

Arcserve® Backup for Windows

Administration Guide

r16.5 SP1

arcserve®

This Documentation, which includes embedded help systems and electronically distributed materials, (hereinafter referred to as the "Documentation") is for your informational purposes only and is subject to change or withdrawal by Arcserve at any time.

This Documentation may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of Arcserve. This Documentation is confidential and proprietary information of Arcserve and may not be disclosed by you or used for any purpose other than as may be permitted in (i) a separate agreement between you and Arcserve governing your use of the Arcserve software to which the Documentation relates; or (ii) a separate confidentiality agreement between you and Arcserve.

Notwithstanding the foregoing, 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.

© 2015 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® Backup
- Arcserve® Unified Data Protection
- Arcserve® Unified Data Protection Agent for Windows
- Arcserve® Unified Data Protection Agent for Linux
- Arcserve® Replication and High Availability

Contact Arcserve

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

<https://www.arcserve.com/support>

With Arcserve Support:

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

You can access other helpful resources appropriate for your Arcserve product.

Documentation Changes

The following documentation updates have been made since the last release of this documentation:

- Updated to include user feedback, enhancements, corrections, and other minor changes to help improve the usability and understanding of the product or the documentation itself.

16.5 SP1 changes:

- Added [NetApp NAS Cluster Protection](#) (see page 941). This section includes information and limitations for supporting the Arcserve Backup NDMP NAS Option and the NetApp Cluster running the Data ONTAP 8.2 version.
- Updated [Backing Up and Recovering Arcserve D2D Data](#) (see page 767). This topic is now updated with the migration of Arcserve D2D and arcserve Unified Data Protection (UDP), where Arcserve D2D is now D2D/UDP.
- Added [Backing Up and Recovering Arcserve UDP Nodes from the Arcserve UDP Recovery Point Server](#) (see page 789). This new topic describes how to back up and recover Arcserve UDP nodes from Arcserve UDP Recovery Point Server.
- The following topics have been updated to accommodate the new feature, Double Tape Copy that is introduced in this release. Double Tape Copy lets you migrate data from a disk staging device to two destinations simultaneously.
 - [Backup Manager Destination Options](#) (see page 149)
 - [How Backup to Disk to Tape Works](#) (see page 201)
 - [How to Use Disk Staging to Manage Backup Data](#) (see page 202)
 - [Tasks You Can Perform Using Disk Staging](#) (see page 206)
 - [How to Configure Arcserve Backup to Perform Disk Staging Backups](#) (see page 209)
 - [Back Up Data Using Disk Staging](#) (see page 222)
 - [Modify Pending Data Migration Jobs](#) (see page 320)
- Updated [Arcserve Backup Services, Components, and Applications that Require Administrative Privileges](#) (see page 527). The following components were removed from this topic:
 - AMSSigUpdater.exe which lets Arcserve Backup update the Antivirus Scan Engine signature.
 - ScanEngInst.exe which is a utility to install the eTrust Threat Management Agent is now removed from the Arcserve Backup Services, Components, and Applications table.
 - UpdaterX64.exe which the x64 version of AMSSigUpdater.exe.

- Removed the following sections from the Administrator Guide:
 - Global Backup Options > **Backup Manager Virus Options**
 - Global Restore Options > **Restore Manager Virus Options**
 - Global Options for Archiving > **Archive Manager Virus Options**
 - How Arcserve Backup Engines Work > **CA Antivirus Maintenance**
- Removed the alert option "Virus Detected" from the following topics:
 - [Specify Alert Options for Disk and Tape Staging Backups](#) (see page 217)
 - [Restore Manager Alert Options](#) (see page 277)
 - [Archive Manager Alert Options](#) (see page 366)
- Updated [Limitations on Performing Raw Backup and Restore Operations](#) (see page 760). "Arcserve Backup does not scan the backup data for viruses." is no longer a limitation and removed from this list.

16.5 Update 2 changes:

- Updated [Agent for Virtual Machines Options](#) (see page 180). This topic now includes the Transport Mode options.
- Updated [How to Find Files That You Want to Restore](#) (see page 258). This topic now includes another way of recovering a virtual machine by specifying a transport mode.

16.5 changes:

- Updated [Specify Local Backup Options](#) (see page 134). This topic now includes information about how to enable the Windows Server 2012 NTFS Data Deduplication function on local backups.
- Updated [Point in Time Restore Options](#) (see page 159). This topic now contains cross-references to how to enable this option.
- Added [NTFS Data Deduplication](#) (see page 159). This new topic describes how you can perform optimized backups for full backups on data-deduplication enabled volumes. This new option is located on the Global Backup Options dialog under the Advanced tab.
- Added [How NTFS Data Deduplication Works](#) (see page 255). This new topic describes how Arcserve Backup implements the Windows Server 2012 NTFS Data Deduplication function.
- Updated How Backup to Disk to Tape Works and [How Backup to Tape to Tape Works](#) (see page 230). These topics now contain More Information links to topics that describe how to leverage Disk Staging and Tape Staging to protect your backup environment.
- Updated [Best Practices - How to Recover a Arcserve Backup Server from a Disaster without Using the Disaster Recovery Option](#) (see page 287). This topic includes a note about having enough space in the System Reserved Volume and the recommended volume size to have.

- Added [Implement Logical Block Protection](#) (see page 438). This new topic describes how to implement Logical Block Protection on stand-alone tape drives.
- Updated [Configure Cloud Connections](#) (see page 478). This topic describes that you cannot create a unique bucket name that ends with a period. This topic also describes the new Proxy Settings button and the Do not use proxy option.
- Added [Arcserve Backup Cannot Perform a Hyper-V Backup Job That Contains Data on SMB Shares](#) (see page 882). This troubleshooting topic helps resolve failed Hyper-V backup jobs that contain disks, configuration files, or both on SMB shares.
- Added [Arcserve Backup Cannot Perform a Hyper-V Writer or SQL VSS Writer Backup Job That Contains Data on SMB Shares](#) (see page 882). This troubleshooting topic helps resolve failed Hyper-V VSS writer and SQL Server VSS writer backup jobs that contain data on SMB shares.
- Added [Updating Passwords Using ca_jobsecmgr Does Not Complete Successfully](#) (see page 871). This new troubleshooting topic describes how to resolve the issue of updating passwords when using the ca_jobsecmgr command.
- Added [Unable to View All Job Logs for Rotation Backups After the Database is Pruned](#) (see page 904). This new troubleshooting topic describes how to configure the database pruning jobs to retain all of the job logs associated with repeating and rotation jobs.

Contents

Chapter 1: Introducing Arcserve Backup 23

Introduction	23
Arcserve Backup Functionality	23
How to Access Arcserve Backup Managers, Wizards, and Utilities.....	25
Arcserve Backup Utilities.....	30
Merge Utility	30
Media Assure & Scan Utility	32
Compare Utility	33
Count Utility	33
Copy Utility.....	33
Purge Utility	34
Report Writer Utility	34
How You Can Manage Jobs Using the cabatch Command	35
User Profile Utility	35
Arcserve Backup Command Line Utilities.....	35
Arcserve Backup Security	36
Arcserve Backup Administrator Profile	36
How Arcserve Backup Equivalence Works	36
System Account.....	37
Equivalency and the System Account	37
How Backup and Restore Operations Function on 64-bit Windows Platforms	38
How to Protect Virtual Machine Environments	39
Arcserve Backup Enterprise Module	40
Enterprise Level Password Management Utility	40
Using Arcserve Backup Documentation	41
Download the Bookshelf.....	41
Configure the Documentation Location.....	42

Chapter 2: Protecting Data Using Arcserve Backup 43

Arcserve Backup Components	44
Central Management	45
Central Job Management	47
Central Job Monitoring	48
Central Database Management	49
Central Logging	50
Central Reporting.....	50

Central Alert Management	51
Central Arcserve Server Administration.....	52
Central Device Management	52
Central License Management	53
Central Job History	55
Locate Information Using Quick Search	76
How Password Management Works	79
Change a Session/Encryption Password	80
Enable Password Management.....	81
How User Profile Management Works.....	81
Roles and Permissions.....	83
How Windows User Authentication Works	89
Configure Windows Security Setting Option.....	90
Open the Manager or Manager Console	90
Log in to Arcserve Backup	92
Add a Windows User	93
Add an Arcserve Backup User	94
Change Passwords from the Home Page	95
Modify Windows User Properties	96
Modify Arcserve Backup User Properties	96
Delete a User.....	97
Add a User to a Role.....	97
Remove a User from a Role.....	98
Using the Audit Log	98
Create an Audit Log Report	102
How Arcserve Backup Processes Backup Data Using Multistreaming	103
Tasks Supported by Multistreaming	105
Multistreaming Support for Local Backup Jobs.....	106
How Arcserve Backup Processes Backup Data Using Multiplexing	106
Tasks Supported by Multiplexing	109
How Arcserve Backup Secures Data.....	109
Encryption and Decryption	110
Federal Information Processing Standards (FIPS)	110
Arcserve Backup and FIPS Compliance	111
Arcserve Backup Data Encryption	112
Effective Media Management.....	115
Configure Devices Using the Device Wizard	116
Configure Device Groups	117
Back Up and Restore Data.....	118
Backup Requirements Plan	118
Add Computers to the Preferred Shares/Machines Tree	119
Backup Media Rotations and Scheduling Options	120

Preflight Checks for Your Backups	126
Start Arcserve D2D	128
Start Arcserve Replication	130

Chapter 3: Backing Up Data **133**

How Arcserve Backup Lets You Back Up Data	133
Specify Local Backup Options.....	134
Submit a Backup Job	136
Backup Manager	137
Options on the Backup Manager Start Tab	139
How to Specify Source Data Using the Classic View and the Group View	140
Backup Manager Destination Options	149
Backup Job Schedules and Rotations	152
Local Backup Options for UNIX and Linux Agents	152
Global Backup Options	154
Backup Manager Alert Options	154
Backup Manager Media Exporting Options	155
Backup Manager Advanced Options	156
Backup Manager Encryption/Compression Options	159
Backup Manager Volume Shadow Copy Service Options	162
Backup Manager Backup Media Options	164
Backup Manager Verification Options	166
Backup Manager Operation Options	167
Backup Manager Pre/Post Options	171
Backup Manager Agent Options	173
Backup Manager Job Log Options	184
Files and Objects that Arcserve Backup Does Not Back Up	184
Skip or Include Database Files in Backups.....	187
Enable Arcserve Backup to Manage Open Files on Remote Computers.....	189
Multiplexing Job Options	189
Specify Multiplexing Options	190
How the Job Status Manager Monitors Multiplexing Jobs	190
Verify Multiplexing Data Integrity.....	191
Using Multiplexing with Microsoft Exchange Backup Jobs	192
Specify Multistreaming Options.....	192
Entire Node Backups	193
Back Up an Entire Node that Contains Database Files.....	193
Create Repeating Backup Jobs	195
Cross-Job Duplicated Source Check	197
Back Up Remote Servers	197
Submit Static Backup Jobs.....	198

Backup Staging Methods	200
How Backup to Disk to Tape Works	201
How to Manage Backup Data Using Tape Staging	229
Archive Staging Methods	241
Archive Disk Staging	241
Archive Tape Staging.....	243
Backing up Multiple Data Mover Servers in a Single Job	246
Back up Multiple Data Mover Servers in a Single Job	247
Back up Multiple Data Mover Servers in a Single Job Using Staging.....	250
Disaster Recovery.....	254
How NTFS Deduplication Works	255

Chapter 4: Restoring Data 257

Restore Manager.....	257
How to Find Files That You Want to Restore	258
How Arcserve Backup Lets You Browse a Large Number of Items in the Restore Manager	262
Duplicate Backup Sessions	265
Smart Restore	265
Export the Restore by Query Results and View the Results in a Spreadsheet	265
Restore Data by Query on UNIX and Linux Platforms	266
Restore Manager Markers	268
Restore Manager Location Options	269
Restore Job Schedules.....	269
Specify Run as Administrator on Windows Server 2008 and 2012 Systems	270
Global Restore Options	271
Restore Manager Backup Media Options	271
Restore Manager Destination Options	272
Restore Manager Operation Options.....	274
Restore Manager Pre/Post Options	276
Restore Manager Job Log Options	277
Restore Manager Alert Options	277
System State Restore Options.....	278
Restoring Data Scenarios	280
Restore Data that was Backed Up Using Staging	280
Restore a Remote Agent on a System without the Disaster Recovery Option	282
Restore Arcserve Backup Member Servers without the Using the Disaster Recovery Option	283
Best Practices - How to Recover a Stand-alone Server from a Disaster Using the Disaster Recovery Option	285
Best Practices - How to Recover a Arcserve Backup Server from a Disaster without Using the Disaster Recovery Option.....	287
How to Restore Data from Cloud Storage.....	299

Chapter 5: Customizing Jobs

301

Job Customization Methods.....	301
Dynamic Job Packaging	302
Static Job Packaging	305
Convert Jobs Submitted Using the Classic View to the Group View	306
Rotation Schemes.....	307
How You Can Manage GFS Rotation Jobs on File System Devices	309
Media Pool Specification.....	312
Backup Method Options	312
How Job Filters Work	312
Filter Options.....	314
Types of Filters	315
Schedule Custom Jobs.....	317
Custom Schedules	318
Tasks You Can Perform Using the Job Status Manager	319
Modify Pending Data Migration Jobs.....	320
Update Multiple Jobs	322
How to Manage Jobs Using the Job Queue Tab	323
View Job Details Using the Activity Log.....	329
Tape Log Tab	332
Job Detail Tab.....	333
Job Log Tab.....	333
How Save Agent and Save Node Information Works	333
Add, Import, and Export Computers Using the Source Tree.....	334
Add Multiple Agents and Nodes Using .csv and .txt Files	337
Export Multiple Agents and Nodes to a Text File.....	338
Filter Nodes by Agent Type	339
Modify the IP Address or Host Name of Agents and Nodes	340
Delete Agents and Nodes from the Source Tree.....	340
How to Use the Job Scheduler Wizard to Schedule Jobs	341
Job Scripts	342
Create a Job Script.....	342
Execute a Job Using a Script.....	342
Job Templates	343
Create Custom Job Templates	343
Save Job Template	344
Open Job Template	345
Windows-Powered NAS and Storage Server 2003 Device Configuration	346
Access Arcserve Backup Through the Windows-powered NAS Device.....	346
Arcserve Backup and Windows-powered NAS Device Configuration	347
Scan Jobs for Cloud Storage	349

Merge Jobs to Cloud Storage	349
How Arcserve Backup Purges Data from Cloud Storage	350

Chapter 6: Archiving Data **351**

Overview of How Arcserve Backup Archives Data	351
Submit Archive Jobs	351
Archive Manager	353
Options on the Archive Manager Start Tab	354
How to Specify Source Data for Archiving Using the Classic View and the Group View	355
Archive Job Schedules	356
Options on the Archive Manager Policies Tab	356
Options on the Archive Manager Destination Tab	358
Global Options for Archiving	359
Archive Manager Media Options	360
Archive Manager Operation Options	362
Archive Manager Pre/Post Options	364
Archive Manager Job Log Options	366
Archive Manager Alert Options	366
Archive Manager Advanced Options	369
Archive Manager Encryption/Compression Options	370

Chapter 7: Managing Devices and Media **373**

Device Management Tools	373
Tape Library Configuration	373
RAID Device Configuration Option	379
Virtual Library Configuration Option	380
Control Devices Using Removable Storage Management	380
Configure Devices Using Enterprise Module Configuration	382
Device Configuration	383
Device Manager	395
Maintenance Tasks	396
Schedule Device Management Jobs	403
Device Management Functions for Libraries	404
Configure VTLs to Function as Arcserve Backup VTLs	418
Media Movement	419
Device Group Configuration Using the Device Manager	419
Universal Serial Bus (USB) Storage Devices	424
Prerequisites for Backing Up to Removable Drives	425
Filter Libraries	426
Removable Drive Support	427
How Arcserve Backup Supports Write Once Read Many (WORM) Media	428

DLTSage Error Handling.....	429
How Arcserve Backup Cures Tape Drive Errors.....	433
How Arcserve Backup Integrates with Secure Key Manager	433
How to Ensure that Arcserve Backup Spans Media in a Single Drive Autoloader.....	435
Media Assure	436
How Uninterrupted Drive Cleaning Works	437
Implement Logical Block Protection	438
How to Optimize Tape Usage.....	439
How Media Pools Work.....	445
Save Sets and Scratch Sets.....	447
Serial Numbers.....	448
GFS Media Pools.....	449
Media Maximization in GFS Rotation Jobs.....	450
Media Pool Manager.....	454
Create Media Pools	455
How You Can Create a Rotation	456
Media Management Administrator (MM Admin).....	457
Media Management and Tape Service	457
Media Management Administrator Terms	458
MM Admin Interface	458
MM Admin Toolbar	459
MM Admin Window.....	459
Schedule Object	460
Reports Object	462
Find Media in Vault Object.....	464
Status Object.....	464
Reset the Status of Vault Processing.....	464
How the Media Management Process Works.....	465
Vault Management	466
Create Schedules.....	468
Modify Schedules.....	468
Delete Tape Volume Movement Schedules	469
How You Can Manage Tape Volumes and VCDs	469
Tape Volume Retention Policies	471
Slot Detail and Status Information	474
Find Specific Media in a Vault	476
How to Manage Cloud Storage	477
Cloud Storage Limitations	477
Configure Cloud Connections.....	478
Create Cloud-Based Devices	480
Browse Cloud Folders.....	483
Format Cloud-Based Device	485

Erase Cloud-Based Device	486
Online and Offline Cloud-Based Device	486
Specify Compression on Cloud-Based Device	487
Migrate Data to Cloud Storage in a Staging Backup Job	487
Configure Throughput to Eucalyptus-Based Cloud Media	489

Chapter 8: Administering the Backup Server 491

How Arcserve Backup Engines Work.....	491
How Engine Status Affects Arcserve Backup Operations	492
Service State Icons	493
Stopping and Starting Arcserve Backup Services	493
Configure Arcserve Backup Engines	499
Job Engine Configuration	500
Tape Engine Configuration	503
Database Engine Configuration.....	514
Alert Configuration.....	518
Additional Server Admin Functions.....	520
Modify the Arcserve Backup System Account	520
Reconfigure Node Tier Assignments	521
Manage Arcserve Backup Component Licenses.....	523
Release Licenses from Servers	525
Configure Multiple Network Interface Cards.....	526
Authentication Levels for Arcserve Backup Services, Components, and Applications	527
Arcserve Backup Services, Components, and Applications that Require Administrative Privileges	527
Arcserve Backup Services, Components, and Applications that Require the Highest Available Privileges	533
Arcserve Backup Domains.....	538
Manage Domain Users and Groups Using the ca_auth Command Line Utility.....	538
Create caroot Equivalence	539
How to Manage Multiple Domains Using the Job Status Manager	539
How to Process Computer Name Changes in an Arcserve Domain	542
Manage User Profiles Using the User Profile Utility.....	554
Add a User Using the User Profile Utility	554
Delete a User Using the User Profile Utility	555
Change a User Password Using the User Profile Utility	555
Assign Roles to Users Using the User Profile Utility.....	556
Suspend Users Using the User Profile Utility	556
Restore the Arcserve Backup Job Queue	557
BSBAB--Manage Arcserve Servers Using the Server Configuration Wizard-AdminW	559
Tasks You Can Perform Using the Server Configuration Wizard	561
Data Migration Limitations in a Arcserve Backup Domain.....	562
Start the Server Configuration Wizard	564

