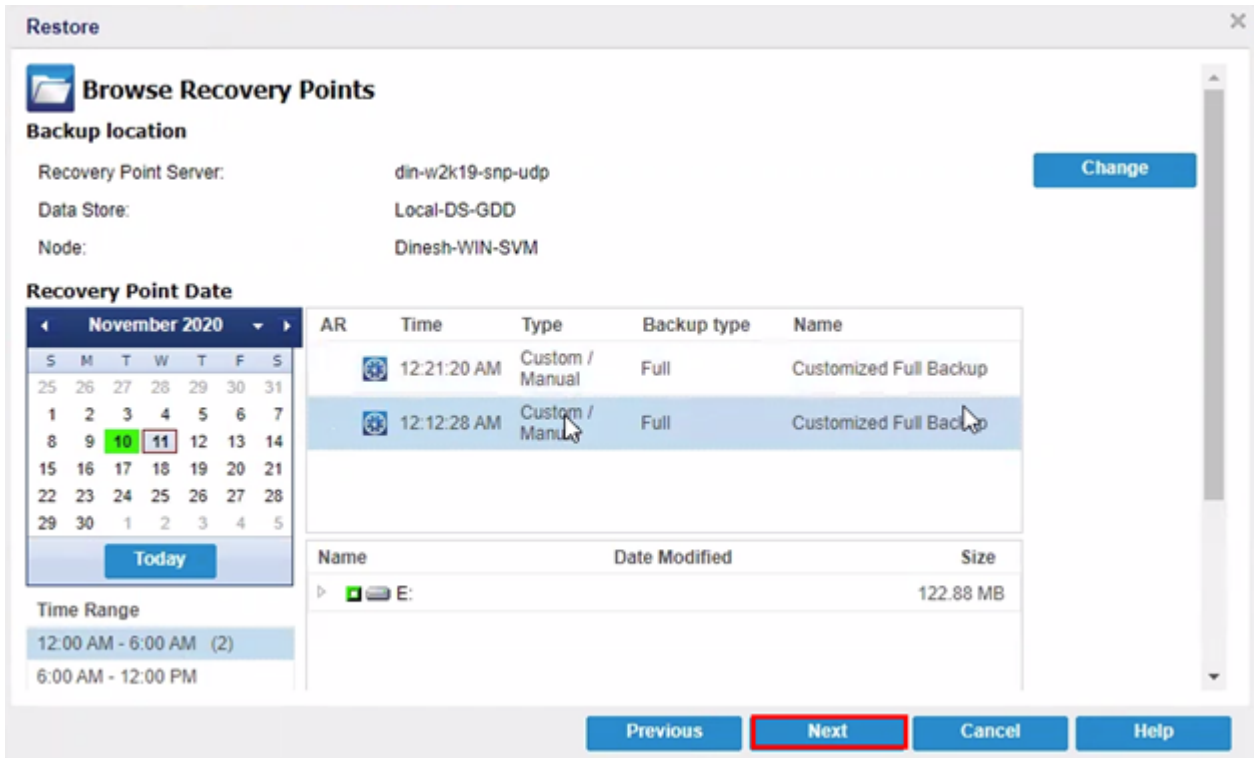
From Arcserve UDP Agent (Windows):

- a. Log into Arcserve UDP Agent (Windows).
- b. From the home page, click **Restore**.
The Restore dialog opens.

2. On the Restore dialog, click the **Browse Recovery Points** option.



3. On the Browse Recovery Points page, do the following:



a. To update the backup location, click **Change**.

The Source dialog opens.

- b. On the Source dialog, select one of the following backup locations, and then click **OK**:

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

Select local disk or shared folder

- Specify or browse to the location where your backup images are stored, and select the appropriate backup source.
- To verify the connection to the specified location, click the green arrow button. If required, enter the Username and Password credentials to access to that source location.

The Select backup location dialog opens.

- Select the folder where the recovery points are stored, and then click **OK**.

The Select backup location dialog closes, and you can see the backup location on the Source dialog.

- 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 then click **Refresh**.

All the agents are listed in the Data Protection Agent column on the Source dialog.

- b. Select the agent from the displayed list.
- c. Click **OK**.

The recovery points are listed in the Browse Recovery Points dialog.

- c. Select the calendar date for the backup image to restore.

All the dates containing recovery points for the specified backup source are highlighted in green.

The corresponding recovery points for that date are displayed, with the time of the backup, the type of backup that was performed (Full, Incremental, or Verify), and the name of the backup.

- d. Select a recovery point to restore.

The backup content (including any applications) for the selected recovery point displays.

Note: If a clock icon appears with a lock symbol, it indicates that the recovery point contains encrypted information and may require a password for restore.

- e. Select the content to restore.
 - ◆ For a volume-level restore, you can specify to restore the entire volume or selected files/folders within the volume.
 - ◆ For an application-level restore, you can specify to restore the entire application or selected components, databases, instances, and so on, within the application.

4. Click **Next**.

The Restore Options dialog opens.

The recovery point and content to restore is specified.

Define the Restore Options

After you specify a recovery point and content to restore, define the copy options for the selected recovery point. This section explains how to define the restore options for the selected recovery point.

Follow these steps:

1. On the Restore Options dialog, select the restore destination.

The screenshot shows the 'Restore' dialog box with the following options:

- Restore options**
 - Destination**
 - Select the restore destination
 - Restore to original location
 - Restore to
 - Use selected network for restore traffic
 -
 - Continue to run job even when unable to connect to the selected destination network
 - Use dedicated ethernet if current machine enables SMB Multichannel
 - Block Level Restore**
 - Perform Block Restore
 - Best for large volumes with many small files.
 - Use of this option will overwrite the target volume.
 - Resolving Conflicts**
 - Specify how to resolve conflicts
 - Directory Structure**
 - Whether to create root directory during restore

At the bottom of the dialog are buttons for **Previous**, **Next**, **Cancel**, and **Help**.

The available destination options are:

Restore to original location

Restores to the original location from where the backup image was captured.

Note: If you performed the recovery point backup using host-based agentless backup, restoring to original location is to restore the file back in to the virtual machine. In this case, a dialog box opens. You may enter the credentials of the hypervisor, and the operating system of the virtual machine.

For VMware VM:

Set Credential for Source vCenter/ESX Server

vCenter/ESX Server Information

vCenter/ESX Server: abc123-vc

Protocol: HTTP HTTPS


Port Number: 443

User Name: hbbaadmin

Password: ●●●●●●●●

VM Settings

VM Name: shuli02-UEFI

VM username: 

VM password:

OK Cancel

Note: To create or write files inside the VM, consider the following requirements for the settings and account permission of virtual machine:

- ◆ VMware Tools is installed and running.
- ◆ Firewall must allow File and Printer Sharing.
- ◆ The account is built-in local administrator, built-in domain administrator, or domain account that is member of the local Administrators group. If other accounts are used, do the following:
 - Disable the UAC remote access. To disable UAC remote access, see [Import Virtual Machine Using Additional Administrative Account](#).
 - Disable UAC in the Local Security Policy by disabling the setting Run all administrator in Admin Approval Mode at secpol.msc -> Local Policies -> Security Options. (Secpol.msc is Microsoft's security policy editor).

Important! Do not attempt to disable the UAC in the User Account Control Settings dialog box that opens from the control panel.

For Hyper-V VM:

The screenshot shows a dialog box titled "Set the credentials for the source Hyper-V Server". It has two main sections: "Hyper-V Server Information" and "VM Settings".

- Hyper-V Server Information:**
 - Hyper-V/Hyper-V Cluster Server: abc123 -hyperv1
 - User Name: administrator
 - Password: [masked with dots]
- VM Settings:**
 - VM Name: abc123-hv102
 - VM username: [empty]
 - VM password: [empty]

At the bottom of the dialog are "OK" and "Cancel" buttons.

Note: To create or write files inside the VM, consider the following requirements for the settings and account permission of virtual machine:


- ◆ Hyper-V integration services are installed and running.
- ◆ Firewall must allow File and Printer Sharing.
- ◆ The account is built-in local administrator, built-in domain administrator, or domain account that is member of the local Administrators group. If other accounts are used, disable the UAC remote access. To disable UAC remote access, see [Import Virtual Machine Using Additional Administrative Account](#).
- ◆ If virtual machine guest OS is Client version Windows (such as Windows 10), you need to manually configure firewall to allow Windows Management Instrumentation (WMI).

Restore to

Restores to the specified location. To restore data to the specified location, do the following:

- ◆ To provide the destination, click **Browse**.
The Select alternative restore location dialog appears.

- ◆ Select the existing folder or create a new folder as needed, and then click **OK**.



- ◆ To verify the connection to the specified location, click the green arrow button. If necessary, enter the Username and Password to gain access to that location.
2. (Optional) To enable the communication between Windows Agent and Recovery Point Server, select the **Use selected network for restore traffic** check box, and then select the network from the drop-down list.

Notes:

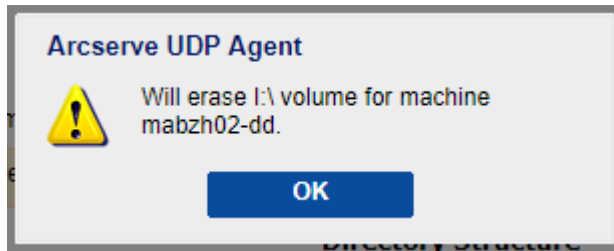
- If the selected backup network is not accessible and to continue the backup job with the available network or with the default network, click the **Continue to run job even when unable to connect to the selected backup network** check box.
- To define the constraint on SMB Multichannel so that the data transfers only through the selected network, select the **Use dedicated ethernet if current machine enables SMB Multichannel** check box.

This option is not available by default. To enable this option, create the *UseDedicatedEthernet* string registry in the following path, and then set the registry value to 1:

SOFTWARE\Arcserve\Unified Data Protection\Engine

3. To improve the throughput when restoring a large volume with many small files, select a volume, and then under Block Level Restore, click the **Perform Block**

Restore check box. Other options get disabled and a notification appears. To continue with the block level restore, click **OK**.



Note: The data gets restored to current UDP Agent machine, and the target volume is overwritten. The target volume is not accessible during the restore job.

4. Specify the Resolving Conflicts option that Arcserve UDP performs if conflicts are encountered during the restore process.

The available options are:

Overwrite existing files

Overwrites (replaces) any existing files, which are at the restore destination. All the objects are restored from the backup files regardless of their current presence on your computer.

Replace active files

Replaces any active files upon reboot. If during the restore attempt Arcserve UDP Agent (Windows) discovers that the existing file is currently in use or being accessed, it will not immediately replace that file, but instead to avoid any problems will delay the replacement of the active files until the next time the machine is rebooted. (The restore occurs immediately, but the replacement of any active files is done during the next reboot).

This option is only available if you select the **Overwrite existing files** option.

Note: If you do not select this option, any active file is skipped from the restore.

Rename files

Creates a new file if the file name already exists. Selecting this option copies the source file to the destination with the same file name but a different extension. Data is then restored to the new file.

Skip existing files

Skips over and not overwrite (replace) any existing files that are located at the restore destination. Only objects that are not currently existing on your machine are restored from the backup files.

Default: Skip existing files.

5. For Directory Structure, if you want to create a root directory during restore, select the **Create root directory** checkbox.

The **Create root directory** checkbox specifies that if a root directory structure exists in the captured backup image, Arcserve UDP recreates that same root directory structure on the restore destination path.

If this option not selected, the file or folder is restored directly to the destination folder.

For example, if during the backup you captured the files "C:\Folder-1\SubFolder2\A.txt" and "C:\Folder1\SubFolder2\B.txt", and during the restore you specified the restore destination as "D:\Restore".

- If you select to restore the "A.txt" and "B.txt" files individually, the destination for the restored files will be "D:\Restore\A.txt" and "D:\Restore\B.txt" (the root directory above the specified file level will not be recreated).
- If you select to restore from the "SubFolder2" level, the destination for the restored files will be "D:\Restore\SubFolder2\A.txt" and "D:\Restore\SubFolder2\B.txt" (the root directory above the specified folder level will not be recreated).

If this option selected, the entire root directory path for the files/folders (including the volume name) is recreated to the destination folder. If the files/folders to be restored are from the same volume name, then the destination root directory path does not include that volume name. However, if the files/folders to be restored are from different volume names, then the destination root directory path does include the volume name.

For example, if during the backup you captured the files "C:\Folder-1\SubFolder2\A.txt", "C:\Folder1\SubFolder2\B.txt", and also E:\Folder-3\SubFolder4\C.txt", and during the restore you specified the restore destination as "D:\Restore".

- If you select to restore just the "A.txt" file, the destination for the restored file will be "D:\Restore\Folder1\SubFolder2\A.txt" (the entire root directory without the volume name will be recreated).
 - If you select to restore both the "A.txt" and "C.txt" files, the destination for the restored files will be "D:\Restore\C\Folder1\SubFolder2\A.txt" and "D:\Restore\E\Folder3\SubFolder4\C.txt" (the entire root directory with the volume name will be recreated).
6. For Recovering ACL, to skip the original permission for the restored files/folders, select the **skip recovering ACL of files / folders** checkbox. Selecting this option lets you inherit the permissions of target folder instead. If you do not select this option,

the original permissions are kept.

7. For **Backup Encryption or Password Protection**, if necessary, specify the backup encryption password when the data you are trying to restore is encrypted.

A password is not required if you are attempting to restore from the same Arcserve UDP Agent (Windows) computer from where the encrypted backup was performed. However, if you are attempting to restore from a different Arcserve UDP Agent (Windows) computer, a password is required.

Note: If a clock icon appears with a lock symbol, it indicates that the recovery point contains encrypted information and may require a password for restore.

8. Click **Next**.

The Restore Summary page opens.

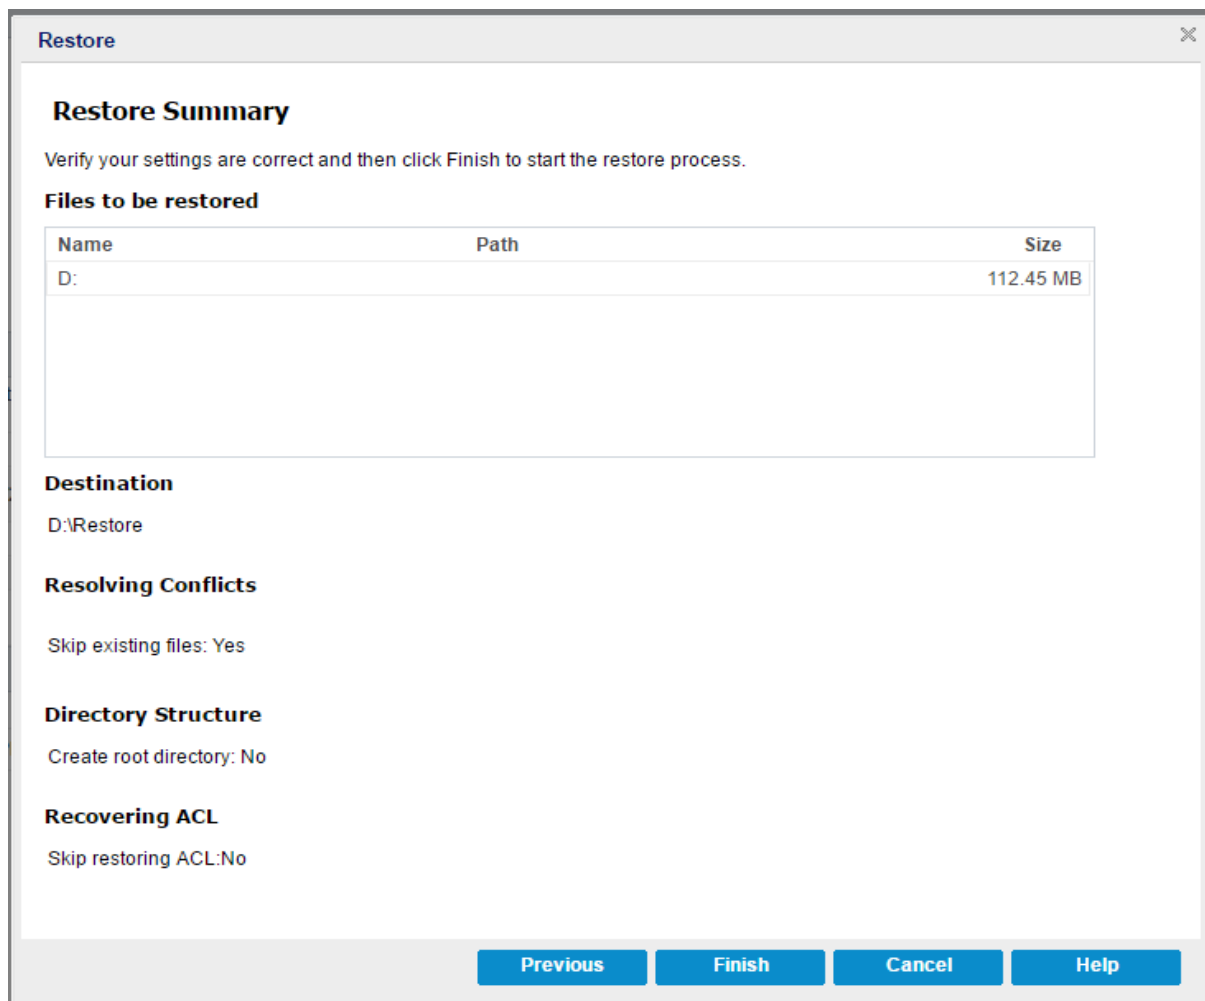
The restore options are defined to restore from a recovery point.

Restore the Recovery Point Content

After you define the restore options, verify that your settings are correct and confirm the restore process. The Restore Summary page helps you review all the restore options that you defined and modify them if necessary.

Follow these steps:

1. On the Restore Summary page, review the displayed information to verify that all the restore options and settings are correct.



2. Do one of the following:
 - ◆ If the summary information is incorrect, click **Previous** and go back to the applicable dialog to change the incorrect setting.
 - ◆ If the summary information is correct, click **Finish** to launch the restore process.

The recovery point content is restored.

Verify that Content is Restored

After the completion of the restore process, verify that content is restored to the specified destination.

Follow these steps:

1. 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.

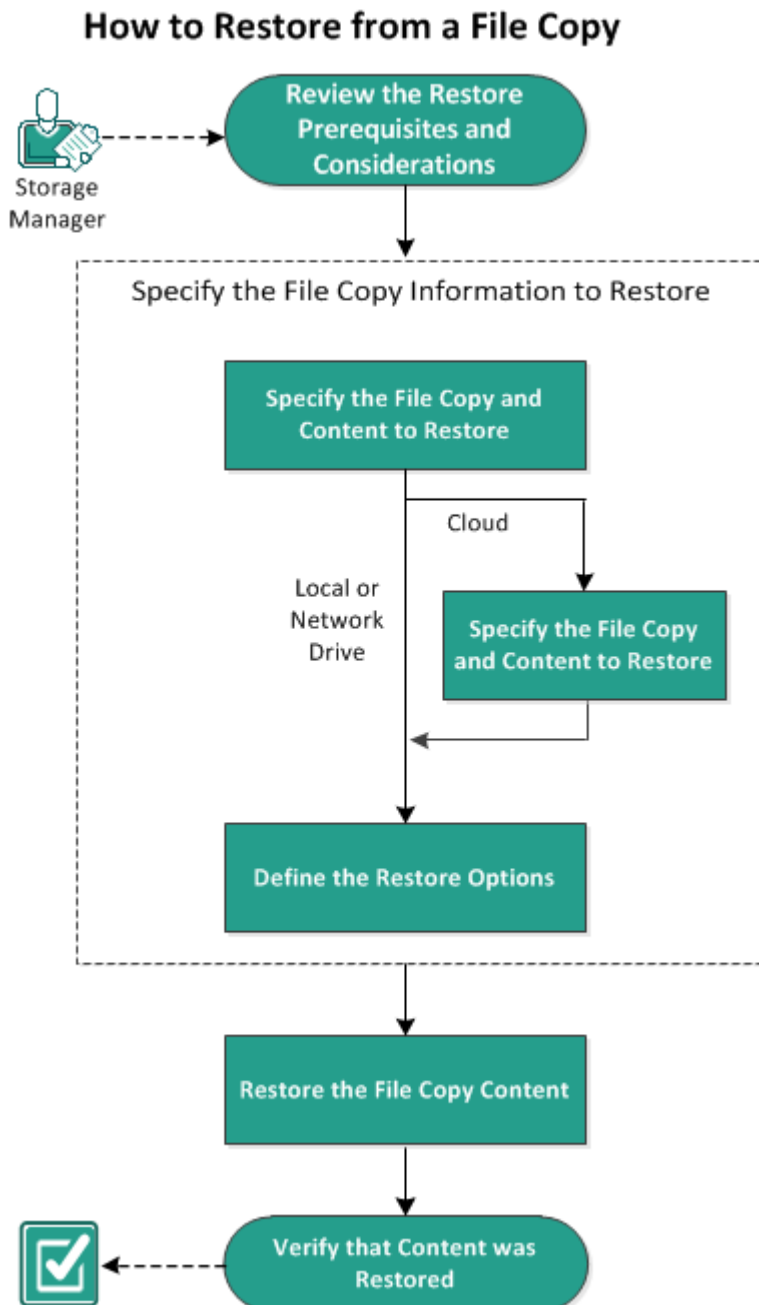
3. Verify the content to confirm the restore job.

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:



Perform the following tasks to restore from a File Copy:

1. [Review the Restore Prerequisites and Considerations](#)

2. [Specify the File Copy Information to Restore](#)
 - a. [Specify the File Copy and Content to Restore](#)
 - [Specify Cloud Configuration for Restore](#)
 - b. [Define the Restore Options](#)
3. [Restore the Recovery Point Content](#)
4. [Verify that Content was Restored](#)

Review the Restore Prerequisites and Considerations

Verify that the following prerequisites exist before performing a restore:

- You have at least one file copy available to restore.
- You have a valid and accessible file copy destination to restore the file copy content from.
- You have a valid and accessible target location to restore the file copy content to.
- Review the [Compatibility Matrix](#) that provides the supported operating systems, databases, and browsers.

Review the following restore considerations:

- Arcserve UDP only allows one restore job to run at the same time. If you attempt to launch a restore job manually, while another restore job is running, an alert message opens informing 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 (Thread-sForArchive).
 - Restore from a File Copy can download more than one file at a time (Thread-sForRestore).
 - Catalog Synchronization uses multiple threads (ThreadForCatalogSync).

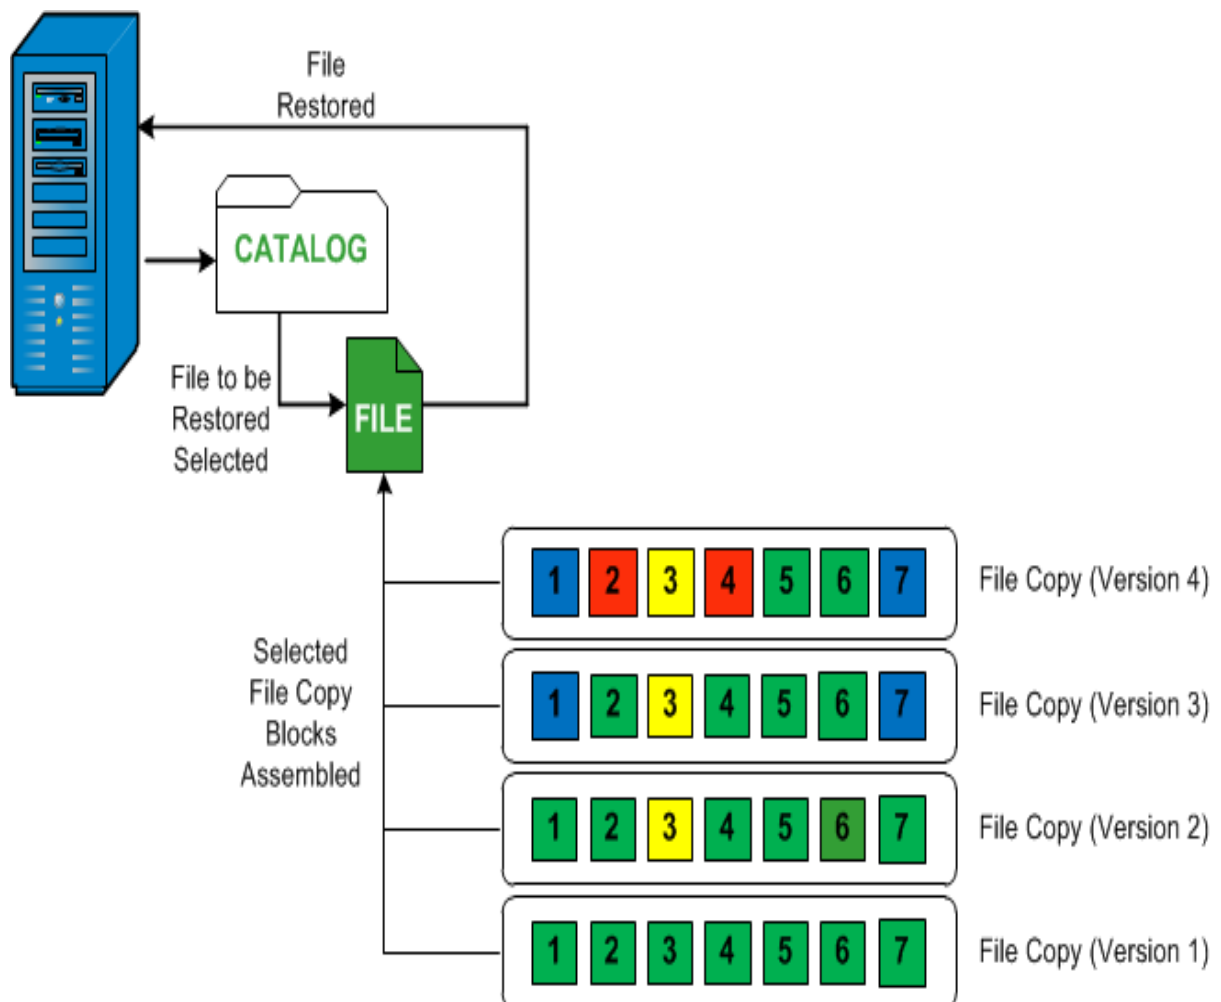
You can change the default File Copy Registry values by modifying the appropriate DWORD value. For more information, see [Configure File Copy Settings to Optimize Performance](#) in the online help.

- (Optional) Understand how the restore process works. For more information, see [How File Level Restores Work](#).

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](#)
2. [Define the Restore Options](#)

Specify the File Copy and Content to Restore

Use the **Browse File Copies** option to restore from a file copy. This restore method allows you to browse the file copied data and specify exactly which file you want to restore.

Follow these steps:

1. Access the restore method selection dialog in one of the following ways:

- ◆ From Arcserve UDP:

- a. Log into Arcserve UDP.
- b. Click the **resources** tab.
- c. Select **All Nodes** in the left pane.
All the added nodes are displayed in the center pane.
- d. In the center pane, select the node and click **Actions**.
- e. Click **Restore** from the **Actions** drop-down list.

The restore method selection dialog opens.

Note: You are automatically logged in to the agent node and the restore method selection dialog is opened from the agent node.


- ◆ From Arcserve UDP Agent (Windows):

- a. Log into Arcserve UDP Agent (Windows).
- b. From the home page, select **Restore**.

The restore method selection dialog opens.


2. Click the **Browse File Copies** option.

The **Restore** dialog opens. The destination that is currently showing in the **Restore From** field is the configured default **File Copy** destination.



- 3. If necessary, you can click **Add** to browse to an alternate location where your file copy images are stored.

The **Destination** dialog opens displaying the available alternate destination options.



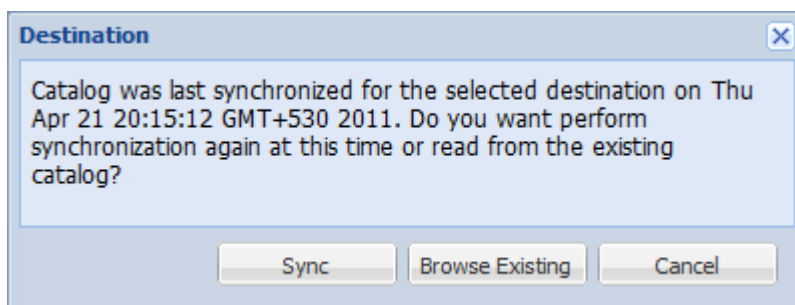
Local or network drive

The **Select a Backup Location** dialog opens, allowing you to browse to and select an alternate local or network drive location.

Cloud

The **Cloud Configuration** dialog opens, allowing you to access and select an alternate cloud location. For more information about this dialog, see *Specify Cloud Configuration for Restore*.

Regardless of whether you selected to restore from **Local or network drive** or from **Cloud**, when you change the destination to an alternate location a pop-up dialog will appear, asking if you want to perform a new catalog synchronization or read from the existing catalog.



- If 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 or file archive cloud location.

Configure the way to access a new cloud storage location

Storage Name	<input type="text" value="Enter a storage name"/>
Storage Service	<input type="text" value="Amazon S3"/> ▼
Bucket Region	<input type="text" value="Select a bucket region"/> ▼
Access Key ID	<input type="text" value="Enter a key ID"/>
Secret Access Key	<input type="text"/>
<input type="checkbox"/> Connect using a proxy server	<input type="button" value="Proxy Settings"/>
Bucket Name	<input type="text" value="Enter a bucket name"/>

Note: Bucket name will be prefixed with 'arcserve-[agent hostname]'

Amazon S3 Storage Enable Reduced Redundancy Storage

The available options are Amazon S3, Amazon S3-compatible, Windows Azure, Windows Azure-compatible, Fujitsu Cloud (Windows Azure), and Eucalyptus-Walrus. (Amazon S3 is the default vendor).

Note: If you are using Eucalyptus-Walrus as your file copy cloud vendor, you will not be able to copy files whose entire path length is greater than 170 characters.

The configuration options for each cloud vendor are similar (with some different terminology), and any differences are described.

1. From the **Browse File Copies** option or the **Find Files/Folders to Restore** option, click Add.

The **Destination** dialog opens.

2. Select **Cloud** and click **Browse**.

The **Cloud Configuration** dialog opens.

3. Enter the following details:

Storage Name

Specifies the name of the cloud storage. This name will be added to Console for identifying the cloud account. Each cloud account must have a unique storage name.

Storage Service

Select of the service from the drop-down list. The configuration option varies depending on the storage service that is selected.

Access Key ID/Account Name/Query ID

Identifies the user who is requesting access this location.

(For this field, Amazon S3 uses Access Key ID, Windows Azure and Fujitsu 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).

Proxy Settings

Specifies the proxy server settings. Select **Connect using a proxy server** to enable this option. If you select this option, you must also include the IP address (or machine name) of the proxy server and the corresponding port number that is used by the proxy server for internet connections. You can also select this option if your proxy server requires authentication. You then must provide the corresponding authentication information (Domain Name\Username and Password) that is required to use the proxy server.

(Proxy capability is not available for Eucalyptus-Walrus).

Bucket Name

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).

4. Click **Test Connection** to verify the connection to the specified cloud location.
5. Click **OK**.

The cloud account is added to the Console.

Define the Restore Options

After you specify the file copy information to restore, define the copy options for the selected file copy and content.

Follow these steps:

1. On the **Restore Options** dialog, select the restore destination.

The screenshot shows the 'Restore Options' dialog box with the following details:

- Destination:** 'Restore to original location' is selected. There is also a 'Restore to' radio button with an empty text box and a 'Browse' button.
- Resolving Conflicts:** 'Skip existing files' is selected. Other options are 'Overwrite existing files', 'Replace active files', and 'Rename files'.
- Directory Structure:** 'Create root directory' is unchecked.
- File Copy Encryption Password:** A password field is present with asterisks.
- Navigation:** 'Previous', 'Next', 'Cancel', and 'Help' buttons are at the bottom.

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:\Folder-1\SubFolder2\A.txt", "C:\Folder1\SubFolder2\B.txt", and also E:\Folder-3\SubFolder4\C.txt" and during the restore you specified to the restore destination as "D:\Restore".

- If you select to restore just the "A.txt" file, the destination for the restored file will be "D:\Restore\Folder1\SubFolder2\A.txt" (the entire root directory without the volume name will be recreated).
- If you select to restore both the "A.txt" and "C.txt" files, the destination for the restored files will be "D:\Restore\C\Folder1\SubFolder2\A.txt" and "D:\Restore\E\Folder3\SubFolder4\C.txt" (the entire root directory with the volume name will be recreated).

4. Specify the encryption password in **File Copy Encryption Password**.
5. Click **Next**.

The **Restore Summary** dialog opens.


The restore options are defined to restore from a file copy.

Restore the File Copy Content

After you define the restore options, verify that your settings are correct and confirm the restore process. **Restore Summary** helps you to review all the restore options that you defined and modify them if necessary.

Follow these steps:

1. On the **Restore Summary** dialog, review the displayed information to verify that all the restore options and settings are correct.



- ◆ If the summary information is incorrect, click **Previous** and go back to the applicable dialog to change the incorrect setting.
- ◆ If the summary information is correct, click **Finish** to launch the restore process.

The file copy content is restored.

Verify that Content was Restored

After the completion of the restore process, verify that content was restored to the specified destination.

Follow these steps:

1. Navigate to the restore destination you specified.

A list of folders appears.

2. Locate the file to which you have restored the content.

For example, If you select to restore the **A.txt** file to the restore destination as "D:\Restore, then navigate to the following location:

D:\Restore\A.txt.

3. Verify the content to confirm the restore job.

The restored content is successfully verified.

How to Restore From a File Archive

Each time Arcserve UDP performs a successful file archive copy job, it archives all files that have changed since the last successful file archive job. This restore method allows you to browse the archived files and specify exactly which file you want to restore.

The file archive restore process is identical to file copy restore.

Perform the following tasks to restore from a File Archive:

1. [Review the Restore Prerequisites and Considerations](#)
2. [Specify the File Copy Information to Restore](#)
 - a. [Specify the File Copy and Content to Restore](#)
 - ◆ [Specify Cloud Configuration for Restore](#)
 - b. [Define the Restore Options](#)
3. [Restore the Recovery Point Content](#)
4. [Verify that Content was Restored](#)

Review the Restore Prerequisites and Considerations

Verify that the following prerequisites exist before performing a restore:

- You have at least one file copy available to restore.
- You have a valid and accessible file copy destination to restore the file copy content from.
- You have a valid and accessible target location to restore the file copy content to.
- Review the [Compatibility Matrix](#) that provides the supported operating systems, databases, and browsers.

Review the following restore considerations:

- Arcserve UDP only allows one restore job to run at the same time. If you attempt to launch a restore job manually, while another restore job is running, an alert message opens informing 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 (Thread-sForArchive).
 - Restore from a File Copy can download more than one file at a time (Thread-sForRestore).
 - Catalog Synchronization uses multiple threads (ThreadForCatalogSync).

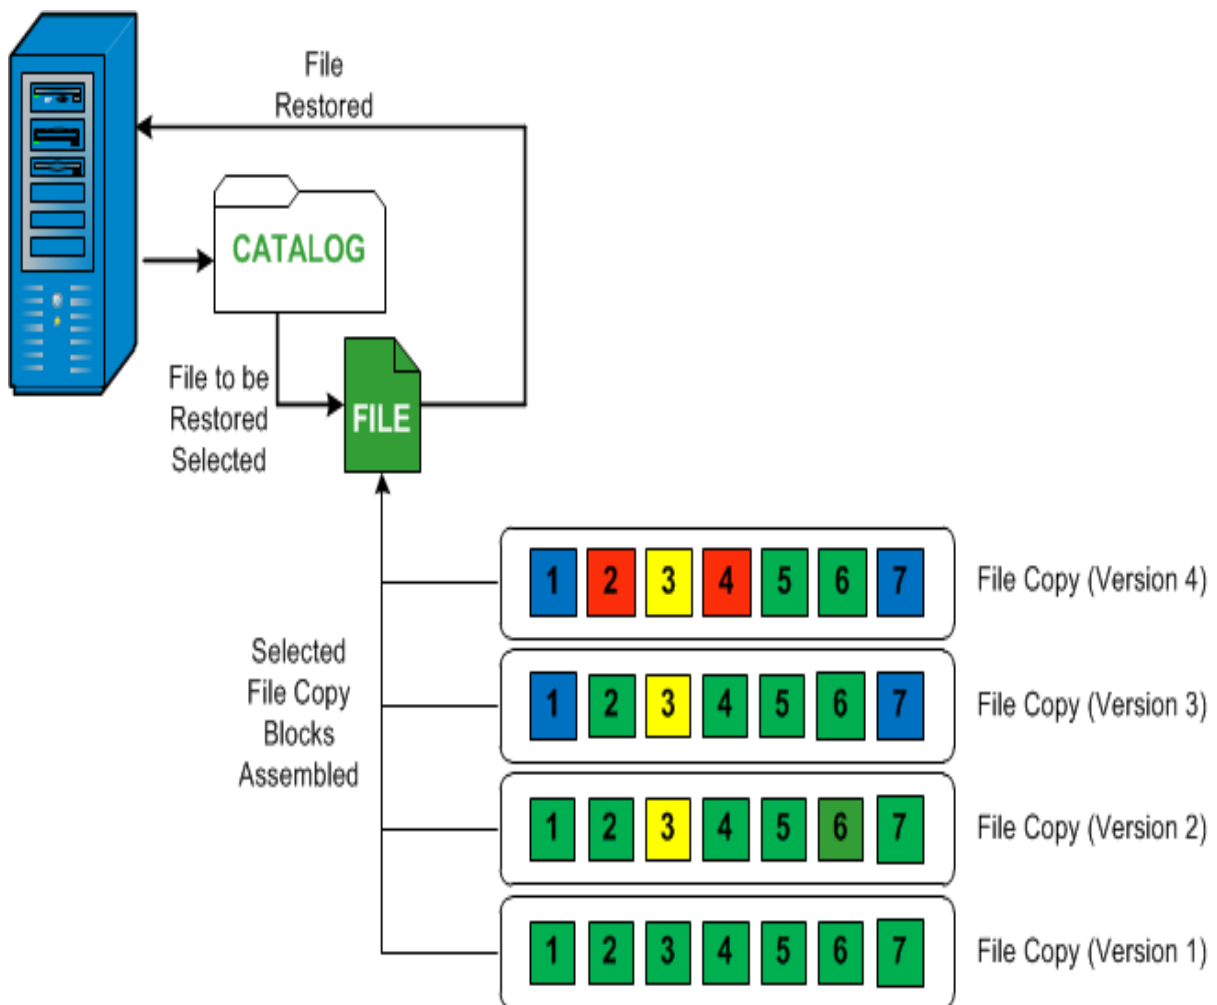
You can change the default File Copy Registry values by modifying the appropriate DWORD value. For more information, see [Configure File Copy Settings to Optimize Performance](#) in the online help.

- (Optional) Understand how the restore process works. For more information, see [How File Level Restores Work](#).

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](#)
2. [Define the Restore Options](#)

Specify the File Copy and Content to Restore

Use the **Browse File Copies** option to restore from a file copy. This restore method allows you to browse the file copied data and specify exactly which file you want to restore.

Follow these steps:

1. Access the restore method selection dialog in one of the following ways:

- ◆ From Arcserve UDP:

- a. Log into Arcserve UDP.
- b. Click the **resources** tab.
- c. Select **All Nodes** in the left pane.
All the added nodes are displayed in the center pane.
- d. In the center pane, select the node and click **Actions**.
- e. Click **Restore** from the **Actions** drop-down list.

The restore method selection dialog opens.

Note: You are automatically logged in to the agent node and the restore method selection dialog is opened from the agent node.

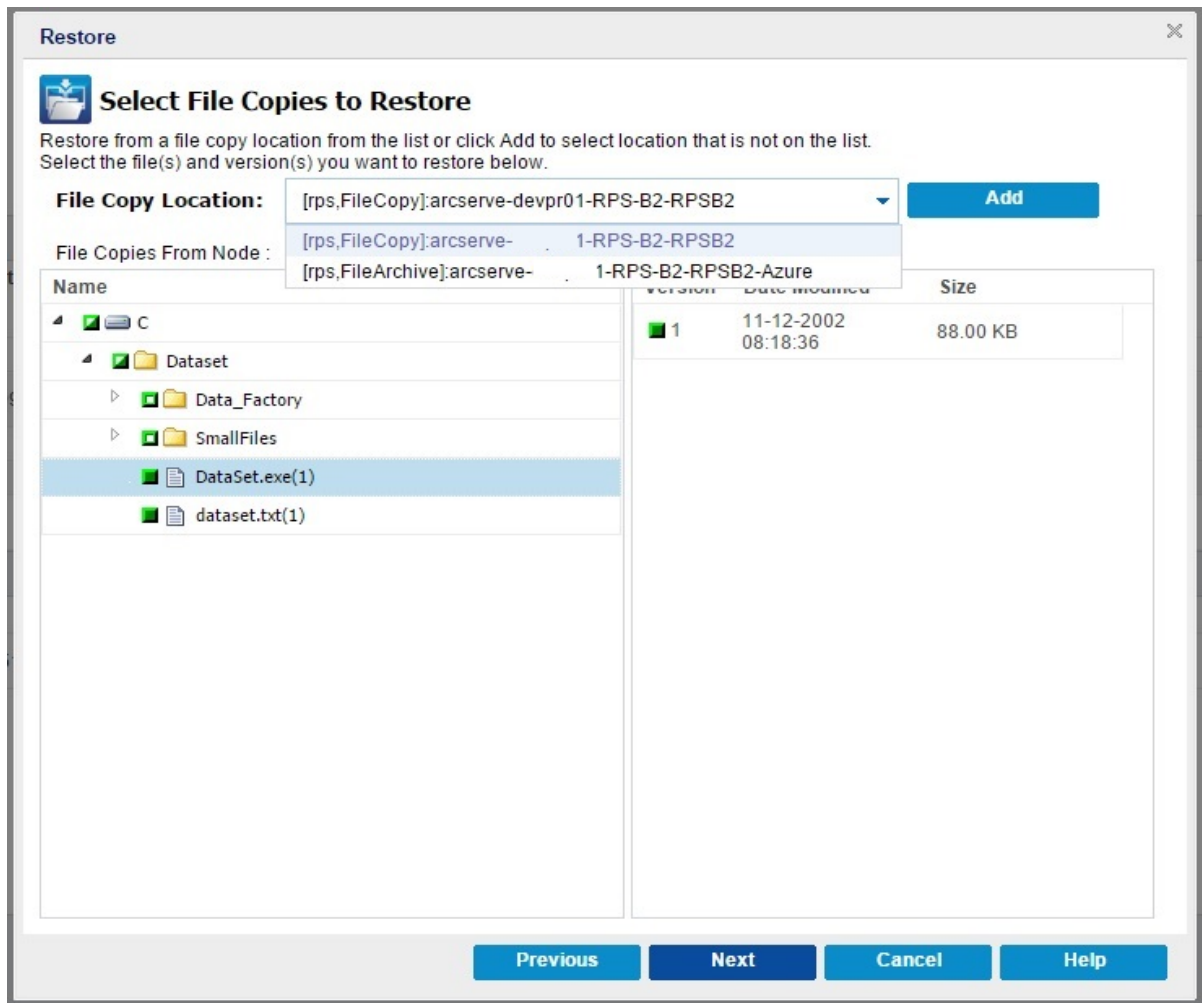
- ◆ From Arcserve UDP Agent (Windows):

- a. Log into Arcserve UDP Agent (Windows).
- b. From the home page, select **Restore**.

The restore method selection dialog opens.


2. Click the **Browse File Copies** option.

The **Restore** dialog opens. The destination that is currently showing in the **Restore From** field is the configured default **File Copy** destination.



- If necessary, you can click **Add** to browse to an alternate location where your file copy images are stored.

The **Destination** dialog opens displaying the available alternate destination options.




Local or network drive

The **Select a Backup Location** dialog opens, allowing you to browse to and select an alternate local or network drive location.

Cloud

The **Cloud Configuration** dialog opens, allowing you to access and select an alternate cloud location. For more information about this dialog, see Specify Cloud Configuration for Restore.

Regardless of whether you selected to restore from **Local or network drive** or from **Cloud**, when you change the destination to an alternate location a pop-up dialog will appear, asking if you want to perform a new catalog synchronization or read from the existing catalog.



- If 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 or file archive cloud location.

Configure the way to access a new cloud storage location

Storage Name	<input type="text" value="Enter a storage name"/>
Storage Service	<input type="text" value="Amazon S3"/> ▼
Bucket Region	<input type="text" value="Select a bucket region"/> ▼
Access Key ID	<input type="text" value="Enter a key ID"/>
Secret Access Key	<input type="text"/>
<input type="checkbox"/> Connect using a proxy server	<input type="button" value="Proxy Settings"/>
Bucket Name	<input type="text" value="Enter a bucket name"/>

Note: Bucket name will be prefixed with 'arcserve-[agent hostname]'

Amazon S3 Storage Enable Reduced Redundancy Storage

The available options are Amazon S3, Amazon S3-compatible, Windows Azure, Windows Azure-compatible, Fujitsu Cloud (Windows Azure), and Eucalyptus-Walrus. (Amazon S3 is the default vendor).

Note: If you are using Eucalyptus-Walrus as your file copy cloud vendor, you will not be able to copy files whose entire path length is greater than 170 characters.

The configuration options for each cloud vendor are similar (with some different terminology), and any differences are described.

1. From the **Browse File Copies** option or the **Find Files/Folders to Restore** option, click Add.

The **Destination** dialog opens.

2. Select **Cloud** and click **Browse**.

The **Cloud Configuration** dialog opens.

3. Enter the following details:

Storage Name

Specifies the name of the cloud storage. This name will be added to Console for identifying the cloud account. Each cloud account must have a unique storage name.

Storage Service

Select of the service from the drop-down list. The configuration option varies depending on the storage service that is selected.

Access Key ID/Account Name/Query ID

Identifies the user who is requesting access this location.

(For this field, Amazon S3 uses Access Key ID, Windows Azure and Fujitsu 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).

Proxy Settings

Specifies the proxy server settings. Select **Connect using a proxy server** to enable this option. If you select this option, you must also include the IP address (or machine name) of the proxy server and the corresponding port number that is used by the proxy server for internet connections. You can also select this option if your proxy server requires authentication. You then must provide the corresponding authentication information (Domain Name\Username and Password) that is required to use the proxy server.

(Proxy capability is not available for Eucalyptus-Walrus).

Bucket Name

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).

4. Click **Test Connection** to verify the connection to the specified cloud location.
5. Click **OK**.


The cloud account is added to the Console.

Define the Restore Options

After you specify the file copy information to restore, define the copy options for the selected file copy and content.

Follow these steps:

1. On the **Restore Options** dialog, select the restore destination.



The available destination options are:

Restore to Original Location

Restores to the original location from where the backup image was captured.

Restore to

Restores to the specified location. You can click the green arrow button to verify the connection to the specified location. If necessary, enter the 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. Specify the encryption password in **File Copy Encryption Password**.
5. Click **Next**.

The **Restore Summary** dialog opens.