Promote a Member Server to a Primary Server.....	565
Demote a Primary Server or Stand-alone Server to a Member Server.....	568
Move a Member Server to a Different Arcserve Backup Domain.....	572
Change the Password for the Arcserve Backup Domain Administrator (caroot) Account.....	573
Repair the Arcserve Backup Configuration.....	574
Repair the Arcserve Database Connection on a Primary Server.....	575
Repair the Arcserve Database Connection on a Member Server.....	576
Register Virtual Clusters as Arcserve Backup Domain Member Servers.....	577
How Arcserve Backup Protects Active Directory Data on Domain Controller Servers.....	579
Back up the Active Directory.....	582
Restore Active Directory Objects.....	583
Microsoft Exchange Server 2010 Mailboxes Appear as Legacy Mailboxes After Recovering the Active Directory.....	589
Reset Microsoft Exchange Server User Passwords After Recovering the Active Directory.....	589
Install and Uninstall Arcserve Backup Server Based Options.....	590
Discovery Configuration.....	591
How the Discovery Service Detects Other Computers.....	592
IP Subnets/Windows Domains Discovery.....	593
Enable Discovery Using TCP/IP Subnet Sweep.....	594
Discovery Configuration for the SAN Option.....	597
Discover Client Agent Systems with Non-default IP Addresses.....	597
Arcserve Backup Maintenance Notifications.....	598
Disable Maintenance Notification Messages.....	600
Enable Maintenance Notification Messages.....	600
Apply Arcserve Backup Component Licenses.....	601
Managing Firewalls.....	602
Allow Arcserve Backup Services and Applications to Communicate Through the Windows Firewall.....	602
How to Configure Your Firewall to Optimize Communication.....	602

Chapter 9: Managing the Database and Reporting **603**

How to Manage the Database and Reports.....	603
Database Manager.....	604
Database Views.....	604
Enable Media Pool Maintenance.....	607
How to Protect the Arcserve Backup Database.....	607
Agent for Arcserve Database.....	608
How the Database Protection Job Works.....	613
How to Back Up the Arcserve Backup Database.....	614
Modify, Create, and Submit a Custom Database Protection Job.....	614
Specify Microsoft SQL Server 2008 Express Backup Options for the Arcserve Backup Database.....	617
Specify Microsoft SQL Server Backup Options for the Arcserve Backup Database.....	619

Start the Arcserve Backup Database Protection Job	623
Access Requirements	624
Delete the Arcserve Backup Database Protection Job	624
Re-create the Arcserve Backup Database Protection Job	625
Re-initialize the Arcserve Backup Database	626
How to Restore the Arcserve Backup Database	628
Recover the Arcserve Backup Database Using Arcserve Database Recovery Wizard	628
Recover the Arcserve Backup Database Using the ca_recoverdb Command	632
Specify Microsoft SQL Server 2008 Express Edition and Microsoft SQL Server Database Restore Options	635
Restore the Arcserve Backup Database (Different Domain)	646
How to Recover the Arcserve Database When the SQL Server Instance Hosting the Arcserve Database is Not Functional	649
How Arcserve Database Recovery Wizard Works	649
How the Catalog Database Works	650
Catalog Browsing	652
Catalog Database Pruning	652
How a Centralized Catalog Database Works	653
Configure the Catalog Database	653
Move the Arcserve Backup Catalog Database to a Different Location	654
Using Microsoft SQL Server as the Arcserve Backup Database	658
Microsoft SQL Server Database Considerations	658
Remote Database Considerations	659
Specify ODBC Communication for Remote Database Configurations	660
How to Calculate the Number of Required SQL Connections	661
How to Enable TCP/IP Communication on Microsoft SQL Server Databases	661
Database Consistency Checks	661
Specify a Arcserve Backup Database Application	662
Configure Microsoft SQL Server as the Arcserve Backup Database	662
Move the Arcserve Backup Database to a Different System or Instance	664
Configure Microsoft SQL Server 2008 Express as the Arcserve Backup Database	667
Arcserve Backup Logs and Reports	667
Activity Log Data	668
Tape Log	668
Job Log	668
Report Manager	669
Report Manager Reports	671
Custom Report Job Scheduling	679
Create Custom Reports Using the Report Writer Utility	680
Report Generation for Multiple Arcserve Backup Servers	681
Arcserve Backup Diagnostic Utility	682
Diagnostic Utility Components	682

Configure Computers Running Windows Vista and Windows 7 Operating Systems to Communicate with the Diagnostic Wizard	683
Create Reports Using the Express Mode Diagnostic Utility	684
Create Reports Using the Advanced Mode Diagnostic Utility	685
View Reports Using the Diagnostic Report Manager	686
Arcserve Backup Infrastructure Visualization	687
Infrastructure Visualization Software Requirements	688
Infrastructure Visualization Operations	688
Infrastructure Visualization Color Scheme	689
Arcserve Backup Infrastructure Visualization Views	689
Dashboard Integration with Infrastructure Visualization	703

Chapter 10: Using the Alert Manager **705**

How the Alert Manager Works	705
Alert Manager Components	707
Set Up Alerts	707
Alert Manager Configuration	709
Ports Option	709
Broadcast Alerts	710
Email Notification	710
Windows Event Log Notification	711
Alert Manager Pager Options	712
Configure SMTP Notification	713
SNMP Notification	713
Trouble Tickets	714
Event Priorities	714
Message Testing	715
Alert Activity Details	715

Chapter 11: Managing Agents Using Central Agent Admin **717**

How the Arcserve Backup Central Agent Admin Works	717
Manage Agents	718
Modify Agents	718
Configure Agent Security	718
Start or Stop Agent Services	719
Start Agent Deployment from the Central Agent Admin	719
Configure Agents	719
Add Computers	721
Add Nodes	721
Manage Agent Logs	723
Configure SRM PKI	725

Configure SRM Exclude Paths	726
Configure Node Tiers.....	727

Chapter 12: Using Deduplication **729**

How Data Deduplication Works	729
How to Plan a Deduplication Installation	731
Deduplication Considerations	732
Supported Functions Matrix	733
Licensing Requirements for Deduplication	735
Create Data Deduplication Devices.....	735
Deduplication Device Group Configuration	738
Device Commands for Data Deduplication Devices	738
Back up Data with Deduplication	738
How Regular Backup Jobs Work with Deduplication	739
Configure Deduplication Groups to Use Staging.....	739
How Staging Jobs Work with Deduplication	740
Back Up Data Using Global Deduplication	750
Recover Deduplicated Data.....	751
Restore Deduplicated Data	751
Scan Jobs with Deduplication.....	753
Merge Jobs with Deduplication	754
GFS Rotation Jobs on Deduplication Devices	754
Deduplication Device Purge	755
Delete Deduplication Backup Sessions	756
Deduplication Reports.....	757

Chapter 13: Raw Backup and Restore of Physical Disks and Volumes **759**

Overview of Raw Backup and Restore	759
Licensing Requirements for Raw Backup of Physical Disks and Volumes	759
How Raw Backup Works.....	760
Supported Functions	760
Limitations on Performing Raw Backup and Restore Operations	760
Naming Convention of Physical Disks and Volumes	762
Enable Raw Backup and Restore	763
Perform Raw Backup of a Physical Disk or Volume.....	763
Entire Node Backup.....	764
Raw Backup Restore.....	764
Restore to an Alternate Location as a File	765
Restore to the Original Location	765
Restore to Another Physical Disk or Volume	766

Chapter 14: Backing Up and Recovering D2D/UDP Data **767**

Specify D2D/UDP Backup Options.....	767
How Arcserve Backup Processes Encrypted D2D/UDP Backup Sessions	768
Submit Backup Jobs of D2D/UDP Nodes	768
Submit Backup Jobs of D2D/UDP Nodes through Proxy Servers	770
Submit Static Backup Jobs Containing D2D/UDP Data.....	772
Submit Dynamic Backup Jobs Containing D2D/UDP Data.....	774
Recover D2D/UDP Data at File Level Granularity.....	775
Recover D2D/UDP Data at Application Level Granularity	777
Recover Files and Folders from D2D/UDP Recovery Points	783
Recover D2D/UDP Data from Raw Sessions.....	785
Configure Arcserve Backup to Generate D2D/UDP Catalog Files.....	787

Chapter 15: Backing Up and Recovering Arcserve UDP Nodes from the Arcserve UDP Recovery Point Server **789**

Submit Backup Jobs of Arcserve UDP Nodes through the Arcserve UDP Recovery Point Server	789
Submit Static Backup Jobs Containing Arcserve UDP Data	791
Submit Dynamic Backup Jobs Containing Arcserve UDP Data	793
Recover Arcserve UDP Data at File Level Granularity	795
Recover Arcserve UDP Data at Application Level Granularity.....	797
Recover Arcserve UDP Data from Raw Sessions	802
Configure Arcserve Backup to Generate Arcserve UDP Catalog Files	805

Chapter 16: Using Synthetic Full Backup **807**

Overview of Synthetic Full Backup.....	807
How Synthetic Full Backup Works.....	808
Submit a Synthetic Full Backup Job	811
Specify the Synthetic Schedule for a Synthetic Full Backup	813
Submit a D2D2T Synthetic Full Backup Job Using a Rotation Scheme	813
Submit a D2D2T Synthetic Full Backup Job Using a GFS Scheme	819
Submit a D2D2T Synthetic Full Backup Job Using a Custom Schedule.....	825
Run a Real Full Backup on the Next Full Backup Day	830
Submit an On-Demand Synthetic Full Backup.....	831
Scan Synthetic Full Backup Session Only.....	831
How the Purge Policy and a DB Pruning Job Works for Synthetic Full Backup	832
How Makeup Jobs Work for Synthetic Full Backup.....	832
Point In Time Restore.....	833
Synthetic Full Backup Considerations	836
Working with Synthetic Full Backup.....	837

Chapter 17: Using Arcserve Backup in a Storage Area Network 841

How to License the Storage Area Network (SAN) Option	841
The SAN Environment	842
How Arcserve Backup Works in a SAN	842
Server Management in a SAN	844
Backup Plans	844
Benefits of Using the Option	844
Terminology	844
Install the SAN Option	845
Operating System Compatibility	845
Installation Prerequisites	845
SAN Option Installation	847
Uninstall the Storage Area Network Option	848
Using the SAN Option.....	848
Storage Area Network (SAN) Configuration.....	848
Create Shared Device Groups	849
Data Backup and Restore in a SAN Environment	850
Device Management	851
Media Management.....	851
Control of Job Runtime	852
Reports and Logs.....	852
Arcserve Virtual Libraries	853

Appendix A: Troubleshooting 855

Log in Problems	855
Unable to Log In After Changing the caroot Password	855
Makeup Jobs Created When the Media is Full.....	857
Unable to Log In to Arcserve Backup After Changing the Computer Name	857
Arcserve Backup Cannot Communicate after Changing the IP Address of a Arcserve Backup Server.....	858
Authentication Problems	863
Authentication Security Settings.....	863
Restricted Users Cannot Access the Activity Log and the Audit Log	865
Authentication Errors Occur When Stopping and Starting the CAportmapper Service.....	868
Credential Problems with Cloud Connection	868
Unable to Open Report Writer.....	869
Unable to Create Reports from Report Manager	870
Updating Passwords Using ca_jobsecmgr Does Not Complete Successfully	871
Backup and Restore Problems	872
Jobs Do Not Start on Schedule	872
Cannot Back Up Open Files	873
Restore Job Fails on Citrix Server	874

Local Restore Data Backed Up with Compression and/or Encryption Failed	875
Arcserve Backup Does Not Restore Data Based on File Access Time	875
GUI Freezes in Active Directory Restore Mode	876
Scheduled Backup Jobs Fail After Changing the Login Credentials for Agent Computers	876
Running a Full Backup (Clear Archive Bit) or Incremental Backup on a Source Previously Backed Up in a SFB Job Causes Data Loss	877
Unable to Run an On-Demand Synthetic Full Backup	878
Restore Job Activity Log Errors.....	879
Arcserve D2D Backup Jobs Fail When Submitted from the Backup Manager	880
Arcserve Backup Does Not Back up Files that were Moved or Renamed	880
Arcserve Backup Cannot Perform a Hyper-V Backup Job That Contains Data on SMB Shares	882
Arcserve Backup Cannot Perform a Hyper-V Writer or SQL VSS Writer Back Up Job That Contains Data on SMB Shares	882
Memory Usage Increases When Restoring Data From a Remote FSD Located on a Windows 2008 SP2 and Windows 2008 R2 Server	882
Restore Job Fails on MAC Agents	883
Media Problems	883
Tape Errors Occur When Backing Up or Restoring Data	883
Arcserve Backup Cannot Detect RSM Controlled Devices on x64 Platforms	885
Arcserve Backup Does Not Detect a Cleaning Tape	886
Hardware Does Not Function as Expected.....	887
Autoloaders and Changers Appear Offline	888
Catalog Database Log Files Consume a Large Amount of Disk Space	888
Unrecognized Vaults Appear in the Media Management Administrator	890
SAN Configuration Problems.....	890
Devices are Not Shared	891
Devices are Not Shared and the Tape Engine is Running.....	891
Shared Devices Appear as Unavailable or Offline	892
Shared IBM Devices Appear as Unavailable or Offline.....	893
Backup Jobs Fail	894
Cluster-based Backup and Restore Problems	895
Prevent Job Failures	895
Back Up MSCS Nodes on Remote Machines	896
Back Up Arcserve Backup Database in a Cluster Environment	897
Job Failure: Media Not Mounted	897
Miscellaneous Problems	898
Discovery Service Does Not Function Properly	898
Arcserve Backup Servers and Agent Servers Cannot Communicate with Each Other	898
SRM PKI Alert is Enabled by Default	900
Job Queue Log Files Consume a Large Amount of Disk Space	902
Unable to View all Job Logs for Rotation Backups after the Database is Pruned	904

Appendix B: Using Arcserve Backup in a Cluster-aware Environment **905**

Install Arcserve Backup into a Cluster-aware Environment	905
Cluster Overview	905
How Failover Works	908
Resource Group	909
Virtual Name and Virtual IP Address	909
Shared Disks	910
Mirrored Disks	911
Quorum Disks	911
Arcserve Backup HA Server to Support of Job Failover	912
Protecting Your Cluster with Arcserve Backup	913
MSCS Protection	915
NEC CLUSTERPRO/ExpressCluster Protection	928
NetApp NAS Cluster Protection	941

Appendix C: Using JIS2004 Unicode Characters with Arcserve Backup **947**

Introduction to JIS2004 Unicode Characters	947
Configuration Requirements for JIS2004 Unicode	947
Platforms Supporting JIS2004 Unicode Characters	948
Tasks You Can Perform Using JIS2004 Unicode Characters with Arcserve Backup	948
Arcserve Backup Applications Supporting JIS2004 Unicode Characters	949
Limitations of Using JIS2004 Unicode Characters with Arcserve Backup	950

Appendix D: Protecting Hyper-V Systems Using the Hyper-V VSS Writer **955**

Overview of Protecting Hyper-V VMs Using the Hyper-V VSS Writer	955
Prerequisite Components for Hyper-V VSS Writer Protection	955
Configure Arcserve Backup to Detect Hyper-V VMs	957
How Back Up Using Saved State Works	958
How Back Up Using Child Partition Snapshot Works	958
Back Up Hyper-V VMs Using the Hyper-V VSS Writer	958
Restore Data to Its Original Location	959

Glossary **961**

Index **965**

Chapter 1: Introducing Arcserve Backup

This section contains the following topics:

[Introduction](#) (see page 23)

[Arcserve Backup Functionality](#) (see page 23)

[Arcserve Backup Utilities](#) (see page 30)

[Arcserve Backup Command Line Utilities](#) (see page 35)

[Arcserve Backup Security](#) (see page 36)

[How Backup and Restore Operations Function on 64-bit Windows Platforms](#) (see page 38)

[How to Protect Virtual Machine Environments](#) (see page 39)

[Arcserve Backup Enterprise Module](#) (see page 40)

[Enterprise Level Password Management Utility](#) (see page 40)

[Using Arcserve Backup Documentation](#) (see page 41)

Introduction

Arcserve Backup is a comprehensive, distributed storage management solution for distributed and multiplatform environments. The application can back up and restore data from all the machines on your network, (including machines running Windows, UNIX, and Linux) using optional client agents. Arcserve Backup also provides media and device management utilities.

Arcserve Backup offers control from one management console. It can support small-scale and large-scale enterprise environments comprising of one machine or many, across different platforms and organizations.

Arcserve Backup Functionality

Arcserve Backup provides the components, functions, and utilities required by network managers to obtain and actively manage network backups.

Start the Arcserve Backup Manager by selecting the Manager icon from the program group. The My First Backup tutorial opens the first time you start the Manager. Subsequently, the Home Page appears, but you can still access the tutorial from the Help menu. From the Home Page, you can start and access any Arcserve Backup function using the following navigational features:

- **Home Page**--Provides news and support that links you to tools you can use to help solve problems with your computer. It also provides links to Quick Start, Configuration, Wizards, and Utilities.

- **Navigation Bar**--Quickly lets you independently access the Managers, Wizards, Utilities, most recently used screens, Arcserve Support, and Get Satisfaction. This area can be found on the left-most portion of the page. It is composed of three parts, Contents, Most Recently Used, and Support buttons. The contents of the Most Recently Used area are saved between sessions. The Navigation Bar displays with all Manager pages unless it is disabled. You can easily show or hide the Navigation Bar by selecting Navigation Bar from the View menu on the Arcserve Backup Home Page. To disable this feature, select View and deselect Navigation Bar. The Navigation Bar can be pinned. It can also be automatically hidden if it is unpinned.
- **Quick Start**--Provides quick links to Arcserve Backup Manager functions.
- **Configuration**--Provides access to Device Configuration, which lets you quickly configure the backup devices on your server and to SAN configuration.

From the Configuration menu you can also access Device Group Configuration, which lets you configure device groups and staging groups.
- **Wizards**--Simplifies the most common tasks of Arcserve Backup. You can access the Device, Create Boot Kit, Job Scheduler, and Diagnostic Wizards.
- **Utilities**--Offers several utilities that you can use to manage your database and media. The utilities are Merge, Scan, Compare, Count, Copy, Purge, User Profile, and Report Writer.

How to Access Arcserve Backup Managers, Wizards, and Utilities

Arcserve Backup managers, wizards, and utilities provide the front-end interfaces used to perform all functions necessary to protect your data. You can access these components from the Navigation Bar on the Manager Console.



The following is a list of the components, the menu from which you can access the component, and the functions they perform:

Quick Start Menu

- **Job Status Manager**--Monitors all pending, completed, and active jobs from the Job Status Manager window. You can schedule pending or completed jobs, submit new jobs, delete jobs, and stop active jobs. Log information is provided for each completed job.
- **Backup Manager**--Backs up data to media. You can schedule and configure backups of your computers and servers. Information about each backup job (such as the path and name of each file, and the media used) is logged in the Arcserve Backup database. Using the Backup Manager you can:
 - Specify the source (data that you want to back up) and the destination (media) for your backup job.
 - Define your backup job to back up data on computers running other operating systems such as UNIX, Linux, and Windows.
 - Use database agents running on your server.

- **Restore Manager**--Restores data that has already been backed up by Arcserve Backup. Using the Restore Manager you can:
 - Find all the versions of the files that were backed up.
 - Specify the source and feedback of the restore job.
 - Define a backup method and specify a backup schedule.
 - Perform a complete or partial restore of your data.
- **Server Admin**--Allows you to modify the Arcserve Backup system account and manage the core Arcserve Backup services: Job Engine, Tape Engine, and Database Engine. The Configuration icon allows you to configure tasks for these services including generating an alert and defining message logging. The Database Engine tab allows you to configure the database pruning job.
- **Dashboard**--Provides you with a snapshot overview of your backup infrastructure and your storage resource management (SRM) environment. For more information, see the *Dashboard User Guide*.
- **Infrastructure Visualization**--Provides you with a visual representation of your Arcserve Backup environment, allowing you to view backup status and explore how servers, nodes, and devices are related.

Monitor & Reports Menu

- **Job Status Manager**--Monitors all pending, completed, and active jobs from the Job Status Manager window. You can schedule pending or completed jobs, submit new jobs, delete jobs, and stop active jobs. Log information is provided for each completed job.
- **Report Manager**--Generates reports from data in the Arcserve Backup database. Various reports include information about backup schedules, media errors, backup devices, media pools, and media status and policies.
- **Report Writer**--Creates custom reports or generate predefined reports based on backup activity for a defined period.
- **Dashboard**--Provides you with a snapshot overview of your backup infrastructure and your storage resource management (SRM) environment. For more information, see the *Dashboard User Guide*.
- **Infrastructure Visualization**--Provides a visual representation of your Arcserve Backup environment. Infrastructure Visualization shows each Arcserve Backup Server in a hierarchical form resembling an organization chart. The mini-map feature acts as a scaled-down version of the current view, allowing you to zoom, pan, and highlight portions of the display.

Protection & Recovery Menu

- **Backup Manager**--Backs up data to media. You can schedule and configure backups of your computers and servers. Information about each backup job (such as the path and name of each file, and the media used) is logged in the Arcserve Backup database.

- **Archive Manager**--Lets you specify archive schedules and policies for the selected Windows, UNIX, Linux, and MAC file system data.
- **Restore Manager**--Restores data that was backed up by Arcserve Backup.
- **Arcserve Replication**--Arcserve Replication is a data protection solution that uses asynchronous real-time replication to provide disaster recovery capabilities. This link is active when you install Arcserve Replication.
- **Arcserve D2D**--Arcserve D2D provides a separate, light-weight solution for tracking changes on a local computer at the block level and then backing up only those changed blocks in an incremental fashion. As a result, Arcserve D2D lets you perform frequent backups (as frequently as every 15 minutes), reducing the size of each incremental backup and providing a more up-to-date backup. Arcserve D2D also provides the capability to restore files/folders and applications, and perform bare metal recovery from a single backup. This link is active when you install Arcserve D2D.

Administration Menu

- **Server Admin**--Allows you to modify the Arcserve Backup system account and manage the core Arcserve Backup services: Job Engine, Tape Engine, and Database Engine. The Configuration icon allows you to configure tasks for these services including generating an alert and defining message logging. The Database Engine tab allows you to configure the database pruning job.
- **Device Manager**--Displays information about your storage devices and media. It also allows you to change a drive's compression mode, and perform media functions such as compression, formatting, erasing, ejecting, and retensioning. Arcserve Backup supports a wide variety of media including 4mm, 8mm, DLT, QIC, Iomega's Zip or Jazz media, PDs, MO, and WORM formats.
- **Device Configuration**--A tool that lets you configure backup devices, such as tape libraries, RAID devices, virtual libraries, disk-based devices (for example, file system devices), and Deduplication devices (DDD). It also lets you enable or disable devices for Removable Storage Management (RSM), and register and unregister UNIX and Linux data mover servers with the primary server.
 - **Maximum number of devices supported:** 255 (includes physical devices, FSDs, and DDDs)
 - **Maximum number of FSDs and DDDs supported:** 255 (only if the number of physical devices configured is 0).
- **Device Wizard**--Displays the devices you have installed on a primary or stand-alone server, and lets you easily format, erase, compress, and eject your storage media.
- **Device Group Configuration**--A tool that lets you easily configure the device groups in your Arcserve Backup environment and select the groups that you will use for the staging of data.
 - **Maximum number of Device Groups supported:** 128

- **Media Pool Manager**--Manages, creates, and maintains logical groupings of media for easy identification of backups, to allow efficient scheduling of the maintenance and recycling of your media. You can design media rotation schemes to suit your particular archive needs.
- **MM Admin**--Provides the tools required to organize tape movement to off-site storage locations and protect, control, and manage media resources.
Note: To use the MM Admin, you must install the Enterprise Module.
- **Database Manager**--Displays information from the Arcserve Backup database, such as the jobs processed by Arcserve Backup, the media used by Arcserve Backup, and the devices you are using with Arcserve Backup.
- **Alert Manager**--Sends messages to people in your organization, using various methods of communication, regarding events that occur during the functioning of Arcserve Backup.
- **User Profile Manager**--Lets you assign roles privileges to Arcserve Backup user accounts.
- **Agent Deployment**--Lets you install and upgrade a collection of Arcserve Backup agents on multiple remote hosts simultaneously.
- **Central Agent Admin**--Lets you view agent logs and event logs, configure agent options and security information, specify debug levels for agent registry values, configure node tiers, and run Agent Deployment.

Utilities Menu

- **Job Scheduler Wizard**--Provides an easy way to package and submit jobs that you would typically submit from the Command Prompt window. In addition to the commands associated with Arcserve Backup, you can use this wizard for virtually any executable.
- **Create Boot Kit Wizard**--Creates and updates precautionary and machine-specific boot kits that allow you to recover your data if a disaster occurs.
Note: The Create Boot Kit wizard is available only if the Arcserve Backup Disaster Recovery Option is installed on your system. The Arcserve Backup Disaster Recovery Option is licensed separately.
- **Diagnostic Wizard**--Gathers and packages various Arcserve Backup system logs, that can be used for troubleshooting.
Note: The Diagnostic Wizard appears only if you install the Diagnostic Utility.
- **Merge Utility**--Lets you take media that contains one or more backup sessions and merge the information from the media into your Arcserve Backup database.
- **Media Assure & Scan Utility**--Lets you collect information about your media backup sessions and helps ensure that the sessions on the media are restorable.

- **Compare Utility**--Lets you compare the contents of a media session to files on a computer.
- **Count Utility**--Lets you count the number of files and directories on a computer.
- **Copy Utility**--Lets you copy files from one location to another.
- **Purge Utility**--Lets you delete files and directories from a computer.

Technical Support

- **Arcserve Backup on the Web**--Links you directly to the website where you find product information.
- **Understanding Your Support**--Links you to Arcserve Support where you can learn more about available support programs.
- **Registering for Support**--Links you directly to the registration form for Arcserve support.
- **Accessing Technical Support**--Links you to release-specific support where you can download software, obtain the latest documentation, or view the supported products matrix.

Note: The Support button on the Navigation Bar also links here.

Quick Reference

- **Readme**--Includes updates and supplements to the documentation and Help system.
- **Release Summary**--Includes new feature and product enhancement summaries.

Product News and Information

For up-to-the-minute product news and information, click the Twitter or Facebook links at the bottom of the home page and follow Arcserve Backup online.

RSS

At the bottom of the screen, the RSS bar scrolls through the latest Arcserve Backup news. Click a headline to link directly to Arcserve.com where you can view the complete story.

You must have Internet access to view RSS news. If you did not use the Arcserve Backup domain account to log on, you can provide credentials manually. Click the Refresh button at the right corner of the RSS bar to access the credentials dialog and refresh the news feed.

Note: After you install Arcserve Backup, the RSS bar communicates (by default) with Arcserve.com. If you do not want the RSS bar to communicate with Arcserve.com, you can stop the communication by closing the RSS bar. To close the RSS bar, click the View menu from any Arcserve Backup manager, and then click News Bar.

Arcserve Backup Utilities

Arcserve Backup offers several utilities that you can use to manage files. You can access the Utilities from the Navigation Bar on the Home Page. These utilities are described in the following sections.

Merge Utility

If you need to restore files to a Arcserve Backup server that you did not use to create the backup, or if you removed information from your Arcserve Backup database that you now need, you can use the Merge utility.

The Merge utility lets you take media that contains one or more backup sessions and merge the information from the media into your Arcserve Backup database. The database information from the media is appended to your existing database files.

Each time you run a backup job, Arcserve Backup records information in its databases about the computers, directories, and files that have been backed up, and the media that were used. This allows Arcserve Backup to locate files whenever you need to restore them. This database information is backed up whenever you back up your Arcserve Backup home directory.

If you have media that has a backup session that is not included in the Arcserve Backup database (for example, the backup was created using Arcserve Backup on a different backup server), you can use the Merge Media option to get the media's information into the database in the Arcserve Backup home directory.

Why would you need to use the Merge utility?

You can use the Merge utility when you need to restore files to a Arcserve Backup server that you did not use to create the backup. You can also use the Merge utility if you pruned (deleted) information from your Arcserve Backup database that you now need.

Merge Utility Options

The Merge utility allows you to merge information from media into the database.

Using the Merge utility you can merge:

- All sessions
- A single session
- A range of sessions, using one of the following types of ranges:
 - Specific start session to a specific end session.
 - Specific start session to the end of the media.

Merge Options:

If you select to merge all sessions, the tape containing Sequence number 1 must be present for this operation to complete successfully.

If the tape containing Sequence number 1 is not present, you will be prompted that the media could not be found and request that you continue (after inserting the proper tape) or cancel the operation.

If you want to merge a session from a different tape other than the one containing Sequence number 1, you can only do so by not selecting to merge all sessions, and instead specify the session number or range of session numbers to be included.

If you want to merge a session that spans more than one tape, you must have the tape on which the session header information is located.

Global Options for the Merge Utility

Arcserve Backup provides several types of global merge options. Use the Merge utility option when you want to restore detailed session information in your Arcserve Backup database.

For example, if a backup was created using Arcserve Backup on a different server, you can use Merge to get the media information into the database in the Arcserve Backup home directory. This will allow you to restore media backed up from another server at the file level. This can be useful if detailed information has been pruned from the database. By default, detailed job information is pruned 30 days after the backup to conserve database space. This can be configured in the Server Admin Manager.

Note: By default, all newly merged session details are preserved for one week (7 days) in the Arcserve Backup database, even if the newly merged session details are older than the prune retention time.

The available global merge options are listed below:

- **Backup Media**--Specify media options for the job such as the media timeout period.
- **Pre/Post**--Run commands or batch files before the job runs and/or after it finishes.
- **Job Log**--Enables you to determine the level of detail you want recorded into the Job Queue Log.
- **Database**--Specify whether you want to record detailed information about the jobs, or job and session-level details only.
- **Alert**--Send messages about events in your operation.

Merge Utility - Database Global Options

The Database tab on the Global Options dialog of the Merge Utility lets you specify the level of details you want to merge.

- **Merge Detail Information**--Lets you merge all details, including job and session information.
- **Merge Session Headers Only**--Lets you merge only header information, such as job and session data.

Media Assure & Scan Utility

The Media Assure & Scan Utility lets you collect information about your media backup sessions. Each source that you specify to back up is saved on media as an individual session. Using the Media Assure & Scan Utility you can scan the following types of sessions:

- Single sessions or an entire media.
- A range of sessions, for example:
 - Specific start session to a specific end session.
 - Specific start session to the end of the media.

The results of the Scan job display in the Job Queue. You would need to do this if you are trying to recover a Arcserve Backup server and you need to identify the most recent backup of the Arcserve Backup database so that you can restore it.

If you would like a Scan job to produce a detailed listing of your media contents, use the Log All activity feature on the scan options tab. You can also use the Media Assure & Scan utility if you want a list of the files that were backed up.

Arcserve Backup provides several types of advanced scan options:

- **Backup Media**--Specifies media options for the job.
- **Operation**--Specifies general options for the job such as to scan files with CRC verification or to enable database recording.
- **Pre/Post**--Runs commands or batch files before or after the job.
- **Job Log**--Determines the level of detail you want recorded in the Job Queue Log.
- **Alert**--Sends messages about events in your operation.

Compare Utility

Compare the contents of a media session to files on a machine. Results of the Compare job can be seen in the Job Queue. You could use this option after a backup to verify that the backup copied all of the files to media without error.

Arcserve Backup provides several types of advanced compare options:

- **Backup Media**--Specify media options for the job.
- **Operation**--Specify whether to enable database recording.
- **Pre/Post**--Run commands or batch files before or after the job.
- **Job Log**--Determine the level of detail you want recorded in the Job Queue Log.
- **Alert**--Send messages about events in your operation.

Count Utility

The Count utility counts the number of files and directories on a machine. Results of the Count job can be seen in the Job Queue. You could use this option after a Copy job to verify that the Copy function copied all of the files from one disk to another without error.

Arcserve Backup provides several types of advanced count options:

- **Operation**--Specify whether to enable database recording.
- **Pre/Post**--Run commands or batch files before and after the job.
- **Job Log**--Determine the level of detail you want recorded in the Job Queue Log.
- **Virus**--Scan files for viruses before they are counted.
- **Alert**--Send messages about events in your operation.

Copy Utility

The Copy utility allows you to copy files from one location to another. For example, you can run a copy job on your local machine to store files and directories on another machine that is going to be backed up to media.

Copy options determine related actions that occur during or after the copy operation:

- **Retry**--Specify when to retry open files and file sharing options.
- **Operation**--Specify operation options and whether or not to enable database recording.
- **Destination**--Specify options for the Directory Structure, File Conflict Resolutions, and VMS File Version.
- **Pre/Post**--Run commands or batch files before or after the job.
- **Job Log**--Determine the detail you want recorded in the Job Queue Log.
- **Virus**--Scans files for viruses before they are copied.
- **Alert**--Send messages about events in your operation.

Purge Utility

The Purge utility allows you to delete files and directories from a machine. Results can be seen in the Job Queue.

Arcserve Backup provides several types of advanced purge options:

- **Operation**--Specify some general options for the job such as to remove directories or enable database recording.
- **Pre/Post**--Run commands or batch files before or after the job.
- **Job Log**--Determine the level of detail to record in the Job Queue Log.
- **Alert**--Send messages about events in your operation.

Report Writer Utility

Create custom reports or generate predefined reports based on backup activity for a defined period. You can specify a query or filter report data. Generated reports can be previewed on screen, printed, and saved in either .csv or .xml format.

How You Can Manage Jobs Using the cabatch Command

The cabatch Utility is a job management tool that allows you perform the following tasks:

- Submit and delete jobs in local or remote Arcserve Backup job queues from the command line.
- Modify the execution times of jobs in the job queues.
- Use job scripts created in the Arcserve Backup Manager or in a text file created using the cabatch Job Information Template in the Arcserve Backup home directory.

For more information on the cabatch utility, see the *Command Line Reference Guide*.

User Profile Utility

The User Profile Utility lets the administrator manage user profiles and provide access to Arcserve Backup.

When you install Arcserve Backup, the caroot user profile is set up, by default, with the administrator group assigned to it. The Administrator group provides control over all Arcserve Backup functions operating within a given Arcserve Backup domain.

Using the User Profile utility Arcserve Backup server administrators can:

- Add a user.
- Delete a user.
- Change a user's password.
- Assign a user to a group.

Arcserve Backup Command Line Utilities

Arcserve Backup offers command line utilities that enable direct control over almost all operations that can be performed by a Arcserve Backup server via the DOS prompt. It provides an alternative method of accessing almost all of the operations available from the Arcserve Backup Manager.

A full description and use of these command line utilities can be found in the *Command Line Reference Guide*.

Arcserve Backup Security

The following sections describe Arcserve Backup security functionality.

This section contains the following topics:

[Arcserve Backup Administrator Profile](#) (see page 36)

[How Arcserve Backup Equivalence Works](#) (see page 36)

[System Account](#) (see page 37)

[Equivalency and the System Account](#) (see page 37)

Arcserve Backup Administrator Profile

Arcserve Backup contains a root-level, super user profile that provides complete control of Arcserve Backup.

This profile, referred to as “caroot,” is set up when you install Arcserve Backup for the first time. You can set the password for caroot during installation or later using the `ca_auth` and `authsetup` Command Line utilities. For information about these utilities, see the *Command Line Reference Guide*.

The caroot password can consist of any combination of alphanumeric and special characters, but may not exceed 15 bytes. A password totaling 15 bytes equates to approximately 7 to 15 characters.

Important! You should not leave the password blank for caroot.

The caroot user profile controls access to only the Arcserve Backup Manager Console and backup-related functions, such as backup, restore, and so on.

How Arcserve Backup Equivalence Works

The User Profile Manager lets you create equivalence to the caroot account for any Windows account. This functionality lets you allow users with Windows accounts log in to Arcserve Backup and access the Arcserve Backup managers and utilities. However, to allow users with Windows accounts to execute Arcserve Backup command line utilities (for example, `ca_backup` and `ca_restore`), you must create caroot equivalence for the Windows accounts using the `ca_auth` command line utility.

Note: Users who have been granted caroot equivalence using the `ca_auth` command line utility can run all the command line utilities but cannot log in to the Arcserve Backup Manager Console.

Creating equivalence has the following advantages:

- **Ease of command line usage**--As you create equivalence for a Windows user, the equivalence performs an implicit login on behalf of the logged-in user whenever a command line function requires authentication. This behavior lets the command line utilities run without requiring the user to enter a user name and password each time a command is submitted.
- **Access Restriction**--Although Windows user accounts with caroot equivalence can run all Arcserve Backup command line utilities, Windows user accounts cannot log in to the Arcserve Backup Manager Console and domain. However, you can use the User Profile Manager to allow Windows users to log in to the Arcserve Backup Manager Console and domain with their Windows user account login information.

Note: In addition to the User Profile Manager, you can add Arcserve Backup user accounts using the `ca_auth` command line utility. For more information about the `ca_auth` utility, see the *Command Line Reference Guide*.