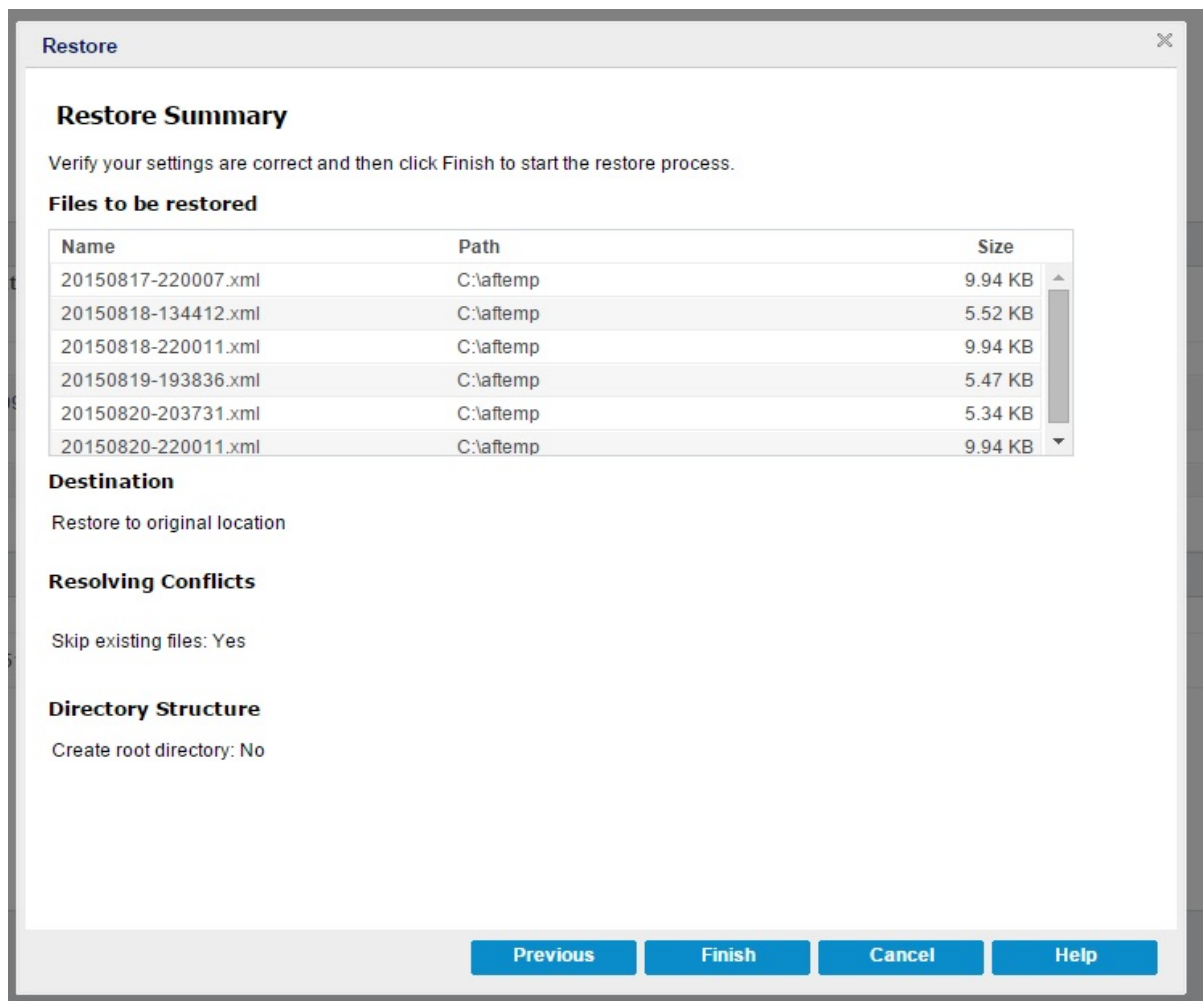
The restore options are defined to restore from a file copy.

Restore the File Copy Content

After you define the restore options, verify that your settings are correct and confirm the restore process. **Restore Summary** helps you to review all the restore options that you defined and modify them if necessary.

Follow these steps:

1. On the **Restore Summary** dialog, review the displayed information to verify that all the restore options and settings are correct.



- ◆ If the summary information is incorrect, click **Previous** and go back to the applicable dialog to change the incorrect setting.
- ◆ If the summary information is correct, click **Finish** to launch the restore process.

The file copy content is restored.

Verify that Content was Restored

After the completion of the restore process, verify that content was restored to the specified destination.

Follow these steps:

1. Navigate to the restore destination you specified.

A list of folders appears.

2. Locate the file to which you have restored the content.

For example, If you select to restore the **A.txt** file to the restore destination as "D:\Restore, then navigate to the following location:

D:\Restore\A.txt.

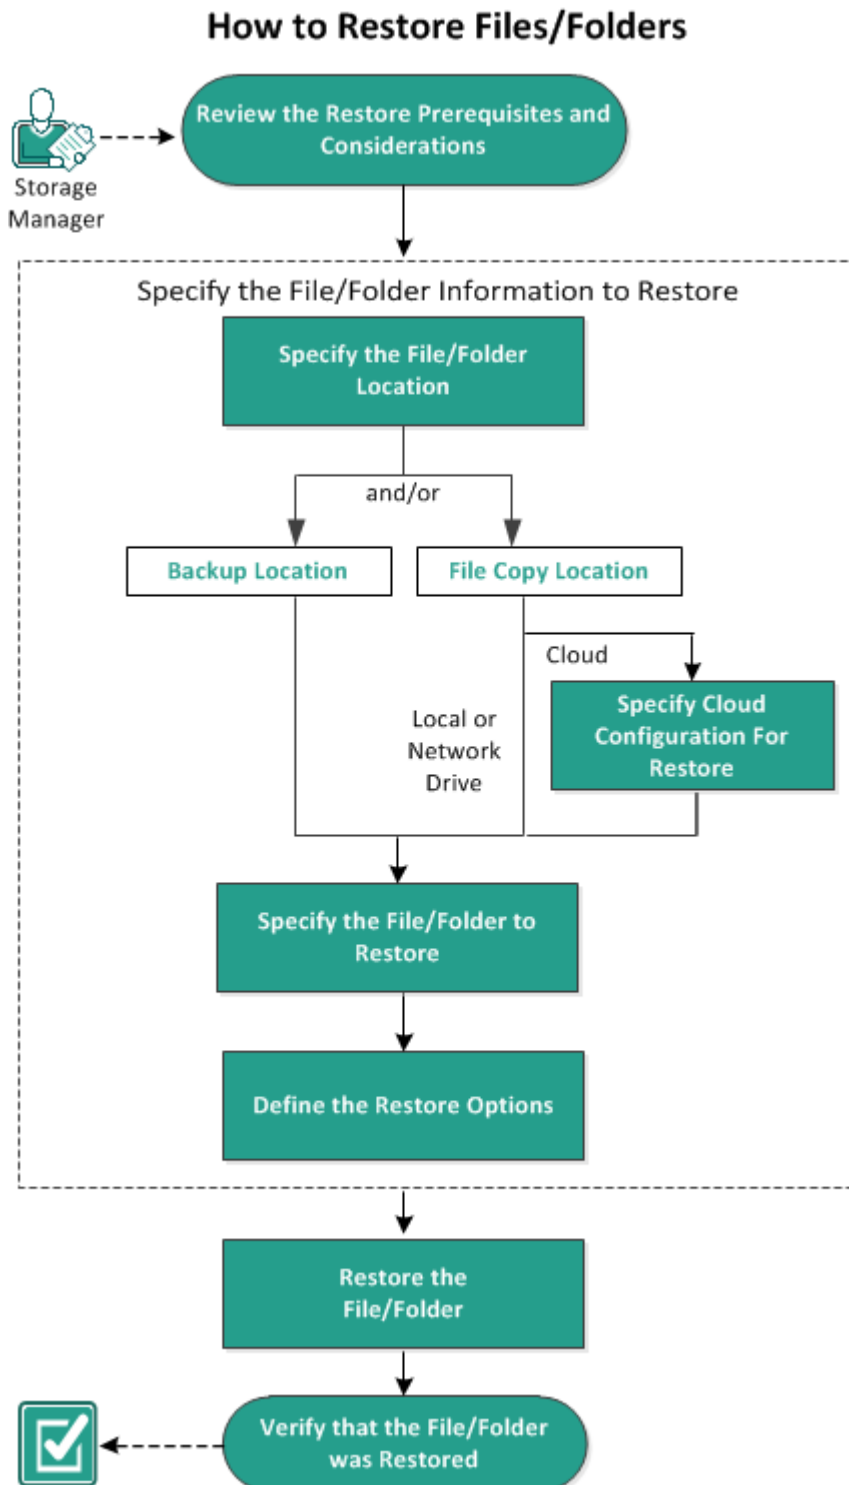
3. Verify the content to confirm the restore job.

The restored content is successfully verified.

How to Restore Files/Folders

Each time Arcserve UDP performs a successful backup, all backed up files/folders are included in the snapshot image of your backup. This restore method allows you to specify exactly which file/folder you want to restore.

The following diagram illustrates the process to restore specific files/folders:



Perform the following tasks to restore files/folders:

1. [Review the Restore Prerequisites and Considerations](#)
2. [Specify the File/Folder Information to Restore](#)
 - a. [Specify the File/Folder Location](#)
 - ◆ [Specify Cloud Configuration for Restore](#)
 - b. [Specify the File/Folder to Restore](#)
 - c. [Define the Restore Options](#)
3. [Restore the File/Folder](#)
4. [Verify that the File/Folder was Restored](#)

Review the Restore Prerequisites and Considerations

Verify that the following prerequisites exist before performing a restore:

- You have at least one backup or file copy version available to restore.
- You have a valid and accessible backup or file copy destination to restore the backup or file copy content from.
- You have a valid and accessible target location to restore the backup or file copy content to.
- Review the [Compatibility Matrix](#) that provides the supported operating systems, databases, and browsers.

Review the following restore considerations:

- For a recovery point without a file system catalog created, to ensure you can browse and select files/folders to restore from the UI, the account/group should be granted access to all the folders/files on all volumes with read/list access before the backup is taken.

The local system (SYSTEM) or built-in administrators group (BUILTIN\Administrators) needs to be added to the ACL of the folders for Arcserve UDP Agent (Windows) to be able to browse a backup without a file system catalog created. Otherwise, Arcserve UDP Agent (Windows) will not be able to browse the folders from the restore UI.

- (Optional) Understand how the restore process works. For more information, see [How File Level Restores Work](#).

Note: The process for restoring from a file copy location is similar to restoring from a backup location.

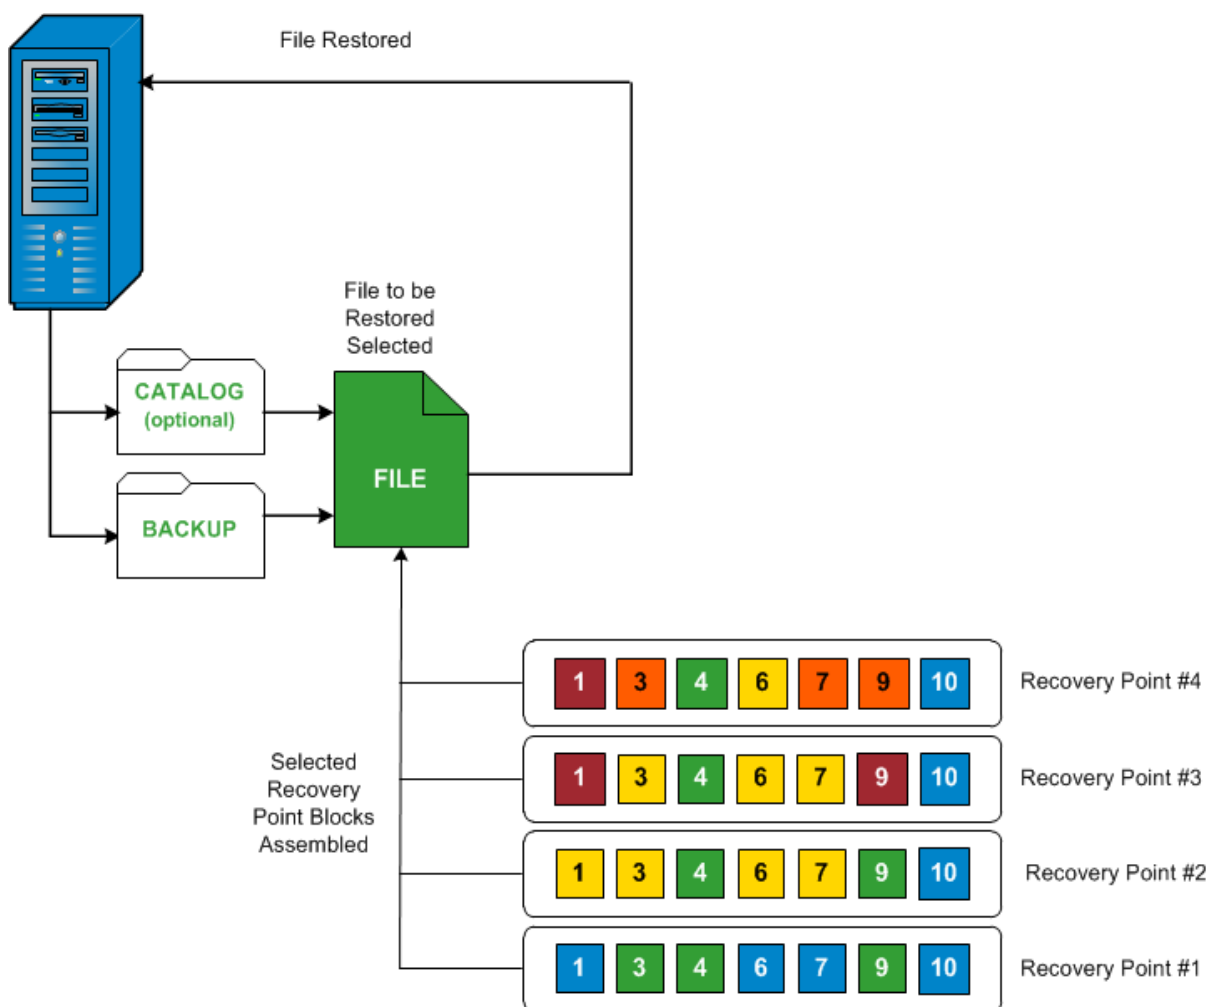
- (Optional) Review the files skipped during restore. For more information, see [Files Skipped During Restore](#).

How File Level Restores Work

During a block-level backup, each backed up file is made up of a collection of blocks that define that particular file. When you need to restore a particular file, you can search your backup and select the file you want to restore and the recovery point you want to restore from. The Arcserve UDP Agent (Windows) then collects the version of the blocks that were used for the recovery point of the specified file, and reassembles and restores the file.

Note: When you specify your backup settings, you have an option to create a file catalog during backup. This file catalog lets you browse the backup sessions faster during restore. If you choose not to create the catalog during backup, it can still be created at a later time.

The following flow diagram shows the process of how Arcserve UDP restores a specific file.



Files Skipped During Restore

During restore by Arcserve D2D some files may be skipped intentionally.

The files and folders in the tables below 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 are skipped when being an important component for Windows or Arcserve D2D.

OS	Folder or Location	File or Folder Name	Remark
All	Root folder of each volume	CAVolTrc.dat	Used by the 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 who 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 contains 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 when you restore to the original or alternate location:

OS	Folder or Location	File or Folder Name	Remark
All	Folder specified in value record under: HKLM\Software\Microsoft\Windows NT\CurrentVersion\WinLogon\SfcDllCache	All files/folders (recursively)	Folder contains a cached dll file

			which is used for System File Checker (SFC) and contents of the system dll cache directory are rebuilt by using SFC.
--	--	--	--

<p>%SystemRoot%\SYSTEM32\dllCache</p>		
<p>Root folder of quorum_device</p>	<p>MSCS*</p>	<p>Used for Microsoft Cluster Server.</p>
<p>%SystemRoot%\SYSTEM32\</p>	<p>perf?00?.dat</p>	<p>Performance data used by the Windows</p>
	<p>perf?00?.bak</p>	<p>performance counter.</p>
	<p>CATROOT*</p>	<p>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.</p>
<p>%SystemRoot%\inetsrv\</p>	<p>metabase.bin</p>	<p>Metabase binary file of earlier IIS versions before 6.0.</p>
<p>File or folder specified in value except "SIS Common Store" under HKLM\SYSTEM\CurrentControlSet\Control\BackupRestore\FilesNotToBackup</p>	<p>All files/folders (recursively)</p>	<p>Do not back up and restore Files and folders . For more inform-</p>

			ation, see the link .
XP W2003	System volume	NTLDR	The main boot loader.
		BOOT.INI	Contains boot configuration (if missing, NTLDR will default to \Windows on the first partition of the first hard drive).
		NTDETECT.COM	Required for booting an NT-based OS. Detects basic hardware information needed for a successful boot.
Vista and later	Root folder of system volume	boot*	Boot folder for Windows.
		bootmgr	Windows boot manager file.
		EFI\Microsoft\Boot*	Used for EFI boot.
	%SystemRoot%\SYSTEM32\	LogFiles\WMI\RTBackup*	Stores ETW trace files (extension .etl) for real time event trace ses-

			sions.
		config\RegBack*	Backup of current registry table.
Win-8 and later	System volume	swapfile.sys	System controller file, normally around 256 MB. It is used by Metro style applications that do not fit the traditional paging characteristics (such as usage pattern, growth, space reservation) of pagefile.sys.
		BOOTNXT	Used to boot from OS, other than Windows 8. Created when enabling the startup options, and updated by Windows.

The Activity log provides the following information:

- Date Time Information: jobxxxx System Files skipped. You can use Bare-Metal Recovery Option (BMR) to restore them.
- Date Time Information: jobxxxx Files or Directories skipped. Which files or directories were skipped can be found in: C:\Program Files\Arcserve\Unified Data Protection\Engine\Logs\Restore-<YYYYMMDD>-<hhmmss>-<Process ID>-<Job ID>.log.

Specify the File/Folder Information to Restore

Arcserve UDP provides you with an option to find and restore a specific file or folder. The goal of running a successful restore job is to identify quickly the data you need and to retrieve it from the appropriate backup location. Each restore job requires a source and destination.

The process involved in restoring by find files/folders is as follows:

1. [Specify the File/Folder Location](#)
 - ◆ [Specify Cloud Configuration for Restore](#)
2. [Specify the File/Folder to Restore](#)
3. [Define the Restore Options](#)

Specify the File/Folder Location

Use the **Find Files/Folders** option to restore files and folders. This restore method allows you to specify exactly which file or folder you want to restore.

Follow these steps:

1. Access the restore method selection dialog in one of the following ways:

From Arcserve UDP:

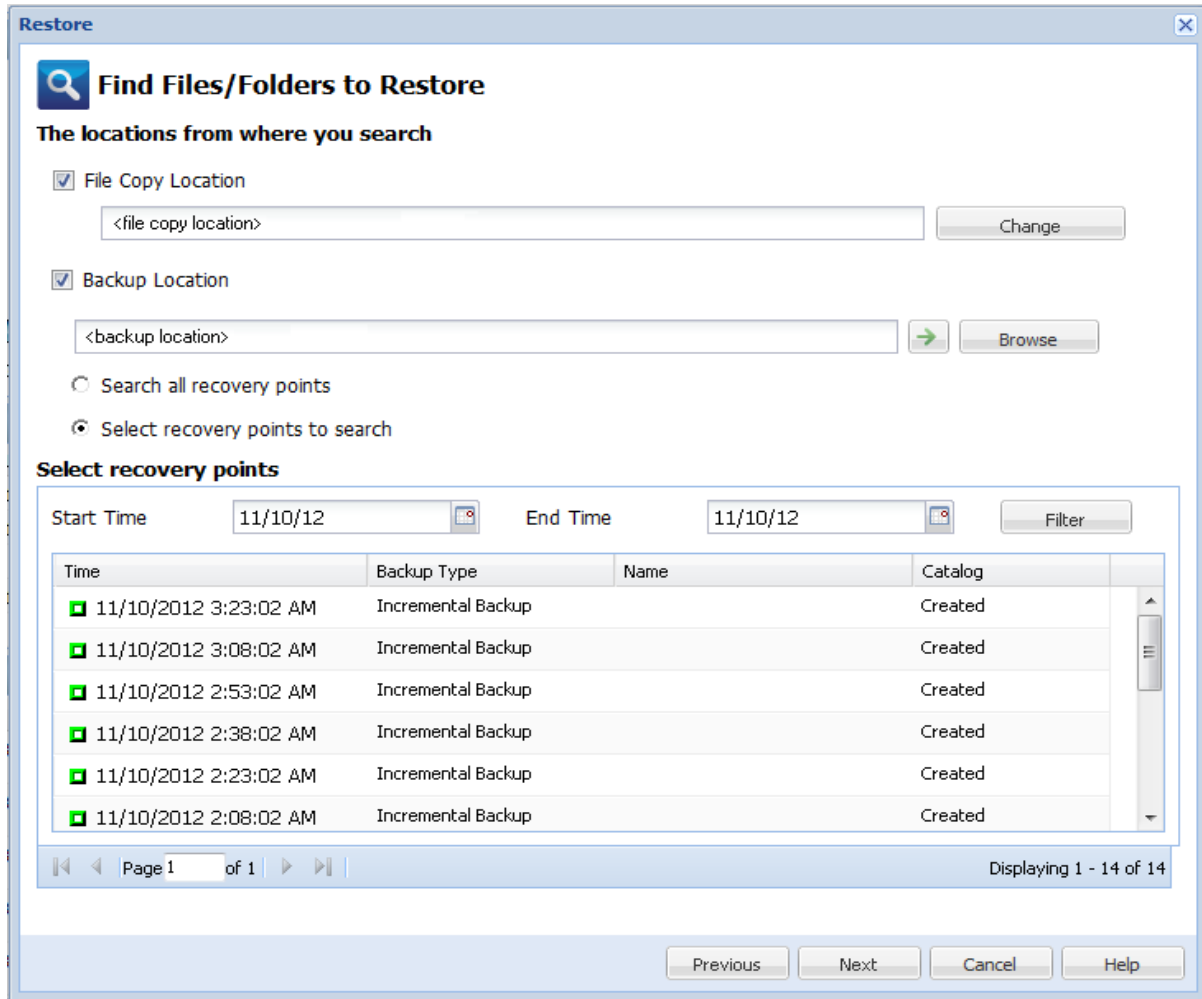
- a. Log into Arcserve UDP.
- b. Click the **resources** tab.
- c. Select **All Nodes** in the left pane.
All the added nodes are displayed in the center pane.
- d. In the center pane, select the node and click **Actions**.
- e. Click **Restore** from the **Actions** drop-down list.

The restore method selection dialog opens.

Note: You are automatically logged in to the agent node and the restore method selection dialog is opened from the agent node.

From Arcserve UDP Agent (Windows):

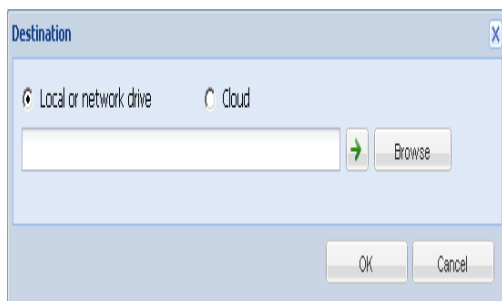
- a. Log into Arcserve UDP Agent (Windows).
 - b. From the home page, select **Restore**.
The restore method selection dialog opens.
2. Click the **Find Files/Folders to Restore** option.
The **Find Files/Folders to Restore** dialog opens.



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**.

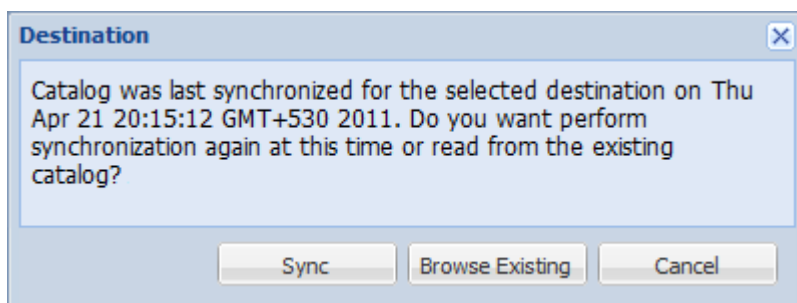
Note: By default, the **Backup Location** and **File Copy Location** fields display the corresponding path used for the most recent backup/file copy destinations.



- ◆ 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](#).

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 Hhttps
 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

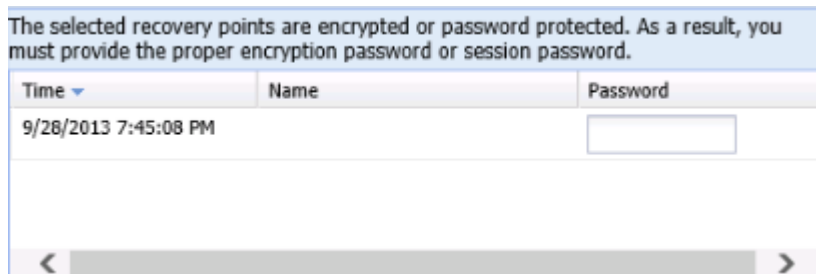
Searches the file or folder in all the recovery points stored in the provided location. You have to specify the file or folder that you want to search on the **Find Files/Folders to Restore** dialog.

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 **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 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.

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.

Note: If you performed the recovery point backup using host-based agentless backup, restoring to original location is to restore the file back in to the virtual machine. In this case, a dialog box opens. You may enter the credentials of the hypervisor, and the operating system of the virtual machine.

For VMware VM:

Set Credential for Source vCenter/ESX Server

vCenter/ESX Server Information

vCenter/ESX Server: abc123-vc

Protocol: HTTP HTTPS


Port Number: 443

User Name: hbbuadmin

Password: ●●●●●●●●

VM Settings

VM Name: shuli02-UEFI

VM username: 

VM password:

OK Cancel

Note: To be able to create or write files inside the VM, consider the following requirements for the settings and account permission of virtual machine:

- VMware Tools is installed and running.
- Firewall must allow File and Printer Sharing.
- The account is built-in local administrator, built-in domain administrator, or domain account that is member of the local Administrators group. If other accounts are used:
 - Disable the UAC remote access. To disable UAC remote access, see [Import Virtual Machine Using Additional Administrative Account](#).
 - Disable UAC in the Local Security Policy by disabling the setting Run all administrator in Admin Approval Mode at secpol.msc -> Local Policies -> Security Options. (Secpol.msc is Microsoft's security policy editor).

Important: Do not attempt to disable the UAC in the User Account Control Settings dialog box that opens from the control panel.

For VMware VM:

Set the credentials for the source Hyper-V Server ✕

Hyper-V Server Information

Hyper-V/Hyper-V Cluster Server: abc123 -hyperv1

User Name: administrator

Password: ●●●●●●

VM Settings

VM Name: abc123-hv102

VM username:

VM password:

OK Cancel

Note: To be able to create or write files inside the VM, consider the following requirements for the settings and account permission of virtual machine:

- Hyper-V integration services are installed and running.
- Firewall must allow File and Printer Sharing.
- The account is built-in local administrator, built-in domain administrator, or domain account that is member of the local Administrators group. If other accounts are used:
Disable the UAC remote access. To disable UAC remote access, see [Import Virtual Machine Using Additional Administrative Account](#).
- If virtual machine guest OS is Client version Windows (such as Windows 10), you need to manually configure firewall to allow Windows Management Instrumentation (WMI).

Restore to

Restores to the specified location. You can click the green arrow button to verify the connection to the specified location. If necessary, enter the 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 file name but a different extension. Data is then restored to the new file.

Skip existing files

Skips over and not overwrite (replace) any existing files that are located at the restore destination. Only objects that are not currently existing on your machine are restored from the backup files.

Default: Skip existing files

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:



Perform the following tasks to restore a virtual machine:

1. [Review the Restore Prerequisites and Considerations](#)
2. [Specify the Virtual Machine Information to Restore](#)
 - a. [Specify the Virtual Machine and the Recovery Point to Restore](#)
 - b. [Define the Restore Options](#)
 - ◆ [Define the Original Location Restore Options](#)
 - ◆ [Define the Alternate Location Restore Options](#)
3. [Restore the Virtual Machine](#)
4. [Verify that the Virtual Machine was Restored](#)

Review the Restore Prerequisites and Considerations

Verify that the following prerequisites exist before performing a restore:

- You have a valid recovery point available to restore from.
- You have a valid and accessible target Virtual Center/ESX or Hyper-V server to recover the virtual machine.
- Review the [Compatibility Matrix](#) that provides the supported operating systems, databases, and browsers.

Review the following restore considerations:

- Arcserve UDP allows multiple restore jobs to run at the same time if the jobs are not for the same VM. If you attempt to launch a restore job, while another restore job is running for the same VM, an alert message informs you that another job is running and requests you to try again later.
- If the Recover VM destination is Windows Server 2008 R2 then the source backup VM should not contain VHDx disks, which are not supported on the Hyper-V server (Windows Server 2008 R2).
- If the Recover VM destination is Windows Server 2008 R2 or Win2012 then the source backup VM's sub-system type should not be generation 2 (which was introduced in Windows Server 2012 R2), and is not supported on the Hyper-V server (Windows Server 2012/2008 R2).

Specify the Virtual Machine Information to Restore

You can recover an entire virtual machine from a recovery point.

The process involved in restoring virtual machine is as follows:

1. [Specify the Virtual Machine and the Recovery Point to Restore](#)
2. [Define the Restore Options](#)
 - ◆ [Define the Original Location Restore Options](#)
 - ◆ [Define the Alternate Location Restore Options](#)

Specify the Virtual Machine and the Recovery Point to Restore

Use the **Recover VM** option to restore a virtual machine that you previously backed up. This method quickly and consistently creates a virtual machine from an Arcserve UDP recovery point on an ESX or Hyper-V server. The recovered virtual machine can then simply be started to complete the recovery process.

Follow these steps:

1. Access the restore method selection dialog in one of the following ways:

From Arcserve UDP:

- a. Log into Arcserve UDP.
- b. Click the **resources** tab.
- c. Select **All Nodes** in the left pane.
All the added nodes are displayed in the center pane.
- d. In the center pane, select the node and click **Actions**.
- e. Click **Restore** from the **Actions** 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 into Arcserve UDP Agent (Windows).
- b. From the home page, select **Restore**.

The restore method selection dialog opens.

2. Click the **Recover VM** option.

The **Recover VM** dialog opens.

Restore
✕

Recover VM

Backup Location

Recovery Point Server: <recovery_point_server_name> Change

Data Store: datastore1

Node: <virtual_machine_name>

Node

Select Node: <virtual_machine_name>

Recovery Point Date

November 2016

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

6:00 AM - 12:00 PM (2)

12:00 PM - 6:00 PM (1)

6:00 PM - 12:00 AM

AR	Time	Type	Backup Type	Name
✔	2:44:01 PM	Custom / Manual	Incremental	Customized Incremental Backup

Name	Date Modified	Size
▶ C:		59.48 GB
▶ E:		9.97 GB
▶ Volume{00972F93-DFF6-4F94-8DAC-		300.00 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.

- Select one of the following options:

Select local disk or shared folder

- 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.

- 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.

- 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**.
- b. All the nodes (agents/virtual machines) are listed in the Node column in the **Source** dialog.
- c. 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. On the Restore Options dialog, select the restore destination.

The screenshot shows the 'Restore' dialog box with the following content:

- Select the restore destination**
 - Restore to original location
 - Restore to an alternative location
- Use selected RPS server network for restore traffic
 - Note: The network is between Proxy Server and Recovery Point Server.
 - Dropdown menu
 - Continue to run job even when unable to connect to the selected destination network
 - Use dedicated ethernet if current machine enables SMB Multichannel
- Use selected Proxy server network for restore traffic
 - Note: The network is between Proxy Server and Hypervisors. To enable the use of the specified network for backup traffic, both backup proxy and the Nutanix Cluster ISCSI Data Service must be in the same network.
 - Dropdown menu
 - Continue to run job even when unable to connect to the selected destination network
 - Use dedicated ethernet if current machine enables SMB Multichannel

The available destination options are:

Restore to original location

Restores the virtual machine to the original location from where the backup image was captured. By default, this option is selected.

For more information, see [Define the Original Location Restore Options](#).

Restore to an alternative location

Restores to the virtual machine to a different location from where the backup image was captured.

For more information, see [Define the Alternate Location Restore Options](#).

2. (Optional) Select the **Use selected RPS server network for restore traffic** check box, and then do the following:
 - a. To enable the communication between Windows Proxy server and Recovery Point Server, select the CIDR network from the drop-down list.

- b. To define the constraint on SMB Multichannel so that the data transfers only through the selected network, select the **Use dedicated ethernet if current machine enables SMB Multichannel** check box.

Note: This option is not available by default. To enable this option, create the *UseDedicatedEthernet* string registry in the following path, and then set the registry value to 1:

SOFTWARE\Arcserve\Unified Data Protection\Engine

3. (Optional) Select the **Use selected Proxy server network for restore traffic** check box, and then do the following:

- a. To enable the communication between Windows Proxy server and Hyper-visor, select the CIDR network from the drop-down list.
- b. To define the constraint on SMB Multichannel so that the data transfers only through the selected network, select the **Use dedicated ethernet if current machine enables SMB Multichannel** check box.

Note: This option is not available by default. To enable this option, create the *UseDedicatedEthernet* string registry in the following path, and then set the registry value to 1:

SOFTWARE\Arcserve\Unified Data Protection\Engine

4. Specify the Resolving Conflicts options that Arcserve UDP performs if conflicts are encountered during the restore process

Overwrite existing Virtual Machine

This option is to specify whether to overwrite the existing virtual machine. By default, this overwrite option is not selected.

Note: For the Overwrite existing Virtual Machine option, an "existing virtual machine" is defined as a VM that has the same VM name and resides in the same ESXi host (for VMware VM), or a VM that has the same VM name and instance UUID and resides in the same Hyper-V host (for Hyper-V VM). For VMware VM, if there is a VM which has the same VM name but resides in a different ESXi host (which is under the same vCenter), the overwrite option does not work. In this case, VM recovery GUI detects that VM and displays an error message and blocks you from proceeding so that a VM is not overwritten by mistake. As a workaround, you need to rename the existing VM or use the "Restore to alternative location" option, and then specify a different VM name.

- ◆ If you select this option, the restore process overwrites (replaces) any existing images of this virtual machine that are at the specified restore destination. The virtual machine image is restored from the backup files

regardless of its current presence on your restore destination.

- ◆ If you do not select this option, VM recovery GUI displays an error message and blocks you from proceeding if the original VM still exists on the original location. As a workaround, you need to rename the existing VM or use the "Restore to alternative location" option, and then specify a different VM name.

Generate new Virtual Machine instance UUID

This option is to specify whether to generate a new instance UUID for the restored VM or keep the original instance UUID.

Note: If you do not select this option, the original instance UUID is set to the restored VM. However, in case the destination vCenter/ESX or Hyper-V host already has a VM with the same instance UUID, new UUID is used instead, and a warning message is displayed in the activity log of VM recovery job.

5. Specify the Post Recovery option.

Power on Virtual Machine

Select whether power is applied to the virtual machine at the end of the restore process. By default, this option is not selected.

Mark as VM Template (available only for VMware VM)

Select whether to convert restored VM to template. If source node is VM when backed up, this option is not selected by default. If source node is template when backed up, this option is selected by default.

The restore options are defined to restore a virtual machine.

Define the Original Location Restore Options

During the Recover VM configuration process, you are required to select the option of where you want to restore the virtual machine to. The available selections are **Restore to the Original Location** and **Restore to an Alternative Location**.

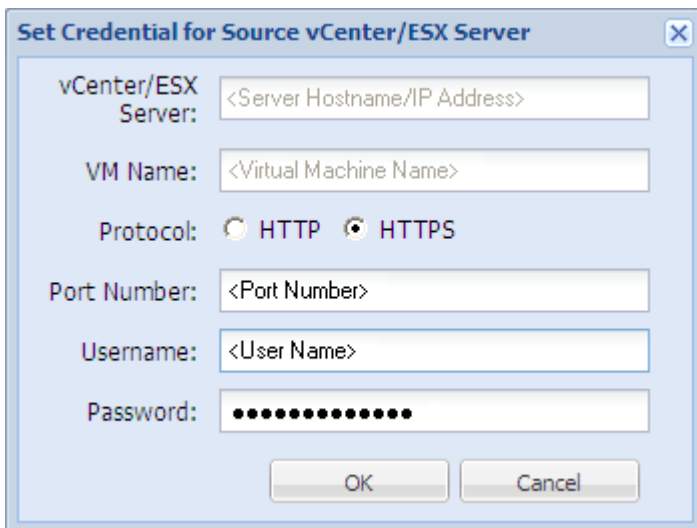
This procedure explains how to restore a virtual machine to the original location.

Follow these steps:

1. From the **Restore Options** dialog, after specifying the **Resolve Conflicts** and **Post Recovery** options, select **Restore to Original Location** and click **Next**.

The appropriate dialog for VMware or Hyper-V is displayed.

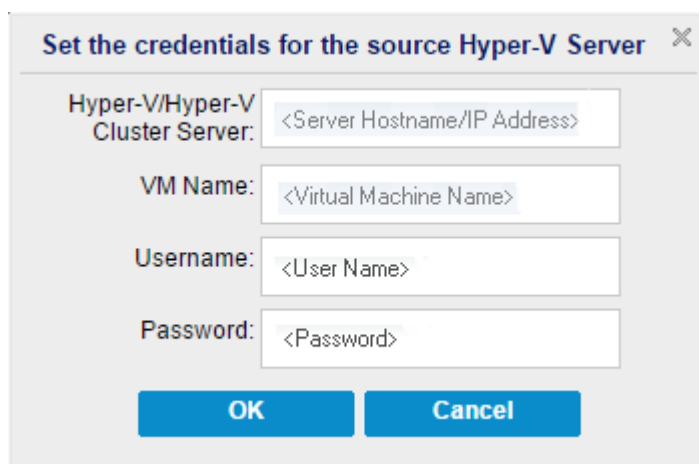
- ◆ For VMware the **Set Credential for Source vCenter/ESX Server** dialog is displayed.



The screenshot shows a dialog box titled "Set Credential for Source vCenter/ESX Server". It contains the following fields and controls:

- vCenter/ESX Server:** A text input field with the placeholder text "<Server Hostname/IP Address>".
- VM Name:** A text input field with the placeholder text "<Virtual Machine Name>".
- Protocol:** Two radio buttons labeled "HTTP" and "HTTPS". The "HTTPS" radio button is selected.
- Port Number:** A text input field with the placeholder text "<Port Number>".
- Username:** A text input field with the placeholder text "<User Name>".
- Password:** A text input field with a masked password represented by a series of black dots.
- At the bottom, there are two buttons: "OK" and "Cancel".

- ◆ For Hyper-V the **Set the credentials for the source Hyper-V Server** dialog is displayed.



The screenshot shows a dialog box titled "Set the credentials for the source Hyper-V Server". It contains the following fields and controls:

- Hyper-V/Hyper-V Cluster Server:** A text input field with the placeholder text "<Server Hostname/IP Address>".
- VM Name:** A text input field with the placeholder text "<Virtual Machine Name>".
- Username:** A text input field with the placeholder text "<User Name>".
- Password:** A text input field with the placeholder text "<Password>".
- At the bottom, there are two buttons: "OK" and "Cancel".

2. Specify the credentials for accessing the virtual machine.

For VMware, complete the following fields.

vCenter/ESX Server

Displays the host name or IP address for the destination vCenter Server or ESX Server system.

Note: You cannot edit this field. You can only view the details.

VM Name

Displays the virtual machine name that you are restoring.

Note: You cannot edit this field. You can only view the details.

Protocol

Specifies the protocol that you want to use for communication with the destination server. The available selections are HTTP and HTTPS.

Port Number

Specifies the port that you want to use for data transfer between the source server and the destination.

Default: 443.

Username

Specifies the user name that has access rights to log in to the vCenter/ESX server where you plan to restore the virtual machine.

Password

Specifies the corresponding password for the User Name.

For Hyper-V, complete the following fields.

Hyper-V/Hyper-V Cluster Server

Displays the host name or IP address for the destination Hyper-V Server or Hyper-V cluster server system.

Note: You cannot edit this field. You can only view the details.

VM Name

Displays the virtual machine name that you are restoring.

Note: You cannot edit this field. You can only view the details.

Username

Specifies the user name that has access rights to log in to the Hyper-V server where you plan to restore the virtual machine. For Hyper-V cluster VM, specify the domain account which has administrative privilege of the cluster.

Password

Specifies the corresponding password for the User Name.

3. Click **OK**.

The **Restore Summary** dialog opens.

The restore options for original location are defined.

Define the Alternate Location Restore Options

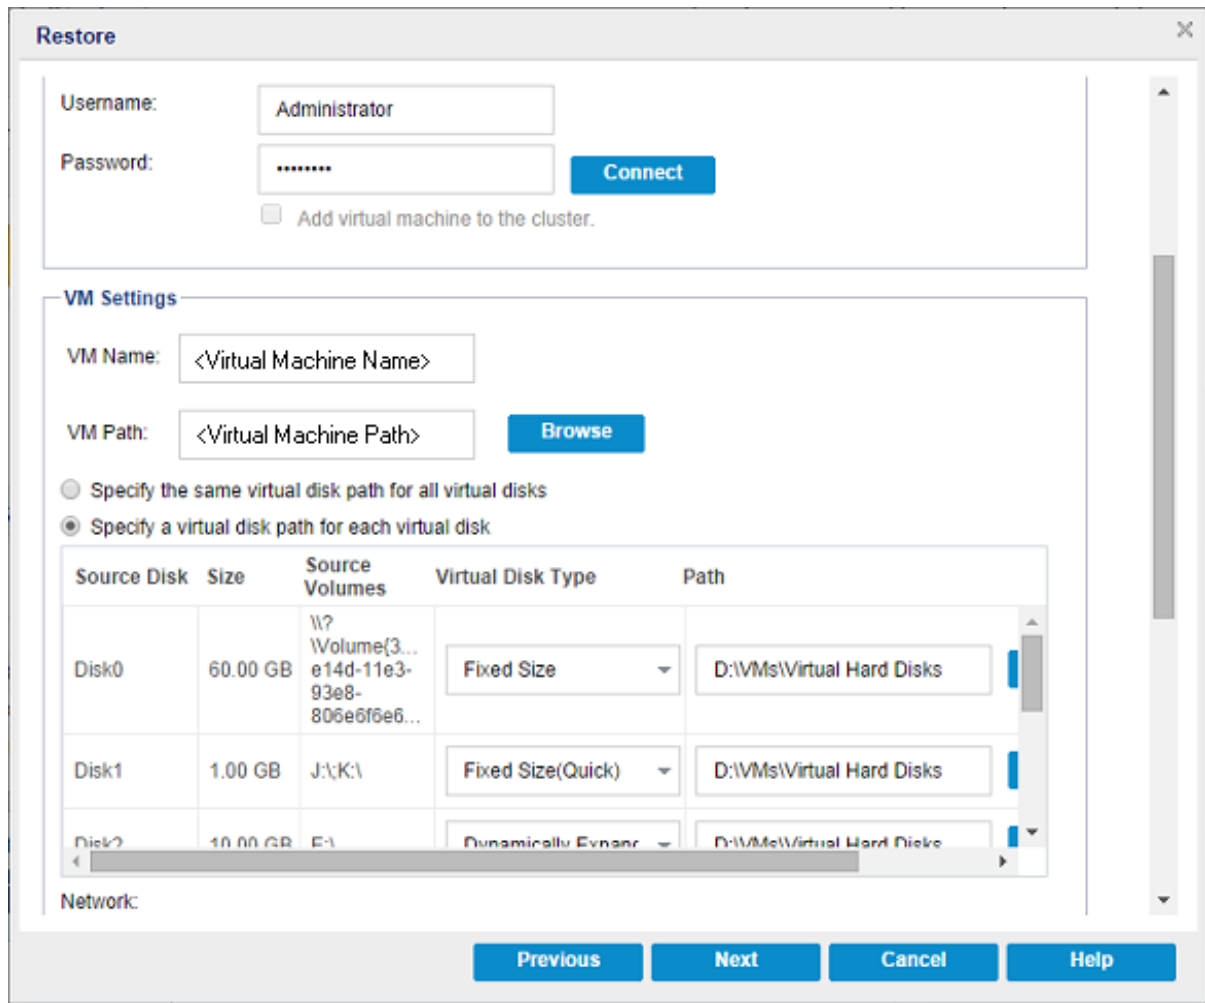
During the Restore VM configuration process, specify where the recovered virtual machine is stored. The available selections are **Restore to the Original Location** and **Restore to an Alternative Location**.

This procedure explains how to restore a virtual machine to alternate location or different data store.

Follow these steps:

1. From the **Restore Options** dialog, after specifying the **Resolve Conflicts** and **Post Recovery** options, select **Restore to an Alternative Location**.
 - ◆ For VMware, the **Restore Options** dialog expands to display additional restore to alternative options.
 - ◆ For Hyper-V, the **Restore Options** dialog expands to display additional restore to alternative options.

If you select the **Specify a virtual disk path for each virtual disk** option, the following dialog appears:



2. Specify the appropriate server Information.

For VMware, enter the following fields:

vCenter/ESX Server

Specifies the host name or IP address for the destination vCenter or ESX server system.

Username

Specifies the user name that has access rights to log into 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: HTTPS.

Note: VMware Virtual Disk Development Kit (VDDK) 6.x.x is in-built with Arcserve UDP 7.0 but VDDK 6.x.x does not support HTTP. Make sure to select HTTPS, unless you manually replace the built-in VDDK 6.x.x with another version of VDDK.

Port Number

Specifies the port that you want to use for data transfer between the source server and the destination.

Default: 443.

For Hyper-V, enter the following fields:

Hyper-V Server

Displays the host name or IP address for the destination Hyper-V Server system.

Username

Specifies the user name that has access rights to log into 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 list contains a listing of all ESX servers that are associated with a vCenter server.

Resource Pool


Selects the **Resource Pool** or **vApp Pool** that you want to use for the virtual machine recovery.

Note: A Resource Pool is a configured collection of CPU and memory resources. A vApp Pool is a collection of one or more virtual machines that can be managed as a single object.

Default: empty.

Click the **Browse Resource Pool** button to display the **Select a Resource Pool** dialog. This dialog contains a listing of all Resource Pools and vApp Pools available for the destination ESX server. Select the pool to use for the virtual machine recovery. You can leave this field blank when you do not

want to assign a Resource Pool or vApp Pool to this virtual machine recovery.



Storage Policy

Specify the VM storage policy that is applied to VM home of restored VM. Select Datastore Default if you do not want to apply the VM storage policy.

Note: If you can see only Datastore Default but actually vCenter has other storage policies defined, then the account used to connect vCenter does not have enough permission to get the storage policy from vCenter. Verify if the account has the privilege Profile-driven storage view at the vCenter level.

VM DataStore

Specify the destination datastore for VM home of restored VM.

Note: By default, only those datastores that are compatible with selected storage policy are listed. If you want to see all datastores, clear selection of the checkbox **Show only compatible datastores for selected storage policy** that is under the Disk Datastore table.

Disk Datastore

Specify Virtual Disk Type, Storage Policy and Target Datastore for each of the virtual disk of the VM, respectively.

- **Virtual Disk Type:** Select one of the following options: Thin, Thick Lazy Zeroed, or Thick Eager Zeroed.

- **Storage Policy:** Select the VM storage policy that is applied to this virtual disk. Select Datastore Default if you do not want to apply the VM storage policy.
- **Target Datastore:** Select the datastore where the virtual disk is restored.

Note: By default only the datastores that are compatible with selected storage policy are listed. If you want to see all datastores, clear selection of the checkbox **Show only compatible datastores for selected storage policy** that is under the Disk Datastore table.

Network

Specifies the vSphere Standard Switch/vSphere Distributed Switch configuration details.

For Hyper-V, enter the following fields.

VM Name

Specifies the virtual machine name that you are restoring.

VM Path

Specifies the destination path (on Hyper-V server) where to save the Hyper-V VM configuration file. The default folder of the VM configuration file for the Hyper-V server is shown by default. You can modify the path directly in the field or click **Browse** to select one.