System Account

The Arcserve Backup services require a valid Windows system account that has Administrator and Backup Operator privileges on the local machine. The services use this account to access local resources, such as the hard drive and the local network.

You are given the option of entering a Windows system account when you first install Arcserve Backup. If you enter a Windows account during installation, Arcserve Backup automatically grants this account Administrator and Backup Operator privileges. If you select Skip during installation, you must enter a valid Windows system account using the Arcserve Backup Administrator and grant it the required privileges manually.

Note: A user in the Backup Operator Group does not have rights to access the Arcserve Backup database. As a result, member servers are not visible to the user in the Backup Manager.

You can change the system account information at any time using the Arcserve Backup Server Admin or the Server Configuration Wizard.

Equivalency and the System Account

Do not confuse the caroot user profile with the Arcserve Backup System Account. The caroot user profile is used to control access to the Arcserve Backup Manager and its related backup functions; the system account provides the security privileges needed by Arcserve Backup services to operate on the local machine.

Although the System Account and the caroot user profile perform different functions, in order for Arcserve Backup to run all of its jobs successfully, you must grant the System Account equivalency to caroot. For example, if the System Account is named BackupAdmin, and the local machine name is BAB01, use the following `ca_auth` command to give the account equivalency to caroot:

- **Local system accounts:**

```
ca_auth -equiv add BAB01\BackupAdmin BAB01 caroot caroot caroot_password
```

- **Domain system accounts:**

```
ca_auth -equiv add DomainName\BackupAdmin BAB01 caroot caroot caroot_password
```

Note: The caroot password can consist of any combination of alphanumeric and special characters, but may not exceed 15 bytes. A password totaling 15 bytes equates to approximately 7 to 15 characters.

For more information on security, see "Administering the Backup Server" or the *Command Line Reference Guide*.

How Backup and Restore Operations Function on 64-bit Windows Platforms

Due to architectural differences between 64-bit and 32-bit Windows platforms, various elements of 64-bit operating systems cannot be accessed by 32-bit applications. These elements include areas of the Windows System Registry, system settings files included in a System State backup, and Volume Shadow Copy Service writers.

To overcome these limitations, and to successfully perform backup and restore operations when the Arcserve Backup server is running a 64-bit version of Windows, you must install the 64-bit version of the Arcserve Backup Client Agent on the Arcserve Backup server.

This configuration lets the 64-bit Client Agent run as a native process on the local Arcserve Backup server, which manifests the capability to perform browse, backup and restore operations on the local file system, System State, System Registry, and Volume Shadow Copy Service writers in the same manner as remote browse, backup, and restore operations using the 32-bit Client Agent for Windows.

For more information about Arcserve Backup Agents and Options supported by 64-bit Windows platforms, see the readme file.

How to Protect Virtual Machine Environments

Use the following methods to protect virtual machine environments using Arcserve Backup:

- **Arcserve Backup Agent for Virtual Machines**--The Agent for Virtual Machines lets you protect environments that rely on virtual machines (VMs) residing in Windows Server Hyper-V systems, VMware ESX/ESXi Host systems, and VMware vCenter Server systems to protect data.

For VMware-based systems, VMware provides a mechanism called Virtual Disk Development Kit (VDDK) that lets you protect the files and data stored in the VMs in VMware ESX Host systems and VMware vCenter Server systems. To integrate Arcserve Backup with VMware VDDK and Windows Server Hyper-V systems, you install and license the Agent for Virtual Machines.

Note: For information about system requirements and supported platforms, see the readme file. For information about installing and configuring the agent, see the *Agent for Virtual Machines Guide*.

- **Scripted Solution for VMware ESX/ESXi Host Systems**--The best method of protecting your VMs and VMware ESX Host systems is to install the Agent for Virtual Machines. However, the scripted solution lets you integrate Arcserve Backup with VMware ESX/ESXi without installing the Agent for Virtual Machines. The scripted solution helps to ensure that your VMs and VMware ESX Host systems are protected as securely as any other server in your environment.

Note: To use the scripted solution, you must install and license the Arcserve Backup Client Agent for Windows.

For information about how to use a scripted solution to protect VMware ESX/ESXi Host systems, see the *Best Practices Guide for VMware ESX Server Backup* on the Arcserve Technical Support website. The best practices guide describes common methods that you can use for data backups on VMs, and any considerations relating to the different methods.

- **Install Arcserve Backup Agents on the VM**--To back up and restore data that resides in your VMs, you can install the Arcserve Backup agents that correspond with the guest operating systems and the applications that are running in your VMs.

Arcserve Backup Enterprise Module

The Arcserve Backup Enterprise Module is a separately-installed component that lets you deploy a number of enhanced features, including the following:

- [Multistreaming](#) (see page 103).
- Disk staging backups and tape staging backups with multistreaming and transmit two or more (up to 32) streams of backup data.

Note: If you do not license the Enterprise Module, Arcserve Backup lets you transmit two streams of backup data for disk staging and tape staging backup jobs. For more information, see [How Backup to Disk to Tape Works](#), and [How Backup to Tape to Tape Works](#) (see page 230).

- [Media Management Option](#) (see page 457).
- [Raw disk backup and restore](#) (see page 759).
- Arcserve Backup for Windows Enterprise Option for VSS Hardware Snap-Shot

Note: For more information, see the *Microsoft Volume Shadow Copy Service Guide*.

In addition, the Arcserve Backup Enterprise Module is a prerequisite component for the following Arcserve Backup options:

- Enterprise Option for IBM 3494
- Enterprise Option for StorageTek ACSLS
- Image Option

Note: For more information about the above options, see the *Enterprise Module Guide*.

Enterprise Level Password Management Utility

When the user password changes, every job in the job queue must be modified to reflect the change. Using the `ca_jobsecmgr` utility you can make global user password changes for all the jobs in the Job Queue for the local Arcserve Backup Server (default).

- Syntax
`ca_jobsecmgr [server arguments] <current security> <new security>`
- Options

For a complete description of the options for this command, see the *Command Line Reference Guide*.

Using Arcserve Backup Documentation

This release of Arcserve Backup contains an online bookshelf, which consists of help systems and PDF files that you access via the Internet. With web-based online help, you can access the latest version of Arcserve Backup documentation without having to install Arcserve Backup product updates. This online version of the bookshelf is the default.

However, if you wish to download and install the bookshelf, you can view the online help and PDF files from a local destination by configuring the documentation to use the help from a local source on the Documentation Configuration dialog.

Download the Bookshelf

By default, the documentation bookshelf is accessed online from Arcserve Support. However, you can download and install the latest version of the Arcserve Backup documentation bookshelf to a local destination.

To install the documentation bookshelf to a local destination:

1. Open the Arcserve Support online bookshelf by clicking Start, Programs, CA, ARCserve Backup, Documentation.

The Arcserve Support documentation bookshelf opens.

2. Click the Download this Bookshelf link to download the bookshelf to your Arcserve Backup server.

Note: You can also click the Download Help link on the Documentation Configuration dialog.

The File Download dialog appears for the file.

Example: Arcserve Backup Backup r16 5-ENU.zip (English version).

3. Click Save.

The Save As dialog appears.

4. Specify or verify the following fields and click Save.

- **Save in**--enter the folder where you want to save the zip file.
- **File name**--verify the file name being saved is correct.
Example: Arcserve Backup Backup r16 5-ENU.zip
- **Save as type**--verify the type is WinZip File or CA Zip.

The Download complete dialog appears and the zip file is saved.

5. Click Open, Open Folder, or Close.

The zip file opens, the folder containing the zip file opens, or the dialog closes.

6. Extract the contents of the zip file to the location specified above where the zip file was saved.

Note: If the extraction utility contains an overwrite existing files option, you should enable the option. However, if the extraction utility prompts you to overwrite the files in the destination directory, click Yes to All.

The documentation bookshelf is downloaded, installed in the local directory, and updated to the latest version.

7. If you wish to view the documentation bookshelf locally, then you must configure the documentation to use the local help as the source on the Documentation Configuration dialog. For more information, see [Configure the Documentation Location](#) (see page 42).

Configure the Documentation Location

The documentation can be configured so that you can view the online help and PDF's from either online or from a local machine. By default, the documentation is accessed online.

To configure the documentation location:

1. Click Start, Programs, CA, Arcserve Backup, Documentation Configuration.

The Documentation Configuration dialog opens.

2. Select one of the following options:

- **Use online Help as source**--select this option to access the documentation from Arcserve Support and receive the latest documentation updates.
- **Use local Help as source**--select this option if you want to use a local version of the downloaded documentation or if you do not have an internet connection.

Important: If you select this option, you must first download the bookshelf to a local directory. Then you must enter a valid help location or click Browse to locate a valid location in the Help Location field. This location must contain the file Bookshelf.html.

3. Click Open Help to test the documentation help location.

If you selected the Use online Help as source option, then the bookshelf will open on Arcserve Support.

If you selected the Use local Help as source option, then the bookshelf will open from the local help location specified above.

4. Click OK.

The documentation is configured.

Chapter 2: Protecting Data Using Arcserve Backup

This section contains the following topics:

[Arcserve Backup Components](#) (see page 44)

[Central Management](#) (see page 45)

[How Password Management Works](#) (see page 79)

[How User Profile Management Works](#) (see page 81)

[How Arcserve Backup Processes Backup Data Using Multistreaming](#) (see page 103)

[How Arcserve Backup Processes Backup Data Using Multiplexing](#) (see page 106)

[How Arcserve Backup Secures Data](#) (see page 109)

[Effective Media Management](#) (see page 115)

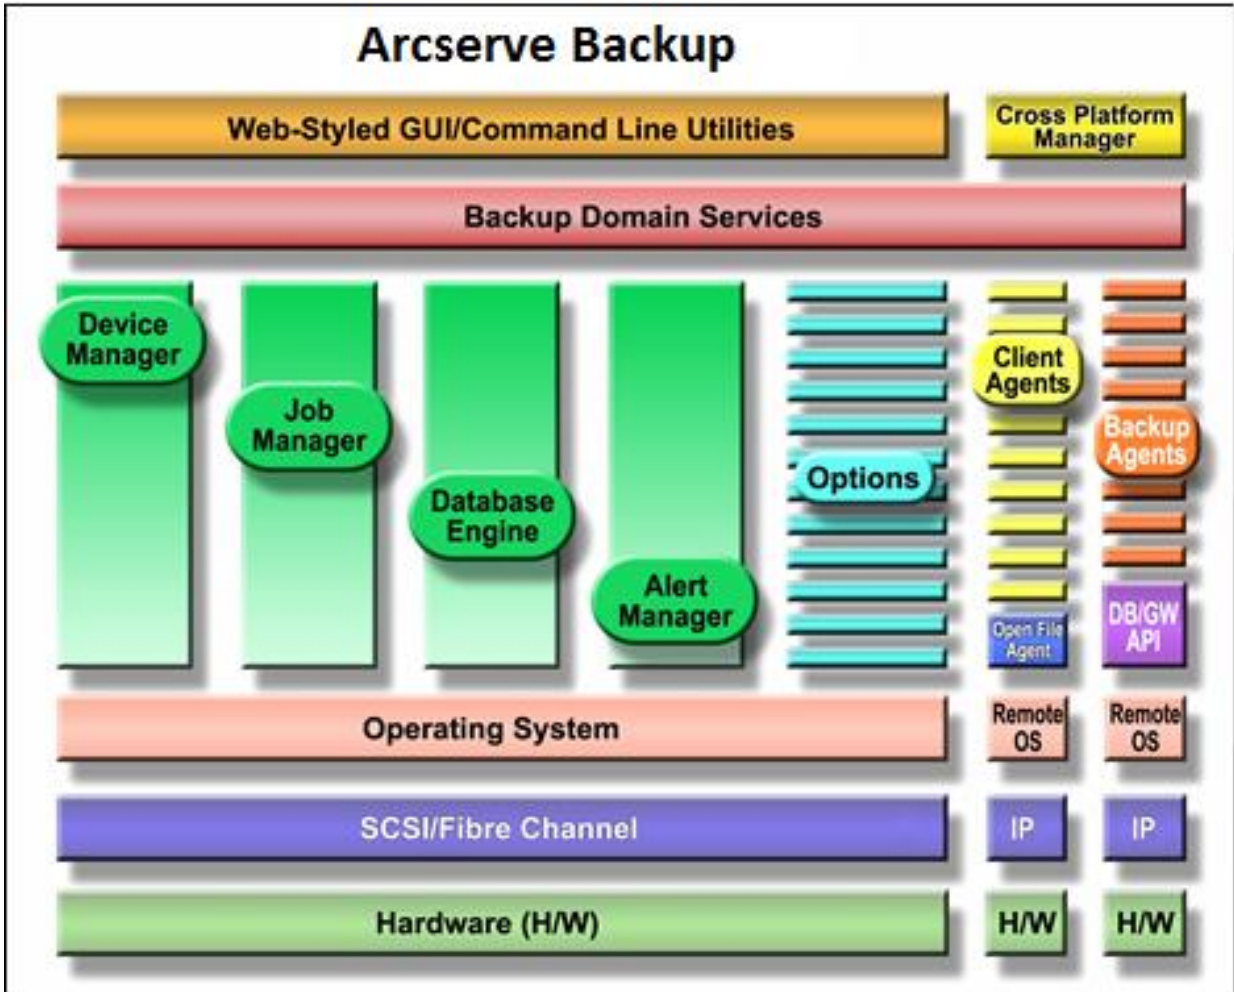
[Back Up and Restore Data](#) (see page 118)

[Start Arcserve D2D](#) (see page 128)

[Start Arcserve Replication](#) (see page 130)


Arcserve Backup Components

Arcserve Backup has a flexible design that allows you to manage and protect your environment. It provides powerful components that work together to accomplish critical administrative tasks seamlessly.



Central Management

The Central Management Option allows you to manage one or more Arcserve servers through a single central system. Within an Arcserve domain, this central system is called the primary server and the other (subordinate) servers are called member servers.



Primary Server

A primary server provides you with a single point to manage the primary server and one or multiple member servers in an Arcserve domain. From the primary server you can centrally manage and monitor jobs that run locally on that primary server and jobs that run remotely on one or more of the member servers in the domain. There can be only one primary server within an Arcserve domain.

Note: You can designate any Arcserve Backup server as the primary server. However, because the primary server is responsible for managing and initializing the shared member servers, you should use your most reliable server as the primary server.

Member Server

A member server executes jobs that are dispatched from the primary server. Within an Arcserve domain, member servers can only belong to one primary server.

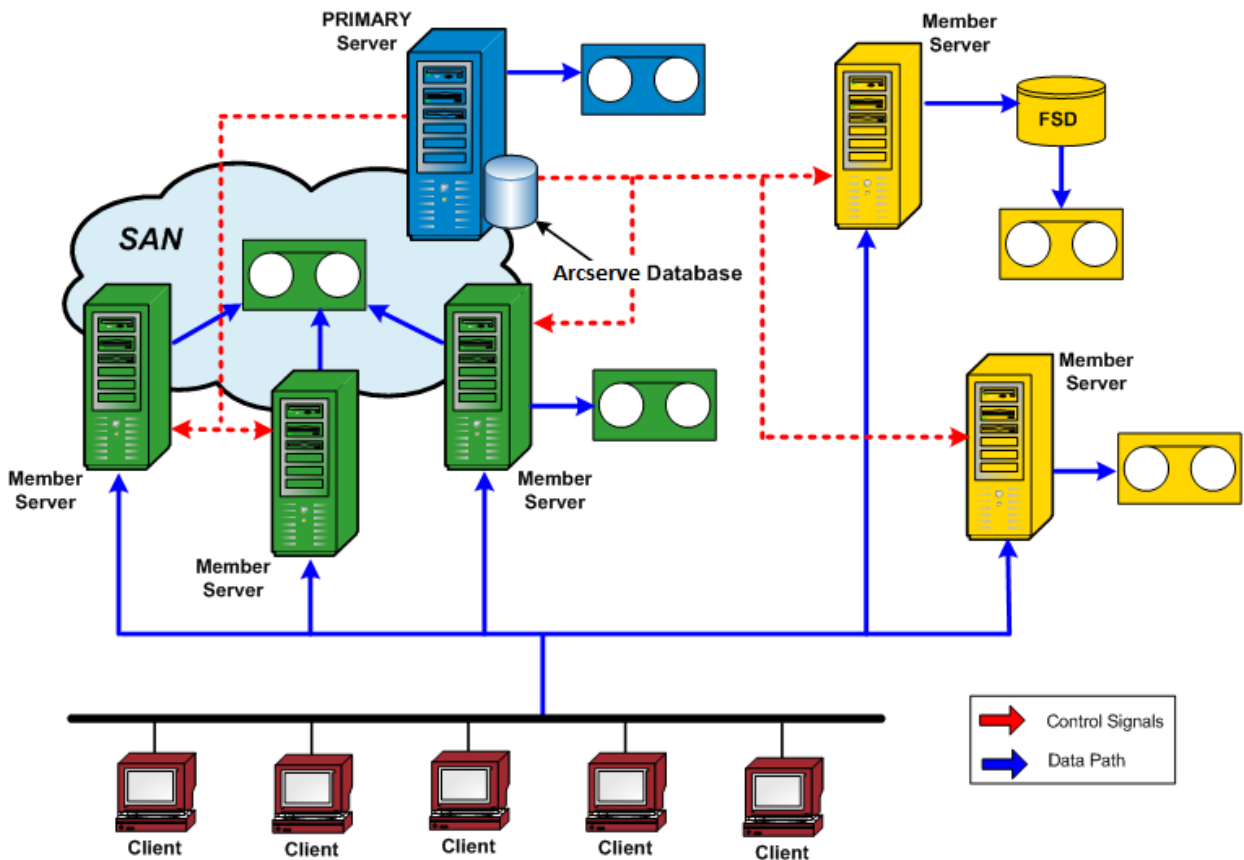
Arcserve Domain

An Arcserve domain is a logical grouping of a primary and one or more member servers that allows easier monitoring and managing of Arcserve Backup servers and users. Within an Arcserve domain, there can only be one primary server and there can be multiple member servers that are controlled by the primary server. An Arcserve domain allows you to manage the domain and select any server from within the domain to perform Arcserve Backup tasks without being required to log in to each server separately.

The Arcserve database (ASDB) can be installed on a primary server or on any remote system in your environment. Be aware that to install the ASDB on a remote system, you must host the ASDB instance using Microsoft SQL Server.

The primary and member servers may or may not be connected through a Storage Area Network (SAN). If the member servers are located on a SAN, the primary server must also be on the SAN.

Note: A SAN environment within an Arcserve domain is an environment where multiple Arcserve servers can share one or more devices (for example, tape libraries).



Central Job Management

Central job management allows you to create, manage, and monitor Arcserve Backup jobs from one central location. Jobs are always submitted on the primary server and can be run either locally on the primary server itself or remotely on any of the associated member servers. With central job management, you can perform job management operations (for example, backup, restore, merge, scan, data migration, tape copy, compare, copy, count, and so on) on all Arcserve servers from the primary server.

All jobs that are scheduled to run on any Arcserve Backup server in the domain will be submitted to the central job queue. This allows you to monitor the job status of all jobs in the domain from the primary server.

To view jobs running from the primary server, select the primary server. To view jobs running from a member server or data mover server, select the member server or a data mover server.

The screenshot displays the Arcserve Backup Central Management console. The left pane shows a tree view of the domain structure:

- Arcserve Backup Domains
 - CURRENT DOMAIN
 - PRIMARY SERVER
 - COMP-001
 - COMP-003
 - COMP-005
 - COMP-007

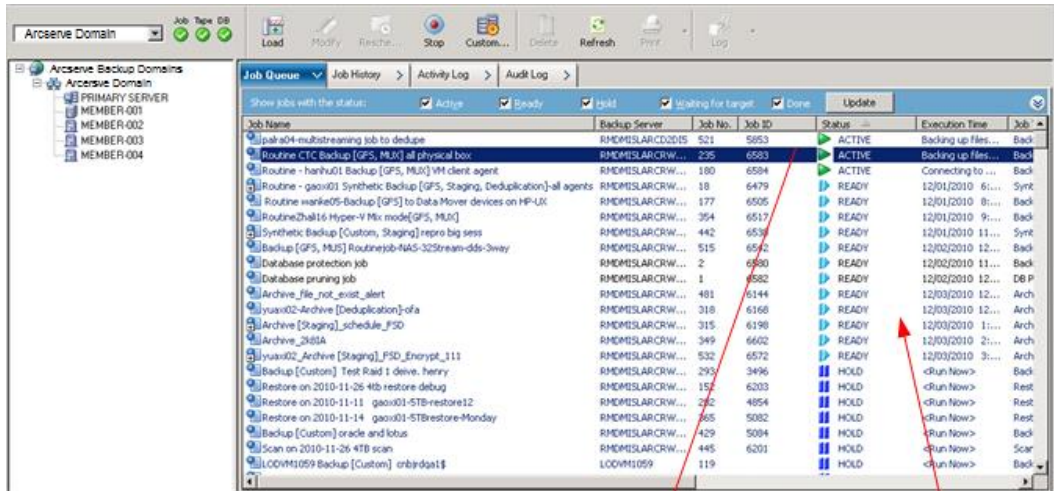
The main pane shows the Job History view for the selected server. The table below represents the data shown in the screenshot:

Job Name	Last Result	MB	Files	Missed	MB/Min.	Time U...	Job ID	Job No
172.0.0.1 (172.0.0.1)								
(2 job execution: 2 Finished, 0 incomplete, 0 failed, 0 canceled, 0 not attempted)								
172.0.0.2 (172.0.0.2)								
(2 job execution: 2 Finished, 0 incomplete, 0 failed, 0 canceled, 0 not attempted)								
Backup [GFS, MUS...]	Finished	3	28	0	.31	00:09:36	6542	5
Backup [GFS, MUS...]	Finished	64,000	81,300	0	216.43	04:58:30	6455	5
172.0.0.3 (172.0.0.3)								
(13 job execution: 6 Finished, 0 incomplete, 7 failed, 0 canceled, 0 not attempted)								
LO00051 Backup [...]	Failed	N/A	N/A	N/A	N/A	00:21:26	1955	1
LO00051 Backup [...]	Failed	N/A	N/A	N/A	N/A	00:21:16	1890	1
LO00051 Backup [...]	Failed	N/A	N/A	N/A	N/A	00:21:04	1808	1
LO00051 Backup [...]	Failed	N/A	N/A	N/A	N/A	00:22:46	1734	1
LO00051 Backup [...]	Failed	N/A	N/A	N/A	N/A	00:21:34	1694	1
LO00051 Backup [...]	Failed	N/A	N/A	N/A	N/A	00:23:12	1629	1
LO00051 Backup [...]	Failed	N/A	N/A	N/A	N/A	00:21:42	1500	1
LO00051 Backup [...]	Finished	0	4	0	.00	00:01:18	1393	1
LO00051 Backup [...]	Finished	0	4	0	.00	00:00:50	1184	1
LO00051 Backup [...]	Finished	0	4	0	.00	00:00:58	968	1
LO00051 Backup [...]	Finished	0	4	0	.00	00:00:46	801	1
LO00051 Backup [...]	Finished	0	6	0	.00	00:00:40	609	1

Central Job Monitoring


Central job monitoring allows you to monitor the progress of all jobs running on any Arcserve server in a domain from the primary server. From the primary server job queue, you can view the real-time status of active jobs within the domain.

Note: Job monitoring is only available for active (running) jobs within the domain. When the job completes, the final status of any job that ran in the domain is displayed in the Job Status Manager.



Job monitoring available for only active jobs.

Job queue for all jobs in the domain.



Central Database Management


Information from all Arcserve Backup servers within a domain is stored in a single central database that can be managed by the primary server. The central database is configured from the primary server and the associated member servers write relevant information into the central database.

Whenever Arcserve Backup performs a backup, all the job, session, and media information from the Arcserve Backup servers is stored in the centralized database. In addition to the database, a central catalog file is also created that contains descriptive information about each session and allows you to select the specific files and directories to be restored without having to query the database itself. The catalog files have been restructured so that they no longer need to be merged into the database to be efficiently searched. When data needs to be restored, Arcserve Backup can quickly browse the content of each session in the catalog file from a single central location to locate the information.

Central Logging

With central logging, Activity Logs and Job Logs for all Arcserve Backup servers in a domain (primary and members) are stored in a central database, allowing you to view the logs from one central location.

Central logging also helps you to perform troubleshooting. You can use the various filters (such as Keywords, Job ID, Job status, Message type, and so on) to isolate the log information to display everything that happened for a specific condition. For example, you can specify to only display the logs for failed jobs, or only display logs that contain a certain keyword in a message or job name, or only display logs for certain job names. Central logging allows you to perform these functions for all Arcserve Backup servers within a domain from one central location.



Central Reporting

With central reporting, you can launch and create scheduled reports for all Arcserve Backup servers in a domain from the primary server. Different reports are generated based on the backup activity stored in the Arcserve Backup database. Central reporting provides the capability to preview a report, print a report, send email, and schedule when to generate a report for all domain servers from the primary server.

For example, from the primary server you can create a report that identifies the agents that failed the most consecutive times, or the agents with the most failed backup attempts, or the agents with the most partial backups. You can find the percentage of successful, incomplete, or failed backup attempts. You can also find the number of errors and warnings generated for the backup job for each agent which helps in determining the agents with most number of errors.

Central Alert Management

With central alerting, alerts are posted from all Arcserve Backup servers in a domain to the primary server. Job level alerts are configured on the primary server and applied to all jobs that are executed on the primary server or any of the associated member servers within the domain.

Central Alerting

Arcserve Backup integrates with Alert Service to deliver alerts via emails, broadcasts, event, pagers and printers. This is accomplished by the central alerting feature. Central Alerting is a mechanism for gathering all alerts from all servers in the domain.

When a job executes, on any server in the domain, the event is recorded in the Arcserve database. The job level alert is forwarded to the Alert Service on the Primary server. The Alert Service then generates a system wide alert according to previously established configurations. For more information on establishing service account configuration, see [Alert Manager Configuration](#) (see page 709).

Central Arcserve Server Administration

Server administration tasks for all Arcserve Backup servers in a domain are performed centrally from the primary server. From the primary server, you can monitor the state of the Arcserve Backup engines (Job Engine, Tape Engine, and Database Engine) for all Arcserve Backup servers in the domain. In addition, you can select an individual server to monitor and manage the state of the engines and services on that server.

DOMAIN: 100-LL-PRIMARY

PRIMARY Server: 100-LL-PRIMARY

MEMBER Servers: 100-3FL-DELL054, 100-LL-BABRW1, 100-LL-BABRW2, 100-LL-BABRW4, 100-LL-BABRW6, 100-LL-BABRW7, 100-LL-BABRW8

Name	Job Engine	Tape Engine	DB Engine
100-3FL-DELL054	Started	Started	Started
100-LL-BABRW1	Started	Started	Started
100-LL-BABRW2	Started	Started	Started
100-LL-BABRW4	Started	Started	Started
100-LL-BABRW6	Started	Started	Started
100-LL-BABRW7	Started	Started	Started
100-LL-BABRW8	Started	Started	Started
100-LL-PRIMARY	Started	Started	Started

Specified server: 100-LL-NPRIMARY (100-LL-PRIMARY)

Name	Status	Up Time (days:hours:minutes)	Description
CA ARCserve Communication Foundation	Started	7 : 20 : 19	Provides data used by CA ARCserve Backup Da...
CA ARCserve Database Engine (ODBC)	Started	7 : 20 : 19	Provides database services for ARCserve Backu...
CA ARCserve Discovery Service	Started	7 : 20 : 19	Enables the discovery of all ARCserve Backup pr...
CA ARCserve Domain Server	Started	7 : 20 : 18	Provides the management of domains and auth...
CA ARCserve Job Engine	Started	7 : 20 : 18	Maintains and executes jobs from the ARCserve...
CA ARCserve Management Service	Started	7 : 20 : 18	Provides remote services for command line utilit...
CA ARCserve Message Engine	Started	7 : 20 : 18	Allows remote management of other ARCserve ...
CA ARCserve Service Controller	Started	7 : 20 : 18	Enables remote start/stop of ARCserve Backup ...
CA ARCserve Tape Engine	Started	7 : 20 : 18	Manages the configuration and operation of bac...

Central Device Management

With central device management, you can manage devices for all Arcserve Backup servers in a domain from the primary server by using the Device Manager. The Device Manager provides information about storage devices that are connected to a server, the media in these devices, and the status of these devices. In addition, the Device Manager also allows you to format, erase, inventory, import, and export media. With central device management, all of these functions can be performed from the primary server for devices connected to the primary server or any of the associated member servers.

Tape Library Auto-Configuration

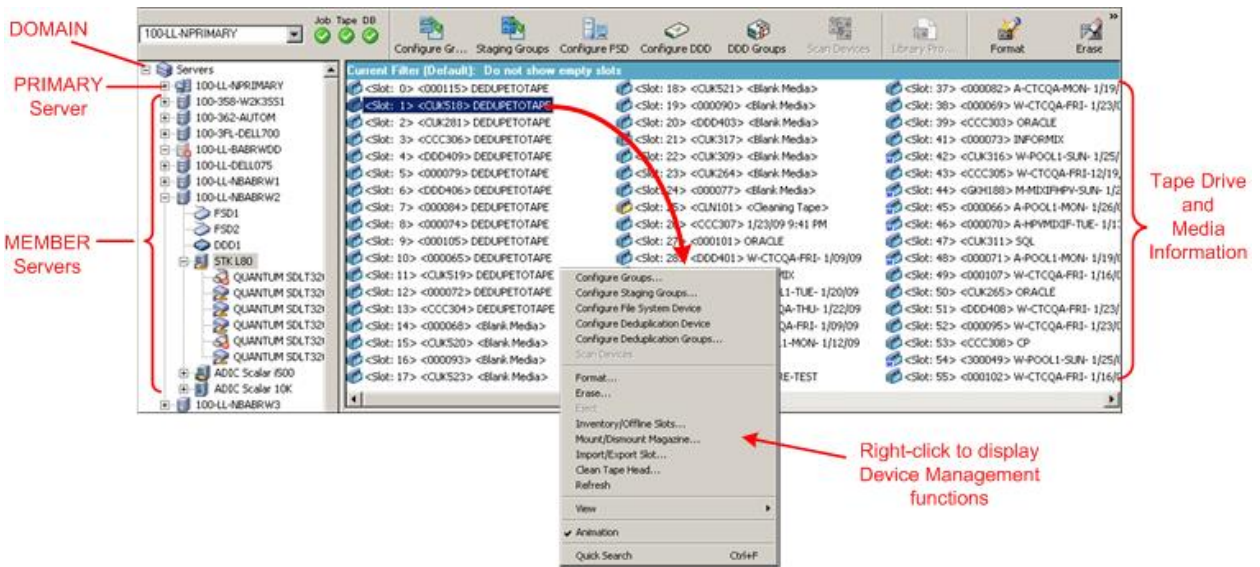
Arcserve Backup will now automatically detect the existence of a tape library and configure it. Therefore you no longer need to run the separate Tape Library Option Setup utility and no longer need to reconfigure a library after replacing bad drives or adding new drives. In addition, library settings can be changed on the fly without stopping the Tape Engine for such tasks as cleaning tapes or specifying cleaning settings.

SAN Auto-Configuration

SAN configuration is now tied to Arcserve Backup domain configuration, eliminating the need to run SAN configuration. Libraries are automatically detected as “shared” on the fly at the Arcserve Backup domain primary server. Domain primary servers can have both SAN and non-SAN domain member servers.

FSD Auto-Configuration

From a central location on the primary server you can create an FSD on any member server without having to stop and start the tape engine.




Central License Management

Arcserve Backup licensing is count-based with licenses for most Arcserve servers within a domain applied centrally on the primary server. Count-based licensing grants a single overall license to the application with a predetermined number of active license rights included in the overall license pool.

Each new user of the application (member server) is granted an active license from the pool on a first-come, first-served basis until the total number of available licenses has been exhausted. If all the active licenses have already been applied and you need to add a license to a different member server, you would first have to manually remove the license from one of the member servers (to reduce the count) and then have the new member server apply for that license (to take up the count).

With central license management, the license allocation is server based. This means that when a license is allocated to a server, central license management will record this allocation and keep this license exclusively used for that server. Future license requests from the same server will always succeed, and requests from other servers will cause a new license to be allocated to the new server. When all available licenses are allocated, license checking places jobs that are running from an Arcserve Member server into a Hold status, and fails jobs associated with a server that is running an Arcserve agent. For all scenarios, when there are no licenses available, you will get an activity log message warning you that the license is a problem.

Through the use of central licensing, you can easily remove license rights to allow other member servers to gain license privileges. From the Server Admin Manager screen on the primary server, you can access the License Management dialog to view the active license counts for each component and also manage which licenses are applied to which servers.



Arcserve Backup licenses are installed on and checked centrally on the Arcserve Backup primary server. However, the following agents must be licensed on the servers where you are installing the agents:

- Arcserve Backup for Windows Agent for Sybase
- Arcserve Backup for Windows Agent for Informix
- Arcserve Backup for Windows Enterprise Option for SAP R/3 for Oracle

More information:

[Manage Arcserve Backup Component Licenses](#) (see page 523)
[Release Licenses from Servers](#) (see page 525)

Central Job History

With central job history, you can view the history of backup jobs on all Arcserve Backup servers within a domain from the primary server. You can view the history based upon the applicable host, source group, or the job itself.

Through central job history, you can locate and review the status of the Arcserve Backup servers that were backed up, the instances (or jobs) for each server, and the volumes (or sessions) for each instance.

You can also view information about the device and the media that were used for the backup job. In addition, central job history is helpful in troubleshooting because any errors or warnings that were generated during each job on any server (primary or member) are also displayed from one central location.

Note: On the Job History tab, the MB/Minute field displays the ratio of megabytes per minute for the entire job. In addition to transferring data from the source location to the destination storage area, a job can include media management activities, pre- and post- scripts, and so on. As a result, the value displayed in the MB/Minute field can be different than the actual throughput. To view the actual throughput for the job, click the Activity Log tab, locate the job, expand Logs for the Master Job, and locate the log entry for Average Throughput.

The screenshot shows the Arcserve Backup software interface. On the left, a tree view displays the domain structure under '100-LL-NPRIMARY'. The main area shows the 'Job History' tab with a table of job entries. Red arrows from external labels point to specific parts of the interface:

- DOMAIN:** Points to the tree view on the left.
- PRIMARY Server:** Points to the '100-LL-NPRIMARY' entry in the tree.
- MEMBER Servers:** Points to a group of servers in the tree, including '100-358-W2K3SS1', '100-362-AUTOM', '100-3PL-DELL700', '100-LL-BABRW00', and '100-LL-DELL075'.
- Server:** Points to a server entry in the job history table, such as '100-362-AUTOM'.
- Instance (Job):** Points to a specific job entry, such as '2009-01-29 04:34:02 agent job Backup ... Finished'.
- Volume (Session):** Points to a session entry, such as '2009-01-27 04:34:10 agent job Backup ... Finished'.
- Summary of job history for server (Host):** Points to a summary row for a server, such as '100-362-DELL046 (9 job execution: 9 finished, 0 incomplete, 0 failed, 0 canceled, 0 not attempted)'.

Job View

The Job view displays all executions of a job. Each execution shows all of the hosts that were backed up. You can also drill down on a host and see the sessions that were backed up.

For each job entry, you also see the following summary information:

	Last Result	MB	Files	Missed	MB/Min...	Time Used	Job ID	Job No.	Session...	Subsession No.	Compression R
Archive - hsa	(0 job execution: 0 finished, 0 incomplete, 0 failed, 0 canceled)										
Archive [Staging] schedule_FSD	(2 job execution: 2 finished, 0 incomplete, 0 failed, 0 canceled)										
2010-12-10 01:57:22	Finished	0	18	0	.00	00:00:18	7293	315			
2010-12-02 22:57:22	Finished	0	18	0	.00	00:00:24	6746	315			
Archive-CSV file job	(0 job execution: 0 finished, 0 incomplete, 0 failed, 0 canceled)										
Archive-fsd	(0 job execution: 0 finished, 0 incomplete, 0 failed, 0 canceled)										
Archive-fsd job1	(0 job execution: 0 finished, 0 incomplete, 0 failed, 0 canceled)										
Archive_StubFile	(0 job execution: 0 finished, 0 incomplete, 0 failed, 0 canceled)										
Archive_yuaxi02_op...	(0 job execution: 0 finished, 0 incomplete, 0 failed, 0 canceled)										
Backup [Custom] or...	(0 job execution: 0 finished, 0 incomplete, 0 failed, 0 canceled)										
Backup [GFS, MUS] R...	(8 job execution: 4 finished, 0 incomplete, 4 failed, 0 canceled)										

- **Number of job executions**--Indicates the number of times the host was supposed to be backed up or backed up by a job.
- **Number of jobs finished**--Number of times the host was backed up successfully.
- **Number of jobs incomplete**--Number of times the host was not completely backed up.
- **Number of jobs failed**--Number of times the backup of the host failed.
- **Number of jobs canceled**--Number of times the backup of the host was canceled.