Note: If you are restoring the virtual machine into Hyper-V cluster and you want the virtual machine to migrate among the cluster nodes, specify the cluster shared volume (CSV) for both- the VM path and the virtual disk path.

Specify the same virtual disk path for all virtual disks

Specify one path (on Hyper-V server) where to save all virtual disks of the VM together. The default folder of the VM disk file for the Hyper-V server is shown by default. You can modify the path directly in the field or click **Browse** to select one.

Note: If you are restoring the virtual machine into Hyper-V cluster and you want the virtual machine to migrate among the cluster nodes, specify the cluster shared volume (CSV) for both- the VM path and the virtual disk path.

Specify a virtual disk path for each virtual disks

Specify the path (on Hyper-V server) for each of the virtual disks of the VM respectively. The default folder of the VM disk file for the Hyper-V server is shown by default. You can modify the path directly in the field or click **Browse** to select one. To assign the virtual disk type, select one of the fol-

lowing 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:

On the **Restore Summary** dialog, review the displayed information to verify that all the restore options and settings are correct.

- ◆ If the summary information is incorrect, click **Previous** and go back to the applicable dialog to change the incorrect setting.
- ◆ If the summary information is correct, click **Finish** to launch the restore process.

The virtual machine is restored.

Verify that the Virtual Machine was Restored

After the completion of the restore process, verify that the virtual machine was restored to the specified destination.

Follow these steps:

1. Navigate to the restore destination you specified.

For example, if you select to restore the virtual machine to the restore destination as original location, then log in to the original vCenter/ESX or Hyper-V Server and check if the virtual machine exists.

If you select to restore the virtual machine to the Alternate location, then log in to the alternate vCenter/ESX or Hyper-V Server provided in the restore options and check if the virtual machine exists.

2. Verify the virtual machine was restored.

The virtual machine is restored successfully.

How to Use Exchange Granular Restore (GRT)

This section provides the following information:

Introduction

The Exchange Granular Restore utility is used to restore Microsoft Exchange email and non-email objects. The utility includes the injection capability for items, such as emails, from offline databases (*.EDB) and log files to the original live Exchange databases, as well as granular data extraction to Personal Storage File (.pst) files.

This utility includes the following key benefits:

- Supports non-email items (for example, Calendar, Contacts, Tasks) and public folders.
- Can work with just a database file as well. Logs are not mandatory, but having them will ensure more recent data available for restore.
- It does not need to generate a catalog and directly restores the mail from the mounted recovery point.
- Takes a minimum amount of time to restore a mailbox level item from a database or user mailbox of any size.
- Supports the command line options to process several databases.

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

Review the Prerequisites and Considerations

Verify that the following prerequisites exist before performing a restore:

- The Exchange Granular Restore utility is available at the following location:
The tool is installed with the Arcserve UDP Agent under the following directory:

X:\Program Files\Arcserve\Unified Data Protection\Engine\Exchange GRT

Note: The tool is installed with the Arcserve UDP Agent.

- The Restore job is set to run from the Exchange machine or HBBU proxy machine.

Note: If you want to run the restore job on any other machine, search the recovery point from the backup destination.

- The database name, Server name, path to database (.edb), and the log files of the user are identified to perform the restore job.

To identify, use the Exchange Management Console (EMC), Exchange Control Panel (ECP), or Exchange Management Shell.

For example:

```
Get-Mailbox -identity "username" | fl Database
```

```
Get-MailboxDatabase -identity "Databasename" | fl Name, Server,  
EdbFilePath,LogFolderPath
```

More information:


[How to Restore Microsoft Exchange Data Using Exchange Granular Restore \(GRT\) Utility](#)

Restoring Microsoft Exchange Data Using Exchange Granular Restore (GRT) Utility


Before you begin, [review the prerequisites and considerations](#).

Perform the following tasks to restore Microsoft Exchange mailbox items, using the Exchange Granular Restore utility:

1. From the Arcserve UDP Agent console, select the [Mount Recovery Point](#) task (recommended) or [restore the Exchange database](#) to the local drive. The Mount Recovery Point dialog opens.



2. Select the recovery point date and click **Mount** for the volume(s) that contain Exchange Database and logs.



Note: If the server that is running the restore job is not the Exchange or HBBU proxy, click **Change** to select the appropriate Recovery Point Server, Data Store, and Exchange Server.

3. Select the drive letter to mount the volume and click **OK**.

Mount Recovery Point

Mount Recovery Point to volume or path

Mount to the following volume: Z: ▼


Mount to the following path (path must be an empty NTFS or ReFS folder):

Backup Encryption or Protection Password

The data you are trying to mount is encrypted. You must provide an encryption password to mount.


Password

4. Launch the Exchange Granular Restore utility from one of the following locations:
Start > All Programs > Arcserve > Unified Data Protection > Arcserve UDP Exchange Granular Restore
or
X:\Program Files\Arcserve\Unified Data Protection\Engine\Exchange GRT\esr.exe
A dialog appears to specify the path for the database and log files.
5. Specify the path to the mounted volume and click **Open**.



The Arcserve UDP Exchange Granular Restore utility opens.

6. Select the user data to restore and click **Export into original mailbox** or **Export into .PST**.



Notes:

- For more details on the supported specifications, features, user options and limitations, see the Exchange Granular Restore user guide (esr.pdf), located at:

%ProgramFiles%\Arcserve\Unified Data Protection\Engine\Exchange GRT or [Exchange Granular Guide](#)

- By default, the utility uses the current user who is logged in to Windows to establish the connection. If the current user does not have permissions to impersonate the selected user, an error message appears in the **Details** pane.

If an error is reported, the recommended action to take is to log in to the machine with an account that has impersonation rights for the selected user or the account of the selected user.

- When the restore job completes, dismount the volume that was used for the recovery.

To dismount the volume, from the Arcserve UDP Agent console, click **Mount Recovery Point** and then click **Dismount**.



How to Restore Microsoft Exchange Data

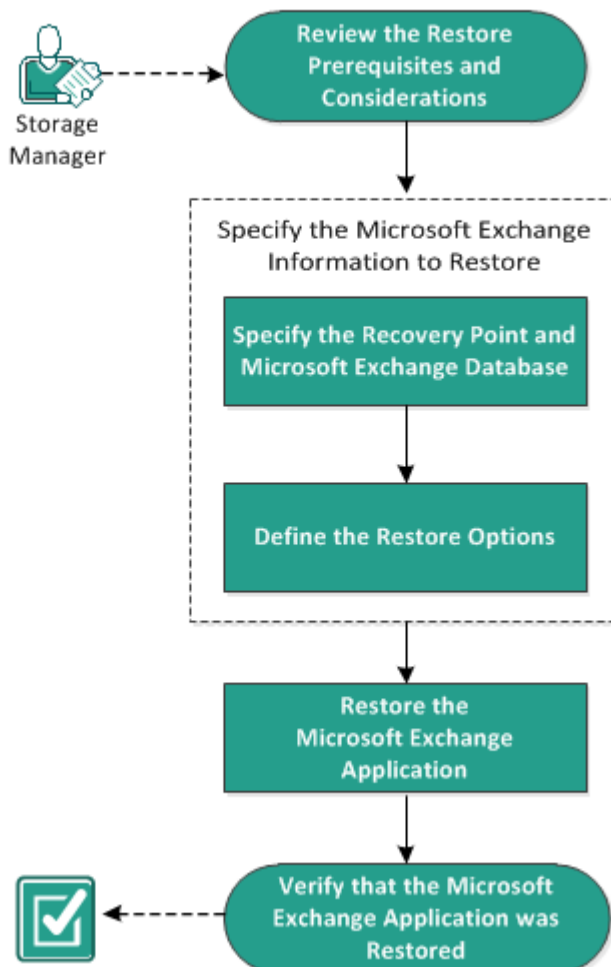
Important! To restore Microsoft Exchange data, it is suggested to use the [Exchange Granular Restore utility](#).

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](#)
2. [Specify the Microsoft Exchange Information to Restore](#)
 - a. [Specify the Recovery Point and Microsoft Exchange Database](#)
 - b. [Define the Restore Options](#)
3. [Restore the Microsoft Exchange Application](#)
4. [Verify that the Microsoft Exchange Application was Restored](#)

Review the Restore Prerequisites and Considerations

Arcserve UDP Agent (Windows) supports the following versions of Microsoft Exchange Server:

- Microsoft Exchange 2010 - Single Server Environment and Database Availability Group (DAG) environment.
- Microsoft Exchange 2013 and 2016 - Single Server Environment and Database Availability Group (DAG) environment.

For Microsoft Exchange Server 2010, 2013, and 2016 DAG environment, Arcserve UDP Agent (Windows) must be installed on all member servers in the DAG group. A backup job can also be performed from any member server for both active and passive database copies, but restore can only be performed to an active database copy.

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, 2013, and 2016.

Mailbox Database Level (Microsoft Exchange 2010, 2013, and 2016)

Specifies if you want to restore a specific Mailbox Database, you can perform a restore at this level.

Mailbox Level (Microsoft Exchange 2010, 2013, and 2016)

Defines if you want to restore a specific Mailbox or mail object.

Verify that the following prerequisites exist before performing a Microsoft Exchange restore:

Database-level restore

- The target machine has the same name and the same version of Microsoft Exchange installed.
- The target database has the same database name and the same storage group name (Microsoft Exchange 200X) and be a part of the same Microsoft Exchange organization.

Granular-level restore

- To restore Microsoft Exchange data, use the [Exchange Granular Restore utility](#).

Specify the Microsoft Exchange Information to Restore

Arcserve UDP Agent (Windows) allows you to not only protect and recover your data, but also helps you to get the Microsoft Exchange Server application that uses that data back up and running. The Microsoft Exchange Server recovery can only be made using the Restore by Recovery Point method.

The process involved in restoring a Microsoft Exchange Application is as follows:

1. [Specify the Recovery Point and Microsoft Exchange Database](#)
2. [Define the Restore Options](#)

Specify the Recovery Point and Microsoft Exchange Database

Use the **Browse Recovery Points** option to restore from a recovery point. When you select a recovery date, all the associated recovery points for that date are displayed. You can then browse and select the Microsoft Exchange database to be restored.

Follow these steps:

1. Access the restore method selection dialog in one of the following ways:

From Arcserve UDP:

- a. Log into Arcserve UDP.
- b. Click the **resources** tab.
- c. Select **All Nodes** in the left pane.

All the added nodes are displayed in the center pane.

- d. In the center pane, select the node and click **Actions**.
- e. Click **Restore** from the **Actions** 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 into Arcserve UDP Agent (Windows).
- b. From the home page, select **Restore**.

The restore method selection dialog opens.

2. Click the **Browse Recovery Points** option.

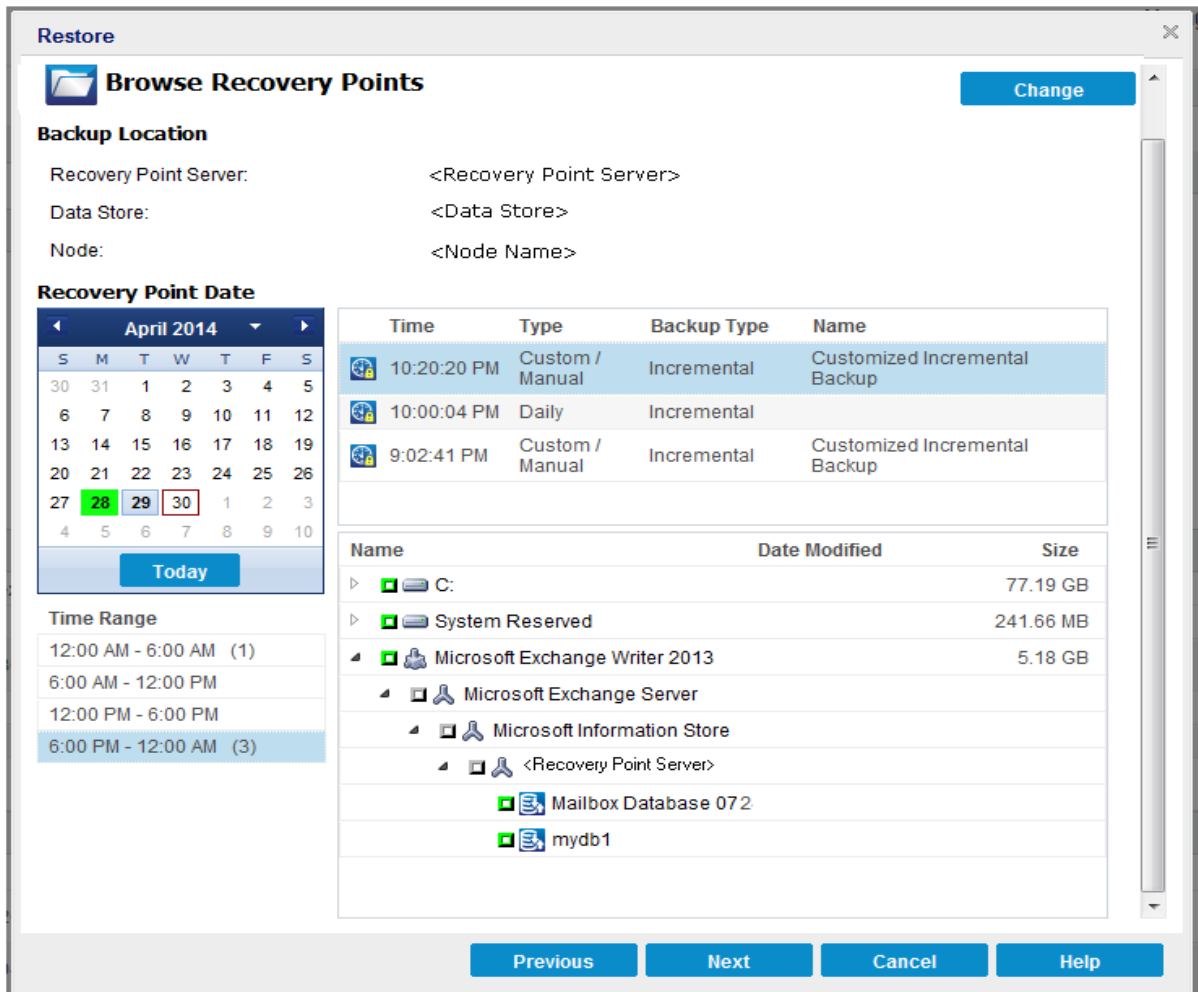
The **Browse Recovery Points** dialog opens.

3. Select the recovery point (date and time) and then select the Microsoft Exchange database to be restored.

The corresponding marker box becomes filled (green) to indicate that the database has been selected for the restore.

Note: If you do not want the transaction log files to be applied after the restore, you must manually delete it before the restore is performed. For more information about manually deleting transaction log files, refer to the Microsoft Exchange

Server documentation.



4. Click **Next**.

The **Restore Options** dialog opens.

Define the Restore Options

After you specify a recovery point and content to restore, define the copy options for the selected recovery point.

Follow these steps:

1. From the **Restore Options** dialog, select the restore destination.

Restore

Restore Options

Destination
Select the restore destination

Restore to original location

Dump file only

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.

2. Select the destination for the restore.

The available options are to restore to the original location of the backup, restore the dump file only, or restore to a Recovery Storage Group/Recovery Mailbox Database.

Restore to original location

Restores to the original location from where the backup image was captured.

Dump file only

Restores the dump files only.

For this option, Arcserve UDP Agent (Windows) will restore the Microsoft Exchange database file to a specified folder, and will not bring it online after recovery. You can then use it to mount on Microsoft Exchange Server manually.

Note: When a Recovery Mailbox Database exists, restore with **Dump file only** option will fail.

Replay log on database

Specifies that when the database files are dumped to the destination folder, you can replay Microsoft Exchange transaction log files and commit them to the database.

Dismount the database before restore and mount the database after restore

Typically before a restore, Microsoft Exchange will perform some checks to help ensure the following:

- The database to be restored is in "Dismounted" status.
- The database is not restored unexpectedly.

To protect a Microsoft Exchange production database from being restored unexpectedly, a switch is added to allow the database to be overwritten during the restore process. Microsoft Exchange will refuse to restore a database if this switch is not set.

For Arcserve UDP Agent (Windows), these two options are controlled by this "Dismount the database before restore and mount the database after restore" option. With this option, Arcserve UDP Agent (Windows) lets you launch the restore process automatically without any manual operations. (You can also specify to dismount/mount database manually).

- If checked, specifies that the recovery process will automatically dismount the Microsoft Exchange database before the restore process and then mount the database after the restore process is completed. In addition, if checked, this option will also allow the Microsoft Exchange database to be overwritten during the restore.
- If unchecked, specifies that the recovery process will not automatically dismount the Microsoft Exchange database before recovery and mount the database after recovery.

The Microsoft Exchange administrator would have to perform some manual operations such as dismount the Microsoft Exchange database, set the Allow Overwrite flag on the database, and mount the Microsoft Exchange database. (The recovery procedure is performed by Exchange during the mounting of the database).

In addition, if unchecked, this option does not allow the Microsoft Exchange database to be overwritten during restore.

Restore to Recovery Database (Microsoft Exchange 2010 and 2013)

Restores the database to a Recovery Database. A Recovery Database is a database that can be used for recovery purposes. You can restore a Microsoft Exchange Mailbox Database from a backup to a Recovery Database and then recover and extract data from it, without affecting the production database that is being accessed by end users.

Before restoring a Microsoft Exchange 2010 or Exchange 2013 database to a Recovery Database, you must first create a Recovery Database.

3. Click **Next**.

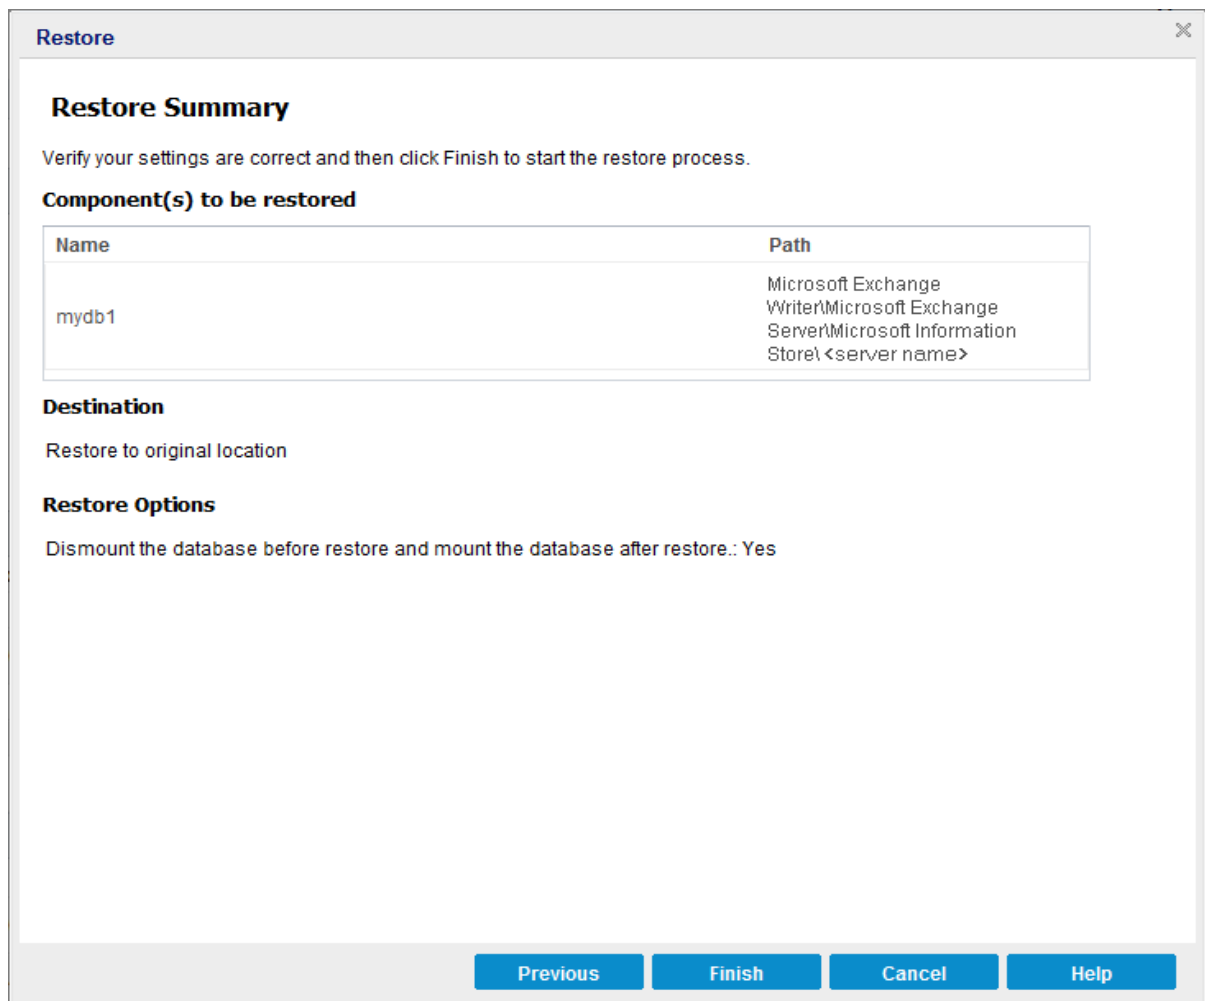
The **Restore Summary** dialog opens.

Restore the Microsoft Exchange Application

After you define the restore options, verify that your settings are correct and confirm the restore process. **Restore Summary** helps you to review all the restore options that you defined and modify them if necessary.

Follow these steps:

1. From the **Restore Summary** dialog, review the displayed information to verify that all the restore options and settings are correct.



- ◆ If the summary information is not correct, click **Previous** and go back to the applicable dialog to change the incorrect setting.
- ◆ If the summary information is correct, click **Next** and then **Finish** to launch the restore process.

The Microsoft Exchange Application is restored.

Verify that the Microsoft Exchange Application was Restored

Follow these steps:


1. Navigate to the Arcserve UDP Agent (Windows) restore destination you specified.
For example, if you select to restore the Microsoft Exchange database to the original location, after the restore is complete, then browse to the physical location to check if the Microsoft Exchange database and logs are restored.
If you select to restore the Microsoft Exchange database to Dump File only location then Arcserve UDP Agent (Windows) will restore the Microsoft Exchange database and logs to a specified location.
2. Verify if the Microsoft Exchange Application was restored and check if the database is mounted and is accessible.
The Microsoft Exchange Application is restored successfully.

How to Restore 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](#)
2. [Specify the Microsoft SQL Server Information to Restore](#)
 - a. [Specify the Recovery Point and Microsoft SQL Server Database](#)
 - b. [Define the Restore Options](#)
3. [Restore the Microsoft SQL Server Application](#)
4. [Verify that the Microsoft SQL Server Application was Restored](#)

Review the Restore Prerequisites and Considerations

Verify that the following prerequisites exist before performing a restore:

- You need Microsoft SQL Server instance before performing a SQL Application restore.
- Review the [Compatibility Matrix](#) that provides the supported operating systems, databases, and browsers.

Review the following restore considerations:

- We cannot restore database across an instance. Restore to alternate location in Arcserve UDP Agent (Windows) means we can restore database and change its database name and file location. For more information, see Microsoft SQL Server Restore to Alternate Location Considerations.
- If the jobs are not for the same VM, Arcserve UDP allows multiple restore jobs to run at the same time. If you attempt to launch a restore job, while another restore job is running for the same VM, an alert message informs you that another job is running and requests you to try again later.
- Arcserve UDP 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 an Arcserve UDP Agent (Windows) restore of a Microsoft SQL Server application to an alternate location, you should consider the following:

If alternate location is on the same machine

For this option, you can either restore a database to a new location (with the same name) or restore with a new name (to the same location):

- **Same Name - New Location**

For example, if Database A is installed in the current SQL Server at "C:\DB_A" and has been backed up. You can use this option and specify "Alternate File Location" to restore Database A to an alternate location such as "D:\Alternate_A.

After the database has been restored, the database file located at the new location "D:\Alternate_A" will then be used.

Important! During restore, if you change the database location but retain the database name, then the previous database will be deleted after the restore is complete. The restored database file will be pointed to the new location.

When you restore to an alternate location, the Instance Name section is unavailable because the Instance Name should always be the same and cannot be changed. As a result, you cannot restore a database to an alternate instance that is currently located on the same MS SQL Server.

- **Same Location - New Name**

For example, if you have two databases (Database A and Database B) installed in the current SQL Server and both have been backed up. You can use this option and specify "New database Name" to restore Database A to same location as Database A_New.

After the databases have been restored, this location will now have three databases (Database A, Database B, and Database A_New).

If alternate location is on the different machine

- The SQL Server installation path must be the same as the path that existed when the backup was performed.

For example, if the backup of the SQL Server is installed at "C:\SQLServer", then the SQL Server on the new Arcserve UDP Agent (Windows) server must also be installed at C:\SQLServer.

- The same instance name for the database that existed when the backup was performed must be installed on Arcserve UDP Agent (Windows) server, otherwise the database associated with that instance will be skipped from the restore.

For example, if the backup of the SQL Server contained "Instance_1" with Database A and Database B and "Instance_2" with Database C, but the Arcserve UDP Agent (Windows) server only has "Instance_1". After the restore is complete, Database A and Database B will be restored, but Database C will not be restored.

- The SQL Server version on the Arcserve UDP Agent (Windows) server must be backwards compatible to the version of the SQL Server used during the backup session.

For example, you can restore an SQL Server 2008 machine to an SQL Server 2010 machine; however, you cannot restore a SQL Server 2010 machine to an SQL Server 2008 machine.

- Restoring a database of 64-bit instance to 32-bit instance is not supported.

Microsoft SQL Server 2012/2014 AAG Restore Considerations

When restoring a Microsoft SQL Server 2012/2014 database that is part of an AlwaysOn Availability Group (AAG), there are some considerations that you should be aware of.

If the MS SQL database is part of the MS SQL 2012/2014 AlwaysOn Availability Group (AAG), and restoring to the original location fails, perform the following tasks:

1. Remove the database to be restored away from the Availability Group. For more information, see [link](#).
2. Share the backup session to Arcserve UDP Agent (Windows) on every Availability Group node and then restore the session by Arcserve UDP Agent (Windows) on every Availability Group node.
3. Add the database back to an Availability Group. For more information, see [link](#).

Specify the Microsoft SQL Server Information to Restore

Arcserve UDP Agent (Windows) allows you to not only protect and recover your data, but also helps you to get the Microsoft SQL Server application that uses that data back up and running. The Microsoft SQL Server recovery can only be made using the Restore by Recovery Point method.

The process involved in restoring a Microsoft SQL Server Application is as follows:

1. [Specify the Recovery Point and Microsoft SQL Server Database](#)
2. [Define the Restore Options](#)

Specify the Recovery Point and Microsoft SQL Server Database

Use the **Browse Recovery Points** option to restore from a recovery point. When you select a recovery date, all the associated recovery points for that date are displayed. You can then browse and select the Microsoft SQL Server database to be restored.

Follow these steps:

1. Access the restore method selection dialog in one of the following ways:

From Arcserve UDP:

- a. Log into Arcserve UDP.
- b. Click the **resources** tab.
- c. Select **All Nodes** in the left pane.
All the added nodes are displayed in the center pane.
- d. In the center pane, select the node and click **Actions**.
- e. Click **Restore** from the server name drop-down menu.

The restore method selection dialog opens.

Note: You are automatically logged in to the agent node and the restore method selection dialog is opened from the agent node.

From Arcserve UDP Agent (Windows):

- a. Log into Arcserve UDP Agent (Windows).
- b. From the home page, select **Restore**.


The restore method selection dialog opens.

2. Click the **Browse 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 doc-

umentation.

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.

The screenshot shows the 'Restore' dialog box with the 'Restore Options' tab selected. Under the 'Destination' section, there are three radio button options: 'Restore to original location', 'Dump file only', and 'Restore to alternative location'. The 'Restore to alternative location' option is selected. Below this, there is a table with columns: Instance Name, Database Name, New Database Name, and Alternative file location. The table contains three rows of data. Below the table, there is a note: 'For SQL writer, the "master" database is not allowed to be renamed.' Underneath that, there is a section for 'Backup Encryption or Protection Password' with a text box for the password. At the bottom of the dialog, there are four buttons: 'Previous', 'Next', 'Cancel', and 'Help'.

Instance Name	Database Name	New Database Name	Alternative file location
MSSQLSERVER	master*		Browse
MSSQLSERVER	model	model_copy	Browse
MSSQLSERVER	msdb	msdb	Browse

2. Select the destination for the restore.

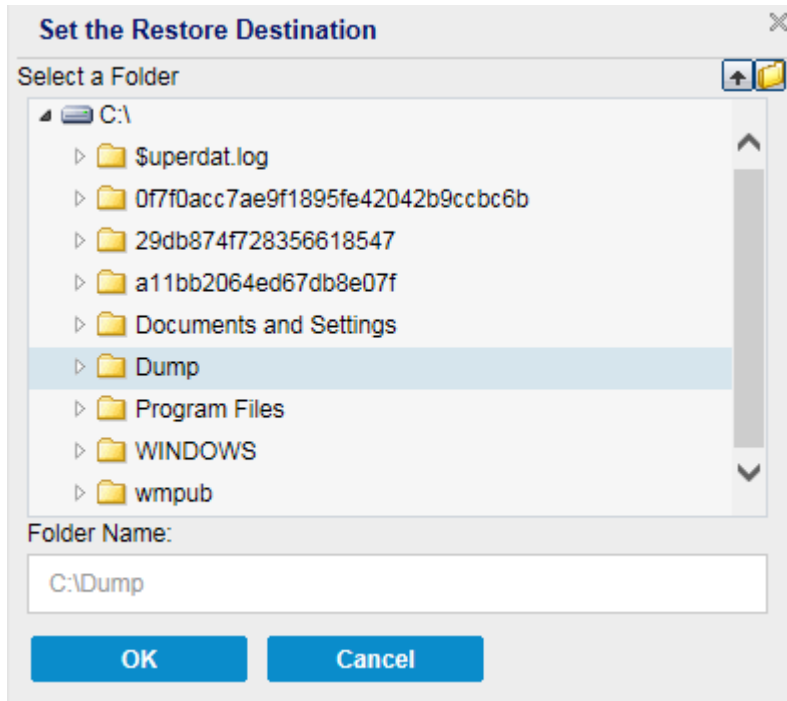
The available options are to restore to the original location of the backup, restore the dump file only, or restore to alternative location.

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](#).

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.

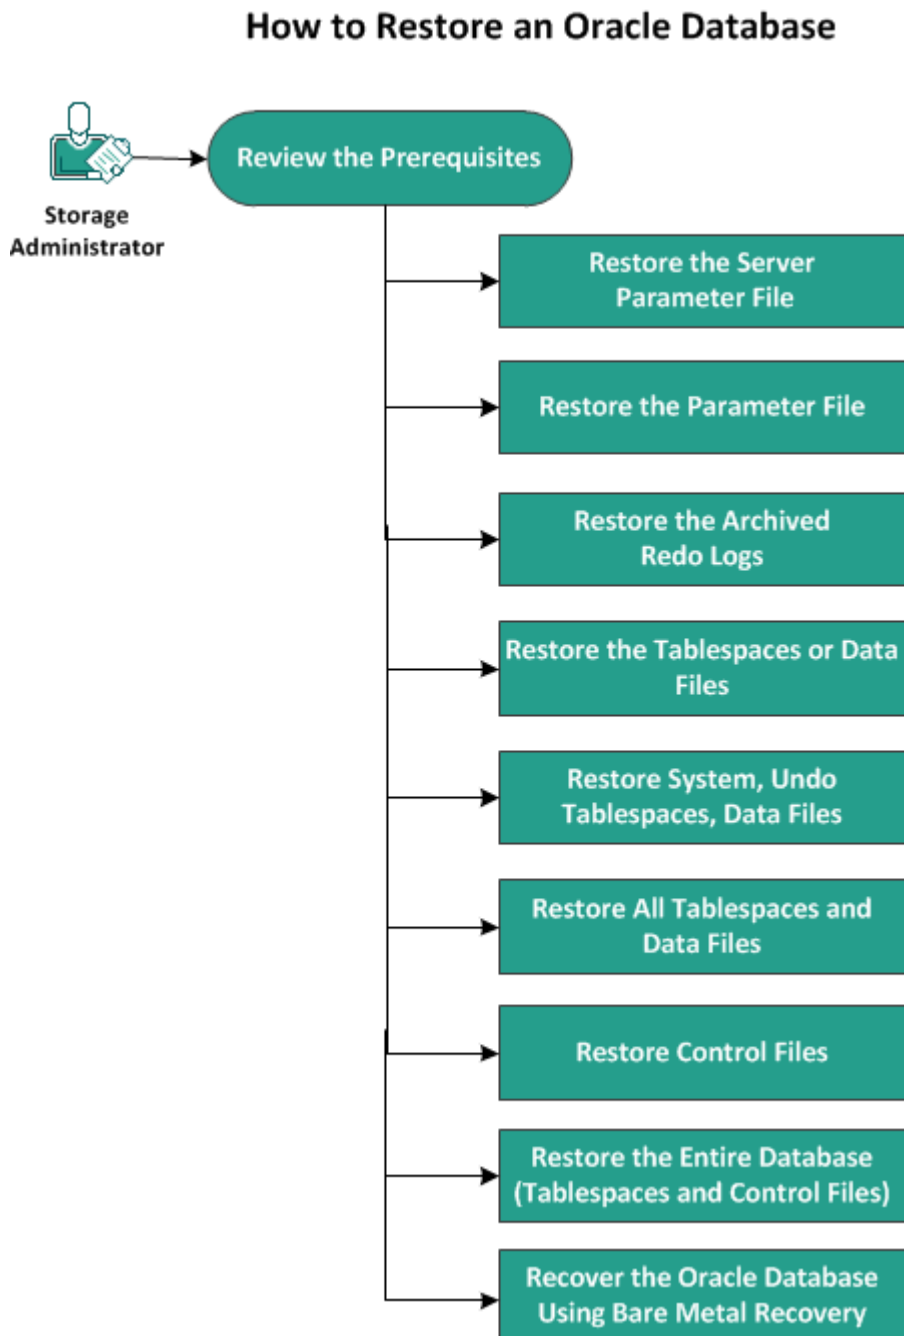
2. Verify if the Microsoft SQL Server Application was restored and check if the database is mounted and is accessible.

The Microsoft SQL Server Application is restored successfully.

How to Restore an Oracle Database

You can restore either certain files and tablespaces or the entire Oracle database using the restore wizard. To restore an Oracle database, locate the files or tablespace on the destination node. Then, you restore the files or tablespace using the restore wizard.

The following diagram illustrates the process to restore an Oracle database:



Perform the following tasks to restore an Oracle database:

- [Review the Prerequisites](#)
- [Restore the Server Parameter File](#)
- [Restore the Parameter File](#)
- [Restore the Archived Redo Logs](#)
- [Restore the Tablespaces or Data Files](#)
- [Restore System, Undo Tablespaces, Data Files](#)
- [Restore All Tablespaces and Data Files](#)
- [Restore Control Files](#)
- [Restore the Entire Database \(Tablespaces and Control Files\)](#)
- [Recover the Oracle Database Using Bare Metal Recovery](#)

Review the Prerequisites and Considerations

Review the following prerequisites before you restore the Oracle database:

- The Oracle VSS writer on the backup node is functioning properly. If the Oracle VSS writer does not function properly, you get a warning message in the Activity Log associated with the backup job.
- You have a valid recovery point.
- To avoid any restore failure problem, you have saved a duplicate copy of your system files before you overwrite the original files.
- Review the [Compatibility Matrix](#) that provides the supported operating systems, databases, and browsers.

Restore the Server Parameter File

The server parameter file is a repository for initialization parameters. Before you restore, you must locate the file. When you locate the files, ensure that the database is in the Open state.

Follow these steps:

1. Log into 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 into 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 into 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 parameter. Before you restore, you must locate the file. When you locate the files, ensure that the database is in the Open state.

Follow these steps:

1. Log into the computer where you want to restore 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 into 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 into 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 into 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 into 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 into 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 into the computer where you want to restore the tablespaces or datafiles.
2. Locate the user tablespaces or datafiles using the following command:

```
SQL> SELECT FILE_NAME, TABLESPACE_NAME FROM DBA_DATA_FILES;
```

3. Change the state of the database to mount, or nomount, or shutdown before you restore the tablespaces or datafiles.

```
SQL> STARTUP MOUNT;
```

```
SQL> STARTUP NOMOUNT;
```

```
SQL> SHUTDOWN IMMEDIATE;
```

4. Log 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 into the destination computer.
7. Navigate to the specific folders and verify that the tablespaces or datafiles are restored.
8. Recover the tablespace or data files.

- ◆ To recover a tablespace, enter the following command at the SQL*Plus prompt screen:

```
SQL> RECOVER TABLESPACE "tablespace_name";
```

- ◆ To recover a data file, enter the following command at the SQL*Plus prompt screen:

```
SQL> RECOVER DATAFILE 'path';
```

Oracle checks for the archive redo log files that it needs to apply and displays the names of the files in a sequence.

9. Enter AUTO in the SQL*Plus prompt screen to apply the files.

Oracle applies the log files to restore the data files. After Oracle finishes applying the redo log file, it displays the following messages:

```
Applying suggested logfile
```

Log applied

After each log is applied, Oracle continues to apply the next redo log file until the recovery is complete.

10. Enter the following command to bring the tablespace online:

```
SQL> ALTER TABLESPACE "tablespace_name" ONLINE;
```

The tablespace is now recovered to the last available log file.

Restore System, or Undo Tablespaces or Data Files

You can restore system, or undo tablespaces or data files. Before you restore, you must locate the file. When you locate the files, ensure that the database is in the Open state.

Follow these steps:

1. Log into the computer where you want to restore system or undo tablespaces or datafiles.

2. Locate the user tablespaces or datafiles using the following command:

```
SQL> SELECT TABLESPACE_NAME, FILE_NAME FROM DBA_DATA_FILES;
```

3. Change the state of the database to mount, or nomount, or shutdown before you restore the tablespaces or datafiles.

```
SQL> STARTUP MOUNT;
```

```
SQL> STARTUP NOMOUNT;
```

```
SQL> SHUTDOWN IMMEDIATE;
```

4. Log into the Arcserve UDP Console.
5. Restore the tablespaces or datafiles using the Restore Wizard. For more information on the restore process, see How to Restore From a Recovery Point.
6. Log into the destination computer.
7. Navigate to the specific folders and verify that the system, or undo tablespaces or datafiles are restored.
8. Recover the tablespace or data files.

- ♦ To recover a tablespace, enter the following command at the SQL*Plus prompt screen:

```
SQL> RECOVER TABLESPACE "tablespace_name";
```

- ♦ To recover a data file, enter the following command at the SQL*Plus prompt screen:

```
SQL> RECOVER DATAFILE 'path';
```

Oracle checks for the archive redo log files that it needs to apply and displays the names of the files in a sequence.

9. Enter AUTO in the SQL*Plus prompt screen to apply the files.

Oracle applies the log files to restore the data files. After Oracle finishes applying the redo log file, it displays the following messages:

```
Applying suggested logfile
```

Log applied

After each log is applied, Oracle continues to apply the next redo log file until the recovery is complete.

10. Enter the following command to bring the tablespace online:

```
SQL> ALTER TABLESPACE "tablespace_name" ONLINE;
```

The tablespace is now recovered to the last available log file.

Restore All Tablespaces and Data Files

You can restore all the tablespaces and data files. Before you restore, you must locate the file. When you locate the files, ensure that the database is in the Open state. If the database is open, use the ALTER TABLESPACE. OFFLINE statement to take the tablespaces or datafiles offline before you begin the restore process.

Follow these steps:

1. Log into the computer where you want to restore the tablespaces or datafiles.
2. Locate the user tablespaces or datafiles using the following command:
SQL> SELECT FILE_NAME, TABLESPACE_NAME FROM DBA_DATA_FILES;
3. Change the state of the database to mount, or nomount, or shutdown before you restore the tablespaces or datafiles.

```
SQL> STARTUP MOUNT;
```

```
SQL> STARTUP NOMOUNT;
```

```
SQL> SHUTDOWN IMMEDIATE;
```

4. Log into the Arcserve UDP Console.
5. Restore the tablespaces or datafiles using the Restore Wizard. For more information on the restore process, see How to Restore From a Recovery Point.
6. Log into 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 into 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 into the Arcserve UDP Console.

5. Restore the tablespaces or datafiles using the Restore Wizard. For more information on the restore process, see How to Restore From a Recovery Point.

6. Log into the destination computer.

7. Navigate to the specific folders and verify that the 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\ O1_MF_1_2_9TKXGGG2_.ARC
```

```
ORA-00280: change 55636 for thread 1 is in sequence #24
```

```
Specify log: {<RET>=suggested | filename | AUTO | CANCEL}
```

14. Specify the full path of the online redo log file and press Enter.

Example: E:\app\Administrator\oradata\orcl\redo01.log

Note: You have to specify the full path multiple times until you get the correct log.

The following messages are displayed:

```
Log applied
```

```
Media recovery complete
```

15. Open the database with the RESETLOGS clause after completing the recovery process.

```
SQL> ALTER DATABASE OPEN RESETLOGS;
```

The lost control files are recovered.

Restore the Entire Database (Tablespaces and Control Files)

You can restore all the entire database (all tablespaces and control files). Before you restore, you must locate the file. When you locate the files, ensure that the database is in the Open state. If the database is open, use the ALTER TABLESPACE.OFFLINE statement to take the tablespaces or datafiles offline before you begin the restore process.

Follow these steps:

1. Log into the computer where you want to restore the tablespaces or datafiles.
2. Locate the user tablespaces or datafiles using the following command:

```
SQL> SELECT TABLESPACE_NAME, FILE_NAME from DBA_DATA_FILES;  
SQL> SHOW PARAMETER CONTROL_FILES;
```
3. Change the state of the database to nomount, or shutdown before you restore the tablespaces or datafiles.

```
SQL> STARTUP NOMOUNT;  
SQL> SHUTDOWN IMMEDIATE;
```
4. Log into the Arcserve UDP Console.
5. Restore the tablespaces or datafiles using the Restore Wizard. For more information on the restore process, see How to Restore From a Recovery Point.
6. Log into the destination computer.
7. Navigate to the specific folders and verify that the tablespaces or datafiles are restored.
8. Recover the database.

```
SQL> RECOVER DATABASE USING BACKUP CONTROLFILE UNTIL CANCEL;
```
9. Apply the prompted archived logs.
Note: If the required archived log is missing, then it implies that a necessary redo record is located in the online redo logs. It occurs because unarchived changes are located in the online logs when the instance failed. You can specify the full path of an online redo log file and press Enter (you may have to try this a few times until you find the correct log).
10. Enter the following command to return the control file information about the redo log of a database:

```
SQL>SELECT * FROM V$LOG;
```

11. (Optional) Enter the following command to see the names of all of the member of a group:

```
SQL>SELECT * FROM V$LOGFILE;
```

Example: After applying the prompted archived logs, you may see the following messages:

```
ORA-00279: change 55636 generated at 24/06/2014 16:59:47 needed for thread 1
```

```
ORA-00289: suggestion e:\app\Administrator\flash_recovery_
area\orcl\ARCHIVELOG\2014_06_24\ O1_MF_1_2_9TKXGGG2_.ARC
```

```
ORA-00280: change 55636 for thread 1 is in sequence #24
```

```
Specify log: {<RET>=suggested | filename | AUTO | CANCEL}
```

12. Specify the full path of the online redo log file and press Enter.

Example: E:\app\Administrator\oradata\orcl\redo01.log

Note: You have to specify the full path multiple times until you get the correct log.

The following messages are displayed:

```
Log applied
```

```
Media recovery complete
```

13. Open the database with the RESETLOGS clause after completing the recovery process.

```
SQL> ALTER DATABASE OPEN RESETLOGS;
```

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 into 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.

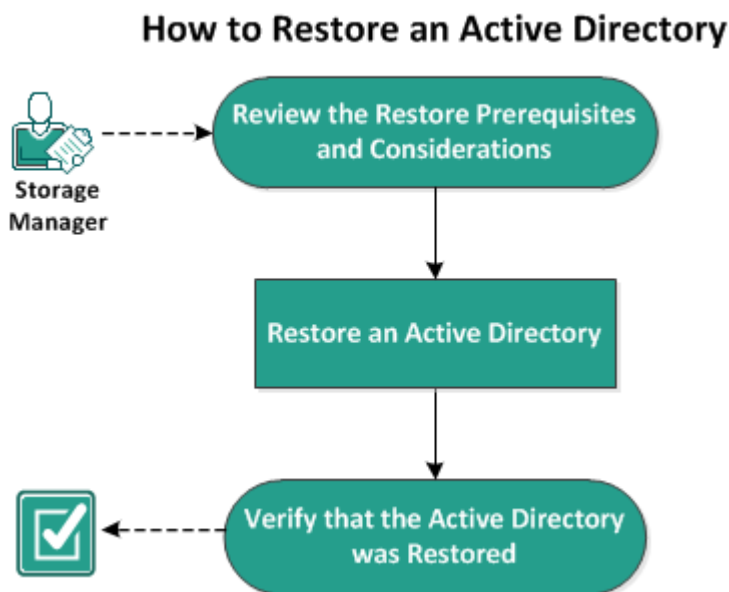
How to Restore an Active Directory

You need to restore a backed up Active Directory session if you have any of the following scenarios:

- You want to recover an attribute of the Active Directory object from any available backed up Active Directory session (not only the last backed up session).
- You want to recover the Active Directory object from any available backed up Active Directory session (not only the last backed up session).
- You want to recover multiple Active Directory attributes or objects from any available backed up Active Directory session (not only the last backed up session).

Important! To perform a granular recovery of an Active Directory, an agent-based backup needs to be performed.

The scenario describes how you can restore an Active Directory.



Perform the following tasks to restore an Active Directory:

1. [Review the Restore Prerequisites and Considerations](#)
2. [Restore an Active Directory](#)
3. [Verify that the Active Directory was Restored](#)

Review the Restore Prerequisites and Considerations

Verify that the following prerequisites exist before performing a restore:

- You have already backed up the volumes that include the Active Directory data-base folder and Log files folder.
- You have the Arcserve UDP Agent (Windows) installed on Domain Controller.
- You have performed an agent-based backup.
- Review the [Compatibility Matrix](#) that provides the supported operating systems, databases, and browsers.

Review the following restore considerations:

- For a recovery point without a file system catalog created, to ensure you can browse and select files/folders to restore, the account/group should be granted access permission to all the folders/files on all volumes with read/list access before the backup is taken.
- You can perform an Active Directory restore only on the Arcserve UDP Agent (Windows).

Restore an Active Directory

After you have installed the Active Directory in different volumes and have performed a backup for both volumes, you may want to restore the volumes with the Active Directory. This scenario describes how you can restore the backed up Active Directory volumes.

Note: Verify that you have completed the prerequisites and backed up Active Directory volumes.

Follow these steps:

1. Access the restore method selection dialog in one of the following ways:

From Arcserve UDP:

- a. Log into Arcserve UDP.
- b. Click the **resources** tab.
- c. Select **All Nodes** in the left pane.
All the added nodes are displayed in the center pane.
- d. In the center pane, select the node and click **Actions**.
- e. Click **Restore** from the **Actions** 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 into Arcserve UDP Agent (Windows).
- b. From the home page, select **Restore**.

The restore method selection dialog opens.

2. From the Restore screen, click Restore Active Directory.

The Restore Active Directory dialog opens.

3. From the Restore Active Directory screen, perform the following steps:

Restore Active Directory

Backup Location

Recovery Point Server: ADTW2K8R2TST Change

Data Store: datastore

Node: ADTW2K8R2DC1

Recovery Point Date

October 2014

Time	Type	Backup Type	Name
3:10:01 AM	Custom / Manual	Full	Customized Incremental Backup

Name	Date Modified	Size
Active Directory		234.02 MB

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

Previous Next Cancel Help

- From the calendar, select Backup date for the Active Directory that you want to restore.
 - From the Time range, select Backup time.
 - From the Restore Active Directory screen, select Backup Job Type and Backup Job Name.
 - From the Name section, select an Active Directory backup session to restore.
4. Click Next.
5. Select the following options to further define the objects, path, and attributes to restore:



- a. From the Object column, select the name of an object. The paths related to the selected object are displayed.
 - b. From the Path column, select a path. The attributes related to the selected path are displayed.

Note: You can use the search icon to browse for the path.
 - c. From the Attribute column, select one or more attributes.
6. Click Next.

The Restore Options screen is displayed.

7. From the Restore Options, select the following objects according to your requirement:
 - a. If the selected object was renamed after backup, click the "Restore with original name of Renamed Objects" option to restore the renamed object.

Note: If you do not select this option, the object will not be restored.
 - b. If the selected object was moved to another container after backup, click the "Restore to original location of Moved Objects" option to restore the moved object.

Note: If you do not select this option, the object will not be restored.

- c. If the selected object was deleted permanently after backup, click the "Restore with the new object ID of Deleted Objects" option to restore the permanently deleted object.

Note: Using this option helps you keep the restored object with the new object ID.

8. Click Next.

The Restore Summary screen is displayed

9. Review the details and perform one of the following action:

- ◆ Click Previous, if you want to modify the details.
- ◆ Click Finish to run restore.

A status message is displayed to inform you when the Restore job is completed. If the restore is unsuccessful, view the logs and try again.

Verify that the Active Directory was Restored

After the completion of the restore process, you can use the Active Directory Users and Computers utility to verify that the Active Directory (object and/or attribute) was restored to the specified destination.

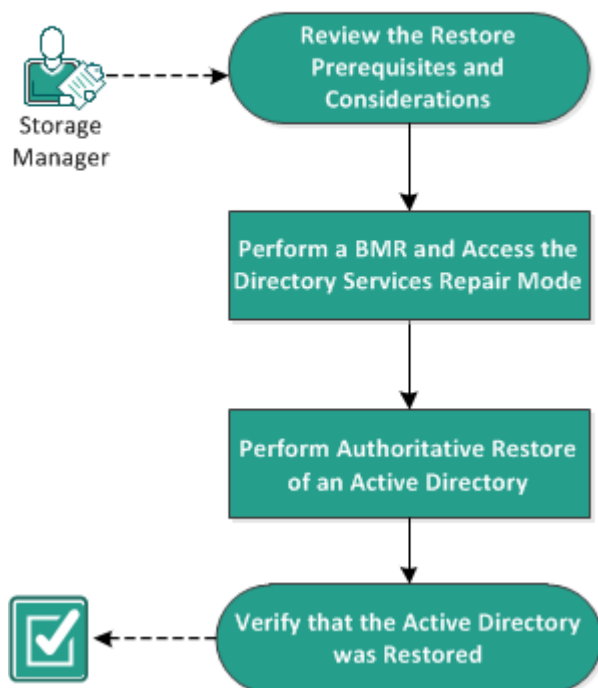
Note: The Active Directory utility is installed automatically with the Active Directory.

How to Perform an Authoritative Restore of an Active Directory after a BMR

When a domain contains more than one domain controller, Active Directory replicates directory objects to all the domain controllers in that domain. The objects contained within a domain can be grouped into Organizational Units (OUs). These OUs can provide a hierarchy structure to a domain and can resemble the organization's structure in managerial or geographical terms. If you inadvertently delete a Active Directory object from a domain controller and want to recover it, you must perform an authoritative restore to return the specified Active Directory object (or container of objects) to its pre-deletion state at the time when it was backed up. For example, you might have to perform an authoritative restore if you inadvertently delete an OU that contains a large number of users.

There are two parts to the authoritative restore process: a non-authoritative restore is performed first by running a BMR, and then an authoritative restore of the deleted Active Directory objects is performed. If you perform only the BMR, the deleted object will not be truly recovered because after the restored Active Directory is updated it will then get replicated back to the pre-restored status by its replication partners, which are also missing the object you wanted to recover.

How to Perform an Authoritative Restore of an Active Directory after a BMR



Complete the following tasks to perform an authoritative restore of an active directory after a BMR:

1. [Review the Restore Prerequisites and Considerations](#)
2. [Perform a BMR and Access the Directory Services Repair Mode](#)
3. [Perform an Authoritative Restore of an Active Directory](#)
4. [Verify that the Active Directory was Restored](#)

Review the Restore Prerequisites and Considerations

Review the following prerequisites and considerations:

- The <"distinguished name"> is the name of the subtree or individual object(s) that is to be marked authoritative. To complete this procedure, you must know the full distinguished name of the object or objects that you want to restore.
- Review the [Compatibility Matrix](#) that provides the supported operating systems, databases, and browsers.

Perform a BMR and Access the Directory Services Repair Mode

There are two parts to the authoritative restore process: a non-authoritative restore is performed first by running a BMR, and then an authoritative restore of the deleted Active Directory objects is performed. As a result, you must make sure that the domain controller that is recovered by the BMR does not automatically reboot before you can access the Directory Services Repair Mode to perform the second part of the process (authoritative restore).

Follow these steps:

1. Perform a BMR for the domain controller that you want to recover, and when the **Bare Metal Recovery Summary of Restore Settings** dialog is displayed, uncheck **Automatically reboot your system after recovery** option.

For more information, see [How to Perform a BMR Using a Backup](#) or [How to Perform a BMR Using a Virtual Standby VM](#).

Important: Do not reboot the domain controller normally after BMR, or else you will lose the opportunity to perform the authoritative restore. The authoritative restore must be performed after you complete the BMR process, but before you fully reboot the recovered domain controller.

2. When the BMR process is completed, initiate a reboot and during the reboot process press **F8** to display the **Advanced Boot Options** screen.

Note: For Windows 8 / Server 2012 and later Operating Systems, accessing Advance Boot Option by pressing F8 at the boot up of machine is disabled by default. To perform BMR and Active Directory server recovery, we provide an option on BMR user interface to help you boot into Advance Boot Option directly without pressing F8. Using this option, you can boot into DSRM mode and perform

Active Directory Authoritative recovery.

arcserve BARE METAL RECOVERY

Bare Metal Recovery(BMR)
- Start Restore Process

This page displays a summary of the disk/volume restore settings that 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, we recommend 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.

Enable Windows F8 boot option helps user perform further recovery or troubleshooting after BMR. For example, press F8 and boot into Active Directory Service Restore mode to perform Active Directory authoritative restore.

Summary of Restore Settings

Restore Item	Status	Progress	Throughput
Restore source volume 'System Reserved' to current destination disk 0	Restoring	24.8%	543.68 MB/Minute
Restore source volume 'C:\' to current destination disk 0	Not Started		
Restore source volume '\\?\Volume{5fc70902-56d8-11e9-80b3-000c292a1b72}\ ...	Not Started		

Automatically reboot your system after recovery.

Do not start Agent service automatically after reboot.

Boot the system to Advanced Boot Options (F8) Menu on the next boot for Windows 8 / Windows Server 2012 and later OS.

Elapsed Time: 00 : 00 : 08
Estimated Time Remaining: 00 : 42 : 00

[24.8%] [64MB/258MB] Restoring basic source volume 'System Reserved' 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 **Advanced Boot Options** screen, select the **Directory Services Repair Mode** and wait for the system to boot up into **Directory Services Repair Mode**.

Perform an Authoritative Restore of an Active Directory

There are two parts to the authoritative restore process: a non-authoritative restore is performed first by running a BMR, and then an authoritative restore of the deleted Active Directory objects is performed.

Follow these steps:

1. Execute **cmd.exe** as an administrator.
2. Run **ntdsutil.exe** to access the Active Directory diagnostic utility.

Note: The ntdsutil.exe is a command-line utility for accessing and managing an Active Directory database.

3. Activate the instance by running the **activate instance <instancename>** command and press Enter. You need to activate the correct instance of ntds to perform maintenance tasks.

The instance name can be retrieved from the ntdsutil.exe by running the "list instances" command. The standard instance of Active Directory is "ntds".

4. Access authoritative restore by running **au r** or **authoritative restore** and press Enter.
5. To restore a subtree or an individual object of an Active Directory, type one of the following commands, and then press Enter.

Note: The "<distinguished name"> is the name of the subtree or object that is to be marked authoritative. To complete this procedure, you must know the full distinguished name of the object or objects that you want to restore.

- To restore a subtree, such as an organizational unit (OU) and all child objects, type: **restore subtree "<distinguished name">**

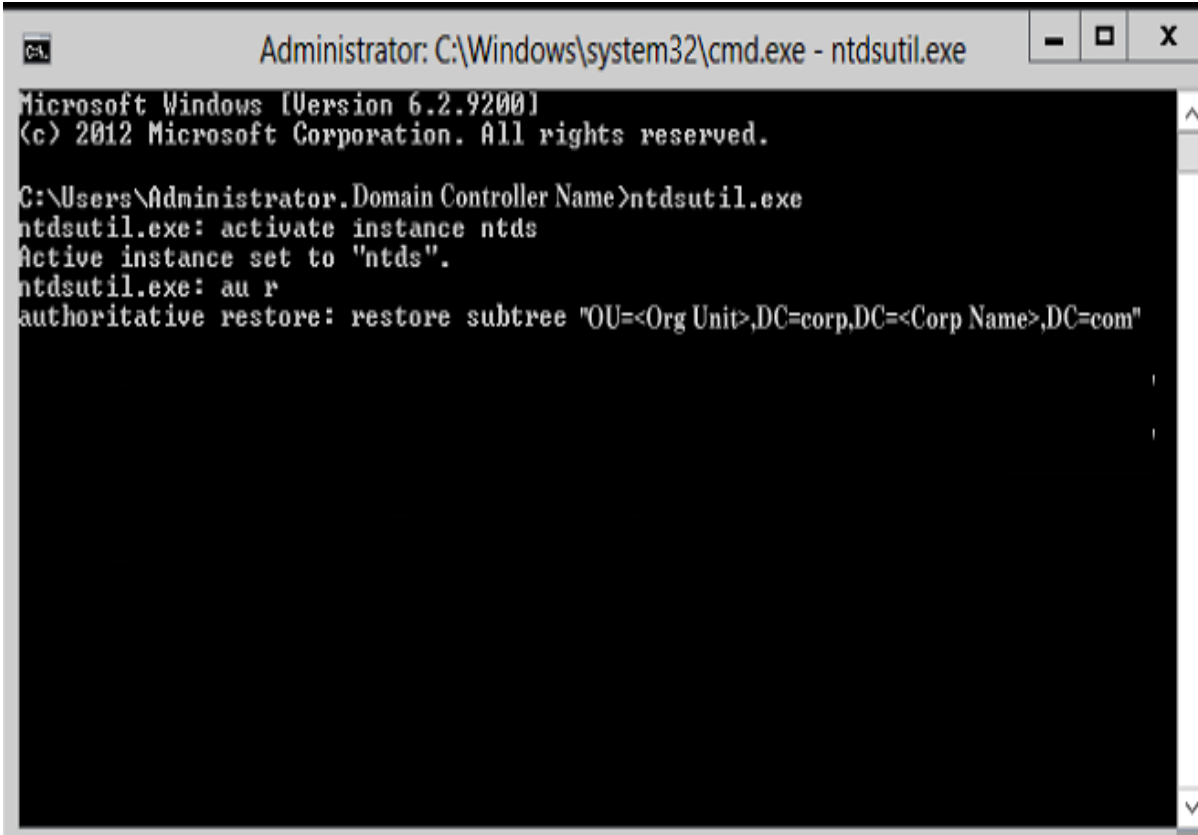
For example: restore subtree "OU=<Organizational Unit>,DC=corp,DC=<Corporate Name>,DC=com"

- To restore a single object or common name (CN), type: **restore object "<distinguished name">**

For example: restore object "CN=<Object Name>,OU=<Organizational Unit>,DC=corp,DC=<Corporate Name>,DC=com"

Note: Always enclose the distinguished name in quotes when there is a space or other special characters within the distinguished name. The most common cause of failure is an incorrectly specified distinguished name or a backup for which the distinguished name does not exist (which occurs if you try to restore a deleted object)

that was created after the backup).



```
Administrator: C:\Windows\system32\cmd.exe - ntdsutil.exe
Microsoft Windows [Version 6.2.9200]
(c) 2012 Microsoft Corporation. All rights reserved.

C:\Users\Administrator.Domain Controller Name>ntdsutil.exe
ntdsutil.exe: activate instance ntds
Active instance set to "ntds".
ntdsutil.exe: au r
authoritative restore: restore subtree "OU=<Org Unit>,DC=corp,DC=<Corp Name>,DC=com"
```

6. From the Authoritative Restore Confirmation dialog, select **Yes** from the pop-up message asking if you are sure you want to perform this Authoritative Restore.
7. Wait for the restore job to complete.
8. At the **authoritative restore** and **ntdsutil** prompts, type quit, and then press Enter.
9. Restart the recovered domain controller in normal operating mode,
10. After the recovered domain controller is started, configure the network settings as necessary (static IP, DNS server etc).
11. From a partner domain controller, access the "Windows Administrative Tools" menu and open **Active Directory Sites and Services**.
12. Run a Replicate job from recovered domain controller. The deleted user is now restored and available from the recovered domain controller and all associated partner domain controllers.

Verify that the Active Directory was Restored

After the completion of the authoritative restore process, verify that the deleted object(s) from the Active Directory were restored to the specified destination.

Follow these steps:

1. For the recovered domain controller, navigate to the Active Directory and verify that the object(s) that were previously deleted are now included.
2. For each domain controller that is associated with the recovered domain controller, navigate to the Active Directory and verify that the object(s) that were previously deleted are now included.


The restored Active Directory is successfully verified.

How to Restore Microsoft Clustered Nodes and Shared Disks

If you have a clustered environment and the clustered nodes and shared disk are not functioning properly, you can easily recover the nodes and disks. You can restore the following items:

- Individual files and folders in a shared disks
- Specific nodes in a cluster
- Entire shared disk
- Entire cluster setup (all clustered nodes and shared disk)

The following diagram illustrates the process to restore clustered nodes and shared disks:



Follow these steps to restore Microsoft clustered nodes and shared disks:

- [Review the Prerequisites](#)
- [Restore Files of a Cluster Shared Disk](#)
- [Restore a Specific Node in a Cluster](#)
- [Restore a Corrupted Cluster Shared Disk](#)
- [Restore the Entire Clustered Nodes and Shared Disk](#)

Review the Prerequisites

Verify that you have completed the following prerequisites:

- You have a valid recovery point for restore.
- You have a valid ISO image for a BMR.
- Review the [Compatibility Matrix](#) that provides the supported operating systems, databases, and browsers.

Restore Files of a Cluster Shared Disk

The shared disk belongs to one of the nodes from the cluster. When you recover any files from the shared disk (not the cluster quorum disk), you need to find the parent node of the shared disk. After you identify the parent node, you can recover files to the parent node from the shared disk.

Note: After a failover happens, you have to browse the recovery point of a different agent to find out the desired recovery point.

Follow these steps:

1. Log into 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 into 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.

Restore from Windows Explorer Using Arcserve UDP Recovery Point View

You can use the Arcserve UDP Recovery Point View from Windows Explorer as an alternate method to restore objects that were backed up. The Arcserve UDP Recovery Point View lets you perform such functions as recover, copy, and paste to quickly and easily restore objects.

From the Arcserve UDP Recovery Point View, you can perform the following restore:

- [Restore a file/folder](#)

Restore a File/Folder Using Arcserve UDP Recovery Point View

For a file/folder restore, you can use the copy and paste function to restore backed up objects (files or folders) from the backup destination to your specific restore destination. (The drag-and-drop method is not supported for a file/folder restore from the Arcserve UDP Recovery Point View).

Note: The path addresses cannot be restored if the length of path is more than 514 characters.

Follow these steps:

1. Access Windows Explorer and from the folders in the left pane, navigate to and select the backup destination.

2. From the right pane, right-click on the backup destination.

A pop-up menu is displayed listing the available options.

3. Select the **Change to Arcserve UDP Recovery Point View** option.

The Windows Explorer View changes to the Arcserve UDP Recovery Point View. When you enter into that directory, all views are logic views from an Arcserve UDP Agent (Windows) user perspective and displays the recovery points stored at that location.

Note: When using the Arcserve UDP Recovery Point View, if you attempt to browse to or locate a recovery point and all recovery points are not displayed, it may be because your backup destinations were split between your local machine and a remote share machine. For the Arcserve UDP Recovery Point View to display all recovery points, the backup locations (for full and corresponding incremental backups) should all be at the same location (either local or remote). If this occurs, it is a limitation only when using the Arcserve UDP Recovery Point View. To remedy this problem, you can use the Arcserve UDP Agent (Windows) restore UI instead to properly display all recovery points, regardless of being at the same or different locations.

4. Select the recovery point that you want to restore from and expand that directory until you can gain access to the file or folder that you want to restore.

Note: From the Arcserve UDP Recovery Point View, the Copy option is only available for file and folder level objects. You cannot use this option to copy volume or machine level objects.

5. Right-click on the file or folder to be restored and select **Copy**. Navigate to the restore location, right-click on the location, and select **Paste**.

Note: Manual operations (such as copy, cut, or paste) for the backup destination folder are not successful if a job is active or a user is browsing recovery points using the Arcserve UDP Recovery Point View.

Note: When restoring backed-up files (with long file names) from the Arcserve UDP Recovery Point View using the copy-and-paste method, the job can fail without any corresponding error or warning. (Windows Explorer has a limitation on the maximum file path length, which can lead to a file copy failure). If this occurs, you can use the installed Arcserve UDP Agent (Windows) UI to perform the restore.

6. When the restore is successfully completed, right-click on the backup destination and select the **Change to Windows Normal View** option.


The Arcserve UDP Recovery Point View changes back to the Windows Explorer View.

Note: During the time that you are in the Arcserve UDP Recovery Point View mode, the merge/purge process for retaining the maximum number of recovery points will be disabled. As a result, Arcserve UDP Agent (Windows) will continue to save recovery points beyond the specified number of recovery points until you exit the Arcserve UDP Recovery Point View and return to the Windows Normal View. When you exit the Arcserve UDP Recovery Point View, all retained recovery points beyond the specified number will then be merged/purged.

How to Copy a Recovery Point

Each time Arcserve UDP Agent (Windows) performs a successful backup, a point-in-time snapshot image of the backup is created. This collection of recovery points lets you locate and specify the exact backup image to copy.

The following diagram illustrates the process to copy a recovery point:



Perform the following tasks to copy a recovery point:

1. [Review the Prerequisites](#)
2. [Configure the Copy Recovery Point Settings](#)
3. [Copy a Recovery Point](#)
 - a. [Select the Recovery Point to Copy](#)
 - b. [Define the Copy Options and Copy the Recovery Point](#)
4. [Verify the Copied Recovery Point](#)

Review the Prerequisites

Review the following prerequisites before copying a recovery point:

- You have at least one full backup available to copy.
- You need a valid destination to copy the recovery point.
- Review the [Compatibility Matrix](#) that provides the supported operating systems, databases, and browsers.

Configure the Copy Recovery Point Settings

Arcserve UDP Agent (Windows) lets you specify the recovery point copy settings. Before you copy a recovery point, configure the copy recovery point settings. For a better understanding about how the options on this dialog can be used to configure your recovery point copy schedule, see [Copy Recovery Points - Example Scenarios](#)

Note: The recovery point copy process is a copy and paste operation only and not a cut and paste operation. As a result, whenever a scheduled copy recovery point job is performed Arcserve UDP Agent (Windows) creates an additional copy of the recovery point to the specified copy destination, while still retaining the original copy of the recovery point at the backup destination that was specified in Backup Settings.

Follow these steps:


1. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Copy Recovery Points** tab. When the **Copy Recovery Points** dialog opens, select **Copy Settings**.

The **Copy Recovery Points** dialog opens.

Notes:

- If the Arcserve UDP Agent (Windows) is being managed by Console, not all settings are available and will be displayed as read-only information.

- When agent is managed by console and not protected in a plan, all the settings are still available except the Preference > Updates panel.



2. Select **Enable Copy Recovery Points**.

When selected, enables the copying of recovery points.

Note: If you do not select this option, no scheduled copying of recovery points is performed.

3. Specify the following recovery point copy schedule settings:

Destination

Specifies the storage location for the copy of the selected recovery point. (Optional) You can click the green arrow button to verify the connection to the specified location.

Note: The maximum length for the specified destination path is 158 characters.

Copy Recovery Points job will run after the specified number of backups are performed

Specifies when the scheduled recovery point copy process is automatically launched. This process is launched based on your selected copy policies and specified number of successful backups (Full, Incremental, and Verify).

Note: Number of successful backup is counted for any custom, daily, weekly monthly backups that are configured.

You can use this setting to control how many times a recovery point copy process is triggered each day. For example, if you schedule to run a backup job every 15 minutes, and copy job after every 4 backups, then it performs 24 recovery point copy jobs each day (1 each hour).

Default: 8

Minimum: 1

Maximum: 1440

Important! If you schedule backup and copy jobs to run at regular intervals and if the copy job is currently running (in active state) when the scheduled time for the backup job time arrives, the backup job fails. (The next backup job will run as scheduled and should be successful if it does not conflict with another copy job). Because the copy operation takes almost same amount of time as performing a full backup, the best practice is not to set a frequent schedule for your recovery point copy jobs.

Specify the number of recovery points to keep

Specifies the number of recovery points that are retained and stored at the specified copy destination. Discards the oldest recovery point, when this number is exceeded.

Note: If you do not have sufficient free space at the target destination, reduce the number of saved recovery points.

Default: 1

Maximum: 1440

4. Select the **Compression** level.

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.

5. If you also want the copied recovery point to be encrypted, specify the following information:

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

Lets you specify and confirm the encryption password being used to encrypt the destination session.

6. Click **Save Settings**.

Your recovery point copy settings are saved.

The copy recovery point settings are successfully configured.

Copy Recovery Points - Example Scenarios

The following example scenarios are provided to give you a better understanding of how the various options can affect your scheduled copying of recovery points.

For this example, assume that you configured your Arcserve UDP Agent (Windows) backup schedule as follows:

- Full Backup - Every 7 days
- Incremental Backup - Every 1 hour
- Verify Backup - Every 3 days

and assume:

- First backup is on Day #1 at 5:00PM (by default, the first backup is always a Full Backup)
- First Incremental Backup will be on Day #1 at 6:00PM (and every hour after)
- Recovery Points retention count is set to 31 (default number)
- Location "D" is configured as the copy destination.

Scenario #1

For this scenario, the Copy Recovery Point settings are as follows:

- Copy after 4 backups
- Retain 1 recovery point

Result:

- At 8:00PM (after the 4th backup), the scheduled copy job will run and consolidate all 4 recovery points into a single recovery point and store it at destination D.
- At 12:00 midnight (after the 8th backup), the next scheduled copy job will run and consolidate all 8 recovery points into a single recovery point and store it at destination D.

The previous recovery point is removed from destination D because the setting is to retain only 1 recovery point at the destination.

Scenario #2

For this scenario, the Copy Recovery Point settings are as follows:

- Copy after 4 backups
- Retain 4 recovery points

Result:

- At 8:00PM (after the 4th backup), the scheduled copy job will run and consolidate all 4 recovery points into a single recovery point (Recovery Point #1) and store it at destination D.
- At 12:00 midnight (after the 8th backup), the next scheduled copy job will run to create Recovery Point #2 and store it at destination D.
- At 4:00AM on Day #2 (after the 12th backup), the next scheduled copy job will run to create Recovery Point #3 and store it at destination D.
- At 8:00AM on Day #2 (after the 16th backup), the next scheduled copy job will run to create Recovery Point #4 and store it at destination D.
- At 12:00 noon on Day #2 (after the 20th backup), the next scheduled copy job will run. A new recovery point will be created and the first recovery point (created after the 8:00PM backup on previous day) is removed from destination D, because the setting is to retain only 4 recovery points at the destination.

Scenario #3

For this scenario, the Copy Recovery Point settings are as follows:

- Copy after 1 backup
- Retain 4 recovery points

Result:

- At 5:00PM (after the 1st backup), the scheduled copy job will run to create a single recovery point (Recovery Point #1) and store it at destination D.
- At 6:00PM (after the 2nd backup), the next scheduled copy job will run to create Recovery Point #2 and store it at destination D.
- At 7:00PM (after the 3rd backup), the next scheduled copy job will run to create Recovery Point #3 and store it at destination D.
- At 8:00PM (after the 4th backup), the next scheduled copy job will run to create Recovery Point #4 and store it at destination D.
- At 9:00PM (after the 5th backup), the next scheduled copy job will run. A new recovery point will be created and the first recovery point (created after the 5:00PM backup) is removed from destination D, because the setting is to retain only 4 recovery points at the destination.

Copy a Recovery Point

When you select a recovery point to copy, all previous backup blocks (full and incremental) are consolidated and captured to recreate a full and most recent backup image.

You can perform the following tasks to protect your backups:

- Copy/export recovery point information to store it safely off-site in the event of a catastrophe.
- Save your recovery points to multiple locations.
- Consolidate your backups to preserve all your recovery points.

The process involved in copying a recovery point is as follows:

1. [Select the Recovery Point to Copy](#)
2. [Define the Copy Options and Copy the Recovery Point](#)

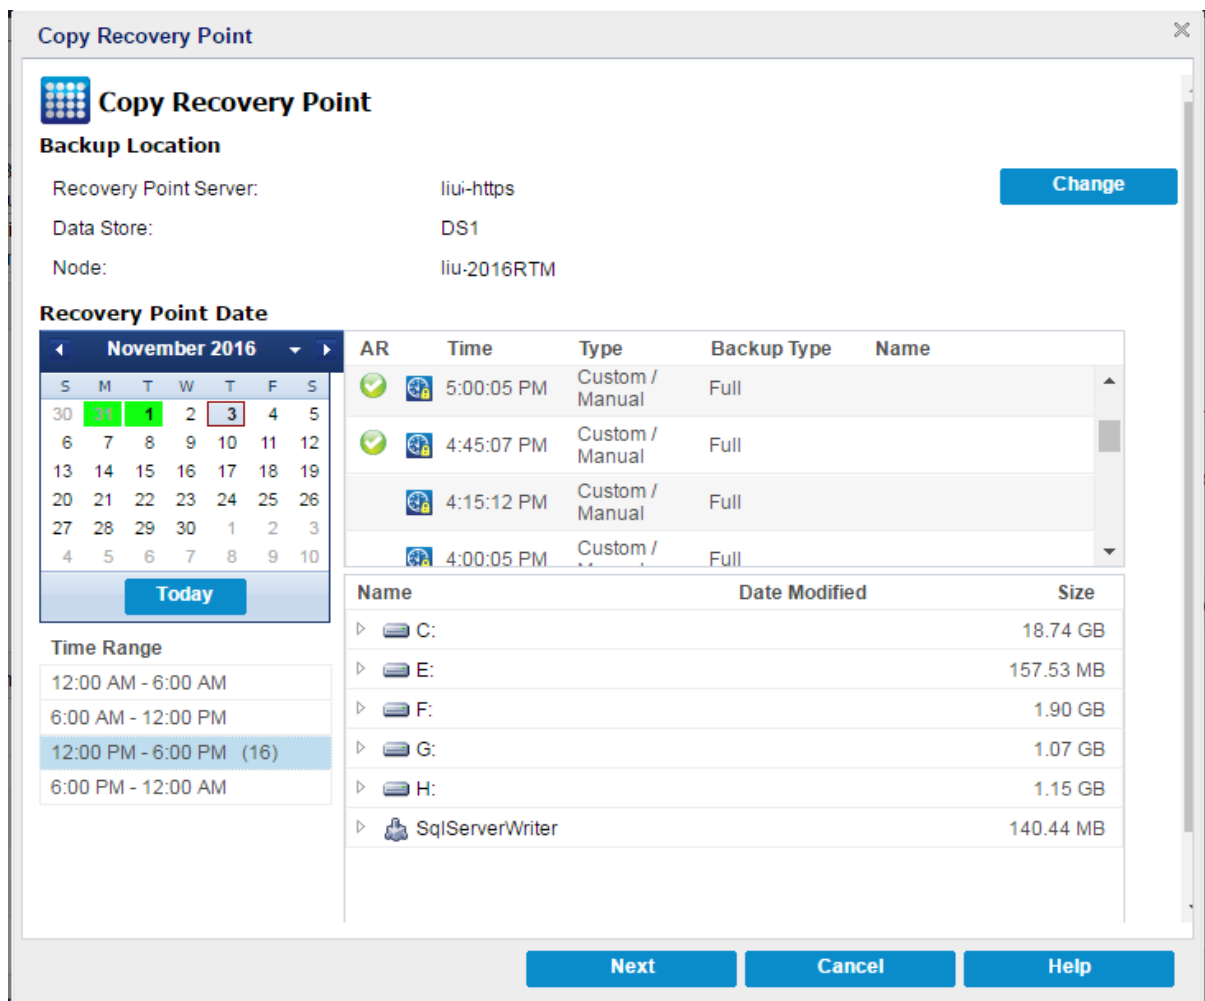
Select the Recovery Point to Copy

Arcserve UDP Agent (Windows) provides a list of available recovery points and lets you select the recovery point to create a copy. You can specify the destination, recovery point date, and time range to copy a recovery point.

Follow these steps:

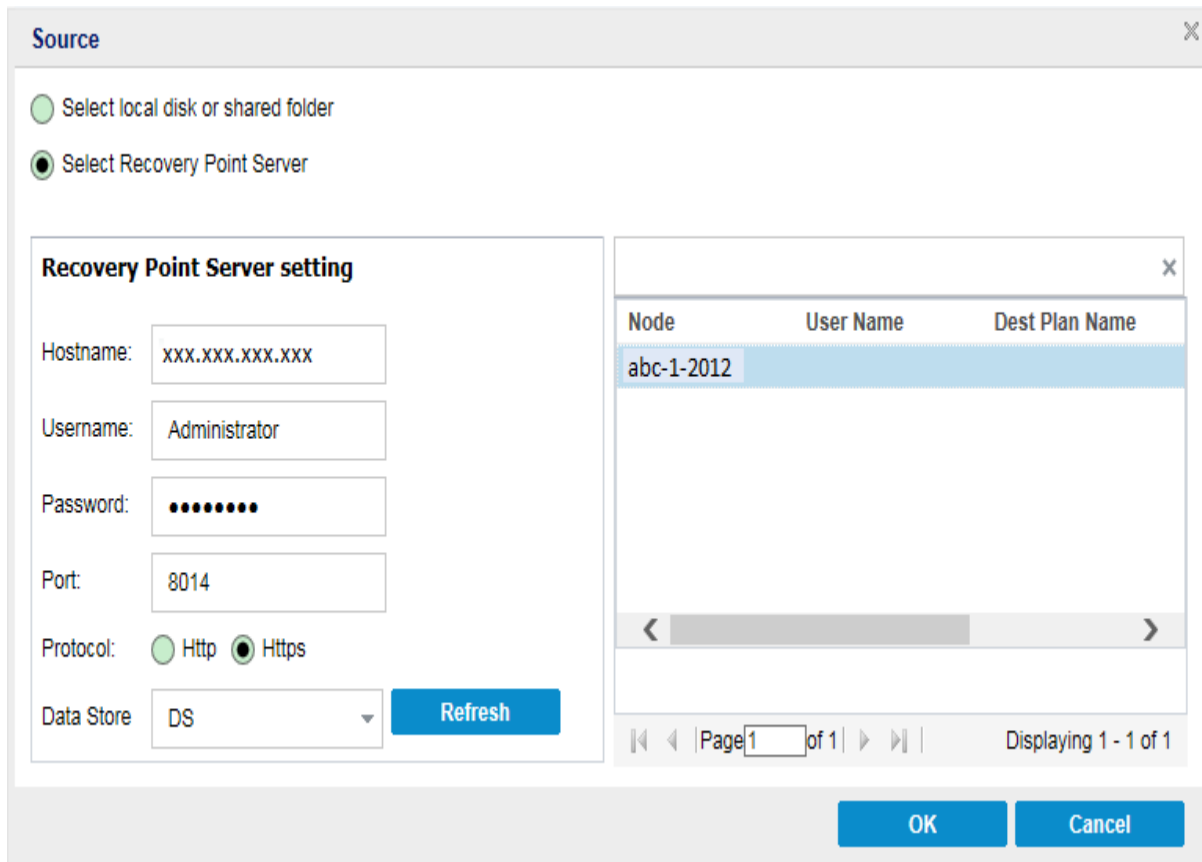
1. On the Arcserve UDP Agent (Windows) home page, select **Copy Recovery Point**.

The **Copy Recovery Point** dialog opens.



2. Click **Change** to change the backup location.

The **Source** dialog opens where you can select the backup location.



3. 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.

Note: All the dates containing recovery points for the specified location are highlighted in green.

4. Select the calendar date for the backup image to copy.

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.

Note: A clock icon with a lock symbol indicates that the recovery point contains encrypted information and can require a password for the restore.

5. Select a **Recovery Point** to copy.

The backup content (including any applications) for the selected recovery point displays.

6. Click **Next**.

The **Copy Options** dialog opens.

The recovery point to copy is specified.

Define the Copy Options and Copy the Recovery Point

After you specify a recovery point to copy, define the copy options to create a copy that combines the previous full and incremental backups for the selected recovery point.

When the backup destination is on a data store of a Recovery Point Server, you have the option to copy a recovery point without providing the password of the recovery point. If the password is not provided, you can still copy the recovery point, but the password, compression, and encryption settings are kept as the source recovery point. If the password is provided, you can change the compression, encryption, and password.

When the backup destination is on a shared or local folder, you must provide the password to copy the recovery point.

Follow these steps:

1. From the **Copy Options** dialog, specify the type of **Copy Options**.

Copy Recovery Point

Copy Options

The recovery point selected to be copied is encrypted or password-protected.

Retain original compression and encryption settings

Select this option if you want to retain the existing encryption and compression settings for the destination recovery point, without providing the session password. The destination recovery point will then use the original protection password, and compression and encryption settings.

Use different compression and encryption settings

Password

Copy Options

Destination

Compression

Encryption Algorithm

Encryption Password

Confirm Encryption Password

Note: This operation will merge all sessions up to, and including, the selected recovery point into a single session, which will then be copied to the specified destination.

Retain original compression and encryption settings

This option is available when you back up to a data store.

Note: When you back up to a shared or local folder, this option is not available.

Use different compression and encryption settings

This option is available when you back up to a data store or a shared or local folder.

Note: When you back up to a shared or local folder, you can only submit a copy recovery point job using this option.

2. Specify the **Copy Options**.

Password

Specify the encrypted password for backup.

Note: This dialog includes two password fields. The upper field is for the password to decrypt the source session, and the lower field is used to encrypt the destination session.

If the recovery point you select to copy was previously encrypted, provide the password.

- If the recovery point to be copied is a backup session of the same computer that is running the copy recovery point job, the encryption password is remembered and this field is automatically populated.
- If the recovery point to be copied is a backup session of another computer, enter the encryption password.

Destination

Specify (or browse to) the storage location of the selected recovery point.

(Optional) You can click the green arrow button to verify the connection to the specified location.

If necessary, enter the Username and Password.

3. If you selected **Use different compression and encryption settings**, select the **Compression** level.

Note: The specified backup compression level has no relation with the copy compression level. For example, in backup destination the compression level can be set to **Standard**; however, when you submit the copy job, the compression can be changed to **No Compression** or **Maximum Compression**.

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.

Note: 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 that is performed after this compression level change is automatically a Full Backup. After performing the Full Backup, performs all future backups (Full, Incremental, or Verify) as scheduled.

4. If you also want the copied recovery point to be encrypted, specify the following information:

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

Lets you specify and confirm the encryption password being used to encrypt the destination session.

Note: When you enable the encryption, specify a new password. You need this password to restore the copied recovery point.

5. Click **Create a Copy**.

A status notification window appears and the copy process for the selected recovery point type is launched immediately.

Note: Arcserve UDP Agent (Windows) allows only one recovery point copy job to run at the same time. If you attempt to launch a recovery point copy job manually while running another scheduled copy job, an alert message opens. The message informs you that another job is running and requests you to try again at a later time.

The recovery point is copied from the backup source to the copy target destination.

Verify the Copied Recovery Point

After you copy a recovery point, verify that the copied recovery point is available at the specified destination.

Follow these steps:

1. Navigate to the Arcserve UDP Agent (Windows) destination you specified.
A list of folders appears.
2. Open the hostname folder, and navigate to the following subordinate folder:
hostname\VStore
3. Open the VStore folder, and navigate to the following session folder:
VStore\S0000000001
4. Locate all files with a D2D extension to verify your copied recovery point at the specified location.

For example, if your computer name is "Department_A" and you copied the recovery point (backup) to "E:\copied_vhd\", navigate to the following location:

E:\copied_vhd\Department_A\VStore\S0000000001.

The copy of your recovery point is successfully verified.

Mount a Recovery Point

Mount Recovery Point provides the ability to mount a recovery point to a drive letter (volume) or an NTFS folder, to view, browse, copy, or open the backup files directly in Windows Explorer.

Note: For Arcserve UDP host-based VM backup, the recovery points are mounted on the backup proxy system.

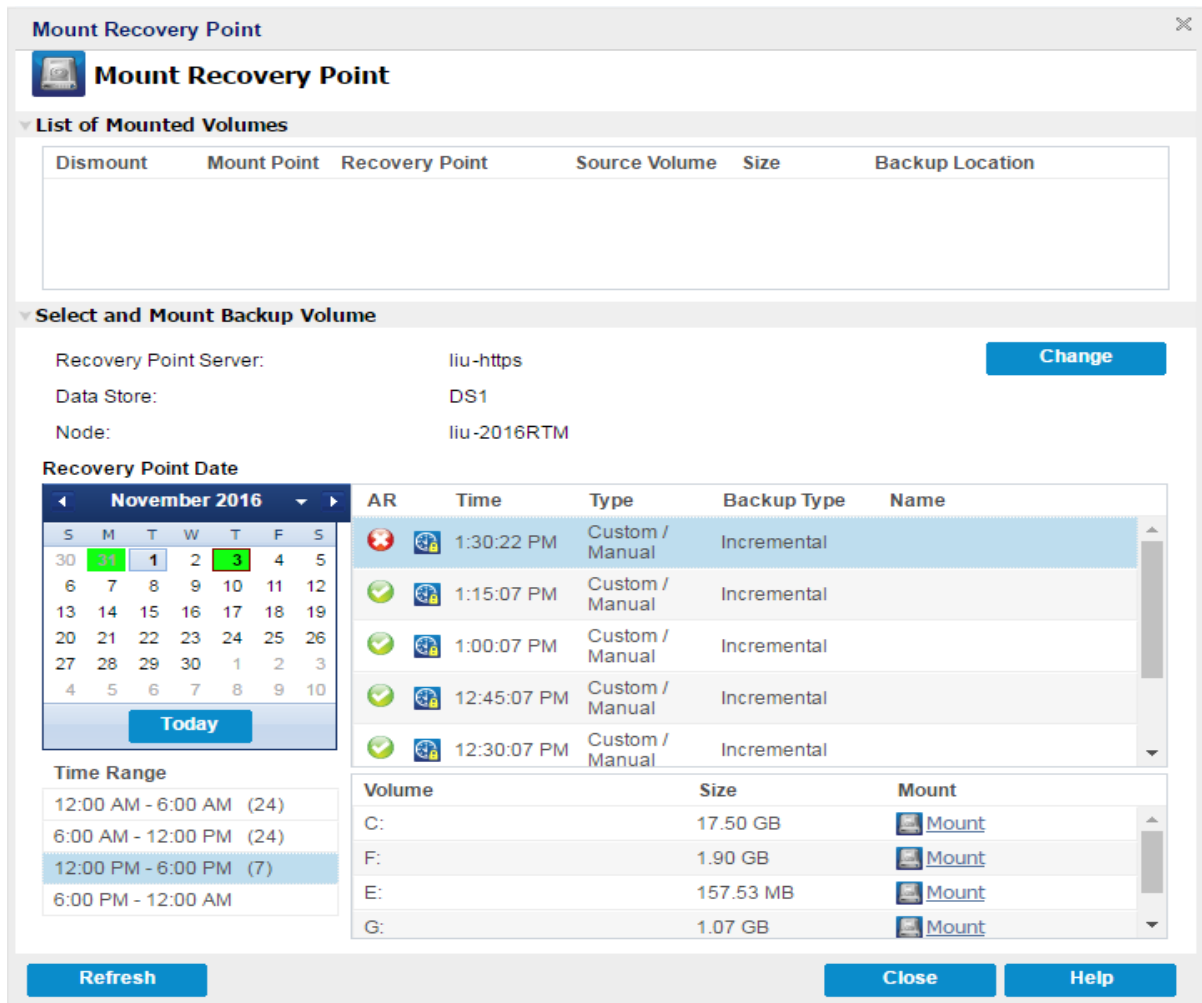
Instead of using the Arcserve UDP Agent (Windows) web interface to find or recover files and folders, you can mount the volumes from a given Arcserve UDP Agent (Windows) backup recovery point to a drive letter, so that you can use Windows Explorer to find or recover any files and folders. The benefit is that Windows Explorer can be more familiar and convenient to use.

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. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Mount Recovery Point**.

The Mount Recovery Point dialog opens.



2. 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 Hhttps
 Data Store:

Node	User Name	Dest Plan Name
abc-1-2012		

Page 1 of 1 | Displaying 1 - 1 of 1

3. 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.

Note: All the dates containing recovery points for the specified location are highlighted in green.

4. Specify the recovery point to mount.

- a. Select the calendar date for the backup image you want to mount.

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.

Note: A clock icon with a lock symbol indicates that the recovery point contains encrypted information and can require a password to mount the recovery point.

- b. Select a recovery point that you want to mount.

The corresponding backup content (including any applications) for that recovery point is displayed.

- c. Locate the Volume that you want to mount and click **Mount**.

You can mount the recovery point to a drive letter (volume) or an empty NTFS folder.

Note: If a volume is mounted already, you cannot mount it again from the same recovery point.

The Mount Recovery Point dialog opens.

5. Select whether you want to mount to a volume or path.
 - ◆ If you mount to a volume, select the volume from the drop-down list.
 - ◆ If you mount to a path, enter or browse for the location.

Important! The path must be to an empty NTFS or ReFS folder.


6. If the selected recovery point was encrypted, provide the encryption password and click **OK**.

The selected backed up volume is mounted and displayed in the List of Mounted Volumes on the Mount Recovery Point dialog. You can now use Windows Explorer to view, browse, copy, or open the backup files.

Note: You cannot delete any backup files from Windows Explorer.

7. When the mounted recovery point is no longer needed, the best practice is to dismount it. Otherwise, the mounted recovery point prevents the Arcserve UDP Agent (Windows) backup from performing a session merge/purge operation.

If Arcserve UDP Agent (Windows) attempts to merge a mounted recovery point, a status alert also displays on the Home Page to inform you the selected recovery point cannot be merged.



Note: If the merge fails and you want to be notified, you can configure Email Alerts in the Preferences Settings to receive an email alert. For more information, see [Specify Email Alert Preferences](#).

- a. To dismount the mount points, select the mount point that you want to dismount and click **Dismount**.

The selected mount point is dismounted and no longer displays in the List of Mounted Volumes on the Mount Recovery Point dialog.

- b. To refresh the list of mount points, click **Refresh**.

The updated list of mount points displays.


Create a VHD File from an Arcserve UDP Agent (Windows) Backup

Each time Arcserve UDP Agent (Windows) performs a successful backup, a point-in-time snapshot image of your backup is also created. Before creating a Virtual Hard Disk (VHD) file from an Arcserve UDP Agent (Windows) backup, you must have at least one Arcserve UDP Agent (Windows) recovery point available.

Follow these steps:


1. From the Arcserve UDP Agent (Windows) home page, select **Copy Recovery Point**.

The Copy Recovery Point dialog opens.



2. Click **Change** to change the backup location.

The Source dialog opens where you can select the backup location.



3. 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.

Note: All the dates containing recovery points for the specified location are highlighted in green.

4. Specify the recovery point to copy.

- a. Select the calendar date for the backup image you want to copy.

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.

- b. Select a recovery point that you want to copy.

The corresponding backup content (including any applications) for that recovery point is displayed.

5. Specify the copy options.

- a. Select **Use different compression and encryption settings** and enter the **Password**.

- b. Select the destination.

You can either specify a location or browse to the location where the copy of your selected recovery point is going to be stored. If necessary, enter the Username and Password.

Note: Verify that you select a location that has sufficient free space available to hold the entire VHD.

- c. Set the level of compression to **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.

6. Click **Create a Copy**.

A status notification window appears and the copy process for the selected recovery point type is launched immediately.

The recovery point image is copied from the backup source to the destination.

7. When the copy process is finished, browse to the destination and navigate to the subordinate folder corresponding to the hostname of the Arcserve UDP Agent (Windows) computer.

8. Open the hostname folder and navigate to the following subordinate folder:

"VStore\S0000000001"

For example, if your computer name is "Department_A" and you copied the recovery point (backup) to "E:\export_vhd\" you would navigate to:

E:\export_vhd\Department_A\VStore\S0000000001

9. Open the "S0000000001" folder to locate all files with a ".vhd" extension.

Each of these files corresponds to an actual physical disk on the source computer which can be used as regular VHD files.

Important! The VHD created by Arcserve UDP Agent (Windows) during the copy process may not boot in the hypervisor because the VHD files may not contain the correct drivers for the VM.

View Logs

The Activity Log contains comprehensive information about all the operations performed by Arcserve UDP Agent (Windows). The log provides an audit trail of every job that is run (with the most recent activities listed first) and can be helpful in troubleshooting any problems that occur.

Note: This task is only available from the Arcserve UDP Agent (Windows) UI and not from the Arcserve UDP Agent (Windows) Monitor.

View Logs

1. From the Arcserve UDP Agent (Windows) home page, select **View Logs**.

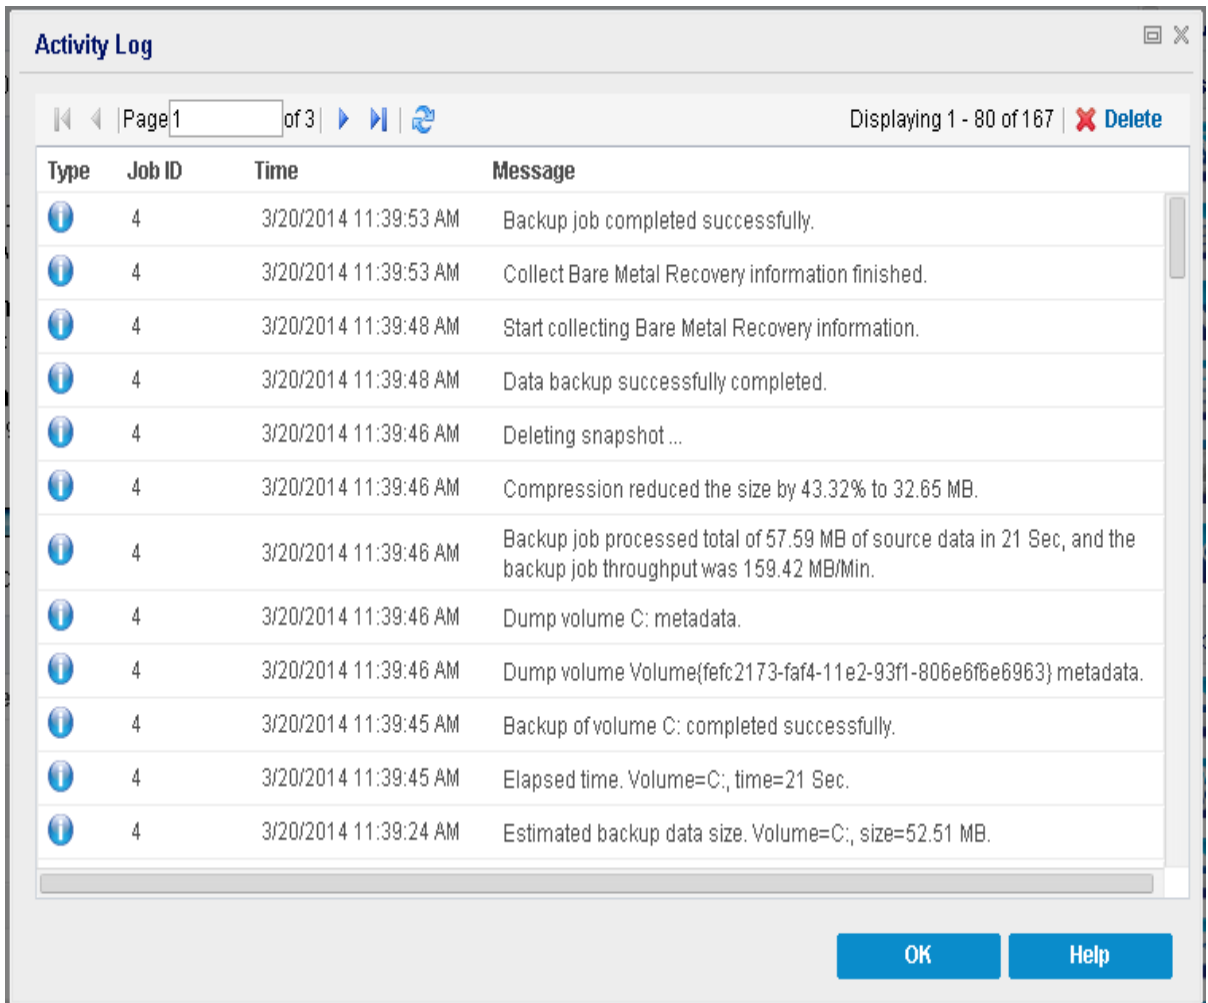
The Arcserve UDP Agent (Windows) Activity Log opens, displaying the following information:

- ◆ Type of Message (Error, Warning, Information)
- ◆ Job ID

Note: The Job ID can be used to easily locate log entries that are related to a specific job and can be helpful when troubleshooting job-related problems.

- ◆ Time that the message was recorded

- ◆ Message indicating the activity performed or the problem encountered.



2. If necessary, you can click the Delete button to purge some or all of the log entries.

The Delete Activity Log dialog opens.

You can then specify to Delete all log records or Delete log records older than a specific date. If you select the "Delete log records older than" option, you can then spe-

cify from the calendar which date will be used as the older than date.

Delete Activity Log

Delete all log records

Delete all log records older than:

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

OK Cancel

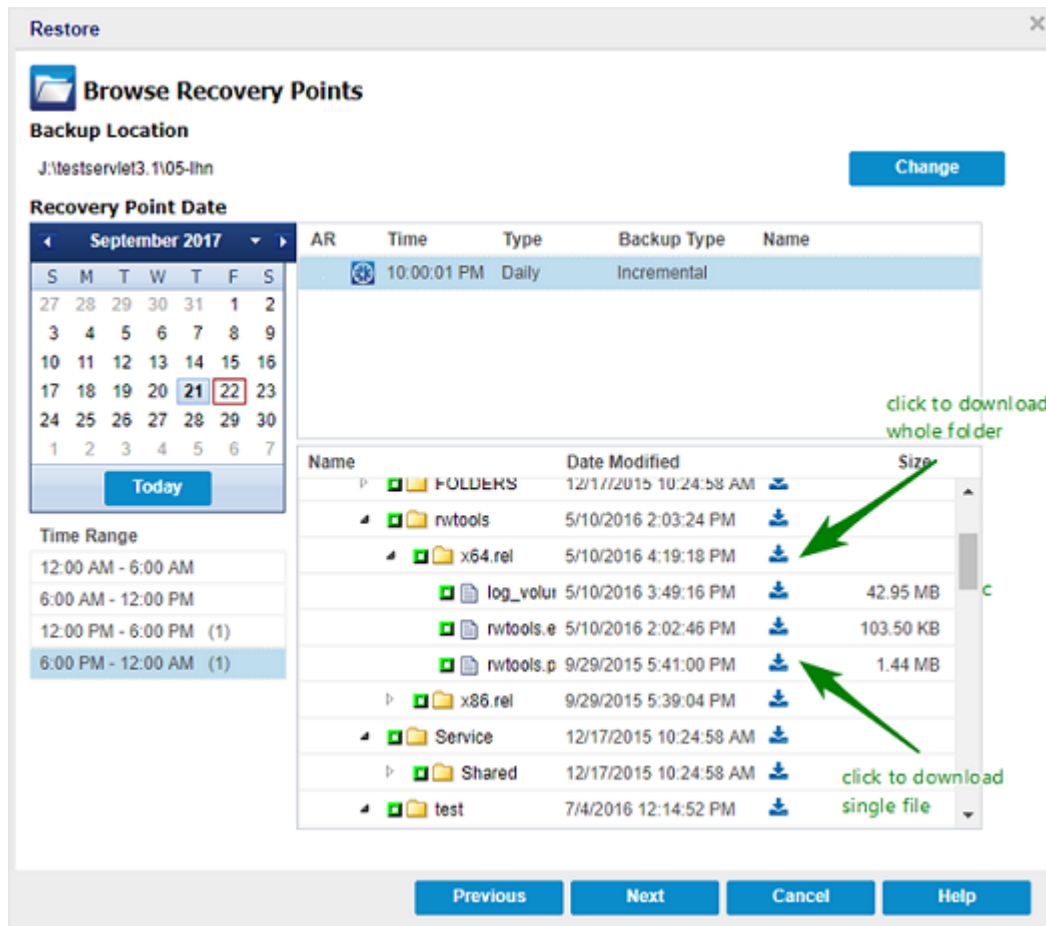
How to Download File/Folders without Restore

Arcserve UDP lets you download a file or complete folder without submitting for restore. From the Restore wizard, the Browse Recovery Points screen lets you directly download any file or a complete folder with all the files. Downloading before restore may help perform a quick check of files to avoid undesired files getting restored.

A single file is downloaded directly in the same format, while a folder is downloaded as a zip file. The zip file has the following name format:

[nodename]_[sessionid]_[timestamp].zip

To download, you simply need to reach the Browse Recovery point screen in the Restore wizard. The below screenshot displays how to perform download of a file or folder:



Considerations for download:

- Downloading or packaging as zip file is not possible for some system file. The agent tomcat service does not have enough privileges to access system file or

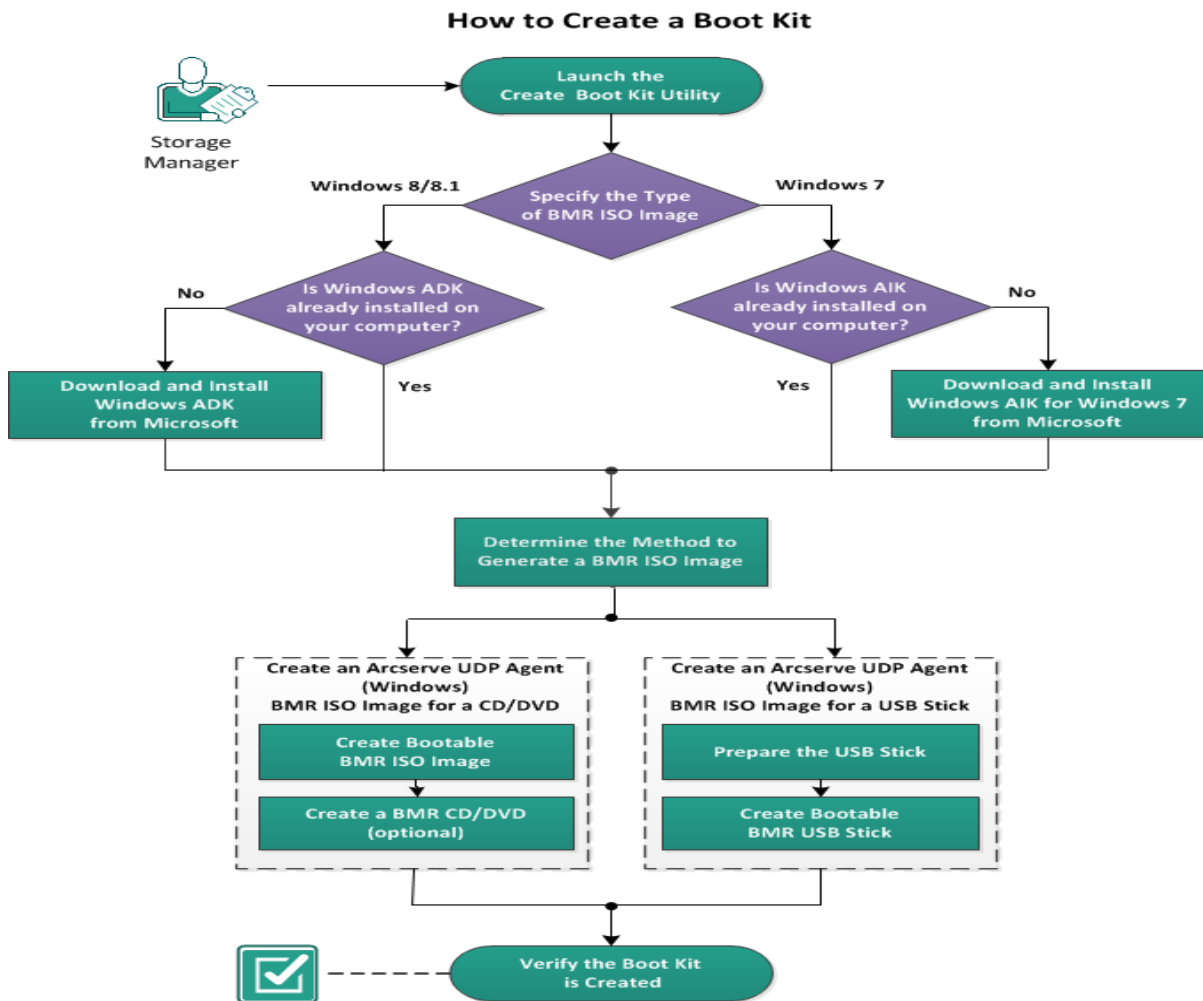
user files of other protected node.

- To avoid excess consumption of Tomcat memory and CPU usage, we recommend submitting a restore job to alternative path while downloading a huge file or folder.
- Using Windows Compressed Folder Tools to browse the downloaded zip files may fail as the tool finds some of the zip entry names too long to browse. We recommend using other zip tools to open the file. For example, WinZip, WinRAR, 7-Zip.
- IE9 user using https in IE9 and agent web service to provide service may not be able to download the files. A known issue from IE9 in downloading resource from a dynamic page through https prevents such download. For more information and solution, click [link](#) for Microsoft article.

How to Create a Boot Kit

The Arcserve UDP Agent (Windows) utilizes a Boot Kit Utility to combine a WinPE (Windows Preinstallation Environment) image and Arcserve UDP Agent (Windows) image to create a BMR ISO image. This ISO image is then burned onto a bootable media. When you perform a bare metal recovery, the Arcserve UDP Agent (Windows) bootable media (CD/DVD or USB stick) is used to initialize the new computer system and allow the bare metal recovery process to begin.

The following diagram illustrates the process to create a boot kit:



Perform the following tasks to create a boot kit:

1. [Launch the Create Boot Kit Utility](#)
2. [Determine the Method to Generate a BMR ISO Image](#)

3. [Create an Arcserve UDP Agent \(Windows\) BMR ISO Image for a CD/DVD](#)
 - a. [Create Bootable BMR ISO Image](#)
 - b. (optional) [Create a BMR CD/DVD](#)
4. [Create an Arcserve UDP Agent \(Windows\) BMR ISO Image for a USB Stick](#)
 - a. [Prepare the USB Stick](#)
 - b. [Create Bootable BMR USB Stick](#)
5. [Verify the Boot Kit is Created](#)

Launch the Create Boot Kit Utility

Arcserve UDP Agent (Windows) provides a Create Boot Kit for Bare Metal Recovery utility to help you generate a WinPE-based ISO image. This ISO image contains all the information needed to perform a bare metal recovery (BMR) if necessary.

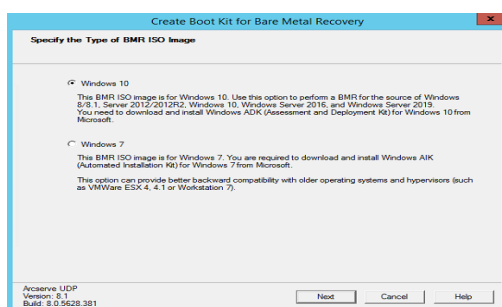
Important! If you upgrade to a newer version or update of Arcserve UDP, you must re-create the BMR ISO using the proper Windows AIK or ADK level to include support for latest features and bug fixes. If you want to perform BMR from an older version of RPS and deduplication data store, you have to use the older version of the BMR ISO.

Follow these steps:

1. You can launch the **Create Boot Kit** utility from the **Advanced** options of the System Tray Monitor or from the Start menu.
2. Specify the type of BMR ISO image to be created (Windows 10, or Windows 7), and then click **Next**.

Once a BMR ISO is created, the ISO file can be used for the same OS level. The following OS levels can use the same ISO:

- ISO created using Windows 7 WAIK – works for Windows 2008, 2008 R2
- ISO created using Windows 10 ADK – works for Windows 10, Windows Server 2016 , Windows server 2019, Windows 8, 8.1, Server 2012 and Server 2012 R2



◆ Windows 10

When launched, the utility immediately checks your computer to determine if the Windows Assessment and Deployment Kit (ADK) is already installed. Windows ADK is a Microsoft tool that lets you deploy Windows operating systems to computers.

Note: You can install Windows 10 ADK on computers running the following operating systems:

- Windows 7
- Windows Server 2008
- Windows Server 2008 R2
- Windows 8
- Windows Server 2012
- Windows 8.1
- Windows 2012 R2
- Windows 10, Windows Server 2016

♦ **Windows 7**

When launched, the utility immediately checks your computer to determine if the Windows Automated Installation Kit (AIK) is already installed. Windows AIK is a Microsoft tool that lets you deploy Windows operating systems to computers.

Note: You can install Windows AIK for Windows 7 on computers running the following operating systems:

- Windows 7
- Windows Server 2008
- Windows Server 2008 R2

3. To create the bootable ISO image, Windows ADK or Windows AIK (as applicable) must be installed on the computer.
- a. If Windows ADK (or AIK) is installed, the utility will advance to the Select the Boot Kit Method screen to let you continue creating the boot kit.
 - b. If Windows ADK (or AIK) is not installed, the corresponding Windows Information screen opens. You need to download and install Windows ADK (or AIK) from the Microsoft Download Center.

Note: For more information on installing Windows ADK (or AIK), see the following websites:

- ♦ [Installing Windows ADK](#)
- ♦ [Installing Windows AIK for Windows 7](#)

You can install Windows ADK (or AIK) using either of the following methods:

- Download the installation media directly from the Microsoft website and install Windows ADK (or AIK) on your computer.

- Click the links on the information screen to open the Microsoft website so that you can download Windows ADK (or AIK) and install it on your computer.

After you install Windows ADK (or AIK), click Next and the utility will advance to the Select the Boot Kit Method screen to let you continue creating the boot kit.


Note: For Windows ADK installation, the following features are required to support creating the boot kit:

- Deployment Tools
- Windows Preinstallation Environment (Windows PE)

Note: For Windows AIK installation, select Windows AIK Setup.

Determine the Method to Generate a BMR ISO Image

The Create Boot Kit utility provides two options for generating an ISO image:



- [Create Bootable BMR ISO Image](#)

This method creates an ISO image that you can then burn onto a CD/DVD for storage. This is the default option. For more information, see [Create an Arcserve UDP Agent \(Windows\) BMR ISO Image for a CD/DVD](#).

- [Create Bootable BMR USB Stick](#)

This method creates an ISO image and burns it directly onto a portable USB stick for storage. For more information, see [Create an Arcserve UDP Agent \(Windows\) BMR ISO Image for a USB Stick](#).

You can then use either of these bootable media 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, it is good practice to create a new ISO image every time you update Arcserve UDP Agent (Windows).

Note: If you are performing a BMR on a virtual machine (VM), then you can also directly attach the ISO image to the VM to start the BMR process without having to first burn it onto a CD/DVD.

Create an Arcserve UDP Agent (Windows) BMR ISO Image for a CD/DVD

The process to create an Arcserve UDP Agent (Windows) BMR ISO image consists of:

- [Create Bootable BMR ISO Image](#)
- [Create a BMR CD/DVD](#)

Create Bootable BMR ISO Image

If you select to create a BMR ISO image, you can then burn this image onto a bootable media (CD or DVD) to initialize the new computer system and allow the bare metal recovery process to begin.

Follow these steps:

1. From the **Select the Boot Kit Method** screen, select **Create Bootable BMR ISO Image** and click **Next**.

The **Select Platform and Destination** dialog opens.

2. Select the applicable platform for the ISO image.

You can select either of the two available options, or both. If you select both platforms, it will result in added time to create the image,

Note: An ISO image that is created from a 32-bit platform should be used to restore a 32-bit server. An ISO image that is created from a 64-bit platform should be used to restore a 64-bit server. If you want to boot a UEFI firmware system, make sure the x64 platform option is selected.

The available options are:

- ◆ BMR ISO image for x86 platform (only).
- ◆ BMR ISO image for x64 platform (only).
- ◆ BMR ISO image for both x86 and x64 platforms.

3. Specify the Destination.

Specify or browse to the location where the BMR ISO image file will be created and stored.

4. Specify the name of the generated BMR ISO image file.
5. After you specify the platform and destination, click **Next**.

The **Select Languages** dialog opens.

6. Select the language for the BMR ISO image. During the BMR procedure, the user interface and keyboard will be integrated with the selected language.

You can select one or more different languages for the BMR ISO image. However, each language selected will result in added time to create the image. The more languages you select, the longer it will take to complete. As a result, you should only select the languages that you actually need.

7. Click **Next**.

The **Specify Drivers** dialog opens.

- Specify the drivers to populate the driver list with drivers to be integrated into the BMR ISO image.

The driver pane is enabled and you can specify any additional drivers that you want to add (or delete) from the BMR ISO image.

Note: When integrating the VirtualBox Host-Only Ethernet Adapter driver into the BMR ISO image, a possible conflict with the Windows ADK components exists. To avoid any conflict, the best practice is not to integrate this driver into the BMR ISO image.

- Include Local Drivers:** Load the local critical device drivers (only oem drivers for NIC, FC, or SCSI) to the driver list. When clicked, the utility checks your computer to determine if there are any critical device drivers that need to be added to the BMR ISO image for this computer. If any critical device drivers are found, they are automatically added to the list.
- Add Driver:** Browse to the drivers you want to be added to the driver list.
- Delete Driver:** Remove any drivers selected from the list that you do not want added to the BMR ISO image.

- Click **Create** to launch the process and create a bootable BMR ISO image.

During the process, the status is displayed.

- When the process is complete a confirmation screen opens to indicate that the BMR ISO image was successfully generated. This screen also displays the location and platform for the image, along with a clickable link to browse to that location.

Create a BMR CD/DVD

After the ISO image is created and saved to the specified destination, you then need to burn this image onto a bootable CD or DVD. You can use this bootable media to initialize the new computer system and allow the bare metal recovery (BMR) process to begin.

To ensure your saved ISO image is always the most recent version:

- You should create a new ISO image every time you update Arcserve UDP Agent (Windows).
- If you saved the ISO image to a remote location, you should burn the CD/DVD only if you need to perform a BMR.
- If you have Arcserve UDP Agent (Windows) installed on multiple computers, you should create a new ISO image (and corresponding CD/DVD) from a known-good computer just prior to performing a BMR so that the image includes all latest Arcserve UDP Agent (Windows) updates.

Create an Arcserve UDP Agent (Windows) BMR ISO Image for a USB Stick

The process to create an Arcserve UDP Agent (Windows) BMR USB stick consists of:

[Prepare the USB Stick](#)

[Create Bootable BMR USB Stick](#)

Prepare the USB Stick

Prior to burning the BMR ISO image onto a USB stick, you must prepare the stick. To create a bootable USB BMR stick, the stick must be made active to enable it to boot a system. You can use the DiskPart command to make the stick active.

Important! If the USB stick needs to be formatted, this process will erase all data currently stored on your USB stick. Verify that there is nothing important on this stick prior to performing this process. If the USB stick was previously formatted, this process will overwrite any files with the same name.

Follow these steps:

1. Open a command prompt (with administrative rights if required by your OS).
2. Type **Diskpart** and press **Enter**.
3. Type **List Disk** and press **Enter**.
A listing of all detected disks is displayed. Determine which of the displayed disks is your USB disk.
4. Select the USB disk by typing **Select Disk <n>** ("n" is the disk number for the USB disk), and press **Enter**.
5. Type **Clean** and press **Enter**.
The system will display "DiskPart succeeded in cleaning the disk."
6. Type **create partition primary** and press **Enter**.
The system will display "succeeded in creating the specified partition".
7. Type **select partition 1** and press **Enter**.
The system will display "Partition 1 is now the selected partition."
8. Type **active** and press **Enter**.
The system will display "DiskPart marked the current partition as active."
9. If necessary, format the USB stick with FAT32 or NTFS file system.
Type **format fs=fat32 quick** or **format fs=ntfs quick**

The USB stick is now prepared and ready for use.