Backup Execution Details for a Selected Job

When you select a specific job execution the following information is displayed:

Group by	Last Result	MB	Files	Missed	MB/Min...	Time Used	Job ID	Job No.	Session...	Subsession No.	Compression R...
Job											
Archive - hsa	(0 job execution: 0 finished, 0 incomplete, 0 failed, 0 canceled)										
Archive [Staging] s...	(1 job execution: 1 finished, 0 incomplete, 0 failed, 0 canceled)										
Archive-CSV file job	(0 job execution: 0 finished, 0 incomplete, 0 failed, 0 canceled)										
Archive-fsd	(0 job execution: 0 finished, 0 incomplete, 0 failed, 0 canceled)										
Archive-fsd job1	(0 job execution: 0 finished, 0 incomplete, 0 failed, 0 canceled)										
Archive_yuaxi02_op...	(0 job execution: 0 finished, 0 incomplete, 0 failed, 0 canceled)										
Backup [Custom] or...	(0 job execution: 0 finished, 0 incomplete, 0 failed, 0 canceled)										
Backup [GFS, MUS] R...	(3 job execution: 3 finished, 0 incomplete, 0 failed, 0 canceled)										
Backup [GFS, MUX] ...	(0 job execution: 0 finished, 0 incomplete, 0 failed, 0 canceled)										
Backup [GFS, MUX] R...	(0 job execution: 0 finished, 0 incomplete, 0 failed, 0 canceled)										
Backup [GFS, Stagin...	(0 job execution: 0 finished, 0 incomplete, 0 failed, 0 canceled)										
Database protection...	(6 job execution: 5 finished, 0 incomplete, 1 failed, 0 canceled)										
hanhu01 Archive	(0 job execution: 0 finished, 0 incomplete, 0 failed, 0 canceled)										

In the top pane, the following information is displayed:

Job Execution Time

The time the job started.

Job Name

The name of the job.

Note: This release of Arcserve Backup does not display blank job names in the Job Status Manager. If you upgraded from a previous ARCserve release, migrated the job history data, and the jobs contained blank job names, the names of the jobs display in the Job Name field in the Job Status Manager in the following format:

<<machine name>> <<job no>>

Last Result

The last result is determined from the following criteria:

- The status is marked as failed if any of the sessions in the job fail.
- The status is marked as incomplete in any of the sessions are incomplete even if some are successful.
- The status is marked as successful only if all sessions are successful.

MB

The amount of data backed up for the job.

Files

The number of files backed up for the job.

Missed

The number of files missed during the backup.

Note: Use Arcserve Backup Agent for Open Files to backup open files to avoid missed files during a backup.

MB/Minute

- At the Job level, MB/Minute indicates the ratio of megabytes and the elapsed time for the entire job, including pre and post scripts, if any, media management activities, and so on. For average master job throughput, refer to the Activity Log.
- At the Host level, MB/Minute indicates the ratio of megabytes and the elapsed time for the entire job, including pre and post scripts, if any, media management activities, and so on for a single host.
- At the Session level, MB/Minute indicates the ratio of megabytes and the elapsed time for a specific volume and its folders, which comprise a single session.

Note: If little or no data is backed up by the backup job, a value of N/A appears in the MB/Minute field.

Time Used

- At the Job level, Time Used indicates the elapsed time for the entire job including pre and post scripts, if any, media management activities, and so on.
- At the Host level, Time Used indicates the elapsed time for the entire job including pre and post scripts, if any, media management activities, and so on for a single host.
- At the Session level, Time Used indicates the elapsed time for the backup of a specific volume and its folders, which comprise a single session.

Job ID

Identifies the specific execution of the job.

Job No.

Identifies the job.

Session Number

Identifies the session number of the session that contains the backup data.

Subsession Number

Identifies the subsession number of the session contains the backup data.

Compression Ratio

The amount of actual data to be stored divided by the amount of data stored after deduplication expressed as a ratio or percentage.

In the bottom pane, the following information is displayed:

Summary

Execution Time

The start time and end time of the selected job.

Total Source Host

The total number of hosts the job attempted to backup.

Total Sessions

The number of sessions that were backed up by the selected job execution.

Total Migrations

The number of sessions migrated in a disk or tape staging job.

Device and Media

Device

The tape drive or file system device used during the backup. Multiple tape drives can also be used for the same host in a single job execution if the job is a multi streaming job.

Media Used

The media that was used during the backup of the host. Multiple media can also be used for the same host in a single job execution if the job is a multi streaming job.

Error and Warning

Displays the errors and warnings that are generated during the back up of a host.

Node Level Details for a Selected Job

When you select the node of an executed job, the following information is displayed:

	Last Result	MB	Files	Missed	MB/Min...	Time Used	Job ID	Job No.	Session...	Subsession No.	Compressi...
[-] Job 001 hardware ...	(9 job execution: 2 finished, 0 incomplete, 7 failed, 0 canceled)										
[-] Job 002 mixif all ser...	(8 job execution: 0 finished, 0 incomplete, 8 failed, 0 canceled)										
[-] Job 003 CP mixed ...	(8 job execution: 8 finished, 0 incomplete, 0 failed, 0 canceled)										
[-] Job 004 agent job B...	(8 job execution: 0 finished, 0 incomplete, 8 failed, 0 canceled)										
[-] Job 005 Backup [Cu...	(8 job execution: 8 finished, 0 incomplete, 0 failed, 0 canceled)										
[-] 2009-01-30 04:46:40	Finished	0	6	0	.00	00:00:54	4961	139			
[-] 2009-01-29 04:46:46	Finished	0	6	0	.00	00:01:28	4804	139			
[-] 2009-01-28 04:47:12	Finished	0	6	0	.00	00:01:26	4686	139			
[-] 2009-01-27 07:39:14	Finished	0	6	0	.00	00:00:54	4549	139			
[-] 2009-01-26 12:01:14	Finished	0	6	0	.00	00:01:02	4424	139			
[-] 2009-01-25 04:46:38	Finished	0	6	0	.00	00:01:00	4368	139			
[-] 2009-01-24 04:46:38	Finished	0	6	0	.00	00:00:50	4319	139			
[-] 2009-01-23 04:46:44	Finished	0	6	0	.00	00:00:38	4188	139			
[-] 100-3FL-DELL700	Finished	0	6	0	N/A	00:00:00	4188	139			
E:\hotfix_11_11	Finished	0	6	0	N/A	00:00:00	4188	139	12		

Detail									
Summary									
Execution Time	2009-01-23 04:46:30----2009-01-23 04:46:30								
Total Sessions	1(1 Finished,0 Failed,0 Cancel,0 Incomplete,0 Other)								
Total Migrations	0(0 Finished,0 Failed,0 Incomplete,0 Pending)								
Device and Media									
Device	HP(Board:4,Bus:0,SCSIID:5,LUN:0)								
Media Used:1	<table border="1"> <thead> <tr> <th>Media Name</th> <th>Barcode</th> <th>SequenceNO</th> <th>RandomID</th> </tr> </thead> <tbody> <tr> <td>DELL700</td> <td></td> <td>1</td> <td>4756</td> </tr> </tbody> </table>	Media Name	Barcode	SequenceNO	RandomID	DELL700		1	4756
Media Name	Barcode	SequenceNO	RandomID						
DELL700		1	4756						
Error and Warning									
No item to display!									

Summary

Execution Time

The start time and end time of the selected node.

Total Sessions

The number of sessions that were backed up for the host.

Total Migrations

The number of sessions migrated in a disk or tape staging job.

Device and Media

Device

The tape drive or file system device used during the backup job. Multiple tape drives can also be used for the same host in a single job execution if the job is a multi streaming job.

Media Used

The media that was used during the backup of the host. Multiple media can also be used for the same host in a single job execution if the job is a multi streaming job.

Error and Warning

Displays the errors and warnings that are generated during the backup of a host.

Session Level Details of a Selected Node

Drilling down even more, when you highlight a session, the following details of the session are displayed:

	Last Result	MB	Files	Missed	MB/Min...	Time Used	Job ID	Job No.	Session...	Subsession No.	Compressio
2009-01-30 04:46:40	Finished	0	6	0	.00	00:00:54	4961	139			
2009-01-29 04:46:46	Finished	0	6	0	.00	00:01:28	4804	139			
2009-01-28 04:47:12	Finished	0	6	0	.00	00:01:26	4686	139			
2009-01-27 07:39:14	Finished	0	6	0	.00	00:00:54	4549	139			
2009-01-26 12:01:14	Finished	0	6	0	.00	00:01:02	4424	139			
2009-01-25 04:46:38	Finished	0	6	0	.00	00:01:00	4368	139			
2009-01-24 04:46:38	Finished	0	6	0	.00	00:00:50	4319	139			
2009-01-23 04:46:44	Finished	0	6	0	.00	00:00:38	4188	139			
100-3FL-DELL700	Finished	0	6	0	N/A	00:00:00	4188	139			
E:\hotfix_11_11	Finished	0	6	0	N/A	00:00:00	4188	139	12		

Detail	
Session Detail	
Execution Time	2009-01-23 04:46:30----2009-01-23 04:46:30
Number	12
Type	NTFS
Path	\\100-3FL-DELL700\E:\hotfix_11_11
Status	Finished
Start time	2009-01-23 04:46:30
End time	2009-01-23 04:46:30
Method	Full
Flags	Catalog
MB	0
Files	6
Missed	0
Device and Media	
Device	HP(Board:4,Bus:0,SCSIID:5,LUN:0)

Session Detail

Execution Time

Indicates the start time and end time of the selected session.

Number

Indicates the session number.

Type

Indicates the type of session backed up.

Path

Indicates the root path of the session.

Status

Indicates the result of the backup session.

Start time

Indicates the start time of the session.

End time

Indicates the end time of the session.

Method

Indicates the type of backup method used for the session.

Flags

Indicates the internal flags created by Arcserve Backup to identify the session.

MB

Indicates the amount of data backed up for the session.

Files

Indicates the number of files backed up for the session.

Missed

Indicates the number of files not backed up during the session.

Device and Media

Device

Indicates the tape drive or file system device used during the backup of the session.

Media Used

Indicates the media that was used during the backup of the session.

Error and Warning

Displays the errors and warnings that are generated during the back up of a session.

Host View

The Host view displays all of the hosts that were backed up and their status each time a job has backed it up. You can also drill down on a host and see the sessions that were backed up. For each host entry, you also see the following summary information:

Number of job execution

Indicates the number of times the host was attempted to be backed up or backed up by a job.

Number of jobs finished

Number of times the host was backed up successfully.

Number of jobs incomplete

Number of times the host was not completely backed up.

Note: To avoid incomplete backups use Arcserve Backup Agent for Open Files to back up open files.

Number of jobs failed

Number of times the backup of the host failed.

Number of jobs canceled

Number of times the backup of the host was canceled.

Backup Execution Details of a Selected Host

When you select a specific job execution the following information is displayed:

Group by	Host	Show history in last	15 Days	Update																																																
<table border="1"> <thead> <tr> <th>Job Name</th> <th>Last Result</th> <th>MB</th> <th>Files</th> <th>Missed</th> <th>MB/Mn...</th> <th>Time U...</th> <th>Job ID</th> <th>Job No.</th> <th>Session No.</th> <th>Subsession No.</th> <th>Compression Ratio</th> </tr> </thead> <tbody> <tr> <td colspan="12">(1 job execution: 1 finished, 0 incomplete, 0 failed, 0 canceled, 0 not attempted)</td> </tr> <tr> <td>A160</td> <td>2011-01-01 17:27:24 Backup [Custom]</td> <td>Finished</td> <td>0</td> <td>1</td> <td>0</td> <td>.00</td> <td>00:00:30</td> <td>248</td> <td>1</td> <td></td> <td>N/A</td> </tr> <tr> <td></td> <td>C:\vs8_remote_... Backup [Custom]</td> <td>Finished</td> <td>0</td> <td>1</td> <td>0</td> <td>N/A</td> <td>00:00:00</td> <td>248</td> <td>1</td> <td>1</td> <td>N/A</td> </tr> </tbody> </table>					Job Name	Last Result	MB	Files	Missed	MB/Mn...	Time U...	Job ID	Job No.	Session No.	Subsession No.	Compression Ratio	(1 job execution: 1 finished, 0 incomplete, 0 failed, 0 canceled, 0 not attempted)												A160	2011-01-01 17:27:24 Backup [Custom]	Finished	0	1	0	.00	00:00:30	248	1		N/A		C:\vs8_remote_... Backup [Custom]	Finished	0	1	0	N/A	00:00:00	248	1	1	N/A
Job Name	Last Result	MB	Files	Missed	MB/Mn...	Time U...	Job ID	Job No.	Session No.	Subsession No.	Compression Ratio																																									
(1 job execution: 1 finished, 0 incomplete, 0 failed, 0 canceled, 0 not attempted)																																																				
A160	2011-01-01 17:27:24 Backup [Custom]	Finished	0	1	0	.00	00:00:30	248	1		N/A																																									
	C:\vs8_remote_... Backup [Custom]	Finished	0	1	0	N/A	00:00:00	248	1	1	N/A																																									
<table border="1"> <thead> <tr> <th colspan="2">Summary</th> </tr> </thead> <tbody> <tr> <td>Total Sessions</td> <td>1(1 Finished,0 Failed,0 Cancel,0 Incomplete,0 Other)</td> </tr> <tr> <td>Total Migrations</td> <td>0(0 Finished,0 Failed,0 Incomplete,0 Pending)</td> </tr> </tbody> </table>					Summary		Total Sessions	1(1 Finished,0 Failed,0 Cancel,0 Incomplete,0 Other)	Total Migrations	0(0 Finished,0 Failed,0 Incomplete,0 Pending)																																										
Summary																																																				
Total Sessions	1(1 Finished,0 Failed,0 Cancel,0 Incomplete,0 Other)																																																			
Total Migrations	0(0 Finished,0 Failed,0 Incomplete,0 Pending)																																																			
<table border="1"> <thead> <tr> <th colspan="2">Device and Media</th> </tr> </thead> <tbody> <tr> <td>Device</td> <td>FSD3(Board:3,Bus:0,SCSIID:2,LUN:0)</td> </tr> <tr> <td>Media Used:1</td> <td>1/01/11 5:27 PM</td> </tr> <tr> <td>Barcode</td> <td></td> </tr> <tr> <td>SequenceNO</td> <td>1</td> </tr> <tr> <td>RandomID</td> <td>E04E</td> </tr> </tbody> </table>					Device and Media		Device	FSD3(Board:3,Bus:0,SCSIID:2,LUN:0)	Media Used:1	1/01/11 5:27 PM	Barcode		SequenceNO	1	RandomID	E04E																																				
Device and Media																																																				
Device	FSD3(Board:3,Bus:0,SCSIID:2,LUN:0)																																																			
Media Used:1	1/01/11 5:27 PM																																																			
Barcode																																																				
SequenceNO	1																																																			
RandomID	E04E																																																			
<table border="1"> <thead> <tr> <th colspan="2">Error and Warning</th> </tr> </thead> <tbody> <tr> <td colspan="2">No item to display!</td> </tr> </tbody> </table>					Error and Warning		No item to display!																																													
Error and Warning																																																				
No item to display!																																																				

In the top pane, the following information is displayed:

Job Name

The name of the job that backed up the host.

Note: This release of Arcserve Backup does not display blank job names in the Job Status Manager. If you upgraded from a previous ARCserve release, migrated the job history data, and the jobs contained blank job names, the names of the jobs display in the Job Name field in the Job Status Manager in the following format:

[<<machine name>>] <<job no>>

Last Result

The last result is determined from the following criteria:

- The status is marked as failed if any of the sessions in the host fail.
- The status is marked as incomplete if any of the sessions are incomplete even if some are successful.
- The status is marked as successful only if all sessions are successful.

MB

The amount of data backed up for the host.

Files

The number of files backed up for the host.

Missed

The number of files missed during the backup job.

Note: Use Arcserve Backup Agent for Open Files to avoid missed files during a backup job.

MB/Minute

- At the Job level, MB/Minute indicates the ratio of megabytes and the elapsed time for the entire job, including pre and post scripts, if any, media management activities, and so on. For average master job throughput, refer to the Activity Log.
- At the Session level, MB/Minute indicates the ratio of megabytes and the elapsed time for a specific volume and its folders, which comprise a single session.

Time Used

- At the Job level, Time Used indicates the elapsed time for the entire job including pre and post scripts, if any, media management activities, and so on.
- At the Session level, Time Used indicates the elapsed time for the backup of a specific volume and its folders, which comprise a single session.

Job ID

Identifies the specific execution of the job.

Job No.

Identifies the job.

Session Number

Identifies the serial number of a session on a media. For example, session number 4, this is the fourth session on the media.

Identifies session number of the session that contains the backup data.

Subsession Number

Identifies the serial number of the logical session in a physical session when a virtual machine (VM) is backed up by the Agent for Virtual Machines. If Arcserve Backup backs up a VM using the Agent for Virtual Machines, Arcserve Backup backs up all volumes and components into a single session. Each volume or component is backed up in a logical session. This logical session contains the serial number, which is the subsession number.

Compression Ratio

The amount of actual data to be stored divided by the amount of data stored after deduplication expressed as a ratio or percentage.

In the bottom pane, the following information is displayed:

Summary**Total Sessions**

The number of sessions that were backed up for the host.

Total Migrations

The number of sessions migrated in a disk or tape staging job.

Device and Media**Device**

The tape drive or file system device used during the backup job. Multiple tape drives can also be used for the same host in a single job execution if the job is a multi streaming job.

Media Used

The media that was used during the backup of the host. Multiple media can also be used for the same host in a single job execution if the job is a multi streaming job.

Error and Warning

Displays the errors and warnings that are generated during the backup of a host.


Session Level Details of a Selected Host

Drilling down even more, when you highlight a session, the following details of the session are displayed:

Job 001 (1 job execution: 1 finished, 0 incomplete, 0 failed, 0 canceled)									
2009-01-21 00:48:50	Finished	24	18	0	30.00	00:00:48	3861	269	
100-362-DELL192	Finished	24	18	0	102.86	00:00:14	3861	269	
sqlldr@NEW	Finished	6	4	0	90.00	00:00:04	3861	269	1
dbasql@NEW...	Finished	4	2	0	120.00	00:00:02	3861	269	2
dbasql@NEW...	Finished	2	2	0	N/A	00:00:00	3861	269	3

Session Detail	
Execution Time	2009-01-21 00:48:34----2009-01-21 00:48:38
Number	1
Type	SQL Server Disaster Recovery Elements
Path	\\100-362-DELL192\sqlldr@NEW
Status	Finished
Start time	2009-01-21 00:48:34
End time	2009-01-21 00:48:38
Method	Clone-Snap
Flags	Agent, Catalog
MB	6
Files	4
Missed	0

Device and Media				
Device	FSD3(Board:2,Bus:0,SCSIID:2,LUN:0)			
Media Used:1	Media Name	Barcode	SequenceNO	RandomID
	1/20/09 8:50 PM		1	8C4B

Error and Warning
 W3558 2009-01-20 20:50:33 Global Agent Options for Microsoft SQL Server are not applied by agents whose version is below r12.5. (Node=100-362-DELL192, Agent Version=r12.0)

Session Detail

Execution Time

Indicates the start time and end time of the selected session.

Number

Indicates the session number.

Type

Indicates the type of session backed up.

Path

Indicates the root path of the session.

Status

Indicates the result of the backup session.

Start time

Indicates the start time of the session.

End time

Indicates the end time of the session.

Method

Identifies the type of backup method used for the session.

Flags

Indicates the internal flags created by Arcserve Backup to identify the session.

MB

Indicates the amount of data backed up for the session.

Files

Indicates the number of files backed up for the session.

Missed

Indicates the number of files not backed up during the session.

Device and Media

Device

Indicates the tape drive or file system device used during the backup of the session.

Media Used

Indicates the media that was used during the backup of the session.

Error and Warning

Displays the errors and warnings that are generated during the back up of a session.

Source Group View

The Source Group view displays the executions of both, application groups and customized source groups. Each execution shows all of the servers selected in the source group that was backed up. You can drill-down on a server and see the details of the sessions that were backed up.

For each source group, you can see the following summary information:

Job Queue	Job History	Activity Log	Audit Log	Tape Log							
Group by: Source Group		Show history in last: 7 Days		Update							
Job Name	Last Result	MB	Files	Missed	MB/Min...	Time U...	Job ID	Job No.	Session...	Subses...	Compression R...
Client Agent	(8 job execution: 2 finished, 0 incomplete, 3 failed, 0 canceled, 3 active, 0 not attempted)										
Client Agent Group Name 1...	(8 job execution: 1 finished, 0 incomplete, 1 failed, 0 canceled, 6 active, 0 not attempted)										
Microsoft SQL Server	(8 job execution: 5 finished, 0 incomplete, 1 failed, 0 canceled, 2 active, 0 not attempted)										
Mixed Group	(8 job execution: 0 finished, 0 incomplete, 6 failed, 0 canceled, 2 active, 0 not attempted)										
SQL Agent Group Name SQ...	(8 job execution: 7 finished, 0 incomplete, 0 failed, 0 canceled, 1 not attempted)										
WinUser Group WinUser G...	(8 job execution: 0 finished, 0 incomplete, 7 failed, 0 canceled, 1 not attempted)										

Number of job execution

Indicates the number of times the source group was supposed to be backed up or backed up by a job.

Number of job finished

Indicates the number of times the source group was backed up successfully.

Number of jobs incomplete

Indicates the number of times the source group was not completely backed up.

Number of jobs failed

Indicates the number of times the backup of the source group failed.

Number of jobs canceled

Indicates the number of times the backup of the source group was canceled.

Number of active jobs

Indicates the number of backup jobs that are still being executed.

Number of jobs not attempted

Indicates the number of backup jobs that are yet to be executed.

Job Level Details for a Selected Host

When you select a specific job, the following details are displayed:

Job Name	Last Result	MB	Files	Missed	MB/Min...	Time U...	Job ID	Job No.	Session...	Subses...	Compl
Group R (1 job execution: 0 finished, 0 incomplete, 0 failed, 0 canceled, 1 active, 0 not attempted)											
LODVM1040											
SERVER-GROUP-1 (2 job execution: 1 finished, 0 incomplete, 0 failed, 0 canceled, 1 not attempted)											
LODVM1040											
2009-11-02 15:3...R2	Not Attemp...	N/A	N/A	N/A	N/A	00:00:00	37	12			
2009-11-02 15:3...R1	Finished	89	51	0	83.44	00:01:04	35	11			
C:\setup R1	Finished	89	51	0	1,335.00	00:00:04	35	11		1	

The following information is displayed in the top pane:

Job Execution Time

Indicates the time when the job execution started.

Job Name

Indicates the name of the job.

Last Result

The last result is determined from the following criteria:

- The status is marked as failed if any of the sessions in the job fail.
- The status is marked as incomplete in any of the sessions are incomplete even if some are successful.
- The status is marked as successful only if all sessions are successful.

MB

Indicates the amount of data backed up for the job.

Files

Indicates the number of files backed up for the job.

Missed

Indicates the number of files missed during the backup.

Note: Use Arcserve Backup Agent for Open Files to back up open files to avoid missed files during a backup.

MB/Minute

- At the Job level, MB/Minute indicates the ratio of megabytes and the elapsed time for the entire job, including pre and post scripts, if any, media management activities, and so on. For average master job throughput, refer to the Activity Log.
- At the Host level, MB/Minute indicates the ratio of megabytes and the elapsed time for the entire job, including pre and post scripts, if any, media management activities, and so on for a single host.
- At the Session level, MB/Minute indicates the ratio of megabytes and the elapsed time for a specific volume and its folders, which comprise a single session.

Note: If little or no data is backed up by the backup job, a value of N/A appears in the MB/Minute field.

Time Used

- At the Job level, Time Used indicates the elapsed time for the entire job including pre and post scripts, if any, media management activities, and so on.
- At the Host level, Time Used indicates the elapsed time for the entire job including pre and post scripts, if any, media management activities, and so on for a single host.
- At the Session level, Time Used indicates the elapsed time for the backup of a specific volume and its folders, which comprise a single session.

Job ID

Identifies the specific execution of the job.

Job No.

Identifies the job.

Compression Ratio

The amount of actual data to be stored divided by the amount of data stored after deduplication expressed as a ratio or percentage.

In the bottom pane, the following information is displayed:

Summary

Total Sessions

Indicates the number of sessions that were backed up by the selected job.

Total Migrations

Indicates the number of sessions migrated in a disk or tape staging job.

Device and Media

Device

Indicates the tape drive or file system device used during the backup. Multiple tape drives can also be used for the same host in a single job execution if the job is a multistreaming job.

Media Used

Indicates the media that was used during the backup of the host. Multiple media can also be used for the same host in a single job execution if the job is a multistreaming job.

Error and Warning

Displays the errors and warnings that are generated during the backup of a host.

Session Level Details for a Selected Group

When you select a specific session, the following details are displayed:

Detail				
Session Detail				
Execution Time	2009-09-26 10:54:24----2009-09-26 11:04:02			
Number	119			
Type	NTFS			
Path	\\L0D0057\C:			
Status	Finished			
Start time	2009-09-26 10:54:24			
End time	2009-09-26 11:04:02			
Method	Full			
Flags	Drive, Agent, Catalog			
MB	12308			
Files	68061			
Missed	0			
Device and Media				
Device	CA(Board:7,Bus:0,SCSIID:0,LUN:4)			
Media Used:1	Media Name	Barcode	SequenceNO	RandomID
	9/25/09 1:16 AM	1100010	1	57CE
Error and Warning				
No item to display!				

Session Detail

Execution Time

Indicates the start time and end time of the selected session.

Number

Indicates the session number.

Type

Identifies the type of session backed up.

Path

Indicates the root path of the session.

Status

Indicates the result of the backup session.

Start time

Indicates the start time of the session.

End time

Indicates the end time of the session.

Method

Indicates the backup method used for the session.

Flags

Indicates the internal flags created by Arcserve Backup to identify the session.

MB

Indicates the amount of data backed up for the session.

Files

Indicates the number of files backed up for the session.

Missed

Indicates the number of files not backed up during the session.

Device and Media

Device

Indicates the tape drive or file system device used during the backup of the session.

Media Used

Indicates the media that was used during the backup of the session.


Error and Warning

Displays the errors and warnings that are generated during the back up of a session.

Filter the Job Queue

Filter options let you refine how you search the Job Queue.

To filter the Job Queue

1. Open the Job Status Manager.
Select the Job Queue tab.
Expand the header bar by clicking .

Choose from the following filter options:

- **Show jobs with the status**--Lets you filter jobs based on the job status.
- **Show done jobs with the results**--Lets you filter done jobs based on the jobs status
- **Keywords**--Lets you filter jobs that contain a particular keyword.
- **In**--Used in conjunction with Keywords, lets you specify whether the keyword is contained within the Backup Server Name or the Job Name.
- **Show jobs owned by other users**--Lets you view all jobs or only the jobs that you own.
- **Show jobs by selected types**--Lets you view jobs based on the type of job. For example, a backup job, a restore job, a migration job, and so on. To specify the types of jobs that you want to view, click Select Types.

2. Click Update.

The filter options are applied.

Filter Job History

Filter options let you refine your job history search.

To filter job history

1. Open the Job Status Manager.

Select the Job History tab.

Click  to expand the header bar.

Choose from the following filter options:

- **Group by**--Specify the type of group to sort by. The options are by job, host, or source group.
- **Show history in last xx days**--Specify the number of days of job history you need. The range is from 1 to 100 days.
- **Show groups with the most recent result**--Specify what type of result history you want to view. You can specify one, all or any combination of options. The options include active, finished, incomplete, failed, and canceled.

Note: The header bar turns yellow when there is a change made to the type of result history you want to view indicating that the advanced filter was used.

- **Keywords**--Specify the keywords that you want to use in the sorting of the job history by Job Name or Host Name.

Note: The header bar turns yellow when a keyword is specified indicating that the advanced filter was used.

- **In**--Specify a Job Name or Host Name. The keywords are used to identify jobs in the chosen category.

2. Click Update.
The filter options are applied.

View Job History

Use the Job History dialog to identify patterns or areas of repeated errors.

To view a job history

1. Open the Job Status Manager.
2. Select the Job History tab.
3. In the Group By drop-down list, select one of the following:
 - Host
 - Job
 - Source Group

Depending upon your selection, the job history appears in either Host View or Job View.

4. Click Update.
The Properties panel displays the job history.

Locate Information Using Quick Search

Locating information manually on the user interface can be time-consuming and tedious, especially when there are many jobs, nodes, tapes, etc. in the system. The Quick Search feature allows you to quickly and easily find the information you need on the user interface. You can use Quick Search to find an item on any tree or list in the Arcserve Backup Manager. For example, you can use Quick Search to locate the following type of information:

- Jobs or logs in the Activity Log
- Jobs in Job History
- Nodes when using Restore by Tree
- Tapes or sessions when using Restore by Session
- Media pools from the Media Pool Manager

Important! Quick Search will only search items displayed on the user interface. It will not search for particular data stored on a tape.

Note: To enhance Quick Search performance, you can filter the Activity Log to reduce the number of items displayed and queried prior to using Quick Search. For more information about filtering the Activity Log, see [Set Activity Log Queries](#) (see page 331).

To locate information using the Quick Search feature

1. Press CTRL+F to open the Quick Search dialog from any tree or list view on the Arcserve Backup user interface.

For a list of where and how you can access the Quick Search feature using **Ctrl+F** from the user interface and a sample of the kind of items that can be searched, see [Quick Search Accessibility](#) (see page 78).

Note: You can also launch Quick Search from the context menu of a tree or list view when you right-click and select Quick Search.

2. Type a keyword to search on or select one from the drop-down list.

If you type a keyword to search on, the search supports the hint during typing feature where the hint is provided from the keyword history. For example, if you type "job" and the keywords "Job 1943" and "Job 2048" have been entered previously, then both "Job 1943" and "Job 2048" will be listed under the text box where you are typing.

Note: The search does not support wildcard characters.

If you click the drop-down arrow, the keyword history displays. By default, the maximum number of keywords remembered is fifty, but this limit can be set in the local registry.

3. (Optional) Click the plus sign icon to expand the Search options field and choose the options that apply.

Match case

Search using case-sensitive capitalization.

Match whole word

Search using whole word matches and not part of a word.

4. Press Enter or click the Search button.

The Search result list displays, showing all the matched items, the number of items found, and the total number of items searched.

Important! The Quick Search feature will only find items that are at the expanded level. Items under a collapsed node will not be found. Therefore, as long as the item is visible on the user interface or can be seen by scrolling, the Quick Search is able to find it.

Note: The search function will work in the background, so you can continue working during long searches without closing the dialog. To stop the search at any time, click the Search button.

5. (Optional) Select or double-click an item from the list of search results.

The search stops if it is still active and the item is selected on the tree or list in the background of the user interface.

6. (Optional) Double-click another item from the results.

The Quick Search dialog remains open.

Note: If the preferred items are not found, you can adjust the keyword and perform another search.

7. To close the Quick Search dialog, press ESC or click the X button to close.

Quick Search Accessibility

The following indicates where and how you can access the Quick Search feature using **Ctrl+F** from the user interface and a sample of the kind of items that can be searched:

Manager	Location	Tree/List View	Searchable Items
Job Status	Arcserve Domain Tree	Tree	Domain, Server
	Job Queue Tab	List	Job
	Job History Tab	List	Job, Host, Job Execution, Session
	Activity Log Tab	List	Job, Log, Message Number, Message
Backup Manager	Source Tab	Tree	Machine, Folder
		List	Machine, Folder, File
	Staging Location Tab	Tree	Server, Group
		List	Server, Group, Media
	Destination Tab	Tree	Server, Group
		List	Server, Group, Media
Restore Manager (By Tree, By Session, By Image, By Backup Media)	Source Tab	Tree	Machine, Session, Folder, Group
		List	Machine, Session, Folder, Group, File, Media
	Destination Tab	Tree	Machine, Folder
		List	Machine, Folder
Device Manager		Tree	Server, Device
		List	Server, Device, Media
Media Pool Manager		Tree	Media Pool
		List	Media Pool, Media
Database Manager		Tree/List	Job, Media, Session, Folder, File, Device

Manager	Location	Tree/List View	Searchable Items
Merge, Media Assure & Scan, Compare, Copy, Count, Purge		Tree/List	Same as Backup and Restore Manager

How Password Management Works

Password Management provides the option to encrypt session passwords during backup and eliminates the need to repeatedly provide passwords. This feature lets you store session passwords in the Arcserve Backup database. During a backup job submission, the passwords are stored in encrypted form and will be automatically used during restore. Along with the session password, information about when to change your password is also stored. If you forget to change your session password, you will receive a reminder through the Activity Log.

Also, part of the enhancement is the ability to restore encrypted tapes on-site without requiring the encryption password. This feature allows other operators to perform different tasks without having to enter the password.

Note: If you run a restore, merge, or compare job for sessions created using an older version of Arcserve Backup, encrypted a tape in a different Arcserve Backup domain, or if Password Management was not enabled during a backup job, you need to provide the session/encryption password manually.

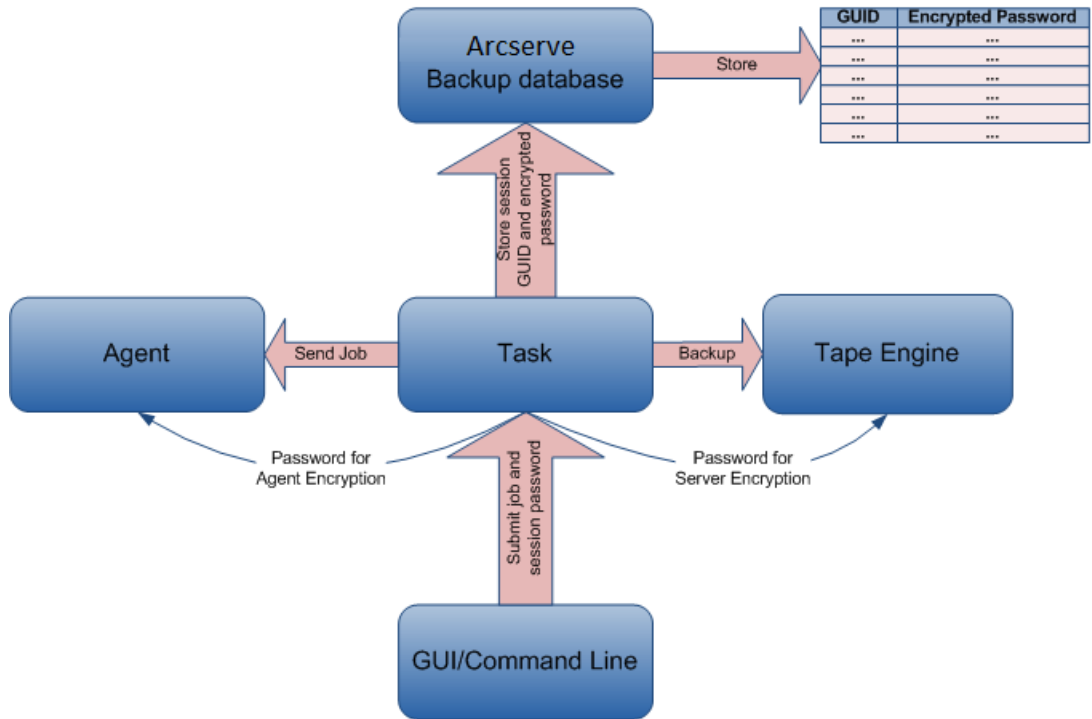
As a backup session is submitted, the session encryption password is saved to the Arcserve Backup database in encrypted format using a random key and the Globally Unique Identifier (GUID) is saved as a binary value. During a restore session, the encrypted password is extracted from the Arcserve Backup database and decrypted. To extract the encrypted password, the session GUID must be known. Depending on how the data was encrypted, either Server Side Encryption or Agent Side Encryption, there are two ways to identify the session GUID.

For Server Side Encryption

The restore session reads the Dummy Session Header from the Tape Engine and if server side encryption was used, the session GUID will be extracted from the Arcserve Backup database.

For Agent Side Encryption

The restore session reads the Session Header from the Tape Engine and extracts the GUID from the Arcserve Backup database.



Change a Session/Encryption Password

An Activity Log warning message is generated seven days in advance of a job session password expiration.

Note: This procedure allows you to change only a Global Option password.

To change a session/encryption password

1. From the Job Status Manager, select the Job Queue tab.
2. Select a job and right-click.
3. From the right-click menu, select Modify Encryption Password.
The Encryption dialog opens.
4. Enter a session/encryption password.

5. Select the option **Save Current Session/Encryption Password into Database**.
6. (Optional) Enter the number of days that must elapse before you need to change the password.

Enable Password Management

When submitting a backup job, you have the option to set a session encryption password.

To enable password management

1. From the Backup Manager, select the Options button on the toolbar.
The Global Options dialog opens.
2. On the Encryption/Compression tab, enter a session encryption/compression password.
3. Select the option **Save Current Session Encryption Password into Database**.
4. (Optional) Enter the number of days that must elapse before you need to change the password.
5. Select the Submit button on the toolbar to submit the backup job and save the session encryption password to the Arcserve Backup database.