```
C:\Windows\System32>diskpart

Microsoft DiskPart version 6.1.7600
Copyright (C) 1999-2008 Microsoft Corporation.
On computer: <computer name>

DISKPART> list disk

   Disk ###  Status              Size               Free               Dyn  Gpt
   -----  -
   Disk 0    Online              465 GB             1024 KB            *
   Disk 1    Online              3745 MB            0 B

DISKPART> select disk 1

Disk 1 is now the selected disk.

DISKPART> clean

DiskPart succeeded in cleaning the disk.

DISKPART> create partition primary

DiskPart succeeded in creating the specified partition.

DISKPART> select partition 1

Partition 1 is now the selected partition.

DISKPART> active

DiskPart marked the current partition as active.

DISKPART> format fs=fat32 quick

   100 percent completed

DiskPart successfully formatted the volume.

DISKPART> exit_
```

Create Bootable BMR USB Stick

If you select to create a bootable BMR (Bare Metal Recovery) USB stick, you can then burn this ISO image directly onto a USB stick to initialize the new computer system and allow the bare metal recovery process to begin.

Follow these steps:

1. If necessary, prepare the USB stick. For more information, see [Prepare the USB Stick](#).
2. From the **Select the Boot Kit Method** screen, select **Create Bootable BMR USB Stick** and click **Next**.

The **Select Platform and Destination** dialog opens.

3. Select the applicable platform for the ISO image.

You can select either of the two available options, or both. If you select both platforms, it will result in added time to create the image,

Note: An ISO image that is created from a 32-bit platform should be used to restore a 32-bit server. An ISO image that is created from a 64-bit platform should be used to restore a 64-bit server. If you want to boot a UEFI firmware system, make sure the x64 platform option is selected.

The available options are:

- ◆ BMR ISO image for x86 platform (only).
- ◆ BMR ISO image for x64 platform (only).
- ◆ BMR ISO image for both x86 and x64 platforms.

4. Specify the USB Drive.

Specify or browse to the drive location where the BMR ISO image file will be created and burned onto the USB stick.

Note: For a USB drive, if you want to boot the UEFI firmware system, you should format the USB drive as a FAT32 file system.

5. Verify that a prepared USB stick is inserted in the specified drive.
6. After you specify the platform and location, click **Next**,

The **Select Languages** dialog opens.

7. Select the language for the generated BMR ISO image. During the BMR procedure, the user interface and keyboard will be integrated with the selected language.

You can select one or more different languages for the BMR ISO image. However, each language selected will result in added time to create the image. The more

languages you select, the longer it will take to complete. As a result, you should only select the languages that you actually need.

8. Click **Next**.

The **Specify Drivers** dialog opens.

9. If necessary, select the Integrate additional drivers option.

The driver pane is enabled and you can specify any additional drivers that you want to add (or delete) from the BMR ISO image.

10. Click **Create** to launch the process and create a bootable BMR ISO image.

During the process, the status is displayed.

11. When the process is complete a confirmation screen opens to indicate that the BMR ISO image was successfully generated and burned onto your USB stick. This screen also displays the location and platform for the image, along with a clickable link to browse to that location.

Verify the Boot Kit is Created

After the BMR ISO image has been successfully created, the Create Boot Kit utility displays a link to connect to the location where the image is saved. Verify the BMR ISO image is saved at that location. By default, the image is saved to the User-Profile folder, with a default image name format consisting of:

BMR_<Platform>_<OS Kernel>_<version>(Build xxx).ISO

Example:

BMR_x86x64_w8.1_Version 5.0 (Build 5.0.1717).ISO

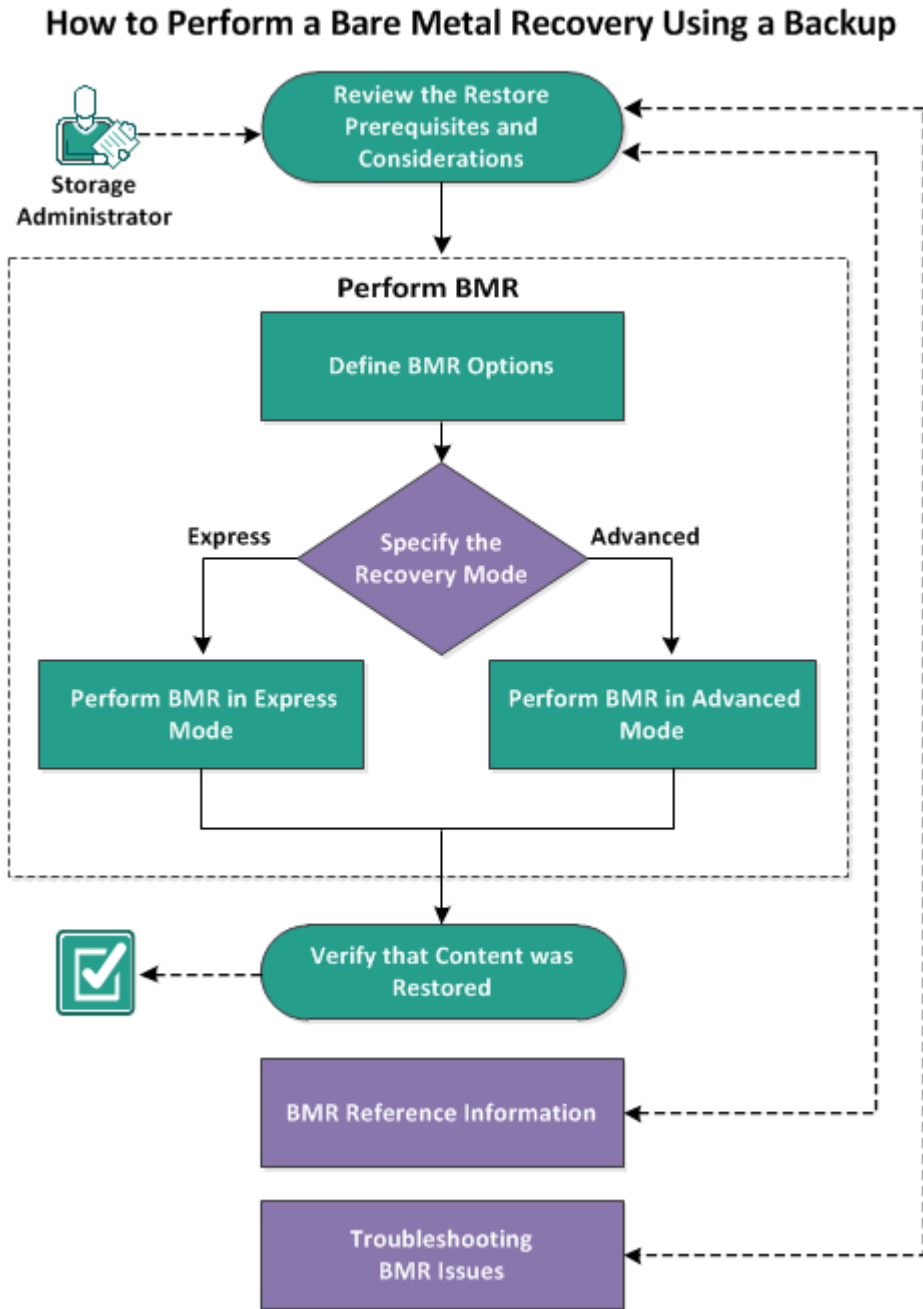
How to Perform a Bare Metal Recovery Using a Backup

Bare Metal Recovery (BMR) is the process of restoring a computer system from "bare metal" including reinstalling the operating system and software applications, and then restoring the data and settings. The BMR process lets you restore a full computer with minimal effort, even to different hardware. BMR is possible because during the block-level backup process, Arcserve UDP Agent (Windows) not only captures the data, but also all information that is related to the following applications:

- Operating system
- Installed applications
- Configuration settings
- Necessary drivers

All relevant information that is necessary to perform a complete rebuild of the computer system from "bare metal" is backed up into a series of blocks and stored on the backup location.

The following diagram illustrates the process for how to perform a BMR using a backup:



Complete the following tasks to perform a BMR using a backup:

1. [Review the BMR Prerequisites and Considerations](#)
2. [Define BMR Options](#)
 - ♦ [Perform BMR in Express Mode](#)
 - ♦ [Perform BMR in Advanced Mode](#)
3. [Verify that the BMR was Successful](#)
4. [BMR Reference Information](#)
5. [Troubleshooting BMR Issues](#)

Review the BMR Prerequisites and Considerations

Verify that the following prerequisites exist before performing a BMR:

- You must have one of the following images:
 - A created BMR ISO image burned onto a CD/DVD
 - A created BMR ISO image burned onto a portable USB stick

Note: Using Arcserve UDP Agent (Windows), you can utilize a Boot Kit Utility to combine a WinPE image and Arcserve UDP Agent (Windows) image to create a BMR ISO image. This ISO image is then burned onto a bootable media. You can then use either of these bootable media (CD/DVD or USB stick) to initialize the new computer system and allow the bare metal recovery process to begin. To ensure your saved image is always the most up-to-date version, create a new ISO image every time you update Arcserve UDP Agent (Windows).

- At least one full backup available.
- At least 2-GB RAM installed on the virtual machine and the source server that you are recovering.
- To recover VMware virtual machines to VMware virtual machines that are configured to behave as physical servers, verify the VMware Tools application is installed on the destination virtual machine.
- Review the [Compatibility Matrix](#) that provides the supported operating systems, databases, and browsers.

Review the following restore considerations:

- Regardless of which method you used to create the Boot Kit image, the BMR process is basically the same.

Note: The BMR process cannot create storage spaces. If the source machine had storage spaces, during BMR you cannot create storage spaces at the destination machine. You can either restore those volumes to regular disks/-volumes or manually create storage spaces before performing the BMR, and then restore the data into those created storage spaces.

- Dynamic disks are restored at the disk level only. If your data is backed up to a local volume on a dynamic disk, you cannot to restore this dynamic disk during BMR. In this scenario, to restore during BMR you must perform one of the following tasks and then perform BMR from the copied Recovery Point:
 - Back up to a volume on another drive.
 - Back up to a remote share.
 - Copy a recovery point to another location.

Note: If you perform BMR with multiple dynamic disks, the BMR may fail because of some unexpected errors (such as fail to boot, unrecognized dynamic volumes, and so on). If this occurs, you should restore only the system disk using BMR, and then after the machine reboot you can restore the other dynamic volumes on a normal environment.

- (Optional) Review the BMR Reference Information. For more information, see the following topics:
 - [How Bare Metal Recovery Works](#)
 - [Operating Systems that Support UEFI/BIOS Conversion](#)
 - [Managing the BMR Operations Menu](#)

Review the following considerations:

- If you upgrade to a newer version or update of Arcserve UDP, you must re-create the BMR ISO using the proper Windows AIK or ADK level to include support for latest features and bug fixes. However, once a BMR ISO is created, the ISO file can be used for the same OS level. The following OS levels can use the same ISO:
 - ISO created using Windows 7 WAIK – works for Windows 2008, 2008 R2
 - ISO create using Windows 8/8.1 ADK – works for Windows 8, 8.1, Server 2012, Server 2012 R2
 - ISO created using Windows 10 ADK – works for Windows 10, Windows Server 2016

Define BMR Options

Prior to initiating the BMR process, you must specify some preliminary BMR options.

Follow these steps:

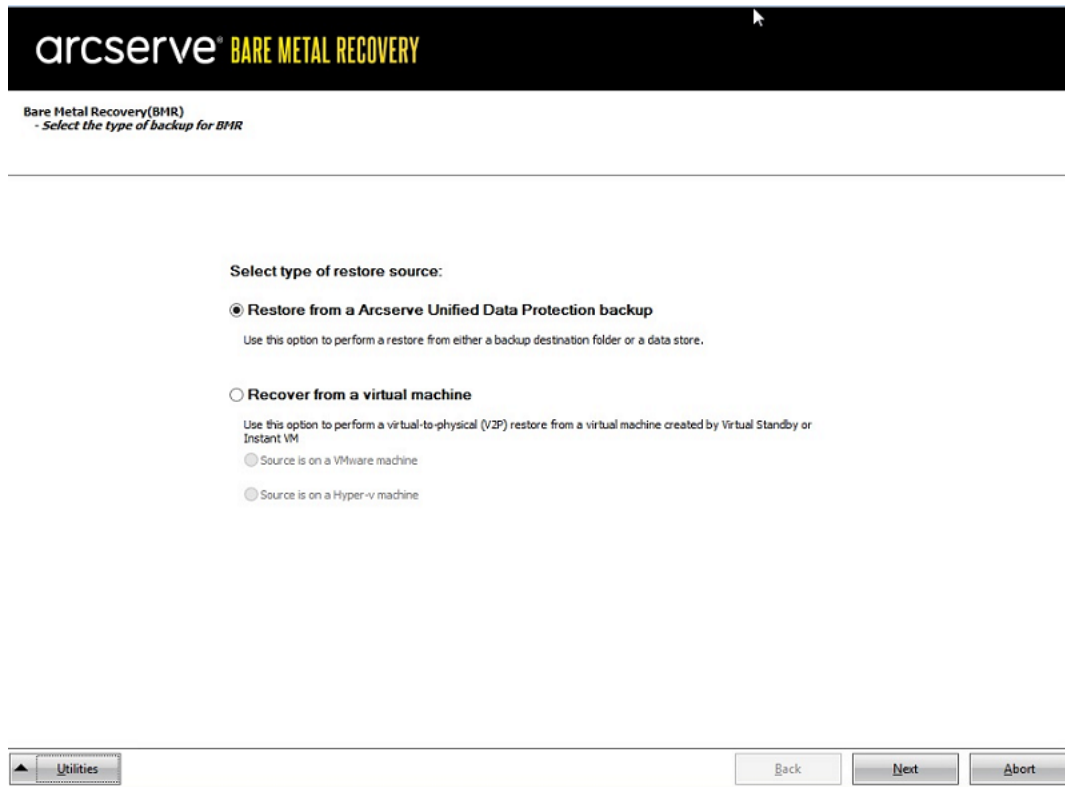
1. Insert the saved Boot Kit image media and boot the computer.
 - If you are using a BMR ISO image burned onto a CD/DVD, insert the saved CD/DVD.
 - If you are using a BMR ISO image burned onto a USB stick, insert the saved USB stick.

The **BIOS Setup Utility** screen is displayed.

2. From the **BIOS Setup Utility** screen, select the CD-ROM Drive option or the USB option to launch the boot process. Select an architecture (x86/x64) and press **Enter** to continue.
3. The Arcserve UDP Agent (Windows) language select screen is displayed. Select a language and click **Next** to continue.



The Bare Metal Recovery process is initiated and the initial BMR wizard screen is displayed.



The BMR wizard screen allows you to select the type of BMR you want to perform:

- **Restore from an Arcserve Unified Data Protection backup**

Use this option to perform a restore from either a backup destination folder or a data store.

This option lets you recover data that was backed up using Arcserve UDP Agent (Windows). This option is used in connection with backup sessions performed with Arcserve UDP Agent (Windows) or with the Arcserve UDP host-based VM backup application.

If you select this option, continue this procedure from here.

- **Recover from a virtual machine**

Use this option to perform a virtual-to-physical (V2P) restore from a virtual standby VM. Virtual-to-physical (V2P) is a term that refers to the migration of an operating system (OS), application programs and data from a virtual machine or disk partition to a computer's main hard disk. The target can be a single computer or multiple computers.

– **Source is on a VMware machine**

Lets you recover data for a machine for which virtual conversion is done to a VMware virtual machine. This option is used in connection with the Arcserve Central Virtual Standby application.

Note: For this option, you can only recover data if the virtual conversion to a VMDK file (for VMware) was performed using Arcserve Central Virtual Standby.

If you select this option, see [Recover using a VMware Virtual Standby VM](#) to continue this procedure.

For more information, see Recover using a VMware Virtual Standby VM in the online help.

– **Source is on a Hyper-V machine**

Lets you recover data for a machine for which virtual conversion is performed to a Hyper-V virtual machine. This option is used in connection with the Arcserve Central Virtual Standby application.

Note: For this option, you can only recover data if the virtual conversion to a VHD file (for Hyper-V) was performed using Arcserve Central Virtual Standby.

If you select this option, see [Recover using a Hyper-V Virtual Standby VM](#) to continue this procedure.

For more information, see Recover using a Hyper-V Virtual Standby VM in the online help.

4. Select **Restore from an Arcserve Unified Data Protection backup** and click **Next**.

The **Select a Recovery Point** wizard screen is displayed.



5. From the **Select a Recovery Point** wizard screen, click **Browse** and select either **Browse from network/local path** or select **Browse from Recovery Point Server**.

a. If you select **Browse** from network/local path, select the machine (or volume) which contains recovery points for your backup image.

Arcserve UDP Agent (Windows) lets you recover from any local drive or from a network share.

- If you recover from a local backup, the BMR wizard automatically detects and displays all volumes containing recovery points.
- If you recover from a remote share, browse to the remote location where the recovery points are stored. If there are multiple machines containing recovery points, all machines are displayed.

You may also need access information (User Name and Password) for the remote machine.

Note: The network must be up and running to browse to remote recovery points. If necessary, you can check/refresh your network configuration information or you can load any missing drivers from the Utilities menu.

- If the BMR module cannot detect any local destination volume, the **Select a Folder** dialog automatically displays. Provide the remote share where the backups are residing.
- If you are restoring from an iSCSI destination, the BMR module may not detect this destination and you need to perform the following:
 1. Click **Utilities**, select **Run** from the pop-up menu, type **cmd**, and then click **OK**.
 2. In the command prompt window, use the following Windows iSCSI commands to set up iSCSI connections:

```
> net start msiscsi  
> iSCSICLI QAddTargetPortal <TargetPortalAddress>  
> iSCSICLI QLoginTarget <TargetName > [CHAP username] [CHAP password]
```

Note: CHAP = Challenge-Handshake Authentication Protocol


For more information about Windows iSCSI command line options, see [link](#).

Note: Extra steps may be needed depending on the iSCSI target software being used. For more information, see the manual of the iSCSI target software.

3. From the BMR screen the disks/volumes connected through the iSCSI disk should be displayed. The iSCSI disk can now be used as the source volume or the backup destination volume.

Note: BMR does not support the case where the OS is installed on an iSCSI disk. Only data disks are supported.

- b. If you select **Browse the Recovery Point Server**, the **Select Agent** dialog displays. Provide the **Recovery Point Server Host Name**, **User Name**, **Password**, **Port**, and **Protocol**. Click **Connect**.




6. Select the folder or Agent Name under Data Store where the recovery points for your backup are stored and click **OK**.

The BMR wizard screen now displays the following information:

- Machine name (in the upper left pane).
- Related backup information (in the upper right pane).
- All the corresponding recovery points (in the lower left pane).
- **Note:** For supported operating systems, you can perform a BMR from a backup performed on a UEFI machine to a BIOS-compatible machine and from a BIOS machine to a UEFI-compatible machine. See [Operating Systems that Support UEFI/BIOS Conversion](#) for a complete listing of firmware conversion supported systems.
- For operating systems that do not support firmware conversion, to perform BMR for a UEFI system, you must boot the computer in UEFI mode. BMR does not support restoring a computer with different firmware. To verify that the boot firmware is UEFI and not BIOS, click **Utilities**, **Abort**.
- For operating systems that do support firmware conversion, after you select a recovery point, if it is detected that the source machine is not the same firmware as your system, you will be asked if you want to convert UEFI to a BIOS-

compatible system or BIOS to UEFI-compatible system.

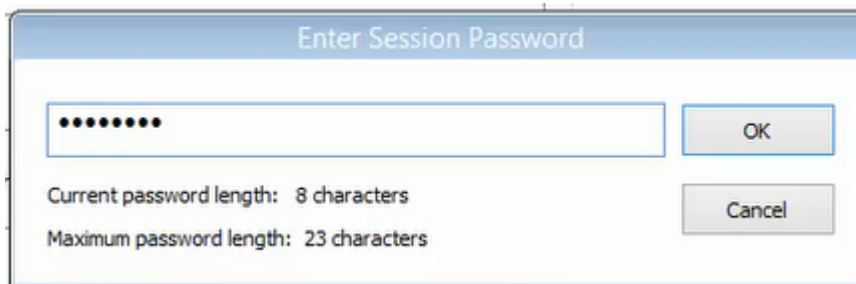


Note: The Arcserve UDP Version 5.0 Update 2 only supports BMR to a smaller disk when the sessions are backed up from Arcserve UDP Version 5.0 Update 2. See the field **Minimum Size Required** for the destination disk size. BMR to a smaller disk is only supported in **Advanced Mode**.

7. Select which recovery point to restore.

The related information for the selected recovery point is displayed (in the lower right pane). This display includes such information as the type of backup that was performed (and saved), the backup destination, and the volumes that were backed up.

If the recovery point contains encrypted sessions (the recovery point clock icon includes a lock), a password required screen appears. Enter the session password and click **OK**.



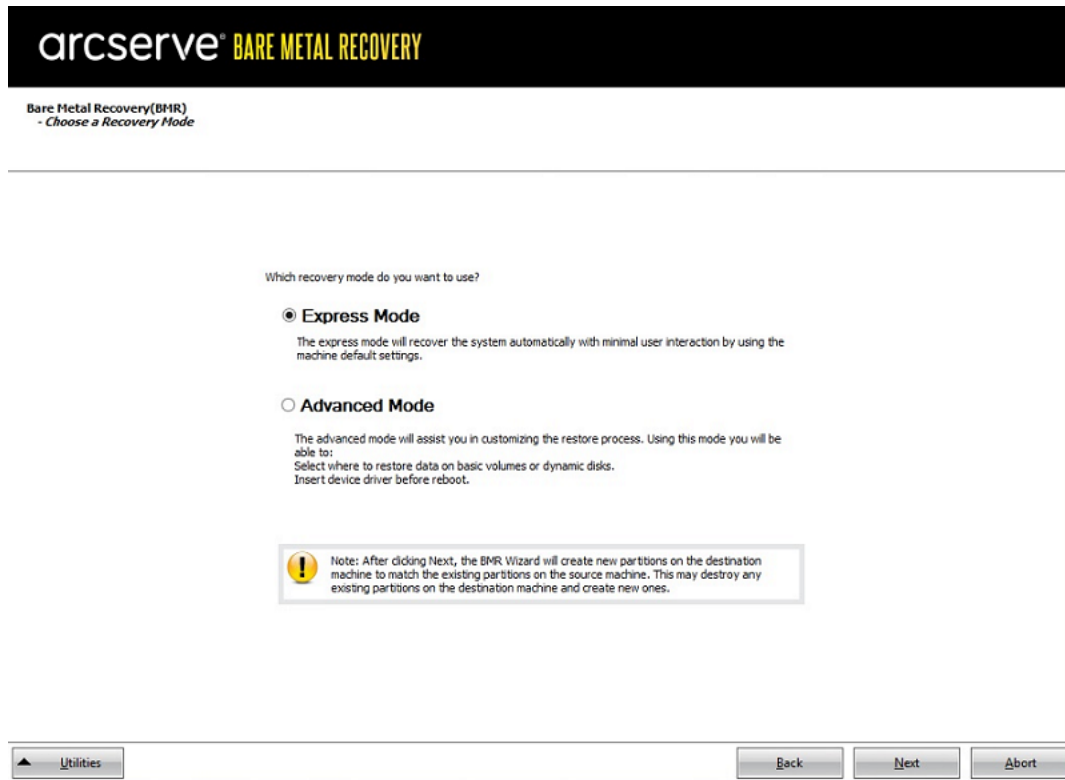
Notes:

If you are restoring from a Arcserve UDP Recovery Point Server, you are asked to provide a session password.

If your machine is a Domain Controller, Arcserve UDP Agent (Windows) supports a 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](#) if you want minimal interaction during the recovery process.
- ◆ Select [Advanced Mode](#) if you want to customize the recovery process.

Default: Express Mode.

Perform BMR in Express Mode

The Express Mode requires minimal interaction during the recovery process.

Follow these steps:

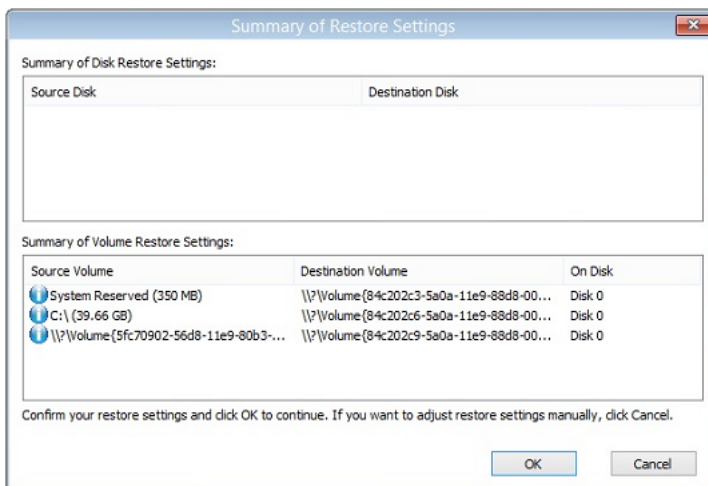
1. From the **Choose a Recovery Mode** dialog, select **Express Mode** and click **Next**.

A confirmation dialog appears.

2. Click **Yes**.

The **Summary of Disk Restore Settings** screen opens, displaying a summary of the volumes that are going to be restored.

Note: On the bottom of restore summary window, the drive letters listed in **Destination Volume** column are automatically generated from the Windows Pre-installation Environment (WinPE). They can be different from the drive letters listed in **Source Volume** column. However, the data is still restored to proper volume even if drive letters are different.



3. 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](#).

- ◆ By default, **Do not start Agent service automatically after reboot** is enabled.
- ◆ If necessary, you can cancel or abort the operation at any time.



Note: Selecting the check box of **Boot system to Advanced Boot Options..** helps when you restore one machine with Active Directory.

4. 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.

5. 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.

6. 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**.
A confirmation dialog appears.
2. Click **Yes**.

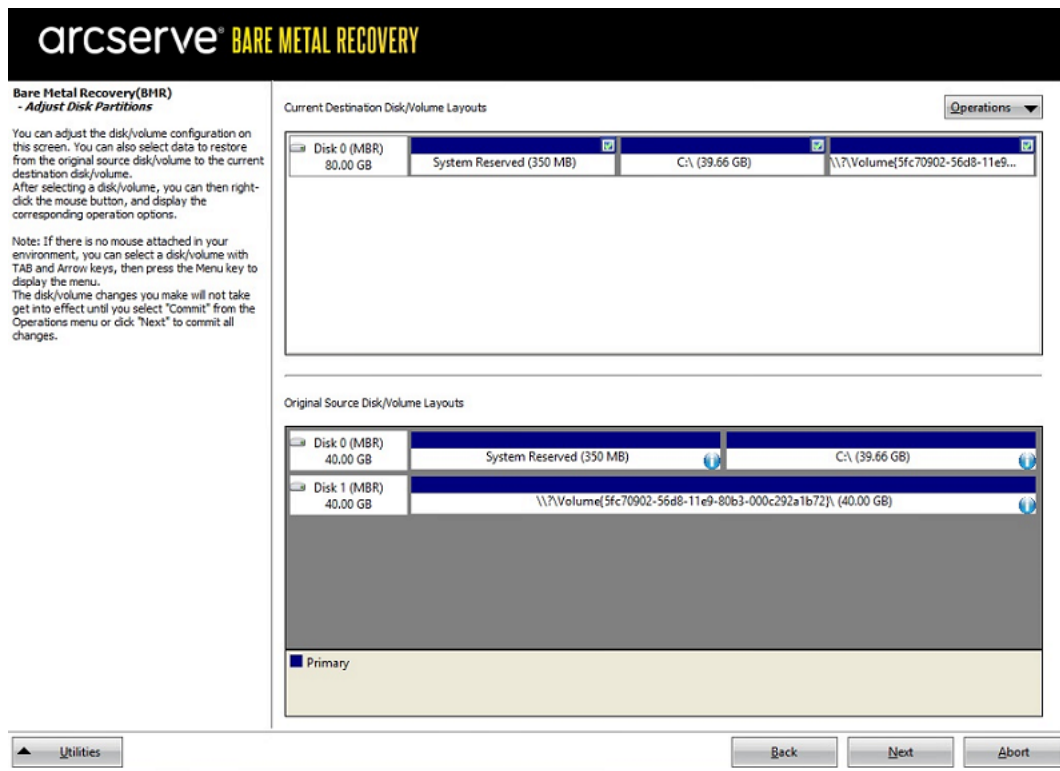
The BMR utility starts locating the machine that is going to be recovered and displays the corresponding disk partition information.

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.

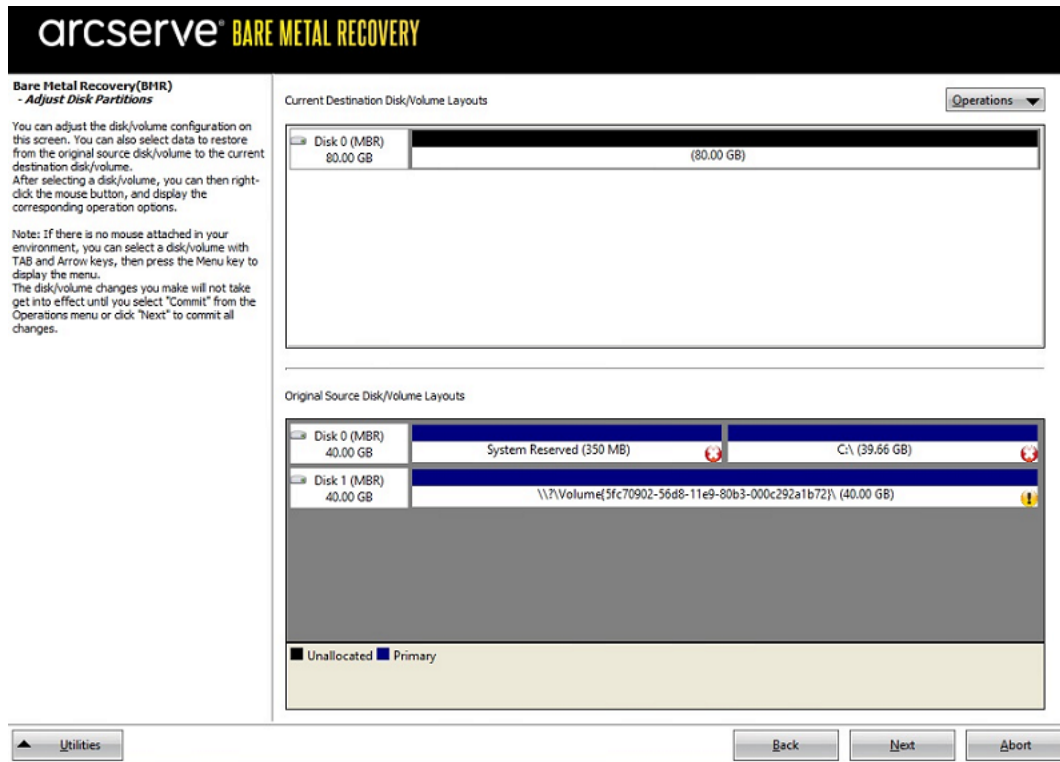
You can create volumes to a smaller disk based on the suggested **Minimum disk space required**. In the example, the original size of the volume is 81568 MB. When you create the volume on the target disk, the suggested minimum size is 22752

MB. In this case, you can create the original volume with a size of 22752 MB.



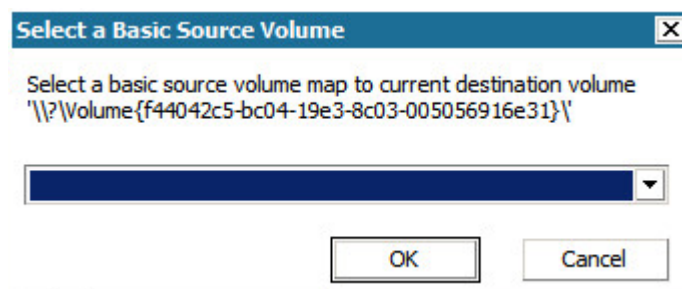
Note: If you perform BMR and you restore the system volume to a disk which is not configured as the boot disk, it will fail to boot the machine after BMR is completed. Ensure that you are restoring the system volume to a properly configured boot disk.

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.



3. If the current disk information you are seeing does not appear correct, you can access the **Utilities** menu and check for missing drivers.
4. 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](#).
5. 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.

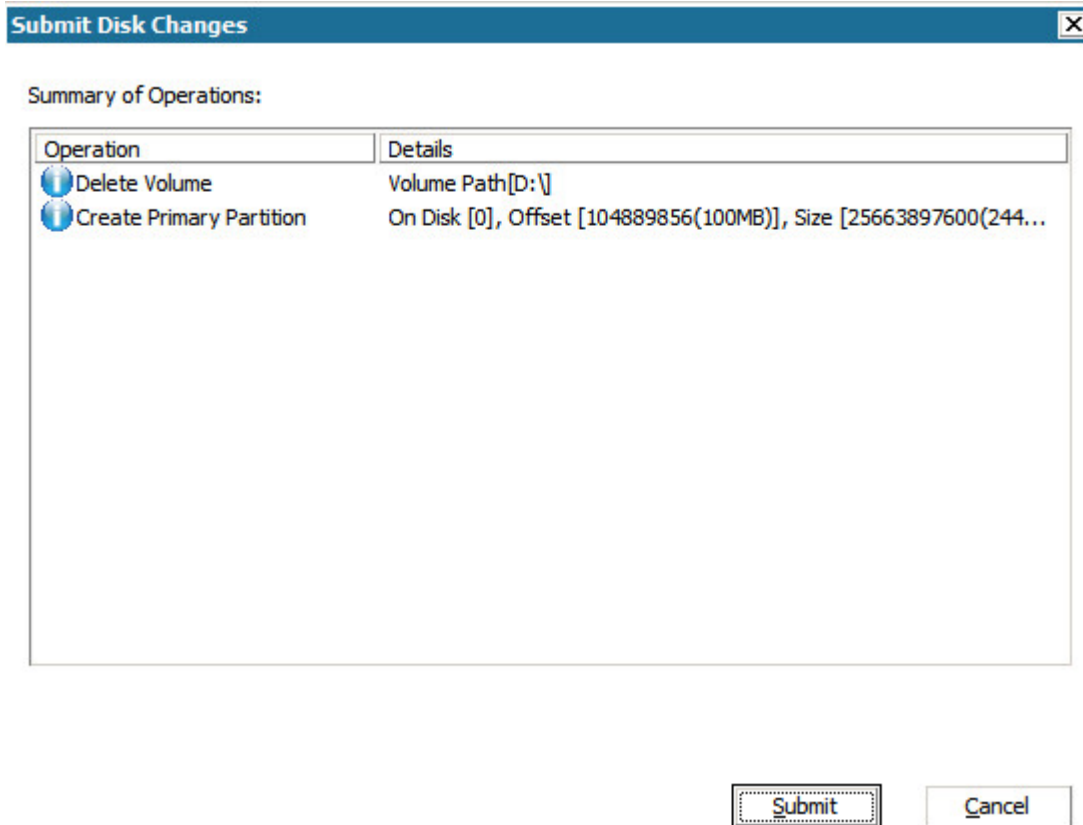
The **Select a Basic Source Volume** dialog opens.



6. From **Select a Basic Source Volume** dialog, click the drop-down menu and select the available source volume to assign to the selected target volume. Click **OK**.
 - On the target volume, a checkmark icon is displayed, indicating that this target volume has been mapped to.

- On the source volume, the red X icon changes to a green icon, indicating that this source volume has been assigned to a target volume.
6. When you are sure all volumes that you want to restore and all volumes containing system information are assigned to a target volume, click **Next**.

The **Submit Disk Changes** screen opens, displaying a summary of the selected operations. For each new volume being created, the corresponding information is displayed.



7. When you have verified the summary information is correct, click **Submit**. (If the information is not correct, click **Cancel**).

Note: All operations to the hard drive do not take effect until you submit it.


On the target machine, the new volumes are created and mapped to the corresponding source machine.

8. When the changes are completed, click **OK**.

The Summary of Disk Restore Settings screen opens, displaying a summary of the volumes that are going to be restored.

Note: On the bottom of restore summary window, the drive letters listed in "Destination Volume" column are automatically generated from the Windows Pre-installation Environment (WinPE). They can be different from the drive letters listed

in "Source Volume" column. However, the data is still restored to proper volume even if drive letters are different.



9. After you have verified that the summary information is correct, click **OK**.


The restore process starts. The BMR wizard screen displays the restore status for each volume.

- Depending upon the size of the volume being restored, this operation can take some time.
- During this process you are restoring, block-by-block whatever you had backed up for that recovery point and creating a replica of the source machine on the target machine.
- By default, the option to reboot your system automatically after recovery is selected. If necessary, you can clear this option and you can reboot manually at a later time.

Important: If you are performing an authoritative restore of an active directory after a BMR, you must uncheck the option **Automatically reboot your system after recovery** and for more information, see [How to Perform an Authoritative Restore of an Active Directory after a BMR](#).

- If necessary, you can select Do not start Agent service automatically after reboot.

- If necessary, you can cancel or abort the operation at any time.



Note: Selecting the check box of "Boot the system to Advanced boot options .. helps when you restore one machine with Active Directory.

10. From the **Utilities** menu, you can access the BMR **Activity Log** and you can use the **Save** option to save the Activity Log.

By default, the Activity Log is saved to the following location:

X:\windows\system32\dr\log.

Note: To avoid getting a Windows-generated error, do not save the Activity Log on your desktop or create a folder on your desktop using the **Save As** option from the BMR **Activity Log** window.

11. If you are restoring to dissimilar hardware (the SCSI/FC adapter which used to connect hard drives could have been changed) and no compatible driver is detected in your original system, a "driver injection" page is displayed to allow you to provide drivers for these devices.

You can browse and select drivers to inject to the recovered system so that even if you are recovering to a machine with dissimilar hardware, you can still bring back the machine after BMR.

12. When the BMR process is completed, a confirmation notification is displayed.

Verify that the BMR was Successful

To verify that the BMR was successful, perform the following tasks:

- Reboot the operating system.
- Verify all systems and applications function correctly.
- Verify all network settings are properly configured.
- Verify the BIOS is configured to boot from the disk on which the boot volume was restored to.
- When the BMR is completed, be aware of the following conditions:
 - The first backup that is performed after the BMR is a Verify Backup.
 - When the machine has been rebooted, you may need to configure the network adapters manually if you restored to dissimilar hardware.

Note: When the machine is rebooting, a Windows Error Recovery screen may be displayed indicating that Windows did not shut down successfully. If this occurs, you can safely ignore this warning and continue to start Windows normally.

- For dynamic disks, if the status of the disk is offline, you can manually change it to online from the disk management UI (accessed by running the Diskmgmt.msc control utility).
- For dynamic disks, if the dynamic volumes are in a failed redundancy status, you can manually resynchronize the volumes from the disk management UI (accessed by running the Diskmgmt.msc control utility).

BMR Reference Information

[How Bare Metal Recovery Works](#)


[Operating Systems that Support UEFI or BIOS Conversion](#)

[Managing the BMR Operations Menu](#)

How Bare Metal Recovery Works

Bare Metal Recovery is the process of restoring a computer system from "bare metal" by reinstalling the operating system and software applications, and then restoring the data and settings. The most common reasons for performing a bare metal recovery are because your hard drive either fails or becomes full and you want to upgrade (migrate) to a larger drive or migrate to newer hardware. Bare metal recovery is possible because during the block-level backup process, Arcserve UDP Agent (Windows) captures not only the data, but also all information related to the operating system, installed applications, configuration settings, necessary drivers, and so on. All relevant information that is necessary to perform a complete rebuild of the computer system from "bare metal" is backed up into a series of blocks and stored on the backup location.

Note: Dynamic disks are restored at disk level only. If your data is backed up to a volume on a dynamic disk, you will not be able to restore this dynamic disk (including all its volumes) during BMR.



When you perform a bare metal recovery, the Arcserve UDP Agent (Windows) boot disk is used to initialize the new computer system and allow the bare metal recovery process to begin. When the bare metal recovery is started, Arcserve UDP Agent (Windows) will prompt you to select or provide a valid location to retrieve these backed up blocks from, as well as the recovery point to be restored. You may also be prompted to provide valid drivers for the new computer system if needed. When this connection and configuration information is provided, Arcserve UDP Agent (Windows) begins to pull the specified backup image from the backup location and restore all backed up blocks to the new computer system (empty blocks will not be restored). After the bare metal recovery image is fully restored to the new computer system, the machine will be back to the state that it was in when the last backup was performed, and Arcserve UDP Agent (Windows) backups will be able to continue as scheduled. (After completion of the BMR, the first backup will be a Verify Backup).

Operating Systems that Support UEFI/BIOS Conversion

If it is detected that the operating system of your source machine is not the same firmware as your system, you will be asked if you want to convert UEFI to a BIOS-compatible system or BIOS to UEFI-compatible system. The following table lists each operating system and the type of conversion supported.

Operating System (OS)	CPU	uEFI to BIOS	BIOS to uEFI
Windows Server 2008	x86	No	No
Windows Server 2008	x64	Yes	Yes
Windows Server 2008 R2	x64	Yes	Yes
Windows 7	x86	No	No
Windows 7	x64	Yes	Yes
Windows 8	x86	No	No
Windows 8	x64	Yes	Yes
Windows Server 2012	x64	Yes	Yes
Windows 8.1	x86	No	No
Windows 8.1	x64	Yes	Yes
Windows 10	x86	No	No
Windows 10	x64	Yes	Yes
Windows Server 2012 R2	x64	Yes	Yes
Windows Server 2016	x64	Yes	Yes
Windows Server 2019	x64	Yes	Yes

Managing the BMR Operations Menu

The BMR Operations menu consists of the following three types of operations:

- Disk Specific Operations
- Volume/Partition Specific Operations
- BMR Specific 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.
- Select the disk header on the BMR UI, click **Operations**, and use the **Clean Disk** operation 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.

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 Bare Metal Recovery Using a Virtual Standby VM or Instant VM

Bare Metal Recovery (BMR) is the process of restoring a computer system from "bare metal" including reinstalling the operating system and software applications, and then restoring the data and settings. The BMR process lets you restore a full computer with minimal effort, even to different hardware. BMR is possible because during the block-level backup process, Arcserve UDP Agent (Windows) not only captures the data, but also all information that is related to the following applications:

- Operating system
- Installed applications
- Configuration settings
- Necessary drivers

All relevant information that is necessary to perform a complete rebuild of the computer system from "bare metal" is backed up into a series of blocks and stored on the backup location.

To perform a BMR from a virtual machine, use one of the following ways:

- Connect to the ESX server directly using the IP address
- Add the correct DNS setting in your BMR machine and resolve the hostname to the IP address

Complete the following tasks to perform a BMR using a virtual standby VM or Instant VM:

1. [Review the BMR Prerequisites and Considerations](#)
2. [Define BMR Options](#)
 - ♦ [Recover Using a Hyper-V Virtual Standby VM or Instant VM](#)
 - ♦ [Recover Using a VMware Virtual Standby VM or Instant VM](#)
 - ♦ [Perform BMR in Express Mode](#)
 - ♦ [Perform BMR in Advanced Mode](#)
3. [Verify that the BMR was Successful](#)
4. [BMR Reference Information](#)
5. [Troubleshooting BMR Issues](#)

Review the BMR Prerequisites and Considerations

Verify that the following prerequisites exist before performing a BMR:

- You must have one of the following images:
 - A created BMR ISO image burned onto a CD/DVD
 - A created BMR ISO image burned onto a portable USB stick
- **Note:** Arcserve UDP Agent (Windows) utilizes a Boot Kit Utility to combine a WinPE image and Arcserve UDP Agent (Windows) image to create a BMR ISO image. This ISO image is then burned onto a bootable media. You can then use either of these bootable media (CD/DVD or USB stick) to initialize the new computer system and allow the bare metal recovery process to begin. To ensure your saved image is always the most up-to-date version, create a new ISO image every time you update Arcserve UDP Agent (Windows).
- At least one full backup available.
- At least 1-GB RAM installed on the virtual machine and the source server that you are recovering.
- To recover VMware virtual machines to VMware virtual machines that are configured to behave as physical servers, verify the VMware Tools application is installed on the destination virtual machine.
- Review the [Compatibility Matrix](#) that provides the supported operating systems, databases, and browsers.

Review the following restore considerations:

- Regardless of which method you used to create the Boot Kit image, the BMR process is basically the same.
 - **Note:** The BMR process cannot create storage spaces. If the source machine had storage spaces, during BMR you cannot create storage spaces at the destination machine. You can either restore those volumes to regular disks/-volumes or manually create storage spaces before performing the BMR, and then restore the data into those created storage spaces.
- Dynamic disks are restored at the disk level only. If your data is backed up to a local volume on a dynamic disk, you cannot to restore this dynamic disk during BMR. In this scenario, to restore during BMR you must perform one of the following tasks and then perform BMR from the copied Recovery Point:
 - Back up to a volume on another drive.
 - Back up to a remote share.
 - Copy a recovery point to another location.

Note: If you perform BMR with multiple dynamic disks, the BMR may fail because of some unexpected errors (such as fail to boot, unrecognized dynamic volumes, and so on). If this occurs, you should restore only the system disk using BMR, and then after the machine reboot you can restore the other dynamic volumes on a normal environment.

- If you attempt to perform a BMR on a Hyper-V VM with a 4 KB disk, add this 4 KB disk to the SCSI controller. If you add it to the IDE controller, the disk will not be detected in the Windows PE system.
- (Optional) Review the BMR Reference Information. For more information, see the following topics:
 - [How Bare Metal Recovery Works](#)
 - [Operating Systems that Support UEFI/BIOS Conversion](#)
 - [Managing the BMR Operations Menu](#)

Review the following considerations:

- If you upgrade to a newer version or update of Arcserve UDP, you must re-create the BMR ISO using the proper Windows AIK or ADK level to include support for latest features and bug fixes. However, once a BMR ISO is created, the ISO file can be used for the same OS level. The following OS levels can use the same ISO:
 - ISO created using Windows 7 WAIK – works for Windows 2008, 2008 R2
 - ISO create using Windows 8/8.1 ADK – works for Windows 8, 8.1, Server 2012, Server 2012 R2
 - ISO created using Windows 10 ADK – works for Windows 10

Define BMR Options

Prior to initiating the BMR process, you must specify some preliminary BMR options.

Follow these steps:

1. Insert the saved Boot Kit image media and boot the computer.
 - ◆ If you are using a BMR ISO image burned onto a CD/DVD, insert the saved CD/DVD.
 - ◆ If you are using a BMR ISO image burned onto a USB stick, insert the saved USB stick.

The **BIOS Setup Utility** screen is displayed.

2. From the **BIOS Setup Utility** screen, select the CD-ROM Drive option or the USB option to launch the boot process. Select an architecture (x86/x64) and press **Enter** to continue.
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.

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 machine

Use this option to perform a virtual-to-physical (V2P) restore from a virtual machine created by Virtual Standby or Instant VM

Source is on a VMware machine

Source is on a Hyper-V machine

The BMR wizard screen allows you to select the type of BMR you want to perform:

- **Restore from an Arcserve Unified Data Protection backup**

Use this option to perform a restore from either a backup destination folder or a data store.

This option lets you recover data that was backed up using Arcserve UDP Agent (Windows). This option is used in connection with backup sessions performed with Arcserve UDP Agent (Windows) or with the Arcserve UDP host-based VM backup application.

For more information, see [How to Perform a Bare Metal Recovery Using a Backup](#) in the online help.

- **Recover from a Virtual Standby VM**

Use this option to perform a virtual-to-physical (V2P) restore from a virtual standby VM or Instant VM. Virtual-to-physical (V2P) is a term that refers to the migration of an operating system (OS), application programs and data from a virtual machine or disk partition to a computer's main hard disk. The target can be a single computer or multiple computers.

– **Source is on a VMware machine**

Lets you recover data for a machine for which virtual conversion is done to a VMware virtual machine. This option is used in connection with the Arcserve Central Virtual Standby or Instant VM application.

Note: For this option, you can only recover data if the virtual conversion to a VMDK file (for VMware) was performed using Arcserve Central Virtual Standby or Instant VM.

If you select this option, see [Recover using a VMware Virtual Standby VM or Instant VM](#) to continue this procedure.

– **Source is on a Hyper-V machine**

Lets you recover data for a machine for which virtual conversion is performed to a Hyper-V virtual machine. This option is used in connection with the Arcserve Central Virtual Standby or Instant VM application.

Note: For this option, you can only recover data if the virtual conversion to a VHD file (for Hyper-V) was performed using Arcserve Central Virtual Standby or Instant VM.

If you select this option, see [Recover using a Hyper-V Virtual Standby VM or Instant VM](#) to continue this procedure.

4. Select **Recover from a Virtual Standby VM**. Then select one of the sources.

- If you select the **Source is on a VMware machine** option, see [Recover using a VMware Virtual Standby VM or Instant VM](#) to continue this procedure.
- If you select the **Source is on a Hyper-V machine** option, see [Recover using a Hyper-V Virtual Standby VM or Instant VM](#) to continue this procedure.

Recover using a Hyper-V Virtual Standby VM or Instant VM

Arcserve UDP Agent (Windows) provides the capability to perform Bare Metal Recovery for virtual-to-physical (V2P) machines. This feature lets you perform virtual-to-physical recovery from the latest state of a standby or instant virtual machine and helps you reduce the loss of your production machine.

Follow these steps:

1. From the select the Type of Bare Metal Recovery (BMR) wizard screen, select the **Recover from a Virtual Standby VM** and select **Source is on a Hyper-V machine** option.

Use this option to perform a virtual-to-physical restore from a virtual standby VM or Instant VM. The term virtual-to-physical refers to the migration of an operating system (OS), application programs and data from a virtual machine or disk partition to a computer's main hard disk. The target can be a single computer or multiple computers.

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 machine


Use this option to perform a virtual-to-physical (V2P) restore from a virtual machine created by Virtual Standby or Instant VM

Source is on a VMware machine

Source is on a Hyper-V machine


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 Arcserve Central Virtual Standby or Instant VM.



4. Select the virtual machine that contains the recovery point snapshots for your backup image.

The backup sessions (recovery point snapshots) for the selected virtual machine are displayed.



5. Select the virtual machine backup session (recovery point snapshot) that you want to recover.

The corresponding details for the selected recovery point snapshot (virtual machine name, backup session name, backed up volumes) are displayed in the right pane.

In addition to selecting one of the listed recovery points, you also have the option to select the **Current State** or the **Latest State** recovery point.

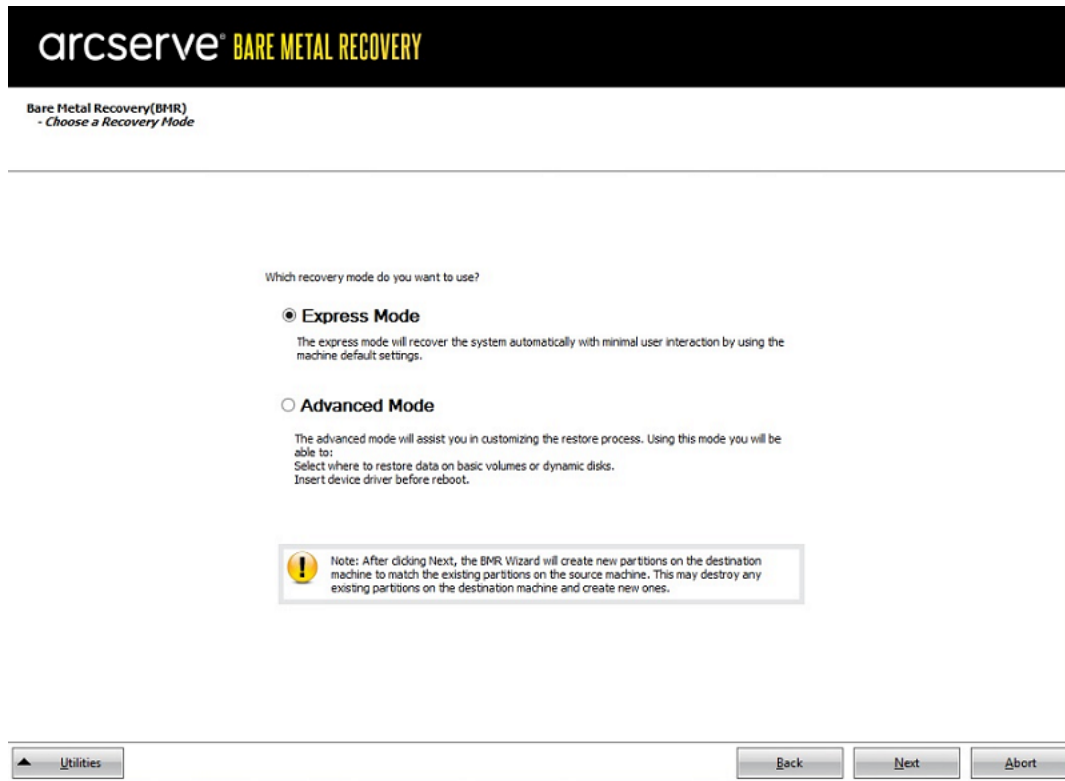
- If the virtual machine that you are recovering from is powered on, the **Current State** recovery point is displayed.

Note: If the virtual machine is powered on, then any data changes in the virtual machine after the BMR process started will not be recovered.

- If the virtual machine that you are recovering from is powered off, the **Latest State** recovery point is displayed.

6. Verify this is the recovery point that you want to restore and click **Next**.

A BMR wizard screen is displayed with the available recovery mode options.



The available options are **Advanced Mode** and **Express Mode**.

- Select **Express Mode** if you want minimal interaction during the recovery process. For more information see, [Perform BMR in Express Mode](#).
- Select **Advanced Mode** if you want to customize the recovery process. For more information, see [Perform BMR in Advanced Mode](#).

Default: Express Mode.

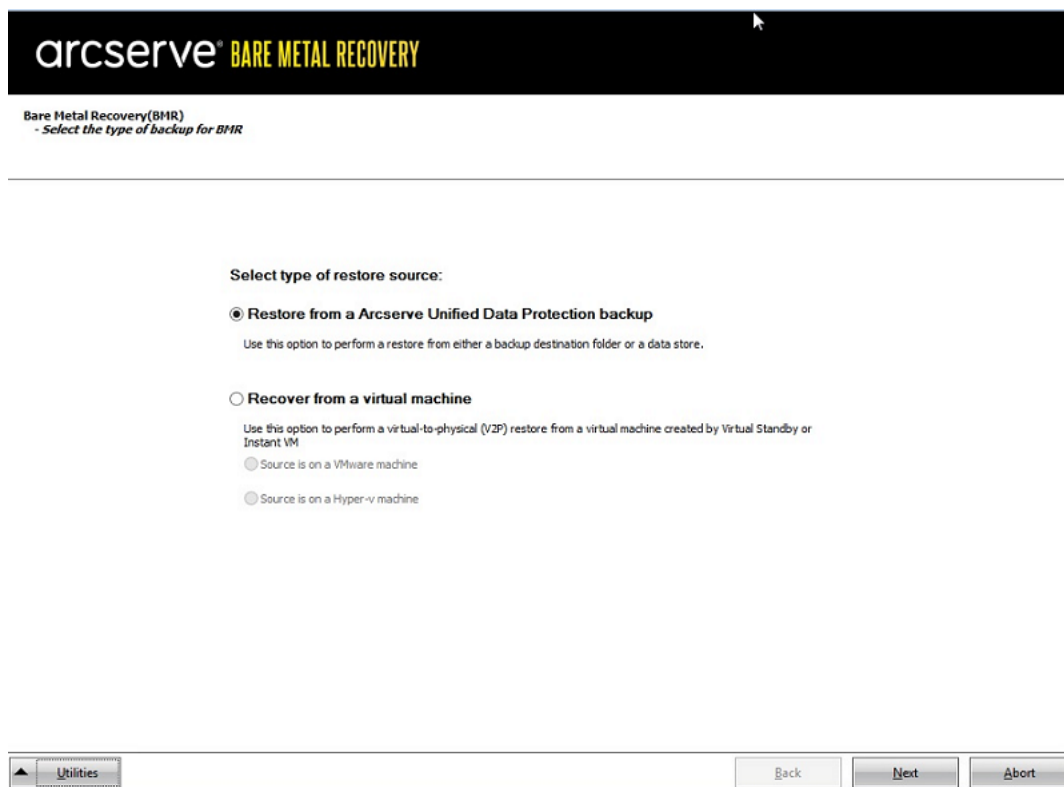
Recover using a VMware Virtual Standby VM or Instant VM

The Arcserve UDP Agent (Windows) provides the capability to perform Bare Metal Recovery for virtual-to-physical (V2P) machines. This feature lets you perform virtual-to-physical recovery from the latest state of a standby virtual machine and helps you reduce the loss of your production machine.

Follow these steps:

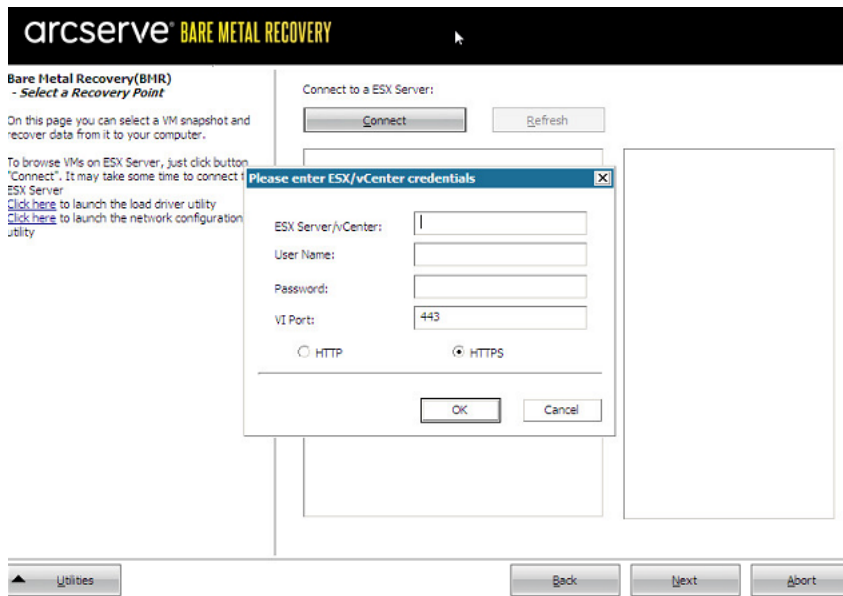
1. From the select the Type of Bare Metal Recovery (BMR) wizard screen, select the **Recover from a virtual machine** and select the **Source is on a VMware machine** option.

Use this option to perform a virtual-to-physical restore from a virtual standby VM or Instant VM. The term virtual-to-physical refers to the migration of an operating system (OS), application programs and data from a virtual machine or disk partition to a computer's main hard disk. The target can be a single computer or multiple computers.



2. Click **Next**.

The **Select a Recovery Point** screen is displayed with the **ESX/VC Credentials** dialog.



3. Enter the credential information and click **OK**.


Note: If you are connecting to a vCenter, you do not need an Administrator permission at the vCenter Server level but you must have an Administrator permission at the Datacenter level. In addition, you must have the following permissions at the vCenter Server level:

- Global, DisableMethods and EnableMethods
- Global, License

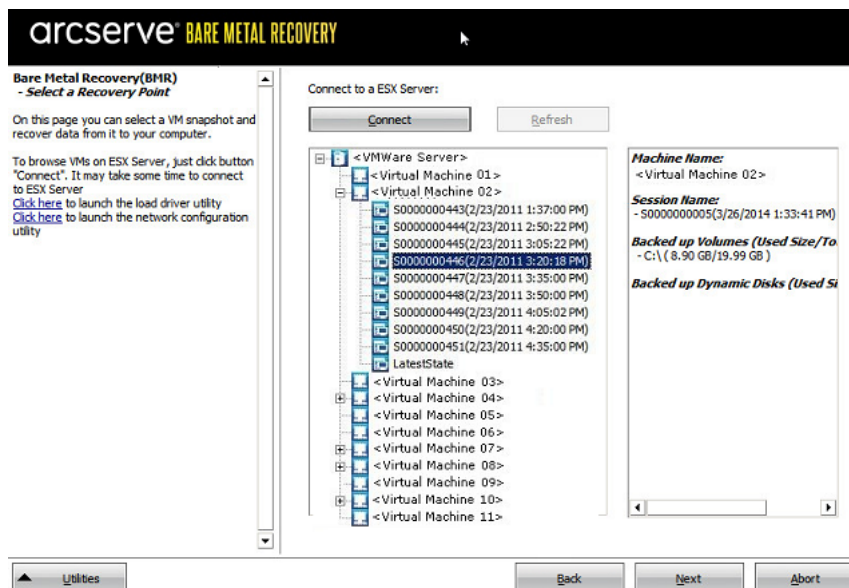
The **Select a Recovery Point** screen is displayed.

The Arcserve 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.

Note: If the virtual machine is powered on, then any data changes in the virtual machine after the BMR process started will not be recovered.

- If the virtual machine that you are recovering from is powered off, the **Latest State** recovery point is displayed.

6. Verify this is the recovery point that you want to restore and click **Next**.

A BMR wizard screen is displayed with the available recovery mode options.



The available options are **Advanced Mode** and **Express Mode**.

- Select **Express Mode** if you want minimal interaction during the recovery process. For more information see, [Perform BMR in Express Mode](#).
- Select **Advanced Mode** if you want to customize the recovery process. For more information, see [Perform BMR in Advanced Mode](#).

Default: Express Mode.

Note: When the virtual machine is located on the VMware ESX(i) server whose edition is 5.0 or 5.1.x, you need to create registry key in the BMR machine. For details, view [link](#).

Create Registry Key in the BMR machine

You can create registry key in the BMR machine. The key is required when the virtual machine is located on the VMware ESX(i) server whose edition is 5.0 or 5.1.x.

Follow these steps:

1. Open the command line console, type *regedit*, and press Enter.
The Windows Registry Editor opens.
2. Locate and click the following registry key:
HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine
3. From the Edit menu, click **New**, and then click String Value.
4. Specify *ESXVersion* as the name for the new entry and then press Enter.
5. Right-click *ESXVersion* and then click **Modify**.
6. Specify *5.1* in the Value data field and then click **OK**.
7. Exit the Registry Editor.

Perform BMR in Express Mode

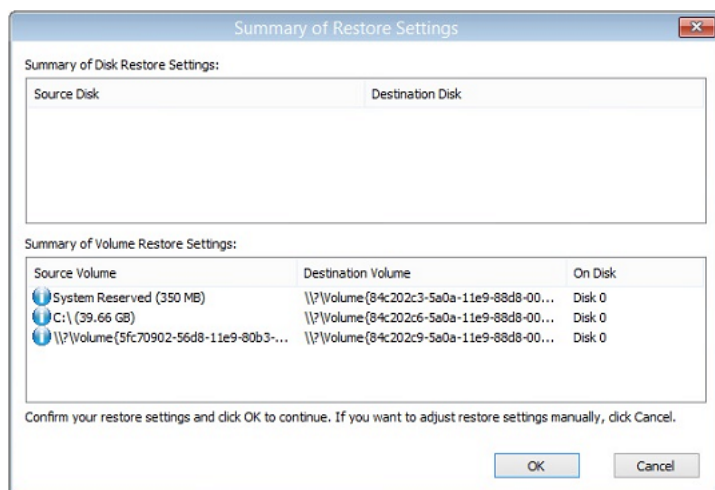
The **Express Mode** requires minimal interaction during the recovery process.

Follow these steps:

1. From the **Choose a Recovery Mode** dialog, select **Express Mode** and click **Next**.

The **Summary of Disk Restore Settings** screen opens, displaying a summary of the volumes that are going to be restored.

Note: On the bottom of restore summary window, the drive letters listed in **Destination Volume** column are automatically generated from the Windows Pre-installation Environment (WinPE). They can be different from the drive letters listed in **Source Volume** column. However, the data is still restored to proper volume even if drive letters are different.




2. After you have verified that the summary information is correct, click **OK**.

The restore process starts. The BMR wizard screen displays the restore status for each volume.

- Depending upon the size of the volume being restored, this operation can take some time.
- During this process you are restoring, block-by-block whatever you had backed up for that recovery point and creating a replica of the source machine on the target machine.
- By default, the option to reboot your system automatically after recovery is selected. If necessary, you can clear this option and you can reboot manually at a later time.

Important: If you are performing an authoritative restore of an active directory after a BMR, you must uncheck the option **Automatically reboot your system after recovery** and for more information, see [How to Perform an Authoritative Restore of an Active Directory after a BMR](#).

- If necessary, you can select Do not start Agent service automatically after reboot.
- If necessary, you can cancel or abort the operation at any time.



3. From the **Utilities** menu, you can access the BMR **Activity Log** and you can use the **Save** option to save the Activity Log.

By default, the Activity Log is saved to the following location:

X:\windows\system32\dr\log.

Note: To avoid getting a Windows-generated error, do not save the Activity Log on your desktop or create a folder on your desktop using the **Save As** option from the BMR Activity Log window.

4. If you are restoring to dissimilar hardware (the SCSI/FC adapter which used to connect hard drives could have been changed) and no compatible driver is detected in your original system, a "driver injection" page is displayed to allow you to provide drivers for these devices.

You can browse and select drivers to inject to the recovered system so that even if you are recovering to a machine with dissimilar hardware, you can still bring back the machine after BMR.

5. When the BMR process is completed, a confirmation notification is displayed.

Perform BMR in Advanced Mode

The **Advanced Mode** lets you customize the recovery process.

Follow these steps:

1. From the **Choose a Recovery Mode** dialog, select **Advanced Mode** and click **Next**.

The BMR utility starts locating the machine that is going to be recovered and displays the corresponding disk partition information.

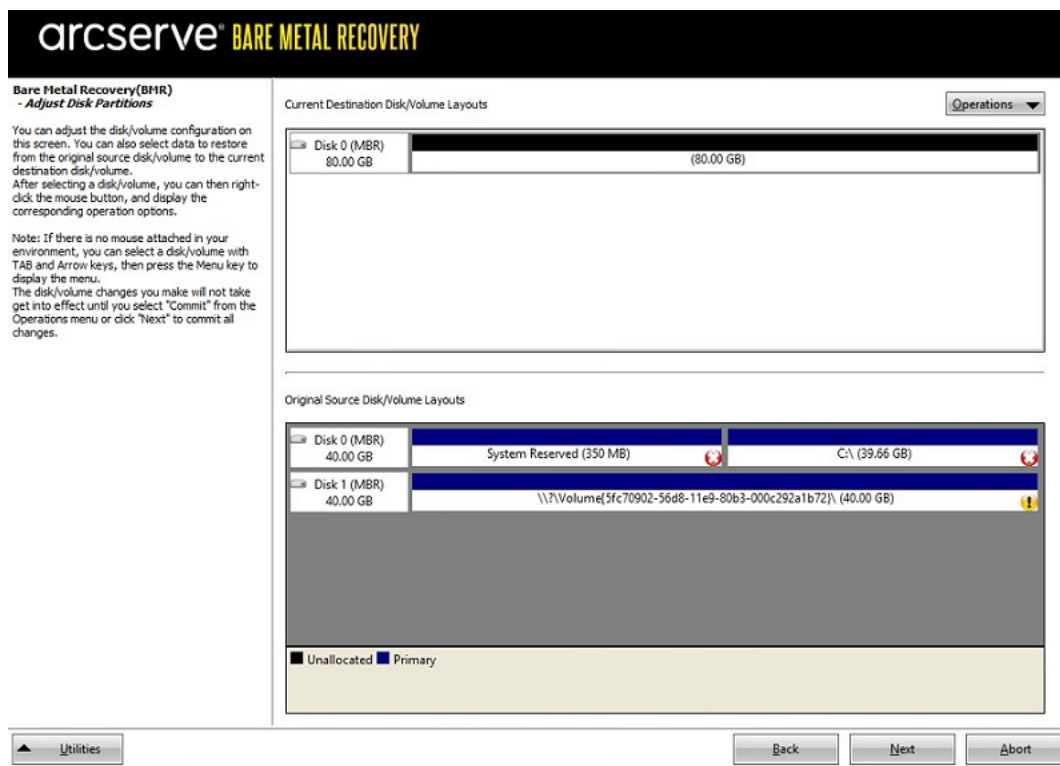
The upper pane shows the disk configuration that you have on the current (target) machine and the lower pane shows the disk partition information that you had on the original (source) machine.

Important! A red X icon displaying for a source volume in the lower pane indicates that this volume contains system information and has not been assigned (mapped) to the target volume. This system information volume from the source disk must be assigned to the target disk and restored during BMR or the reboot fails.

Note: If you perform BMR and you restore the system volume to a disk which is not configured as the boot disk, it will fail to boot the machine after BMR is completed. Ensure that you are restoring the system volume to a properly configured boot disk.

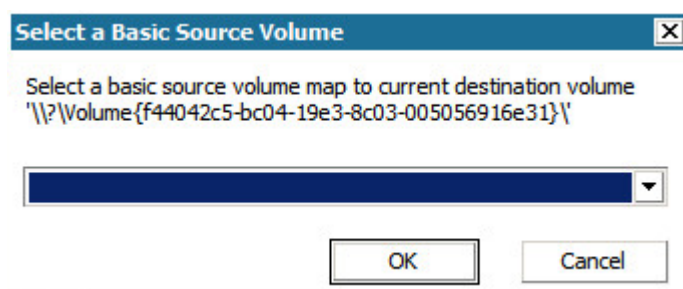
Note: When restoring to another disk/volume, the capacity of new disk/volume must be the same size or larger than original disk/volume. In addition, disk resizing

is for basic disks only, and not for dynamic disks.



2. If the current disk information you are seeing does not appear correct, you can access the **Utilities** menu and check for missing drivers.
3. If necessary, on the target disk/volume pane you can click the **Operations** drop-down menu to display the available options. For more information about these options, see [Managing the BMR Operations Menu](#).
4. Click on each target volume and from the pop-up menu, select the **Map Volume From** option to assign a source volume to this target volume.


The **Select a Basic Source Volume** dialog opens.



5. From **Select a Basic Source Volume** dialog, click the drop-down menu and select the available source volume to assign to the selected target volume. Click **OK**.
 - On the target volume, a checkmark icon is displayed, indicating that this target volume has been mapped to.

- On the source volume, the red X icon changes to a green icon, indicating that this source volume has been assigned to a target volume.
6. When you are sure all volumes that you want to restore and all volumes containing system information are assigned to a target volume, click **Next**.

The Submit Disk Changes screen opens, displaying a summary of the selected operations. For each new volume being created, the corresponding information is displayed.



7. When you have verified the summary information is correct, click **Submit**. (If the information is not correct, click **Cancel**).

Note: All operations to the hard drive do not take effect until you submit it.


On the target machine, the new volumes are created and mapped to the corresponding source machine.

8. When the changes are completed, click **OK**.

The Summary of Disk Restore Settings screen opens, displaying a summary of the volumes that are going to be restored.

Note: On the bottom of restore summary window, the drive letters listed in "Destination Volume" column are automatically generated from the Windows Pre-installation Environment (WinPE). They can be different from the drive letters listed

in "Source Volume" column. However, the data is still restored to proper volume even if drive letters are different.



9. After you have verified that the summary information is correct, click **OK**.


The restore process starts. The BMR wizard screen displays the restore status for each volume.

- ◆ Depending upon the size of the volume being restored, this operation can take some time.
- ◆ During this process you are restoring, block-by-block whatever you had backed up for that recovery point and creating a replica of the source machine on the target machine.
- ◆ By default, the option to reboot your system automatically after recovery is selected. If necessary, you can clear this option and you can reboot manually at a later time.

Important: If you are performing an authoritative restore of an active directory after a BMR, you must 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.



- 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.

- If you are restoring to dissimilar hardware (the SCSI/FC adapter which used to connect hard drives could have been changed) and no compatible driver is detected in your original system, a "driver injection" page is displayed to allow you to provide drivers for these devices.

You can browse and select drivers to inject to the recovered system so that even if you are recovering to a machine with dissimilar hardware, you can still bring back the machine after BMR.

- When the BMR process is completed, a confirmation notification is displayed.

Verify that the BMR was Successful

To verify that the BMR was successful, perform the following tasks:

- Reboot the operating system.
- Verify all systems and applications function correctly.
- Verify all network settings are properly configured.
- Verify the BIOS is configured to boot from the disk on which the boot volume was restored to.
- When the BMR is completed, be aware of the following conditions:
 - The first backup that is performed after the BMR is a Verify Backup.
 - When the machine has been rebooted, you may need to configure the network adapters manually if you restored to dissimilar hardware.

Note: When the machine is rebooting, a Windows Error Recovery screen may be displayed indicating that Windows did not shut down successfully. If this occurs, you can safely ignore this warning and continue to start Windows normally.

- For dynamic disks, if the status of the disk is offline, you can manually change it to online from the disk management UI (accessed by running the Diskmgmt.msc control utility).
- For dynamic disks, if the dynamic volumes are in a failed redundancy status, you can manually resynchronize the volumes from the disk management UI (accessed by running the Diskmgmt.msc control utility).

BMR Reference Information

[How Bare Metal Recovery Works](#)


[Operating Systems that Support UEFI or BIOS Conversion](#)

[Managing the BMR Operations Menu](#)

How Bare Metal Recovery Works

Bare Metal Recovery is the process of restoring a computer system from "bare metal" by reinstalling the operating system and software applications, and then restoring the data and settings. The most common reasons for performing a bare metal recovery are because your hard drive either fails or becomes full and you want to upgrade (migrate) to a larger drive or migrate to newer hardware. Bare metal recovery is possible because during the block-level backup process, Arcserve UDP Agent (Windows) captures not only the data, but also all information related to the operating system, installed applications, configuration settings, necessary drivers, and so on. All relevant information that is necessary to perform a complete rebuild of the computer system from "bare metal" is backed up into a series of blocks and stored on the backup location.

Note: Dynamic disks are restored at disk level only. If your data is backed up to a volume on a dynamic disk, you will not be able to restore this dynamic disk (including all its volumes) during BMR.



When you perform a bare metal recovery, the Arcserve UDP Agent (Windows) boot disk is used to initialize the new computer system and allow the bare metal recovery process to begin. When the bare metal recovery is started, Arcserve UDP Agent (Windows) will prompt you to select or provide a valid location to retrieve these backed up blocks from, as well as the recovery point to be restored. You may also be prompted to provide valid drivers for the new computer system if needed. When this connection and configuration information is provided, Arcserve UDP Agent (Windows) begins to pull the specified backup image from the backup location and restore all backed up blocks to the new computer system (empty blocks will not be restored). After the bare metal recovery image is fully restored to the new computer system, the machine will be back to the state that it was in when the last backup was performed, and Arcserve UDP Agent (Windows) backups will be able to continue as scheduled. (After completion of the BMR, the first backup will be a Verify Backup).

Operating Systems that Support UEFI/BIOS Conversion

If it is detected that the operating system of your source machine is not the same firmware as your system, you will be asked if you want to convert UEFI to a BIOS-compatible system or BIOS to UEFI-compatible system. The following table lists each operating system and the type of conversion supported.

Operating System (OS)	CPU	uEFI to BIOS	BIOS to uEFI
Windows Server 2008	x86	No	No
Windows Server 2008	x64	Yes	Yes
Windows Server 2008 R2	x64	Yes	Yes
Windows 7	x86	No	No
Windows 7	x64	Yes	Yes
Windows 8	x86	No	No
Windows 8	x64	Yes	Yes
Windows Server 2012	x64	Yes	Yes
Windows 8.1	x86	No	No
Windows 8.1	x64	Yes	Yes
Windows 10	x86	No	No
Windows 10	x64	Yes	Yes
Windows Server 2012 R2	x64	Yes	Yes
Windows Server 2016	x64	Yes	Yes
Windows Server 2019	x64	Yes	Yes

Managing the BMR Operations Menu

The BMR Operations menu consists of the following three types of operations:

- Disk Specific Operations
- Volume/Partition Specific Operations
- BMR Specific 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.
- Select the disk header on the BMR UI, click **Operations**, and use the **Clean Disk** operation 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.

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.

Using the PowerShell Interface

This section contains the following topics:

How to use the PowerShell Interface

Arcserve UDP provides PowerShell capabilities that allows you to submit a backup job, perform a restore, and recover VM from the command line. The PowerShell interface is named as UDPPowerCLI.ps1.

- ◆ [Review the Prerequisite](#)
- ◆ [Using the PowerShell Interface for Arcserve UDP](#)
- ◆ [PowerShell Syntax and Parameters](#)
- ◆ [PowerShell Examples](#)

Review the Prerequisite

Review the following prerequisites before using the PowerShell interface:

- You must have Windows 2008 R2 Server or higher versions.
- You must have PowerShell 3 or higher version installed on your server.

Using the PowerShell Interface for Arcserve UDP

The PowerShell utility is bundled with the Arcserve UDP installation file. When you install Arcserve UDP, typically the file gets installed at the following location:

C:\Program Files\Arcserve\Unified Data Protection

In such cases, on Console, UDPPowerCLI.ps1 gets installed at the following location:

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

On RPS or Agent, UDPPowerCLI.ps1 gets installed at the following location:

C:\Program Files\Arcserve\Unified Data Protection\Engine\PowerCLI

Refer the following items to help you use the PowerShell interface:

- Update the PowerShell execution policy to allow the scripts to run. For example, update the execution policy to **Set-ExecutionPolicy RemoteSigned**.

Note: For more information about changing the execution policy, see the Microsoft [website](#).

- Run the following PowerShell command to get the detailed help messages and examples for the scripts:

```
On Console:
```

```
Get-Help 'C:\Program Files\Arcserve\Unified Data Protection\Management\PowerCLI\UDPPowerCLI.ps1' -full
```

On RPS or Agent:

```
Get-Help 'C:\Program Files\Arcserve\Unified Data Protection\Engine\PowerCLI\UDPPowerCLI.ps1' -full
```

PowerShell Syntax and Parameters

SYNTAX 1

```
UDPPowerCLI.ps1 -Command <CreatePswFile> -Password <String> -PasswordFile <string> [<CommonParameters>]
```

SYNTAX 2

```
UDPPowerCLI.ps1 -Command <Backup> [-UDPConsoleServerName <String>] [-UDPConsoleProtocol <{http|https}>] [-UDPConsolePort <int>] [-UDPConsoleUserName [<String>]] [-UDPConsolePassword <String>] [-UDPConsolePasswordFile <String>] [-UDPConsoleDomainName <String>] -planName <String> -nodeName <String> [-backupJobType <String>] [-jobDescription <String>] [-waitJobFinish <String String>] [-timeOut <int>] [-agentBasedJob <{true|false} String>] [-backupScheduleType <String>] [<CommonParameters>]
```

SYNTAX 3

```
UDPPowerCLI.ps1 -Command <Restore> [-UDPConsoleServerName <String>] [-UDPConsoleProtocol <String>] [-UDPConsolePort <int>] [-UDPConsoleUserName <String>] [-UDPConsolePassword <String>] [-UDPConsolePasswordFile <String>] [-UDPConsoleDomainName <String>] [-UDPAgentServerName <String>] [-UDPAgentProtocol <String>] [-UDPAgentPort <int>] [-UDPAgentUserName <String>] [-UDPAgentPassword <String>] [-UDPAgentPasswordFile <String>] [-UDPAgentDomainName <String>] [-RestoreDirectoryPath <String>] [-RestoreFilePath <String>] [-BackupSessionNumber <int>] [-VmName <String>] -RestoreDestination <String> [-RestoreDestinationUserName <String>] [-RestoreDestinationPassword <String>] [-CreateRootFolder <String>] [-ChangeFileName <String>] [-ReplaceActiveFilesFlag <String>] [-OverwriteExistFiles <String>] [<CommonParameters>]
```

SYNTAX 4

```
UDPPowerCLI.ps1 -command <RecoverVM> [-UDPConsoleServerName <String>] [-UDPConsoleProtocol <String>] [-UDPConsolePort <int>] [-UDPConsoleUserName <String>] [-UDPConsolePassword <String>] [-UDPConsolePasswordFile <String>] [-UDPConsoleDomainName <String>] [-UDPAgentServerName <String>] [-UDPAgentProtocol <String>] [-UDPAgentPort <int>] [-UDPAgentUserName <String>] [-UDPAgentPassword <String>] [-UDPAgentDomainName <String>] [-UDPAgentPasswordFile <String>] [-BackupSessionNumber <int>] -RecoverVmName <String> [-OverwriteExistingVM <String>] [-PoweronVM <String>] [<CommonParameters>]
```

DESCRIPTION

A utility to connect to the Arcserve UDP Console service, and submit backup and restore jobs.

PARAMETERS

-Command <String>

Specifies the command that is used. Currently, the following strings are supported:

- CreatePswFile
- Backup
- Restore
- RecoverVM

Required? **true**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPConsoleServerName <String>

Specifies the DNS name of the UDP server (the server where you have installed the Console) to which you want to connect. If this value is not specified, then the cmdlet uses the default value, the DNS name of the local machine.

Required? **false**

Position? **named**

Default value **\$env:COMPUTERNAME**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPConsolePort <int>

Specifies the port number you want to use for the connection. If this value is not specified, then the cmdlet uses the default value, 8015.

Required? **false**

Position? **named**

Default value **8015**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPConsoleProtocol <String>

Specifies the protocol on the server that you want to use for the connection. The protocol can be either http or https. If this value is not specified, then the cmdlet uses the default value, http.

Required? **false**

Position? **named**

Default value **http**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPConsoleUserName <String>

Specifies the user name you want to use for connecting to the UDP server. If the user name is not specified, then the cmdlet uses the user name currently used to log into the system.

Required? **false**

Position? **named**

Default value **\$env:UserName**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPConsolePassword <String>

Specifies the password you want to use for connecting to the UDP server.

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-passwordFile <String>

Specifies to generate the password file.

Required? **true**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPConsolePasswordFile <String>

Specifies the UDP password file you want to use for connecting to the UDP server.

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPAgentServerName <String>

Specifies the DNS name of the UDP agent server to which you want to connect for restore.

Required? **false**

Position? **named**

Default value **\$env:COMPUTERNAME**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPAgentProtocol <String>

Specifies the internet protocol that you want to use to connect to the UDP agent server. It can be either http or https. If this value is not specified, then the cmdlet uses the default value, http.

Required? **false**

Position? **named**

Default value **http**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPAgentPort <int>

Specifies the port number that you want to use to connect to the UDP agent server. If this value is not specified, then the cmdlet uses the default value, 8014.

Required? **false**

Position? **named**

Default value **8014**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPAgentUserName <String>

Specifies the user name that you want to use to connect to the UDP agent server. If the user name is not specified, then the cmdlet uses the user name currently used to log into the system.

Required? **false**

Position? **named**

Default value **\$env:UserName**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPAgentPassword <String>

Specifies the password that you want to use to connect to the UDP agent server.

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPAgentPasswordFile <String>

Specifies the UDP agent password file that you want to use to connect to the UDP agent server.

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPAgentDomainName <String>

Specifies the domain name where the specified UDP agent user is located.

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-NodeName <String>

Specifies the name of node that you want to back up.

Required? **true**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-RestoreFilePath <String>

Specifies the file that you want to restore.

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-RestoreDirectoryPath <String>

Specifies the directory that you want to restore.

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-BackupSessionNumber <int>

Specifies the session number to use for the restore job.

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-VmName <String>

Specifies the host name of a virtual machine for restoring file or directory from its backup session.

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-RestoreDestination <String>

Specifies the directory path where the files will be restored.

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-RestoreDestinationUserName <String>

Specifies the user name of the destination machine where you want to restore data. The user name belongs to the user who can log in to the destination machine.

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-RestoreDestinationPassword <String>

Specifies the password that you will use to log into the destination machine.

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-CreateRootFolder <String>

Specifies that if a root directory structure exists in the captured backup image, Arcserve UDP recreates that same root directory structure on the restore destination path. When this option is not selected, the file or folder is restored directly to the destination folder. You can use any one of the following strings:

- True
- False

Required? **false**

Position? **named**

Default value **False**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-ChangeFileName <String>

Creates a new file if the file name already exists. Selecting this option copies the source file to the destination with the same filename but a different extension. Data is then restored to the new file. You can use any one of the following strings:

- True
- False

Required? **false**

Position? **named**

Default value **False**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-ReplaceActiveFilesFlag <String>

Replaces any active files after reboot. If during the restore attempt Arcserve 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 when the **OverwriteExistingFiles** parameter is True. You can use any one of the following strings:

- True
- False

Required? **false**

Position? **named**

Default value **False**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-OverwriteExistingFiles <String>

Overwrites (replaces) any existing files, which are at the restore destination. All objects are restored from the backup files regardless of their current presence on your computer. You can use any one of the following strings:

- True
- False

Required? **false**

Position? **named**

Default value **False**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-UDPConsoleDomainName <String>

Specifies name of domain where the specified user is located. If this value is not specified, then the cmdlet uses the domain name of local machine; or the DNS name of local machine if it is not in a domain.

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-PlanName <String>

Specifies the plan name that defines the backup job setting.

Required? **true**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-BackupJobType <String>

Specifies the type of the backup job. One of the following values can be used: Full (indicates a Full backup), Incr (indicates an Incremental backup), or Rsyn

(indicates a Resync backup). If you do not provide any value, then the cmdlet uses the default value, **Incr**. The following strings are supported:

- Full
- Incr
- Rsyn

Required? **false**

Position? **named**

Default value **Incr**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-JobDescription <String>

Specifies the description for the backup job.

Required? **true**

Position? **named**

Default value **PowerCLIJo**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-RecoverVmName <String>

Specifies the host name of the virtual machine that you want to recover.

Required? **true**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

-OverwriteExistingVM <String>

Specifies that if the value is true, the restore job overwrites the existing virtual machine. The default value is false. You can use any one of the following strings:

- True
- False

Required? **false**

Position? **named**

Default value **False**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-PoweronVM <String>

Specifies that if the value is true, the virtual machine is powered on after it is recovered. The default value is false. You can use any one of the following strings:

- True
- False

Required? **false**

Position? **named**

Default value **False**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-waitJobFinish <{true|false} String>

Specifies that if the value is true, the command waits for further instructions until the backup job is complete. The default value is false. You can use any one of the following strings:

- True
- False

Required? **false**

Position? **named**

Default value **False**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-timeOut <int>

Specifies the maximum waiting time (in seconds) for the backup job to complete.

Required? **false**

Position? **named**

Default value **600**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-agentBasedJob <String String>

Specifies that if true then for two nodes with the same node name, the cmdlet lets the node that has the agent-based task to submit the backup job. The default value is False. You can use any one of the following strings:

- True
- False

Required? **false**

Position? **named**

Default value **False**

Accept pipeline input? **false**

Accept wildcard characters? **false**

-backupScheduleType <String>

Specifies schedule backup job, submits the specified schedule backup job immediately, and runs only once. The following strings are supported:

- Daily
- Weekly
- Monthly

Required? **false**

Position? **named**

Default value

Accept pipeline input? **false**

Accept wildcard characters? **false**

<CommonParameters>

This cmdlet supports the common parameters such as **Verbose**, **Debug**, **ErrorAction**, **ErrorVariable**, **WarningAction**, **WarningVariable**, **OutBuffer**, and **OutVariable**. For more information, see [about_CommonParameters](#).

INPUTS

OUTPUTS

- 0 or 1

If job is successfully submitted successfully, the command returns 0; otherwise returns 1.

PowerShell Examples

Example 1

```
C:\PS>UDPPowerCLI.ps1 -Command CreatePswFile -password myPlainPassword -passwordFile myPasswordFile
```

Description

The command encrypts plain passwords that are present in the password file.

Example 2

```
C:\PS>UDPPowerCLI.ps1 -Command Backup -UDPConsoleUserName myUsr -UDPConsolePassword myPsw -PlanName myPlan
```

Description

On the local server, the command connects to the UDP Console service with HTTP protocol over port 8015, and then submits an Incremental backup job for the plan named *myplan*.

Example 3

```
C:\PS>UDPPowerCLI.ps1 -Command Backup -UDPConsoleUserName myUsr -UDPConsolePasswordFile myUDPPasswordFile -NodeName myNodeName
```

Description

On the local server, the command connects to the UDP Console service with HTTP protocol over port 8015, and then submits an Incremental backup job for the node named *myNodeName*.

Example 4

```
C:\PS>UDPPowerCLI.ps1 -Command Backup -UDPConsoleServerName myServer -UDPConsoleProtocol https -UDPConsolePort 8018 -UDPConsoleUserName myUsr -UDPConsolePassword myPsw -UDPConsoleDomainName myDomain -PlanName myPlan -BackupJobType Full -JobDescription myJob
```

Description

The command connects to the UDP Console service on the server named *myServer* with HTTPS protocol over port 8018, and then submits a Full backup job for the plan named *myPlan*, and set the job description as *myJob*.

Example 5

```
C:\PS>UDPPowerCLI.ps1 -Command Backup -UDPAgentServerName yourUDPAgentServer -UDPAgentPasswordFile myUDPAgentPasswordFile -jobType Incr
```

Description

The command connects to the UDP agent service on the server named *yourUDPAgentServer* with HTTP protocol over port 8014, and then submits an Incremental backup job for *yourUDPAgentServer*.

Example 6

```
C:\PS>UDPPowerCLI.ps1 -Cmd Backup -Svr myServer -Ptc https -Prt 8018 -Usr myUsr -Psw myPsw -Dmn myDomain -Pln myPlan -Jbt Full -Jbd myJob
```

Description

The command shortens the Parameter name.

Example 7

```
C:\PS>UDPPowerCLI.ps1 -Command restore -UDPAgentServerName yourUDPAgentServer -UDPAgentPasswordFile myUDPAgentPasswordFile -RestoreDirectoryPath 'c:\Test' -BackupSessionNumber 1
```

Description

The command connects to the server named *yourUDPAgentServer* using the user name of the environment, the default HTTP protocol, and port 8014. It verifies the backup session number is 1 from the *yourUDPAgentServer* backup configuration and then restores the directory to the original location, with the restore option selected as Overwrite existing files.

Example 8

```
C:\PS>UDPPowerCLI.ps1 -Command restore -UDPAgentServerName yourUDPAgentServer -UDPAgentUserName UDPAgentUsername -UDPAgentPasswordFile myUDPAgentPasswordFile -UDPAgentProtocol 'https' -UDPAgentPort 8018 -UDPAgentDomainName UDPAgentdomainName -BackupSessionNumber 1 -RestoreFilePath 'C:\1.txt' -RestoreDestination 'C:\restore' -RestoreDestinationUserName remoteAccessUser -RestoreDestinationPassword remoteAccessPsw -CreateBaseFolder 'true'
```

Description

The command connects to the server named *yourUDPAgentServer* using the HTTPS protocol and port 8018. It verifies the backup session number is 1 from the *yourUDPAgentServer* backup configuration and then restores the 1.txt file to an alternate location, with the restore option selected as Overwrite existing file and create root directory.

Example 9

```
C:\PS>UDPPowerCLI.ps1 -Command restore -UDPAgentServerName yourUDPAgentServer -UDPAgentPasswordFile myUDPAgentPasswordFile -RestoreDirectoryPath 'c:\Test' -BackupSessionNumber 1 -RestoreDestination 'C:\restore' -RestoreDestinationUserName remoteAccessUser -RestoreDestinationPassword
```

```
remoteAccessPsw -servername yourUDPServer -vmname sourceVMName -UDPConsolePasswordFile myUDPPasswordFile -domainname yourUDPDomainName -OverwriteExistFiles 'true' -CreateRootFolder 'true'
```

Description

The command connects to the server named *yourUDPAgentServer* using the user name of the environment, the default HTTP protocol, and port 8014. Then, it connects to the UDP server using the default port 8015 and protocol HTTP to check the backup session number is 1. Lastly, it restores the directory to an alternate location, with the restore option selected as Overwrite existing file and create root directory.

Example 10

```
C:\PS>UDPPowerCLI.ps1 -Command RecoverVM -UDPAgentServerName yourUDPAgentServer -UDPAgentPasswordFile myUDPAgentPasswordFile -BackupSessionNumber 1 -UDPConsoleServerName yourUDPServer -recovervmname sourceVMName -UDPConsolePasswordFile myUDPPasswordFile -UDPConsoleDomainName yourUDPDomainName -OverwriteExistingVM 'true' -PoweronVM 'true'
```

Description

The command connects to the server named *yourUDPAgentServer* using the user name of the environment, the default HTTP protocol, and port 8014. Then, it connects to the UDP server using the default port 8015 and protocol HTTP to check the backup session number is 1. Lastly, it recovers the VM to the original location, with the recover VM option selected as Overwrite existing vm and power on vm after recovered.

Example 11

```
C:\PS>UDPPowerCLI.ps1 -Command Backup -UDPAgentServerName myServer -UDPAgentPassword myPassword -UDPAgentDomainName myDomainName -UDPAgentUserName myPassword -backupJobType 'incremental' -backupScheduleType 'weekly' -jobDescription 'PowerCLIJob'
```

Description

The command submits weekly backup job on the UDP Agent immediately and runs only one time.

Example 12

```
C:\PS>UDPPowerCLI.ps1 -Command Backup -UDPConsoleServerName myServer -UDPConsolePasswordFile myPasswordFile -UDPConsoleDomainName myDomainName -nodeName myNodeName -UDPConsoleUserName myAdmin -backupJobType 'incremental' -jobDescription 'PowerCLIJob' -waitJobFinish 'true' -timeout 600 -agentBasedJob 'true'
```

Description

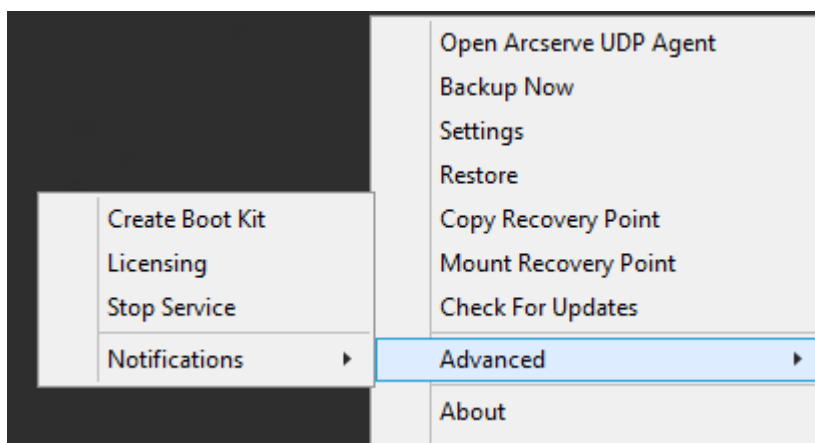
The command submits the backup job and sets the timeout in seconds to wait for the job to complete.

Add Arcserve UDP Agent (Windows) Licensing

Arcserve UDP Agent (Windows) requires you to license your product to receive authorized and uninterrupted access to the related components.

Arcserve UDP Agent (Windows) will function for a period of 30 days after you begin using it. Then, apply an appropriate license key to continue using it.

To add an Arcserve UDP Agent (Windows) license, access the Arcserve UDP Agent (Windows) Monitor Advanced options.



Note: For Windows Core Operating Systems (Windows Server 2008/R2, 2012/R2 Core edition), run the "ArcserveLicense.exe" file and provide the proper license key information. The ArcserveLicense.exe file is located in the following directory:
C:\Program Files\CA\SharedComponents\CA_LIC

Follow these steps:

Note: Perform this operation locally on the computers running Arcserve UDP Agent (Windows) software.

1. Access the Arcserve UDP Agent (Windows) Monitor, click the **Advanced** option, and select **Licensing**.

The License Verification Entry dialog opens, displaying all valid Arcserve licensed products.

Note: If no Arcserve products were previously licensed, the field of this dialog is empty.

2. Enter the 25-digit license key and then click **Add**.

The component is licensed.

3. Identify the next component to license and repeat Step 2.
4. Click **OK** to accept the key after all components are defined as a licensed product.
All components you specified are licensed.

The license key information is stored in the Arcserve.olf file on each of the computers that are running your Arcserve software.

Change Server Communication Protocol

By default, Arcserve UDP Agent (Windows) uses the Hypertext Transfer Protocol Secure (HTTPS) for communication among all of its components. In addition, if you do not need this extra level of security, you can easily change the protocol being used back to HTTP.

Note: After you change the protocol from HTTP to HTTPS or from HTTPS to HTTP, restart the browser and reconnect to Arcserve UDP Agent (Windows).

Follow these steps:

1. To change the protocol from HTTP to HTTPS, launch the **changeToHttps.bat** utility tool from the following default location:

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

Note: The location of the BIN folder can vary depending upon your Arcserve UDP Agent (Windows) installation path.

When the protocol has been successfully changed, the following message displays:

"Protocol has been changed to HTTPS. Use https://localhost:8014 to access Arcserve UDP Agent (Windows) system."

Note: When the protocol is changed to HTTPS, a warning displays in the web browser due to a self-signed security certificate. The message asks you to either:

- ◆ Ignore the warning and proceed. or
- ◆ Add that certificate to the browser to prevent that warning from coming back in the future.

2. To change the protocol from HTTPS to HTTP, launch the **changeToHttp.bat** utility tool from the following default location:

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

Note: The location of the BIN folder can vary depending upon your Arcserve UDP Agent (Windows) installation path.

When the protocol has been successfully changed, the following message displays:

"Protocol has been changed to HTTP. Use http://localhost:8014 to access Arcserve UDP Agent (Windows) system."

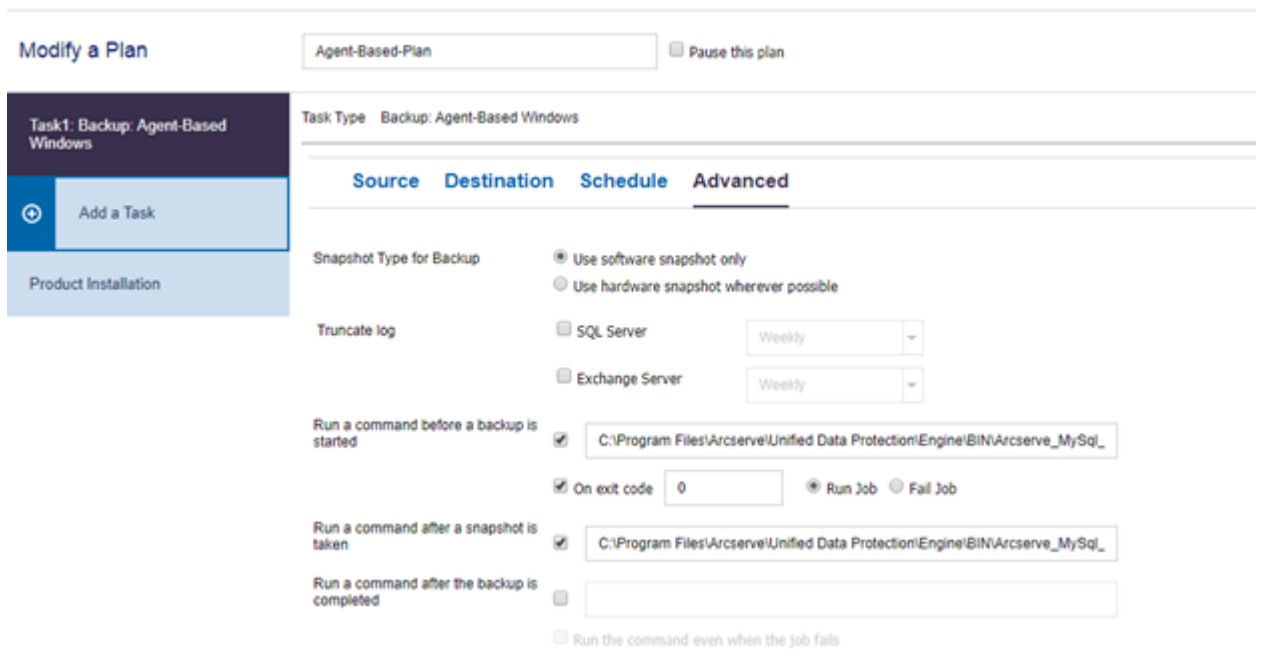
Use Scripts to Backup and Restore MySQL Database

The following [scripts](#) are available to perform backup of the MySQL database. When running the scripts, you do not have to stop your database to perform a backup.

- **Arcserve_MySql_PreBackup_script.bat:** This script closes all open tables, and it locks all the tables for all the databases with a global read lock.
- **Arcserve_MySql_PostSnapshot_script.bat:** This script releases all the locks.
- **Arcserve-MySQL-pre-post-snapshot-conf.bat:** This script helps to capture database details -hostname, DB username, DB password, Port.

To use the scripts, follow these steps:

1. Extract [UDP-MySQL-Windows-scripts.zip](#) provides seven files.
2. Place all the files in BIN folder of the agent installation folder. The default agent installation location is C:\Program Files\Arcserve\Unified Data Protection\Engine\.
3. Provide MySQL database details (hostname, DB username, DB password, Port) under Arcserve-MySQL-pre-post-snapshot-conf.bat
4. Configure agent-based plan from UDP Console and select the MySQL node as source.



5. Check the Activity log for pre and post script execution status. The ArcserveMySQLPrepost.log and ArcserveMySQLReadLock.log files are available in C:\Program Files\Arcserve\Unified Data Protection\Engine\Logs.

Restore MySQL Database

1. Stop MySQL Service.
2. To restore to the original location, do the following:
 - c. Delete files and directories from the current MySQL\data\ folder.
 - d. Restore database folder from recovery point to MySQL\data\ folder.
5. Start MySQL Service.

Note: The steps provided here help to restore the entire MySQL server data, not just the single database.

Modify Arcserve-MySQL-pre-post-snapshot-conf.bat

Edit the following two configurable parameters, which are available in the Arcserve-MySQL-pre-post-snapshot-conf.bat file.

- set AMSQLREADLOCKTIMEOUT=25 // time-out parameter in minute for acquiring the read lock on MySQL database server
- set AMSQLREADLOCKRETRY=3 // Retry count to acquire Read Lock <in pre-script> and Remove the Read lock <in pos-script>

Note: If read-lock is not successful in “AMSQLREADLOCKTIMEOUT” minutes, the backup fails.

Use Scripts to Backup and Restore PostgreSQL Database

The following [scripts](#) are available to perform backup of the PostgreSQL database. When running the scripts, you do not have to stop your database to perform a backup.

- postgresql_pre_script.bat: This script puts the database into backup mode.
- postgresql_post_snapshot_script.bat: This script removes the database from backup mode.
- postgresql_pre_post_conf.bat: This is a configuration file where PostgreSQL variables might need to be updated.

Prerequisites

Before you begin the backup, make sure to do the following:

- WAL level is set to archive (or hot_standby)
- archive_mode is set to on
- archive_command has to be set to specify the archive location

Note: To apply the settings, reboot the server after configuring these settings in the postgresql.conf file.

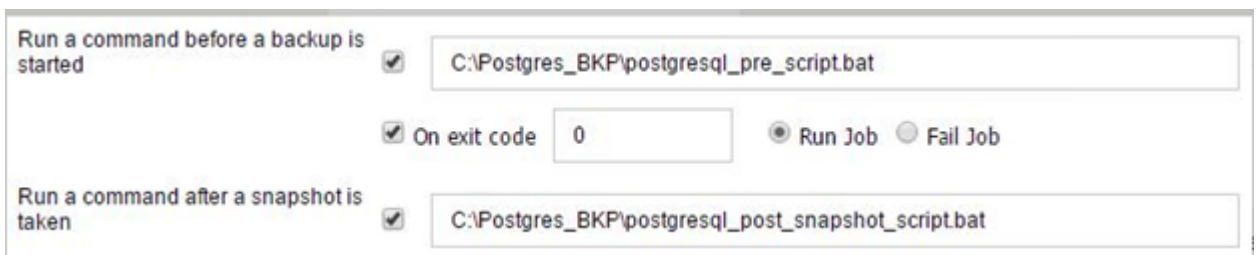
The following commands help to check the status of the archive mode after reboot:

- show archive_mode
- show archive_command
- show WAL level

Apply Scripts

Follow these steps:

1. Extract the [PostgreSQL_UDP_Windows_Scripts.zip](#), which contains the following three files: postgresql_pre_script.bat, postgresql_post_snapshot_script.bat, postgresql_pre_post_conf.bat.
2. Create a folder on a node where PostgreSQL database is running, for example: C:\PostgresBKP, and then copy all three files.
3. Make sure to check the postgresql_pre_post_conf.bat for all values set against the variables and make modifications for any required changes as per your environment.
4. Configure the plan from UDP Console and select the PostgreSQL node as source.



The screenshot shows a configuration window with two sections. The first section is titled 'Run a command before a backup is started' and has a checked checkbox. To its right is a text box containing the path 'C:\Postgres_BKP\postgresql_pre_script.bat'. Below this is another checked checkbox labeled 'On exit code' followed by a text box containing the value '0'. To the right of the 'On exit code' section are two radio buttons: 'Run Job' (which is selected) and 'Fail Job'. The second section is titled 'Run a command after a snapshot is taken' and also has a checked checkbox. To its right is a text box containing the path 'C:\Postgres_BKP\postgresql_post_snapshot_script.bat'.

5. Confirm the backup status. To know the status of PostgreSQL backup, check for the postgresql_pre_post_backup.txt file, which gets created under the directory 'C:\Program Files\Arcserve\Unified Data Protection\Engine\Logs'.

Restore PostgreSQL Database

1. Stop the database server.
2. To restore to the original location, do the following:
 - a. Delete files and directories from the current /data folder.
 - b. Perform a restore of entire /data folder.
3. Delete the files from the following folders after completion of restore from /data folder:
 - pg_dynshmem/
 - pg_notify/
 - pg_serial/
 - pg_snapshots/
 - pg_stat_tmp/
 - pg_subtrans/
 - pg_internal.init
4. Go to the folder, which is configured for WAL Archiving, and do the following:
 - a. Delete the files present in the restored pg_wal directory which contains the information related to transactions performed during the backup.
 - b. Now, copy files from the user defined archived location to the pg_wal folder, for data consistency and point-in-time recovery.
5. Start the Database server.

Restore to Alternate location on the Same Server:

Follow these steps:

1. Stop the database server.
2. Run PGDATA configuring it to the " new_data_directory_path".
3. Initialize the newly created DB using `initdb` cmd.
4. Delete files and directories from the current /data folder.
5. Perform a restore of entire /data folder.
6. Delete the files from the following folders after completion of restore from /data folder:

- pg_dynshmem/
- pg_notify/
- pg_serial/
- pg_snapshots/
- pg_stat_tmp/
- pg_subtrans/
- pg_internal.init

7. Go to the folder, which is configured for WAL Archiving, and do the following:
 - a. Delete the files present in the restored pg_wal directory, which contains the information related to transactions performed during the backup.
 - b. Now, copy files from the user defined archived location to the pg_wal folder, for data consistency and point-in-time recovery.
8. Start the Database server.

Note: Make sure that the database startup is performed in the session where the PGDATA gets updated.

Sample postgresql_pre_post_conf.bat

```
PG_BIN_PATH=C:"Program Files"\PostgreSQL\12\bin\  
PG_DATA_DIR=C:"Program Files"\PostgreSQL\12\data\  
PG_USERNAME=postgres  
PGPASSWORD=postgres  
Set PG_PORT=5432
```

Chapter 6: Troubleshooting Arcserve UDP Agent (Windows)

This section contains the following topics:

Troubleshooting Overview	696
Arcserve UDP Agent Service could not be started because of port conflict	697
Reboot Not Required After Agent Deployment	700
Unable to Connect to Cloud	701
Unable to Change Destination to Removable Device	702
Unable to display Arcserve UDP Agent (Windows) UI in Firefox	704
Settings Disabled when Opening Agent UI	705
Unable to Open the SQL database in SQL Management Studio from Mounted Volume	706
Recovery of SQL Server Databases to Original Location fails	707
Login Link Does not Work at Arcserve UDP Agent Home	708
Troubleshooting Installation Issues	709
Troubleshooting Update Issues	716
Troubleshooting Uninstall Issues	720
Troubleshooting User Interface Issues	723
Troubleshooting Backup Issues	726
Troubleshooting BMR Issues	733
Troubleshooting Merge Issues	740
Troubleshooting Exchange Issues	744

Troubleshooting Overview


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.

Arcserve UDP Agent Service could not be started because of port conflict

The port that Arcserve UDP Agent (Windows) uses can conflict with the default port that Tomcat uses. This conflict causes Tomcat to fail when Arcserve UDP Agent (Windows) is started before it. To remedy this problem, you can change the Tomcat default port as follows:

1. Access the Arcserve UDP Agent (Windows) Monitor, click the **Advanced** option, and select **Stop Service**.

The Arcserve UDP Agent Service is stopped.



2. To edit/configure the Tomcat behavior, open the Tomcat server.xml file located in the following path:

C:\Program Files\Arcserve\Unified Data Protection\Engine\TOMCAT\conf

3. Locate the <Server> tag inside the server.xml file.

```
<?xml version="1.0" encoding="UTF-8"?>
<Server port="18005" shutdown="SHUTDOWN">
  <Listener className="org.apache.catalina.core.JasperListener"/>
  <Listener className="org.apache.catalina.core.JreMemoryLeakPreventionListener"/>
  <Listener className="org.apache.catalina.core.ThreadLocalLeakPreventionListener"/>
  <Service name="Catalina">
    <Connector connectionTimeout="180000" port="8014" protocol="HTTP/1.1"/>
    <Engine defaultHost="localhost" name="Catalina">
      <Host appBase="webapps" autoDeploy="false" deployOnStartup="false" deployXML="false"
        <Context debug="0" docBase="C:\Program Files\Arcserve\Unified Data Protection
        <Context debug="0" docBase="C:\Program Files\Arcserve\Unified Data Protection
        <Context debug="0" docBase="C:\Program Files\Arcserve\Unified Data Protection
      </Host>
    </Engine>
  </Service>
</Server>
```

4. Edit the <Server> tag as follows:

From:

```
<Server port="18005" shutdown="SHUTDOWN">
```

To:

```
<Server port="18006" shutdown="SHUTDOWN">
```

```
<?xml version="1.0" encoding="UTF-8"?>
<Server port="18006" shutdown="SHUTDOWN">
  <Listener className="org.apache.catalina.core.JasperListener"/>
  <Listener className="org.apache.catalina.core.JreMemoryLeakPreventionListener"/>
  <Listener className="org.apache.catalina.core.ThreadLocalLeakPreventionListener"/>
  <Service name="Catalina">
    <Connector connectionTimeout="180000" port="8014" protocol="HTTP/1.1"/>
    <Engine defaultHost="localhost" name="Catalina">
      <Host appBase="webapps" autoDeploy="false" deployOnStartup="false" deployXML="false"
        <Context debug="0" docBase="C:\Program Files\Arcserve\Unified Data Protection
        <Context debug="0" docBase="C:\Program Files\Arcserve\Unified Data Protection
        <Context debug="0" docBase="C:\Program Files\Arcserve\Unified Data Protection
      </Host>
    </Engine>
  </Service>
</Server>
```

5. Save and close the server.xml file.

The command to shut down Tomcat has now been configured so that the server can receive it on the named port (18006).

6. Access the Arcserve UDP Agent (Windows) Monitor, click the **Advanced** option, and select **Start Service**.

The Arcserve UDP Agent Service is started.

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.

Reboot Not Required After Agent Deployment

Symptom

After deployment, the deploy Agent on Windows 2008 x86 platform displays the following message:

Deployment completed successfully but requires a reboot.

Solution

You do not need to reboot. You need to manually start agent web services for backup to work. Reboot is required only to restore data directly into the system. Without reboot, you can use alternative system to restore data.

Note: Valid only for systems that have “UMDF/KMDF” versions lower than 1.9. For example, Server 2008 and lower.

Unable to Connect to Cloud

If you are attempting to file copy to cloud, but cannot connect your machine to the cloud server, perform the following troubleshooting procedure:

1. From the File Copy Settings Destination dialog, click the Configure button to display the Cloud Configuration dialog and verify the following are correct:
 - Proxy credentials (Username and password)
 - Proxy Server IP address and corresponding port number
 - Access key and Secret key for access to the specified proxy server
 - Vendor URL address for the specified cloud provider
2. To eliminate any potential clock skew error, verify that your machine has the correct time zone set and the clock is in sync with the global time.
3. Resubmit the file copy job.


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 Change Destination to Removable Device

If you are attempting to configure your backup destination setting to a removable device, and you are unable to save the setting without a failure, it may be caused by the local security policy settings. If this occurs, perform the following troubleshooting procedure:

Note: This problem may occur in multiple cases, such as when you are attempting to browse to a destination or saving a backup destination setting.

1. Run GPEDIT.msc to open the "Local Group Policy Editor" dialog.
2. Select Computer Configuration -> Windows Settings -> Security Settings -> Local Policies -> Audit Policy.
3. Double click "Audit object access" to access the "Audit object access Properties" dialog.



4. On the Local Security Settings tab, uncheck the Success and Failure options, and click Apply to save the settings.
5. Reboot the server or run 'GPupdate/force'. (If you run 'GPupdate/force' then after the command is completed, log off and then log back on).

6. If the trouble continues, the computer is most likely a part of a domain. You will need to find out which group policy has the setting enabled by running the following in an administrative command-prompt:

```
gpresult /H C:\gpresult.html
```

7. Open the file C:\gpresult.html and navigate to the following section:

Computer Details -> Settings -> Policies -> Windows Settings -> Security Settings -> Local Policies\Audit Policy -> Audit Object Access.

Note: For Windows 7 operating systems, the location of Audit object access in the file C:\gpresult.html varies slightly by replacing "Computer Details -> Settings" with "Computer Configuration" in the navigation path.

8. The group policy is located under the "Winning GPO" column. Edit that group policy, and then reboot the server.

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 display Arcserve UDP Agent (Windows) UI in Firefox

If you are using Firefox to connect to the local Arcserve UDP Agent (Windows) server, where the browser and Arcserve UDP Agent (Windows) are both on the same machine, certain proxy settings may cause the Arcserve UDP Agent (Windows) UI not to be displayed.

If this condition occurs, connect to the loopback address 127.0.0.1 or use the host-name in Firefox instead of using localhost.

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.

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.

Unable to Open the SQL database in SQL Management Studio from Mounted Volume

Symptom

Unable to Open the SQL database in SQL Management Studio from mounted volume.

Solution

If the database could not be attached in SQL Management Studio, use the following registry to make the mounted volume writable, then dismount and mount the volume, and attach the database again.

Registry:

Forcewritable "=dword:00000001

under

[...\Engine\AFStorHBAmgmt]

Recovery of SQL Server Databases to Original Location fails

Symptom

Recovery of SQL Server databases to original location fails with error "sqlwriter is in bad status".

When recovering database to original location it is expected to have the original SQL Server instance service up and running.

Solution

Start the SQL Server instance service and try the restore job again.

Login Link Does not Work at Arcserve UDP Agent Home

- Agent login fails while using the Microsoft Edge browser.

Symptom

When logging in from the Arcserve UDP Agent home page, the link **Login with current Windows Credentials** does not work.

Solution

The feature is not supported at Microsoft Edge. Switch to other browser. On other browsers, when you face the problem, use the following workaround:

- For Internet Explorer or Chrome, verify if the URL is added into the **Local Intranet** list. To add URL, navigate to "Internet Options > Security > Local Intranet > Sites > Advanced settings.
- For Firefox, modify configuration to enable IWA.

For details, click the [link](#).

- Agent login fails with Http error code 500.

Symptom

When logging in from the Arcserve UDP Agent home page, the link **Login with current Windows Credentials** does not work.

Solution

Follow these steps:

1. Open the file from the location: *C:\Program Files\Arcserve\Unified Data Protection\Common\JRE\lib\net.properties*.
2. Replace the folder *C:\Program Files* to your installation location.
3. Modify the entry to the following one:
jdk.http.ntlm.transparentAuth=allHosts
4. Restart the Arcserve UDP Agent service to login again.

Troubleshooting Installation 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.

Unable to install/uninstall Arcserve UDP Agent (Windows) if a previous attempt was interrupted

If during an attempt to install or uninstall Arcserve UDP Agent (Windows), the install/uninstall process was interrupted, you may not be able to successfully continue and complete the process.

For example, any of the following conditions could cause a partial install/uninstall state:

- Your computer is shut down in middle of install/uninstall process.
- You encounter a power outage during install/uninstall and there is no Uninterruptible Power Supply (UPS).

To resolve this problem, perform the following steps:

1. Enter "**regedit**" in the **Run** dialog and click **OK** to open **Registry Editor**.
2. Locate and delete the following entry:
"HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine"
3. Use the search option in the **Registry Editor** to locate and delete all occurrences of the following string:
 - ♦ [Arcserve UDP Agent (Windows) for x86]: {CAAD8AEA-A455-4A9F-9B48-C3838976646A}
 - ♦ [Arcserve UDP Agent (Windows) for x64]: {CAAD1E08-FC33-462F-B5F8-DE9B765F2C1E}
4. Use the search option in the **Registry Editor** to locate and delete all occurrences of the string "Arcserve UDP Agent " under the following key:
 - HKEY_CLASSES_ROOT\Installer\Products
 - HKEY_LOCAL_MACHINE\SOFTWARE\Classes\Installer\Products
 - HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Installer\UserData\S-1-5-18\Products
 - HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Uninstall
5. From the command line, delete the service by entering the following commands:
sc delete ShProvd
sc delete CASAD2DWebSvc
6. Run the command line to remove additional setup files.
 - ♦ x86 operating system:

```
"%ProgramFiles%\Arcserve\SharedComponents\Arcserve Unified Data Protection\Setup\uninstall.exe" /q
```

- ◆ x64 operating system:

```
"%ProgramFiles(x86)%\Arcserve\SharedComponents\Arcserve Unified Data Protection\Setup\uninstall.exe" /q
```

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.

Windows failed to start after Arcserve UDP Agent (Windows) is installed

If Windows failed to start with the following error after Arcserve UDP Agent (Windows) has recently been installed, it can be caused by a Windows internal failure.

File: ARCFlashVolDrv.sys

Status: 0xc0000098

Info: Windows failed to load because a required file is missing, or corrupt.

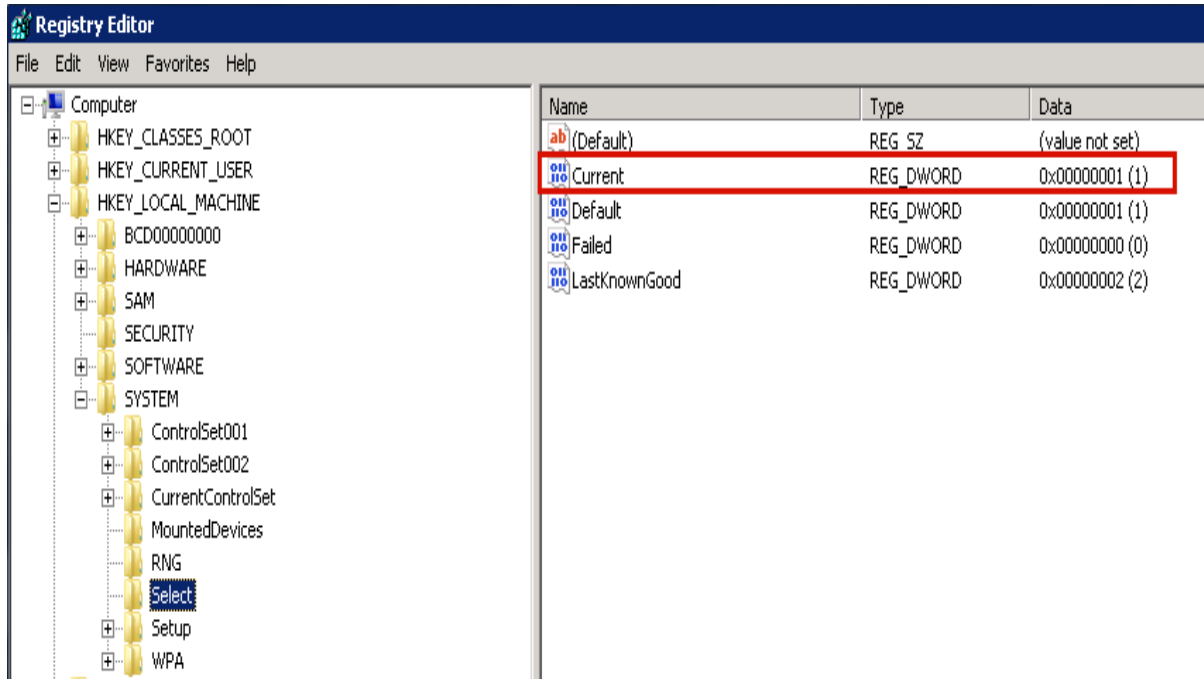
The probable causes for this problem are:

- User temporary folder is not writable
- Insufficient privilege
- Windows update database corrupted

Important! This procedure contains information about modifying the registry. Before you modify the registry, make sure to create a backup of the registry and ensure that you understand how to restore the registry if a problem occurs. For more information about how to back up, restore, and edit the registry, see the relevant Microsoft Knowledge Base [articles](#).

To resolve this problem, perform the following steps to uninstall the driver:

1. Use the Create Boot Kit for Bare Metal Recovery utility to create the BMR ISO image if you do not already have it. For more information, see [How to Create a Boot Kit](#) in the online help.
2. Click Run from the Utilities menu.
3. Enter "regedit" in the Run dialog and click OK to open Registry Editor.
4. Select HKEY_LOCAL_MACHINE and click "Load Hive..." from the File menu in the Registry Editor.
5. Locate SYSTEM file under the %systemroot%\system32\config directory on your system and click open.
6. Enter a name for the hive to be loaded.
7. From the Registry Editor, check the "Current" entry under "HKEY_LOCAL_MACHINE\SYSTEM>Select".



8. Depending on the Current value that is displayed, delete the corresponding entries under the new hive that was just loaded:

For example:

- ◆ If the Current Value is **1**, then delete the following entries:

HKEY_LOCAL_MACHINE\%your_hive_name%\ControlSet001\Services\ARCFlyVolDrv

HKEY_LOCAL_MACHINE\%your_hive_name%\ControlSet001\Services\Eventlog\System\ARCFlyVolDrv

- ◆ If the Current Value is **2**, then delete the following entries:

HKEY_LOCAL_MACHINE\%your_hive_name%\ControlSet002\Services\ARCFlyVolDrv

HKEY_LOCAL_MACHINE\%your_hive_name%\ControlSet002\Services\Eventlog\System\ARCFlyVolDrv

9. Depending on the Current value that is displayed, delete the corresponding value "ARCFlyVolDrv" for the following registry keys:

Important! The "LowerFilters" registry key may also contain other Windows driver names. Be sure to delete only the "ARCFlyVolDrv" value from the list. Do not delete the entire registry key or any other driver names from within the key.

For example:

- ◆ If the Current Value is **1**, then delete the following entries:

HKEY_LOCAL_MACHINE\%your_hive_name%\ControlSet001\Control\Class\{533C5B84-EC70-11D2-9505-00C04F79DEAF}\LowerFilters

HKEY_LOCAL_MACHINE\%your_hive_name%\ControlSet001\Control\Class\{71A27CDD-812A-11D0-BEC7-08002BE2092F}\LowerFilters

- ◆ If the Current Value is **2**, then delete the following entries:

HKEY_LOCAL_MACHINE\%your_hive_name%\ControlSet002\Control\Class\{533C5B84-EC70-11D2-9505-00C04F79DEAF}\LowerFilters

HKEY_LOCAL_MACHINE\%your_hive_name%\ControlSet002\Control\Class\{71A27CDD-812A-11D0-BEC7-08002BE2092F}\LowerFilters

10. Click "Unload Hive..." from the File menu in the Registry Editor.
11. Perform the following troubleshooting steps:
 - a. Verify that the user account has administrator privileges on this machine.
 - b. Verify that the user account has write permissions on the following temp folders:
 - ◆ %windir%/temp
 - ◆ %temp%
 - c. For Microsoft Windows Vista and Microsoft Windows 2008 and later, download and run the [Microsoft System Update Readiness Tool](#). This tool helps to fix any inconsistencies or corruption in the installed updates and system files.
 - d. Determine if there are any pending Windows updates or reboots and take the necessary action. Perform one of the following tasks to show Windows Update related information for the computer:
 - ◆ Click Start, All Programs, Windows Update.
 - ◆ Access [Update](#).
 - e. If there are problems installing several Windows Updates, then examine why the updates cannot be installed on this computer before proceeding to the next step.
12. Reinstall the ARCFlyVolDrv driver by running "ARCFlyVolDrvINSTALL.exe -i -output=c:\install.log" after the machine reboot.
 - ◆ ARCFlyVolDrvINSTALL.exe is located on Arcserve UDP Agent_Home\bin\Driver.
 - ◆ Arcserve UDP Agent_Home is located on the Arcserve UDP Agent (Windows) install path.

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.

Troubleshooting Update 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.

- [Unable to Access Arcserve UDP Agent \(Windows\) After Reboot](#)
- [Unable to Connect to the Arcserve Download Server to Download Updates](#)
- [Failed to Download Arcserve UDP Agent \(Windows\) Updates](#)

Unable to Access Arcserve UDP Agent (Windows) After Reboot

If you are not able to access the Arcserve UDP Agent (Windows) UI, perform the following troubleshooting procedure:

1. From the **Add or Remove Programs** dialog, click the **Add/Remove Windows Components** option to access the **Windows Components Wizard** screen and remove the **Internet Explorer Enhanced Security Configuration** component.
2. Add the host name URL to the **Trusted Sites** in Internet Explorer.
3. Adjust the security level in Internet Explorer.

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 Connect to the Arcserve Download Server to Download Updates

If you are not able to connect to the Arcserve download server to download Arcserve UDP Agent (Windows) updates, follow these steps:

1. From the Arcserve UDP Agent (Windows) home page, click **View Logs**, and verify the error message.
2. Verify that you have a good network connection.
3. Open command line and ping the downloads.arcserve.com server.

Perform *one* of the following to establish connection with the download server:

- ◆ From the Arcserve UDP Agent (Windows) home page, select **Settings**, then **Preferences**, and click **Updates and Download Server**. Click on the proxy settings and verify that the default option **Use browser proxy settings** (for IE and Chrome only) is selected.
 - ◆ From the Arcserve UDP Agent (Windows) home page, select **Settings**, then **Preferences**, and click **Updates and Download Server**. Click on the proxy settings and select **Configure Proxy Settings** and enter the valid proxy server name, port number and credentials and click **OK**.
4. Click **Test Connection** to verify that the connection is established.

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 Download Arcserve UDP Agent (Windows) Updates

If you are not able to download Arcserve UDP Agent (Windows) updates, follow these steps:

1. From the Arcserve UDP Agent (Windows) home page, click **View Logs** and read the error message.
2. Verify that you have a good network connection.
3. Verify that there is enough disk space.
4. From the Arcserve UDP (Windows) installation home path, access the update Log file ("<Product Home>\Update Manager\logs\ARCUpdate.log").
5. Check the log entries for detailed error messages.

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.

Troubleshooting Uninstall 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.

Unable to install/uninstall Arcserve UDP Agent (Windows) if a previous attempt was interrupted

If during an attempt to install or uninstall Arcserve UDP Agent (Windows), the install/uninstall process was interrupted, you may not be able to successfully continue and complete the process.

For example, any of the following conditions could cause a partial install/uninstall state:

- Your computer is shut down in middle of install/uninstall process.
- You encounter a power outage during install/uninstall and there is no Uninterruptible Power Supply (UPS).

To resolve this problem, perform the following steps:

1. Enter "**regedit**" in the **Run** dialog and click **OK** to open **Registry Editor**.
2. Locate and delete the following entry:
"HKEY_LOCAL_MACHINE\SOFTWARE\Arcserve\Unified Data Protection\Engine"
3. Use the search option in the **Registry Editor** to locate and delete all occurrences of the following string:
 - ♦ [Arcserve UDP Agent (Windows) for x86]: {CAAD8AEA-A455-4A9F-9B48-C3838976646A}
 - ♦ [Arcserve UDP Agent (Windows) for x64]: {CAAD1E08-FC33-462F-B5F8-DE9B765F2C1E}
4. Use the search option in the **Registry Editor** to locate and delete all occurrences of the string "Arcserve UDP Agent " under the following key:
 - HKEY_CLASSES_ROOT\Installer\Products
 - HKEY_LOCAL_MACHINE\SOFTWARE\Classes\Installer\Products
 - HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Installer\UserData\S-1-5-18\Products
 - HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Uninstall
5. From the command line, delete the service by entering the following commands:


```
sc delete ShProvd
sc delete CASAD2DWebSvc
```
6. Run the command line to remove additional setup files.
 - ♦ x86 operating system:

```
"%ProgramFiles%\Arcserve\SharedComponents\Arcserve Unified Data Protection\Setup\uninstall.exe" /q
```

- ◆ x64 operating system:

```
"%ProgramFiles(x86)%\Arcserve\SharedComponents\Arcserve Unified Data Protection\Setup\uninstall.exe" /q
```

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.

Troubleshooting User Interface 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. 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.

Unable to display Arcserve UDP Agent (Windows) home page in IE

If you are using an Internet Explorer (IE) web browser to access the Arcserve UDP Agent (Windows) home page and it does not display, it may be that the Arcserve UDP Agent (Windows) website is not included as a "Trusted Site" in your IE browser.

If this condition occurs, add this website as a Trusted Site in your IE browser. For more information about adding a website as a Trusted Site, see [Security zones: adding or removing websites](#).

If the problem persists, use [Live Chat](#) to contact Arcserve support. Live Chat lets you optimize intelligent conversation between you and the Technical Support team, allowing you to address your concerns and questions immediately, while still maintaining access to the product.

Job Monitor data speed displays a 0 or some other abnormal value

Symptom

Windows Performance Counters are disabled.

Solution

From the Registry Editor, delete or enable the following registry keys on all Windows versions:

- Perflib

Path: HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib

Name: "Disable Performance Counters"

Type: DWORD

Value: Set to 0 to enable performance counter.

- Performance

Path: HKLM\SYSTEM\CurrentControlSet\Services\PerfProc\Performance

Name: "Disable Performance Counters"

Type: DWORD

Value: Set to 0 to enable performance counter.

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.

Troubleshooting Backup 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.

Note: If you convert a basic disk to a dynamic disk, and then you restart the server. When you perform an Incremental backup, the backup will be as large as a full backup for that disk. The reason is, when you change the disk from basic to dynamic, Arcserve UDP considers the dynamic disks as a new disk and performs a full backup for the first time. From the next backup, the backup job will be an Incremental backup.

- [SQL Server backup failed due to Out of memory error](#)
- [Backup sessions for Arcserve UDP Agent \(Windows\) do not include any Microsoft SQL database information](#)
- [Catalog Job fails when backing up a large number of files because of less space](#)
- [Catalog Job fails when backing up a large number of files on Windows 2003 x86 machine](#)
- [Failed to create snapshot for selected volumes](#)
- [Unable to change backup destination folder to Arcserve UDP Recovery Point View](#)

SQL Server backup failed due to "out of memory" error

This is caused by a Microsoft known issue: Volume Shadow Copy Service (VSS) cannot create a volume snapshot even when VSS has sufficient memory space.

To resolve this problem, apply the Microsoft [patch](#).

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 sessions do not include Microsoft SQL database information

After upgrading from a previous release, the backup sessions of Arcserve UDP Agent (Windows) do not include any Microsoft SQL database information. This may be caused by the SQL server not starting automatically in a virtual environment. If this occurs, verify that the SQL database is in a good state and retry the backup.

If the problem persists, you can change the startup type of the SQL server to "Automatic (Delayed Start)".

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.

Catalog Job fails Due to Less Space when Backing up Large Number of Files

If you are attempting to back up a large number of files and the catalog generation job fails because there is not enough available space in the Arcserve UDP Agent (Windows) home folder, perform the following task to create a new temp location:

Important! Verify that this new location contains enough free space to hold all of your catalog temporary data.

1. Within the Arcserve UDP Agent (Windows) home folder, access the **Configuration** folder. (The Arcserve UDP Agent (Windows) home folder is located on the Arcserve UDP Agent (Windows) install path).

Program Files\Arcserve\Unified Data Protection\Engine\Configuration

2. Within the **Configuration** folder, create a **switch.ini** file. (File name is case sensitive).
3. Within the new **switch.ini** file, add the following content:

```
[CatalogMgrDll.DLL]
```

```
Common.TmpPath4Catalog="I:\catalogtemp"
```

4. Run the backup job again.

The catalog generation part of the job will now go to the newly created temp folder.

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 create snapshot for selected volumes

If a volume does not have enough disk space, the backup job can fail with the error message "Failed to create snapshot for selected volumes". If the backup job fails, you can perform either task:

- Free up some disk space on the volumes being backed up.
- Reconfigure the **Volume Shadow Copy** settings to save shadow copy to a volume with sufficient free disk space.

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 change backup destination folder to Arcserve UDP Recovery Point View

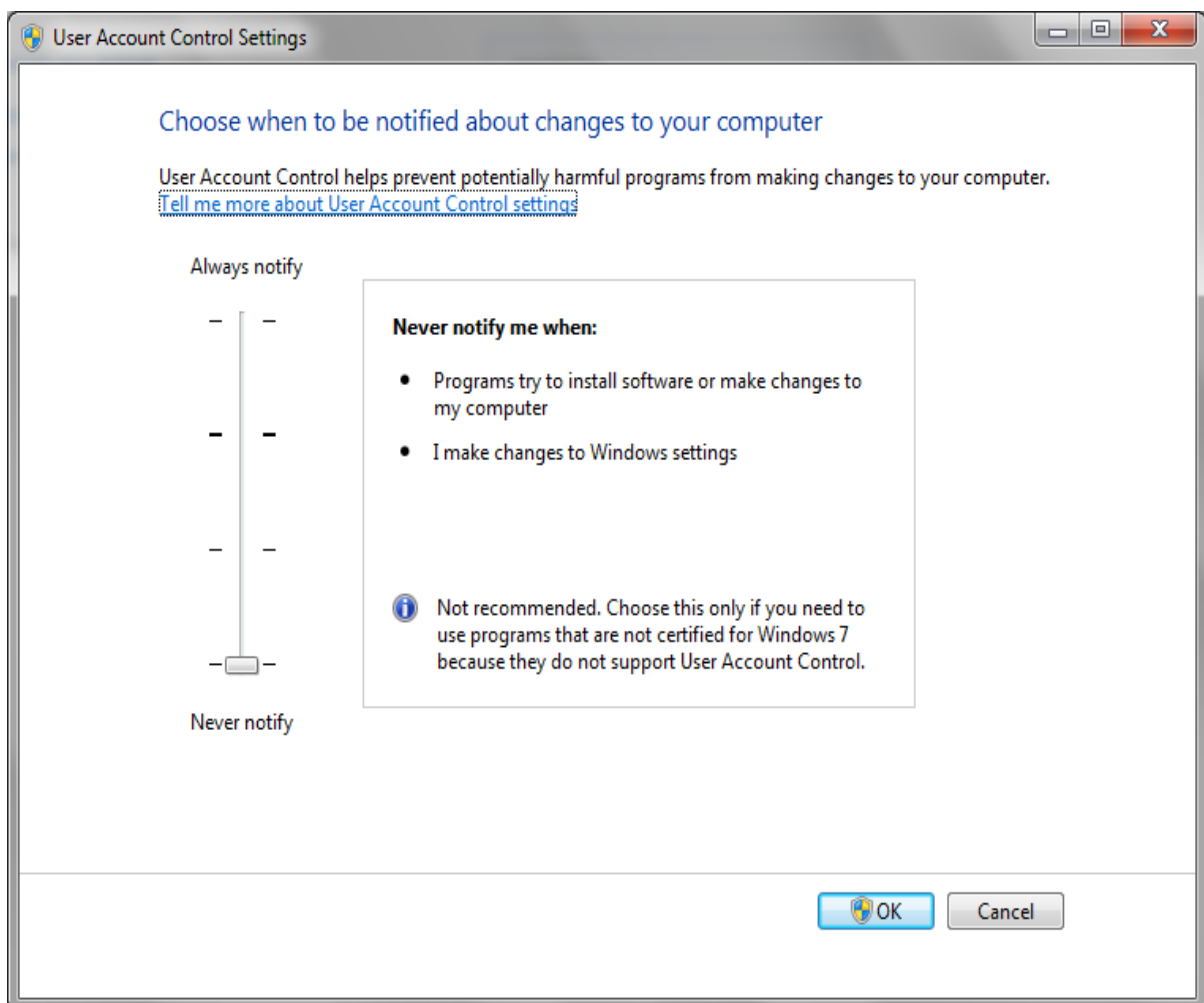
On Windows Vista and later operating systems, if you create an account which belongs to a local administrator group, and from this new account you attempt to change the Arcserve UDP Agent (Windows) backup destination folder to Arcserve UDP Recovery Point View, the folder view cannot be changed and no error message is displayed. This can happen when the **User Account Control** is enabled.

If this condition occurs, you can either disable the **User Account Control** or you can grant Modify privileges to the created Windows account.

To disable the User Account Control, perform the following task:

1. From the Windows **Control Panel**, select **User Accounts**, **User Accounts**, and then **Change User Account Control Settings**.

The **User Account Control Settings** dialog displays.




2. For the **Choose when to be notified about changes to your computer** option, drag the slide bar to the bottom (Never notify).
3. When you disable the **User Account Control**, reboot your computer.

To grant Modify privileges to the created Windows account, perform the following task:

1. From the **Windows Explorer** view, navigate to the specified backup destination.
2. Right-click on the backup destination folder, select **Properties**, and click the **Security** tab.
3. Click **Edit** and Add a user for this destination folder.

The **Permissions** dialog is displayed.



4. For this user, check the **Modify** permissions option to allow control specifically to this user and add it to the folder security list.

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.

Troubleshooting BMR Issues

When a problem is detected, Arcserve UDP Agent (Windows) generates a message to help you identify and resolve the problem. These messages are contained in the Arcserve UDP Agent (Windows) **Activity Log**, which is accessed from the **View Logs** option on the home page UI. In addition, when an incorrect action is attempted, Arcserve UDP Agent (Windows) generally displays a pop-up message to help you identify and quickly resolve the problem.

Slow throughput performance during BMR

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.

Troubleshooting Merge 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.

Merge Session is Skipped

If the oldest recovery point is skipped in the merge operation, perform the following troubleshooting procedure after you submit a new backup when the specified recovery point count limit is exceeded:

1. Open the Mount Recovery Point dialog to see if you have any recovery points mounted. If any recovery points are mounted, dismount them.
2. Open Windows Explorer and switch to the backup destination to see if the session is under the Arcserve UDP Recovery Point View. If it is, change to the Windows Explorer View.
3. Verify if any file copy jobs are running.

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.

Merge Job Failed when Configured to Retain Recovery Sets

Symptom

Possible network failure or busy network.

Solution

Perform any of the following to fix this problem:

- Run a new backup job which will trigger a merge job after the backup is completed.
- Access the Backup Settings dialog and save the Retention Setting again.
- Restart the 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.

Merge Job Fails After Being Paused by a Restore Job

If a merge job is running and you perform another job at the same time, the merge job will automatically pause. After the other job is finished, if you attempt to resume the merge job, then the merge job will fail. This may be caused by the session not being released after the other job is finished, and the unreleased session could not be merged. If a mounted session was not cleanly dismounted, the session lock may not disappear and as a result the session will not be released after the job is finished. If this occurs, run the following command to force a clean session dismount:

```
"%caarcflash_home%\bin\driver\AFMntDrvInstall.exe" -stop
```

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.

Troubleshooting Exchange 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.

Fail to Restore Exchange Database in DAG Node to Original Location

If you want to restore the Exchange database (DB), including DAG DB or local DB in the node, in an Exchange DAG environment, ensure that the status of exchange services with Startup Type as Automatic are running for all the nodes for DAG.

If the problem persists, use [Live Chat](#) to contact Arcserve support. Live Chat lets you optimize intelligent conversation between you and the Technical Support team, allowing you to address your concerns and questions immediately, while still maintaining access to the product.

Restore Job Fails During Dump Exchange Database

Valid on Windows operating system

Symptom

The restore job fails for the following two scenarios:

- When you dump Exchange database to file only and the **Replay log on database** option is selected in host-based agentless backup session. Also, the Exchange writer is not installed on the proxy server.
- You use an agent (Agent A), which does not have the Exchange writer installed, to browse recovery points backed up by another agent (Agent B). You want to dump Exchange database to file only and you have selected the **Replay log on database** option from Agent A.

Solution

To resolve this issue, do not select the **Replay log on database** option.

Unable to Connect across the Domain Live Mailbox from Exchange GRT utility

Symptom

In two scenarios, the proxy cannot connect live mailbox in Exchange server and the following error message is displayed:

Could not resolve domain name . Try using server's IP address instead.

- When proxy is in Workgroup or does not share the same domain with Exchange server then performing HBBU backup and open Exchange GRT utility in proxy cannot connect live mailbox in Exchange server even after using the IP address.
- Connect live mailbox failed across domain in Exchange GRT utility even after using the IP address. The two domains do not have the same windows version. For example, the first one is Windows 2008 and the other one is Windows 2012.

Solution

Add one item in hosts file and save at the following location:

`C:\Windows\System32\drivers\etc`

For example:

`102.54.94.97 DesExchangeServer.domain.com`

APPENDIX: Frequently Asked Questions (FAQ)

This section contains the following topics:

File Copy Related FAQ	750
Encryption Related FAQ	763
Exchange Granular Restore FAQ	767
Service Related FAQ	770
Updates Related FAQ	772

File Copy Related FAQ

The following Frequently Asked Questions are related to the File Copy feature.

Can I restore data if I lose the encryption password?

No. To restore encrypted data that was backed up, you must provide the proper encryption password.

What is the maximum file size that can be backed up/restored?

You do not face any restriction on the file size that can be backed up or restored using Arcserve UDP Agent (Windows) (for example large Outlook PST files, CAD files, video broadcast files).

What is not deleted during a File Copy – Delete Source job?

Yes. The Arcserve UDP Agent (Windows) will exclude deletion of all system state files and application files and folders during a File Copy – Delete Source job. The Arcserve UDP Agent (Windows) only supports Microsoft Exchange and SQL Server and the list of application files is obtained by querying the VSS writers.

Does a File Copy job copy data directly from the local source disks?

A File Copy job will mount the Arcserve UDP Agent (Windows) backup disks and it will then copy the data. It does not actually read from the local source disks.

What is the maximum file size that can be stored on Amazon S3 cloud?

There is no maximum file size that can be stored on an Amazon S3 cloud location.

For any file size less than 64K, will Arcserve UDP Agent (Windows) copy the entire file?

Yes, the granularity limit for block-level incremental backups is set for 64K. The minimum size for a block-level incremental (BLI) backup is 64K.

Can a File Copy job and a Backup run simultaneously?

Yes. The Arcserve UDP Agent (Windows) allows both jobs to run at the same time.

During a File Copy job, will the stub files be copied again?

No. During a File Copy job, Arcserve UDP Agent (Windows) will ignore the stub files and not copy them again.

Does every File Copy job initiate a VSS snapshot like a regular Arcserve UDP Agent (Windows) Backup job?

No. The VSS snapshot is performed only during a Backup job and not during a File Copy job.

Will a File Copy stored on an Amazon S3 cloud location be open source archive format?

No. A File Copy that is stored on an Amazon S3 cloud location will be proprietary only format.

If a File Copy – Delete Source job deletes files, will I be able to perform a BMR from the file copy destination?

No. You just need to perform a restore from the file copy destination. The files that are deleted are only deleted from the source, and not from the recovery point. The recovery points contain the full volume information necessary to perform a complete BMR.

For a File Copy job, is the Delete Source option enabled by default?

No. This option is selected by you when you add a task or set backup settings.

Encryption Related FAQ

The following Frequently Asked Questions are related to the Encryption feature.

If I change the encryption type or the encryption password and the maximum number of recovery points are then reached, what happens?

The image consolidation during backups will continue as usual for images with the older password. When the remaining oldest image is the last Full Backup with the old password, that Full Backup will be deleted.

If I enter a new encryption password, will the old encryption password be asked for first?

No. The Arcserve UDP Agent (Windows) will immediately apply the new password and no longer request the old password.

What happens to data encrypted either using Windows or a third-party encryption system?

- For Windows Encryption File System (EFS) encryption, Arcserve UDP Agent (Windows) will write in encrypted format used in the EFS format.
- For third-party-encryption, it depends on the technology. If volume encryption is enabled or locked, Arcserve UDP Agent (Windows) will not be able to read it and will generate an error.

Exchange Granular Restore FAQ

The following Frequently Asked Questions are related to the Exchange Granular-Level Restore feature.

Can Exchange search attachments in email?

Yes, you can search using keywords of subject, from, to, within date range of send / received time, and find the email attachment as well as the content in the attachments.

Can I restore a mailbox without overwriting the existing data?

Yes, you can restore an entire mailbox and it will not overwrite the existing data in the mailbox store.

Service Related FAQ

The following Frequently Asked Questions are related to the services:

How do I use a different account to start the Arcserve UDP Agent Service?

If you want to change the account used to start the **Arcserve UDP Agent Service**, you must create an account belonging to the local administrators group, and ensure the **Replace a process level token** user right is assigned to this account. For more information, see the Microsoft documentation on how to assign this user right to an account.

Updates Related FAQ

The following Frequently Asked Questions are related to the Updates feature:

Can I use scripted information for specifying Updates proxy settings?

Yes. You can select the "Use Browser Settings" option on the Proxy Settings dialog to inherit the browser proxy settings (accessed from the Updates Preferences).

Can I use a workstation node as an Updates staging server?

Yes. Your workstation node can be used as a staging server for downloading Arcserve UDP Agent (Windows) updates.

Can I manage/operate Updates together or do I need to configure each node separately (one by one)?

No. You must configure each node individually for Updates.

Does an Updates staging server need a separate Arcserve UDP Agent (Windows) license if I am not using any Arcserve UDP Agent (Windows) functions on same staging server?

No. If you are not using Arcserve UDP Agent (Windows) for any function other than just as an Updates staging server, you do not need to have a separate Arcserve UDP Agent (Windows) license for the staging server.

Can I continue to replicate my recovery points backed up in my local RPS server to the remote managed RPS server, after Upgrade?

Question:

Due to production issues, I have not yet upgraded my Arcserve UDP Console, Recovery Point Server installations and Arcserve UDP agents. They are still running Arcserve UDP Version 5.0 Update 1.

However, I upgraded my remote Recovery Point Server to Update 2 as I could manage some downtime for this server. Can I continue to replicate my recovery points backed up in my local RPS server to the remote managed RPS server?

Answer:

No. It has been observed with some basic testing that such configurations should not have a problem and you should be able to continue replicating data to the remote managed RPS server running Update 2. However, it is strongly recommended that you upgrade all your source nodes running Update 1 to Update 2.

Can I continue to replicate backups from my production systems running Update 2 to a remotely managed RPS server running Update 1, after upgrade?

Question:

I have upgraded all of my source nodes including Arcserve UDP Console, RPS servers and Arcserve UDP agent nodes to Update 2 but my destination RPS node is still running Update 1.

Can I continue to replicate backups from my production systems running Update 2 to a remotely managed RPS server running Update 1?

Answer:

No. This configuration is an unsupported one. Update 2 contains several new updates and enhancements. Given that the destination is still at Update 1, you cannot run replication of recovery points backed up using Update 2 to a server that has an older update. The Replication will try to connect for 10 minutes and show the "Preparing" status. After 10 minutes the replication will stop, and the corresponding job log will have an error entry with the following text:

"The specified timeout period expired while communicating with the web service on the destination server."

This is not a network related issue, but an indication of a destination RPS that is not upgraded to Update 2 yet. It is strongly recommended to upgrade the destination to Update 2 as well to ensure seamless functioning of the system as all units will now be at Update 2 level.

APPENDIX: Using the RDX Cleaner Utilities

This section contains the following topics:

What are the RDX Cleaner Utilities?	780
How to Execute the RDX Cleaner Utility	781
How to Execute the RDX Force Cleaner Utility	785

What are the RDX Cleaner Utilities?

RDX is a removable hard disk drive storage system, containing a docking station (RDX dock) and storage media (removable disk cartridge). RDX disk technology combines the strengths of hard disk drive and tape cartridge data storage, allowing you to back up data like a tape drive, with the instant access of a hard drive. It allows for shorter backup windows and quicker restorations. Indirectly, these utilities help to rotate the RDX media based on the backup schedule, in order to maximize the use of the RDX media.

- The **RDX Cleaner** utility is a tool that helps to purge or clear the current backup RDX media destination, if it does not contain the latest Full Backup. It relies upon a verification process which ensures no full backup exists before it clears the content.

[How to Execute the RDX Cleaner Utility](#)

- The **RDX Force Cleaner** utility is a similar tool that helps purge the current backup RDX media destination, but does not rely on any verification process before it forcefully clears the content. You should only use the RDX Force Cleaner utility when you need to clean all backup sessions at the destination. This utility will perform a complete cleanup of the destination without checking for any existing condition or criteria.

[How to Execute the RDX Force Cleaner Utility](#)

Notes: If the RDX Storage device is configured as Backup destination for incremental backups, then consider the following points:

- If planning to change the RDX cartridges every week, then configure the Backup Settings - >Schedule->Recovery point Retentions as 7.
- If planning to change the RDX cartridges every 5 days (excluding Saturday and Sunday), then configure Backup Settings - >Schedule->Recovery point Retentions as 5.
- If planning to change the RDX cartridge daily, then configure Backup Settings - >Schedule->Recovery point Retentions as 1.
- Always keep the Recover Point Retention number more than the cartridges available in the RDX Storage device.

How to Execute the RDX Cleaner Utility

Before you can use the RDX Cleaner utility, you must download a copy of the utility from the [file transfer site](#).

Follow these steps:

1. Download the appropriate RDX Cleaner utility from the [file transfer site](#):
 - ◆ X64 Platform - RDXCleanerX64.exe
 - ◆ X86 Platform - RDXCleanerX86.exe

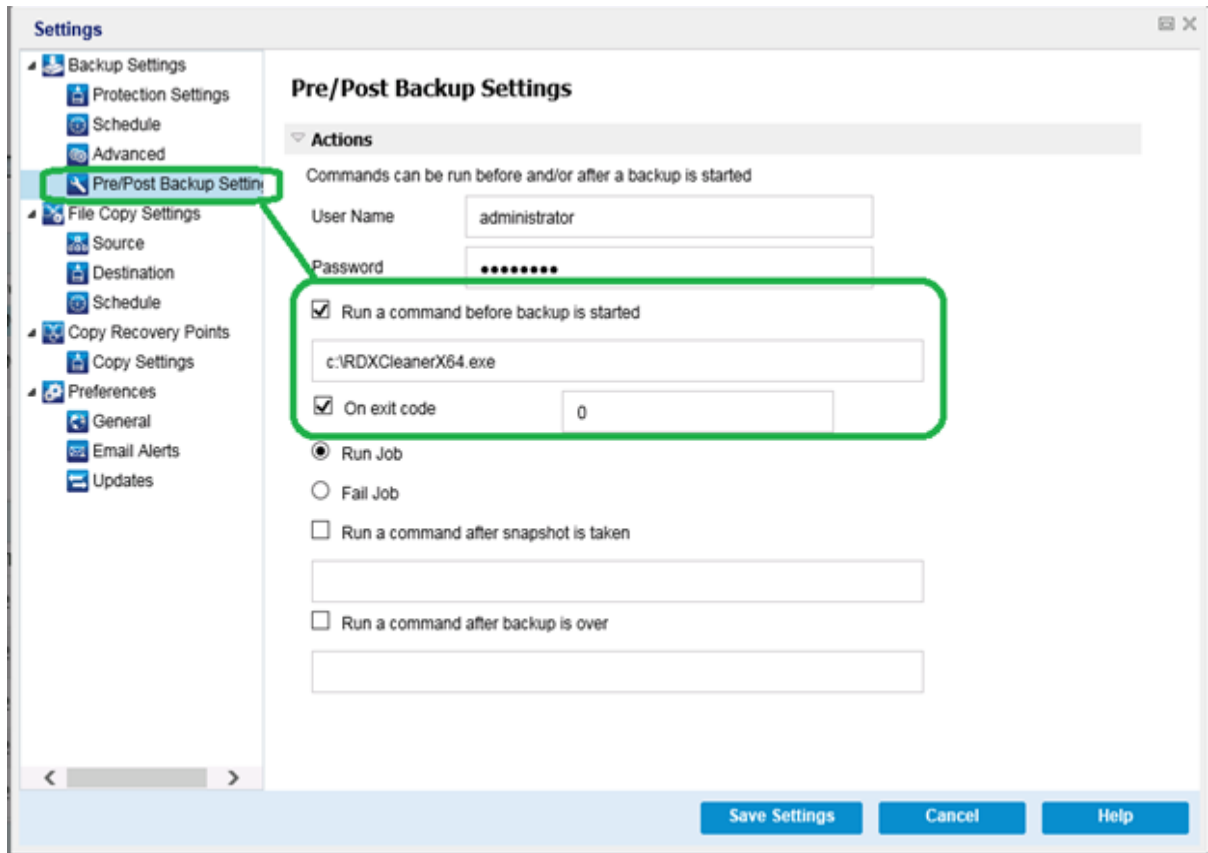
MD5: 8accdc7f14fc30f61e6533b9e16a5758
2. Copy the appropriate version of the RDX Cleaner utility to your local machine (for example C:\) or to any location you specify.
3. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Backup Settings** tab. When the **Backup Settings** dialog opens, select **Pre/Post Backup**.

The **Pre/Post Backup Settings** dialog opens.

4. In the **Actions** section, specify your pre/post backup setting options:
 - a. Select the **Run a command before backup is started** check box.
 - b. Enter the path to the location where you downloaded the RDX Cleaner utility in the command field. For example:
 - ◆ C:\RDXCleanerX64.exe
 - ◆ C:\RDXCleanerX86.exe
 - c. Select the **On exitcode** check box and enter a zero in the On exit code field.

Note: The exit code corresponds to the completion status of the RDX Cleaner command. A zero (0) exit code specifies to run the backup job only when the RDX Cleaner utility successfully completes the deletion of the backup destination content.

- d. Select **Run Job**.



- 5. Click **Save Settings**.

Your pre/post backup settings are saved.

Note: For information about running this utility, see [Post Cleaning Verification \(RDX Cleaner\)](#).

Post Cleaning Verification (RDX Cleaner)

When the RDX Cleaner utility runs, verify the following:

- It creates a new log folder **ClearRDXMediaLogs** in the following location:
C:\Program Files\Arcserve\Unified Data Protection\Engine\Logs

Each time the utility runs, a log file is created with the current time stamp, using the format: **YYYY-MM-DD_HH-MM-SS.txt**

- It clears all the contents of the backup destination folder, except the following files:
 - BackupDestination.ico
 - NodeInfo
 - BackupDev.sig
 - desktop.ini

Before clearing the content of the destination folder, the utility will temporarily move these files to the following folder:

C:\Program Files\Arcserve\Unified Data Protection\Engine\Logs\ClearRDXMediaLogs

After the backup destination has been cleared, the RDX Cleaner utility then moves these files back to the destination folder.

- After the RDX Cleaner utility runs, one of the following codes will be returned:
 - 0 - If one of the following occurs:
 - If the backup destination has the latest Full Backup, then its contents will not be cleared and the backup will be run as submitted.
 - If the backup destination does not have the latest Full Backup, then the content of this destination will be cleared and if the content is deleted successfully a "0" is returned. Because all the content in this destination has been deleted, this backup job will automatically be converted to a Full Backup, regardless of which type was submitted.
 - -1 - Deletion of the backup destination content failed.
 - -2 - Cannot preserve some important files of the backup destination before clearing it.
 - -3 - Current backup destination is not accessible.

Note: The exit code corresponds to the completion status of the RDX Cleaner command. If the exit code is not zero (0), you should check the corresponding log files in the following folder for more detailed information about the reason

for the failure of this cleanup attempt:

C:\Program Files\Arcserve\Unified Data Protection\Engine\Logs\ClearRDXMediaLogs

How to Execute the RDX Force Cleaner Utility

Before you can use the RDX Force Cleaner utility, you must download a copy of the utility from the [file transfer site](#).

Follow these steps:

1. Download the appropriate RDX Force Cleaner utility from the [file transfer site](#):
 - ◆ X64 Platform - RDXForceCleanX64.exe
 - ◆ X86 Platform - RDXForceCleanX86.exe

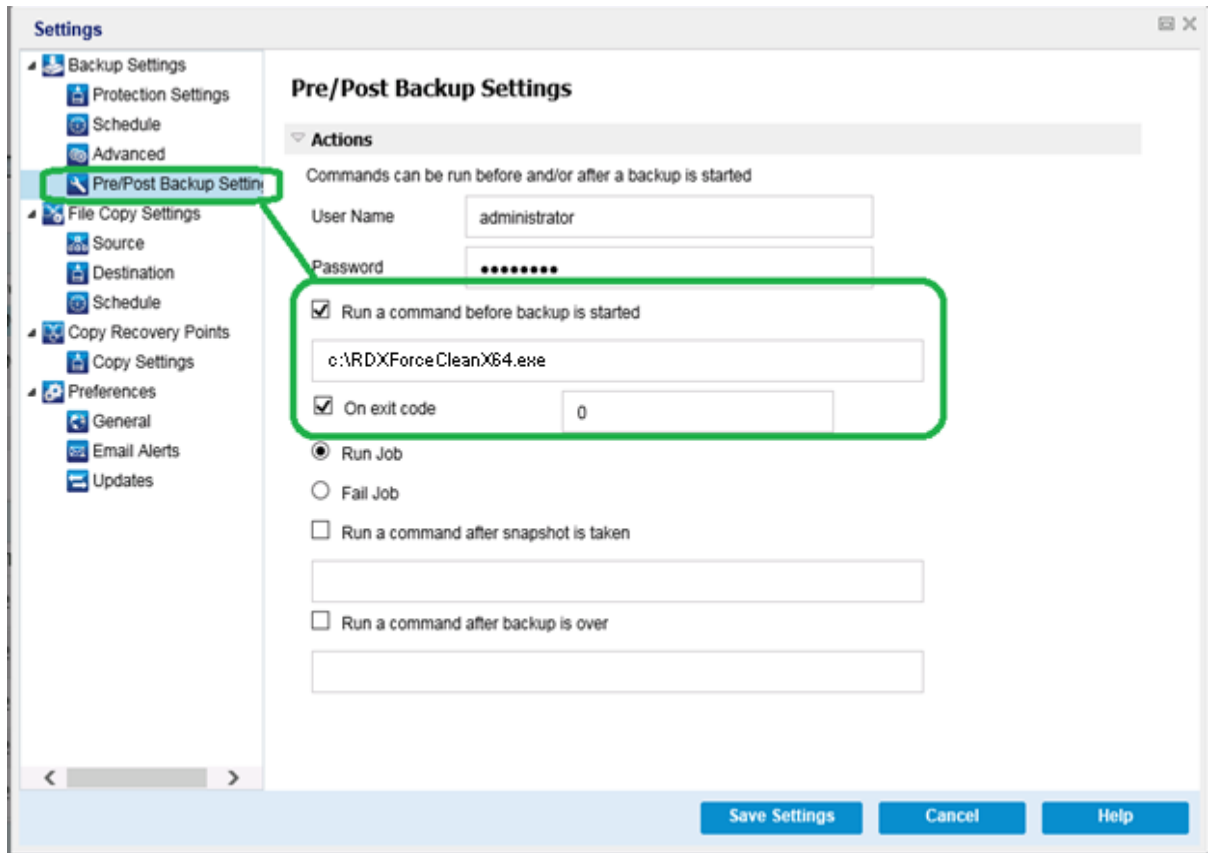
MD5: acd110c67e967f9acfe13f2b0a509d6f
2. Copy the appropriate version of the RDX Force Cleaner utility to your local machine (for example C:\) or to any location you specify.
3. From the Arcserve UDP Agent (Windows) home page (or Arcserve UDP Agent (Windows) Monitor), select **Settings** from the taskbar and then select the **Backup Settings** tab. When the **Backup Settings** dialog opens, select **Pre/Post Backup**.

The **Pre/Post Backup Settings** dialog opens.

4. In the **Actions** section, specify your pre/post backup setting options:
 - a. Select the **Run a command before backup is started** check box.
 - b. Enter the path to the location where you downloaded the RDX Force Cleaner utility in the command field. For example:
 - ◆ C:\RDXForceCleanX64.exe
 - ◆ C:\RDXForceCleanX86.exe
 - c. Select the **On exitcode** check box and enter a zero in the On exit code field.

Note: The exit code corresponds to the completion status of the RDX Force Cleaner command. A zero (0) exit code specifies to run the backup job only when the RDX Force Cleaner utility successfully completes the deletion of the backup destination content.

- d. Select **Run Job**.



- 5. Click **Save Settings**.

Your pre/post backup settings are saved.

Note: For information about running this utility, see [Post Cleaning Verification \(RDX Force Cleaner\)](#).

Post Cleaning Verification (RDX Force Cleaner)

When the RDX Force Cleaner utility runs, verify the following:

- It creates a new log folder **ClearRDXMediaLogs** in the following location:

C:\Program Files\Arcserve\Unified Data Protection\Engine\Logs

Each time the utility runs, a log file is created with the current time stamp, using the format: **YYYY-MM-DD_HH-MM-SS.txt**

- It clears all the contents of the backup destination folder, except the following files:
 - BackupDestination.ico
 - NodeInfo
 - BackupDev.sig
 - desktop.ini

Before clearing the content of the destination folder, the utility will temporarily move these files to the following folder:

C:\Program Files\Arcserve\Unified Data Protection\Engine\Logs\ClearRDXMediaLogs

After the backup destination has been cleared, the RDX Force Cleaner utility then moves these files back to the destination folder.

- After the RDX Force Cleaner utility runs, one of the following codes will be returned:
 - 0 - Deletion of all backup contents was successful.
 - -1 - Deletion of the backup destination content failed.
 - -2 - Cannot preserve some important files of the backup destination before clearing it.
 - -3 - Current backup destination is not accessible.

Note: The exit code corresponds to the completion status of the RDX Force Cleaner command. If the exit code is not zero (0), you should check the corresponding log files in the following folder for more detailed information about the reason for the failure of this cleanup attempt:

C:\Program Files\Arcserve\Unified Data Protection\Engine\Logs\ClearRDXMediaLogs

APPENDIX: Arcserve UDP Terms and Definitions

Agent-Based Backup	790
Compression	790
configuration	791
Dashboard	791
Destination	791
Data Store	791
Discovered Nodes	791
Encryption	791
Host-Based Agentless Backup	792
HOTADD Transport Mode	793
Job	793
NBD Transport Mode	793
NBDSSL Transport Mode	793
Nodes	793
Plan	793
Protected Nodes	793
Recent Event	794
Recovery Point	794
Recovery Point Server	794
Replicate	794
Resources	794
SAN Transport Mode	794
Systems	794
Tasks	794
Unprotected nodes	795

Agent-Based Backup

An Agent-Based backup is a method to back up data using an agent component. The agent is installed on the source node.

Compression

Compression is used for backups. Compression is often selected to decrease disk space usage, but also has an inverse impact on your backup speed due to the increased CPU usage.

The available options are:

No Compression

This option has the lowest CPU usage (fastest speed), but also has the highest disk space usage for your backup image.

Standard Compression

Some compression is performed. This option provides a good balance between CPU usage and disk space usage. This is the default setting.

Maximum Compression

Maximum compression is performed. This option provides the highest CPU usage (lowest speed), but also has the lowest disk space usage for your backup image.

Notes:

- If your backup image contains uncompressible data (such as JPG images, ZIP files, and so on), you may need to allocate additional storage space to handle such data. As a result, if you select any compression option and have uncompressible data in your backup, it could result in an increase in disk space usage.
- If you change the compression level from No Compression to either Standard Compression or Maximum Compression, or if you change from either Standard Compression or Maximum Compression to No Compression, the first backup performed after this compression level change is automatically a Full Backup. After the Full Backup is performed, all future backups (Full, Incremental, or Verify) are performed as scheduled.

This option is available only for the local or remote share destinations. You cannot change the compression setting if the Arcserve Unified Data Protection agent is backed up to data store.

- If your destination does not have sufficient free space, you may consider increasing the Compression setting of the backup. This option is available only for the local or remote share destinations. You cannot change the compression setting if the Arcserve Unified Data Protection agent is backed up to data store.

configuration

A tab on the Arcserve UDP Console to define configuration parameters such as email alerts, database settings, and installation preferences.

Dashboard

A tab on the Arcserve UDP Console that lets you 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.

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).

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.

Discovered Nodes

Discovered nodes are physical or virtual systems that are added to the Arcserve UDP Console by discovering them from active directory or vCenter/ESX server, importing from a file, or manually adding them using its IP address.

Encryption

The Arcserve Unified Data Protection solution provides encryption feature for data.

When the backup destination is a recovery point server, the available encryptions are No Encryption and Encrypt data with AES-256. You can set this to create a data store. When the backup destination is the local or remote share, the available encrypt format options are No Encryption, AES-128, AES-192, and AES-256. You can

set the option while creating a plan to backup to local or share folder, or set this from backup setting for standalone Arcserve Unified Data Protection Agent.

Encryption settings

- a. Select the type of encryption algorithm that you want to use for backups.

Data encryption is the translation of data into a form that is unintelligible without a deciphering mechanism. The Arcserve Unified Data Protection solution uses secure, AES (Advanced Encryption Standard) encryption algorithms to achieve maximum security and privacy of your specified data.

- b. When an encryption algorithm is selected, provide (and confirm) an encryption password.

- ◆ The encryption password is limited to a maximum of 23 characters.
- ◆ A full backup and all its related incremental and verify backups must use same password to encrypt data.
- ◆ If the encryption password for an incremental or verify backup is changed, a full backup must be performed. This means after changing encryption password, the first backup will be full, despite the original backup type.

For example, if you change the encryption password and submit a customized incremental or verify backup manually, it automatically converts to a full backup.

Note: This option is available only for the local or remote share destinations. You cannot disable the encryption setting if the Arcserve Unified Data Protection agent is backed up to data store.

- c. The Arcserve Unified Data Protection solution has encryption password and session password.

- ◆ The encryption password is required for data store.
- ◆ The session password is required for node.
- ◆ If the data store is encrypted, then session password is mandatory. If the data store is not encrypted, the session password is optional.

A password is not required when you are attempting to restore to the computer from which the backup was performed. However, when you attempt to restore to a different computer, a password is required.

Host-Based Agentless Backup

A Host-Based Agentless backup is a method to back up data without using an agent component on the source machine.

HOTADD Transport Mode

The HOTADD transport mode is a data transport method that lets you back up virtual machines configured with SCSI disks. For more information, see the Virtual Disk API Programming Guide on the VMware website.

Job

A job is an Arcserve UDP action to back up, restore, create virtual standby, or replicate nodes.

NBD Transport Mode

Network Block Device (NBD) transport mode, also referred to as LAN transport mode, uses the Network File Copy (NFC) protocol to communicate. Various VDDK and VCB operations use one connection for each virtual disk that it accesses on each ESX/ESXi Server host when using NBD.

NBDSSL Transport Mode

Network Block Device Secure Sockets Layer (NBDSSL) transport mode uses the Network File Copy (NFC) protocol to communicate. NBDSSL transfers encrypted data using TCP/IP communication networks.

Nodes

A node is a physical or virtual system that Arcserve UDP protects. Arcserve UDP can protect physical nodes and virtual machines in a vCenter/ESX or Microsoft Hyper-V server.

Plan

A plan is a group of tasks to manage backup, replication, and creation of virtual standby machines. A plan consists of a single or multiple tasks. Tasks are a set of activities to define the source, destination, schedule, and advanced parameters.

Protected Nodes

Protected nodes are the nodes that have scheduled backup plans to back up data on regular intervals.

Recent Event

Recent Events are the jobs that are still running or jobs that were recently completed.

Recovery Point

A recovery point is a point in time backup snapshot of a node. A recovery point is created when you back up a node. Recovery points are stored on the backup destination.

Recovery Point Server

A recovery point server is a destination node where you install the server. You can create data stores in a recovery point server.

Replicate

Replicate is a task that duplicates the recovery points from one server to another server.

Resources

resources is a tab on the Arcserve UDP Console. From the **resources** tab, you can manage source nodes, destinations, and plans.

SAN Transport Mode

The SAN (Storage Area Network) transport mode lets you transfer backup data from proxy systems connected to the SAN to storage devices.

Systems

Systems are all type of nodes, devices, and virtual machines that can be managed by Arcserve Unified Data Protection. This includes physical, virtual, Linux, and standby virtual machines.

Tasks

A task is a set of activities to define various parameters to back up, replicate, and create virtual standby machines. These parameters include source, destination,

schedule, and some advanced parameters. Each task is associated with a plan. You can have more than one task in a plan.

Unprotected nodes

Unprotected nodes are the nodes that are added to Arcserve Unified Data Protection but a plan is not assigned. When a plan is not assigned, you cannot back up data and the node remains unprotected.