How User Profile Management Works

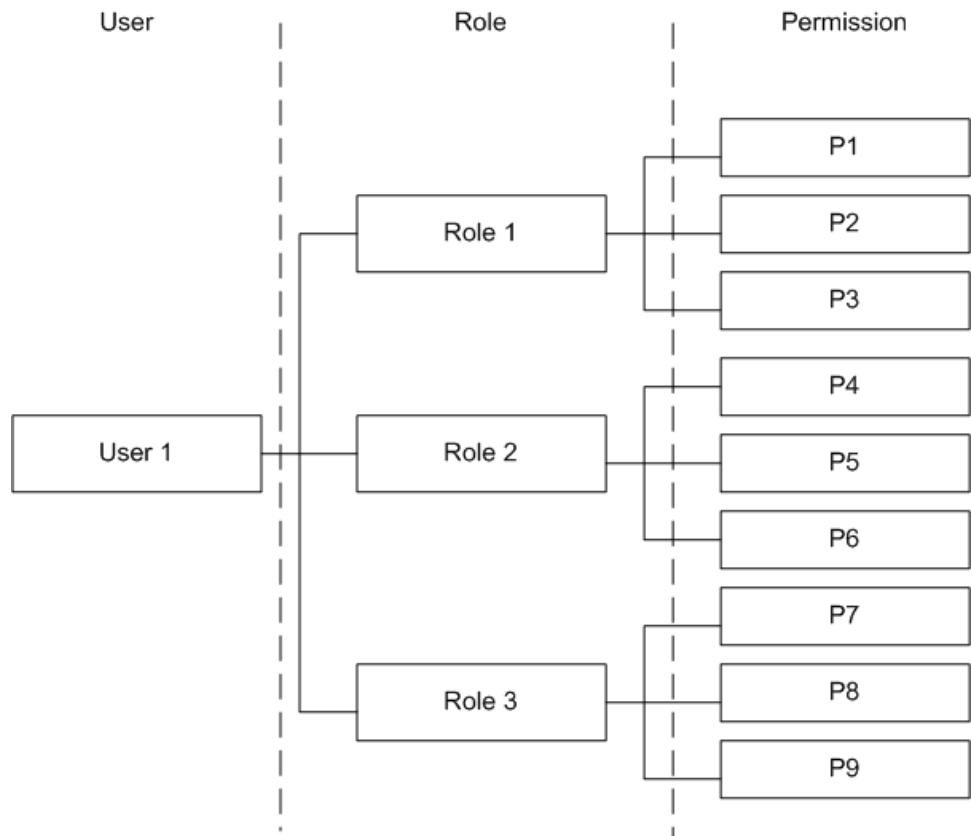
User Profile Management allows you to have different users access Arcserve Backup with different privileges. Using the User Profile Manager, you can assign individual users different roles with specific permissions. For example, you can have three users: one having an administrator role, one having a backup operator role, and one having a restore operator role. The ability to assign permissions based on role decreases the chances of an unauthorized user accessing the Arcserve Backup domain.

When you install Arcserve Backup, the caroot user profile is set up, with the Administrator group assigned to it by default. The Administrator group provides control over all Arcserve Backup functions operating within a given Arcserve Backup domain.

Using the User Profile Manager, Arcserve Backup supports the following management functions for users and roles:

- Add a user.
- Delete a user.
- Change a user's password.
- Assign a user to a role.
- Delete a user from a role.
- Assign a role to a user.
- Delete a role from a user.

You can assign a user multiple roles, providing the user with a variety of permissions. The following diagram illustrates a user with multiple roles:



Roles and Permissions

You can assign a user multiple roles and each role consists of specific permission set. All users must have at least one role assigned to them. Some of the roles have very restricted permissions. For example, the Report Operator can only view and create reports, while the Arcserve Backup Administrator can perform all operations.

Note: Only the Administrator role can add or delete users.

Permission	Role						
	Admin	Backup	Restore	Device	Monitor	Report	Tape
Submit Jobs	X	X	X				
Tape and Device	X	X	X	X	X		X
Job	X	X	X	X	X		X
Log	X	X	X	X	X		X
Reports	X	X				X	X
Service	X	X	X	X	X		X
Media Pools	X	X	X	X	X		
Database	X	X	X				
Dashboard	X				X	X	
MMO	X	X					
Other	X	X	X	X	X	X	X

Submit Job Options Permission Details

The following table outlines the Submit Job Options permission details and identifies corresponding roles:

Submit Job Options	Role		
	Admin	Backup	Restore
Backup	X	X	
Restore	X		X
Compare	X	X	X
Scan	X	X	X

Submit Job Options	Role		
	Admin	Backup	Restore
Merge	X	X	X
Generic*	X		
Count	X		
Purge	X		
Copy	X		

* **Note:** This is a job created using the Job Scheduler Wizard.

Tape and Device Operations Permission Details

The following table outlines the Tape and Device Operations permission details and identifies corresponding roles:

Tape and Device Operations	Role					
	Admin	Backup	Restore	Device	Monitor	Tape
View	X	X	X	X	X	X
Format/Erase	X	X		X		
Config FSD	X	X		X		
Config Groups	X	X		X		
Config Staging Groups	X	X		X		
Config DDD	X	X		X		
DDD Groups	X	X		X		
Retention	X			X		X
Compression	X			X		X
Eject	X			X		X
Rebuild	X			X		X
Mount	X			X		X
Import/Export	X			X		X
Clean	X			X		X
Rescan	X			X		X
View Properties	X			X		X

Tape and Device Operations	Role					
	Admin	Backup	Restore	Device	Monitor	Tape
Set Auto Clean	X			X		X

Job Operations Permission Details

The following table outlines the Job Operation permission details and identifies corresponding roles:

Note: The Backup Operator and the Restore Operator can only modify, reschedule, run, stop or delete jobs that they submit. If the Ownership Checking Exemption Privilege permission is checked the Backup Operator and the Restore Operator can operate jobs submitted by any operator.

Job Operations	Role					
	Admin	Backup	Restore	Device	Monitor	Tape
View all Status	X	X	X	X	X	X
Add	X	X	X			
Modify	X	X	X			
Reschedule	X	X	X			
Run/Stop	X	X	X			
Delete	X	X	X			
Modify Username	X	X	X			
Run PFC	X	X				
View current job status	X	X	X			
Modify password	X	X				

Log Operations Permission Details

The following table outlines the Log Operations permission details and identifies corresponding roles:

Tape and Device Operations	Role					
	Admin	Backup	Restore	Device	Monitor	Tape
View Job History	X	X	X	X	X	X
View Activity Log	X	X	X	X	X	X

Tape and Device Operations	Role					
	Admin	Backup	Restore	Device	Monitor	Tape
Delete Activity Log	X					
View Tape Log	X	X	X	X	X	X
Delete Tape Log	X					
View Audit Log	X	X	X	X	X	X
Delete Audit Log	X					

Report Operations Permission Details

The following table outlines the Role Operations permission details and identifies corresponding roles:

Report Operations	Role			
	Admin	Backup	Report	Tape
View/Create	X	X	X	X
Design	X	X	X	X

Service Operations Permission Detail

The following table outlines the Service Operations permission details and identifies corresponding roles:

Service Operations	Role					
	Admin	Backup	Restore	Device	Monitor	Tape
View Service Status	X	X	X	X	X	X
Set Auto Start Engine	X					
Adjust Arcserve Backup System Account	X					
Config System engines	X					
Start/Stop all services	X					
Start/Stop specified services	X					
Start/Stop all services in domain	X					

Service Operations	Role					
	Admin	Backup	Restore	Device	Monitor	Tape
Add/View Licenses	X					
Manage Licenses	X					
View service status	X					
Install/Uninstall options	X					

Media Pool Operations Permission Detail

The following table outlines the Media Pool Operations permission details and identifies corresponding roles:

Media Pool Operations	Role					
	Admin	Backup	Restore	Device	Monitor	Tape
View Media Pools	X	X	X	X	X	X
Create new media pools	X	X				
Delete media pools	X	X				
Move Media between Scratch Set and Save Set	X	X				
Assign Media to Scratch Set and Save Set	X	X				
Remove Media from Scratch Set and Save Set	X	x				

Database Operations Permission Detail

The following table outlines the Database Operations permission details and identifies corresponding roles:

Database Operations	Role			
	Admin	Backup	Restore	Monitor
View media/sessions in database	X	X	X	X
View jobs in database	X	X	X	X
View devices in database	X	X	X	X
Delete media/sessions in database	X			

Database Operations	Role			
	Admin	Backup	Restore	Monitor
Delete jobs in database	X			
Delete devices in database	X			

MMO Operations Permission Detail

The following table outlines the MMO Operations permission details and identifies corresponding roles:

MMO Operations	Role	
	Admin	Backup
All MMO operations	X	X

Other Operations Permission Detail

The following table outlines the Other Operations permission details and identifies corresponding roles:

Other Operations	Role						
	Admin	Backup	Restore	Device	Monitor	Report	Tape
Configure Alert Manager	X	X	X	X	X	X	X
View Alert Manager	X	X	X	X	X	X	X
Use Diagnostic Manager	X	X	X	X	X	X	X

Extended Permissions

The User Profile Manager includes the following extended permissions:

- **Security Administrator**--The Security Administrator permission is only selectable if Arcserve Backup Administrator role is selected. In order to perform user management task the Security Administrator must be selected.
- **Ownership Checking Exemption Privilege**--The Ownership Checking Exemption Privilege is only selectable if the Backup or Restore Operator role is assigned to a user. If the Ownership Checking Exemption Privilege permission is checked the Backup Operator and the Restore Operator can operate jobs submitted by any operator.

How Windows User Authentication Works

Arcserve Backup Windows User Authentication simplifies Arcserve Backup user management. It lets Windows users log in to the Arcserve Backup domain with their Windows user account login information.

Users can log in to Arcserve Backup after the Arcserve Backup administrator has added the user to the Arcserve Backup database. Any valid Windows user can be added as a user from the User Profile Manager.

Window User Authentication is a two step process. The process is as follows:

- The user must be authenticated in the Windows domain.
- The user must be authenticated in the Arcserve Backup database.

Note: To help ensure that a Windows user account with a blank password can log in to Arcserve Backup successfully, you must configure a [Windows Security Setting Option](#) (see page 90).

To help ensure that users that are logged in to Arcserve Backup with a Windows user account that has Arcserve Backup administrative privileges can access database related activities (for example, view the Activity Log, view the Audit Log, monitor jobs, and so on), the Windows user account must be configured as follows:

- The Windows user account must be able to log in to Microsoft SQL Server or Microsoft SQL Server Express databases.
- The Windows user account must be assigned the Microsoft SQL Server or Microsoft SQL Server Express SysAdmin role.

Configure Windows Security Setting Option

To make sure a Windows user with a blank password will not fail to log on to Arcserve Backup, you must configure a Windows Security Setting Option.

To configure windows security setting option

1. From the Start menu select Control Panel.
The Control Panel opens.
2. Select Administrative Tools.
The Administrative Tools dialog opens.
3. Select Local Security Policy.
The Local Security Policy dialog opens.
4. Double-click the Accounts: Limit local account use of blank passwords to console logon only option.
5. Select Disable and click OK.
The Windows Security Setting Option is configured to accept blank passwords.

Open the Manager or Manager Console

The Manager Console is an interface that lets you administer backup and restore operations in your environment. With the Manager Console, you can log in to and administer local and remote Arcserve Backup servers and domains.

This release of Arcserve Backup provides you with a redesigned Manager Console. If you are running an older release of Arcserve Backup in your environment, you must log in to the system running the older release using the previous version of the Manager.

To open the Manager or Manager Console

- Do one of the following:
 - To access a server running this release of Arcserve Backup, click the Windows Start button, point to Programs, CA, ARCserve Backup, and click Manager.
 - To access an Arcserve server running a previous release, browse to the following file:

C:\Programs Files\CA\ARCserve Backup\ARCserveMgr.exe

Double-click ARCserveMgr.exe.

- If you installed a previous Arcserve Backup release in the default installation directory, and used the installation process to upgrade Arcserve Backup, you can open the Manager by clicking the Windows Start button, select Programs, CA, ARCserve Backup, and click Manager.

The Default Server Information page appears.

- To change the default server or specify a different server, select a server from the Arcserve Backup Primary Server list. If the target server does not appear in the drop-down list, you can input the host name or IP address of the server in the Arcserve Backup Primary Server list.

- To change the user, select either Arcserve Backup Authentication or Windows Authentication and specify a user name and password.

By default, Arcserve Backup does not remember your security information. To save the user name and password information that you entered for this server, you must explicitly select Remember the security information. If you do not save this information, Arcserve Backup prompts you to provide Arcserve Backup security credentials the first time you open managers, wizards, and so on, and you must provide a Arcserve Backup user name and password.

4. Enter caroot in the User Name field, the appropriate password in the Password field, and click OK.

The first time you log in to Arcserve Backup, a tutorial, called My First Backup, appears. This tutorial lets you become familiar with the basics of backing up and restoring data in a controlled and directed way. This tutorial appears automatically only the first time you log in. However, you can access My First Backup from the Help menu.

Log in to Arcserve Backup

When you open the Arcserve Backup Manager Console, you must log in to Arcserve Backup. The first time you log in to Arcserve Backup, you can log in as caroot, which has administrator privileges, and provide the appropriate password in the password field. Optionally, you can log in to Arcserve Backup using the Windows account that was provided when you installed Arcserve Backup, or with any Windows administrative account associated with the computer that you are logging in to.

After you log in, you can change the password for the caroot user and add new users. You can also add new users using the command line utility, `ca_auth.exe`. For information about `ca_auth.exe`, see the *Command Line Reference Guide*.

Note: The caroot password can consist of any combination of alphanumeric and special characters, but may not exceed 15 bytes. A password totaling 15 bytes equates to approximately 7 to 15 characters.

To log in to Arcserve Backup

1. Open the Arcserve Backup Manager Console.

To open the Manager Console, click Start on the toolbar, select Programs, CA, ARCserve Backup, and click Manager.

The Default Server Information page appears.

- To change the default server or specify a different server, select a server from the Arcserve Backup Primary Server list. If the target server does not appear in the drop-down list, you can input the host name or IP address of the server in the Arcserve Backup Primary Server list.

- To change the user, select either Arcserve Backup Authentication or Windows Authentication and specify a user name and password.

By default, Arcserve Backup does not remember your security information. To save the user name and password information that you entered for this server, you must explicitly select Remember the security information. If you do not save this information, Arcserve Backup prompts you to provide Arcserve Backup security credentials the first time you open managers, wizards, and so on, and you must provide a Arcserve Backup user name and password.

- Enter caroot in the User Name field, the appropriate password in the Password field, and click OK.

The first time you log in to Arcserve Backup, a tutorial, called My First Backup, appears. This tutorial lets you become familiar with the basics of backing up and restoring data in a controlled and directed way. This tutorial appears automatically only the first time you log in. However, you can access My First Backup from the Help menu.

Add a Windows User

Before users can be assigned a role, you must add them to the Arcserve Backup database. A Windows user can log in using their standard Windows log in.

To add a Windows user

- From the Arcserve Backup Manager, select Administration, User Profile.

The User Profile manager opens.

2. Click the Add User button on the toolbar.
The Add User Dialog opens.
3. On the General tab, select Windows Authentication and enter the following information:
 - **Username**--Enter the new user name in the following format: domain\username. Alternatively, you can click the Search button for a list of available users.
 - **Description**--(Optional) In the description box, enter information about the user account.
4. Select the Roles tab, assign a role for the user.
Note: A minimum of one role must be assigned to a user when the user is created.
5. Click OK.
The Windows user is added to the Arcserve Backup database.

Be aware of the following:

- Any valid Windows user can be added to Arcserve Backup from the User Profile Manager.
- Only the Arcserve Backup Administrator with a Security Administrator role can add a user.
- Note: To make sure a Windows user with a blank password will not fail to log on to Arcserve Backup, you must configure a [Windows Security Setting Option](#) (see page 90).

Add an Arcserve Backup User

Before users can be assigned a role, you must add them to the Arcserve Backup database.

To add a Arcserve Backup user

1. From the Arcserve Backup Manager, select Administration, User Profile.
The User Profile manager opens.
2. Click the Add User button on the toolbar.
The Add User Dialog opens.

3. On the General tab, select Arcserve Backup Authentication and enter the following information:
 - **Username**--Enter the new user name. You cannot use the "\" character.
 - **Password**--You must enter and confirm the password.
 - **Description**--(Optional) In the description box, enter information about the user account.
4. Select the Roles tab, assign a role for the user.

Note: A minimum of one role must be assigned to a user when the user is created.
5. Click OK.

The Windows user is added to the Arcserve Backup database.

Be aware of the following:

- Any valid Windows user can be added to Arcserve Backup from the User Profile Manager.
- Only the Arcserve Backup Administrator with a Security Administrator role can add a user.

Change Passwords from the Home Page

All users can change their password from the Arcserve Backup Home Page or from the User Profile Manager.

To change passwords from the Home Page

1. Open the Arcserve Backup Manager.

The Arcserve Backup Home Page GUI opens.
2. From the menu, select File and Change Password.

The Change User Password dialog opens.
3. Enter the password fields provided and click OK.

If the password credentials are met, your password is successfully changed and a message dialog displays.
4. Click OK on the message dialog and restart all Arcserve Backup Managers that are connected to the server.

Modify Windows User Properties

When necessary, you can change a user's assigned roles or suspend or activate the user account.

To modify Windows user properties

1. From the Arcserve Backup Manager, select Administration, User Profile.
The User Profile manager opens.
2. Click the Properties button on the toolbar.
The User Properties Dialog opens.
3. On the General tab, choose from the following properties:
 - **Description**--Lets you add information in the description field about the user.
 - **Status**--Lets you specify if the user account is active or suspended.
4. On the Role tab, add or delete the roles assigned to the user.

Modify Arcserve Backup User Properties

When necessary, you can change a user's assigned roles or suspend or activate the user account.

To modify Arcserve Backup user properties

1. From the Arcserve Backup Manager, select Administration, User Profile.
The User Profile manager opens.
2. Click the Properties button on the toolbar.
The User Properties Dialog opens.
3. On the General tab, choose from the following properties:
 - **Description**--Add information in the description field about the user.
 - **Status**--Specify if the user account is active or suspended.
 - **Password**--Specify a new password.
 - **Confirm Password**--Re-enter the new password.
4. On the Role tab, add or delete the roles assigned to the user.

Delete a User

Use the following steps when you want to delete a user from Arcserve Backup.

To delete a user

1. From the Arcserve Backup Manager, select Administration, User Profile.
The User Profile manager opens.
2. Select the user you want to delete and click the Delete button on the toolbar.
The delete confirmation box appears.
3. Select Yes.
The user is deleted

Be aware of the following:

- Only the Arcserve Backup Administrator with a Security Administrator role can delete a user.
- The Arcserve Backup user caroot cannot be deleted.
- You cannot delete the current user.

Add a User to a Role

You can add users to a particular role.

Note: Before a user can be assigned a role, you must add them to the Arcserve Backup domain.

To add a user to a role

1. From the Arcserve Backup Manager, select Administration, User Profile.
The User Profile manager opens.
2. From the Security tree, select Roles.
3. Highlight a role and click the Properties button on the toolbar.
The Role Properties dialog opens.
4. Select the Users tab and click Add.
The Select Users dialog opens.
5. Select a user and click OK. Alternately, you can double click a user.
The Role Properties dialog re-opens.
6. Click OK.
The user is added to the role.

Remove a User from a Role

You can remove users from a particular role.

To remove a user from a role

1. From the Arcserve Backup Manager, select Administration, User Profile.
The User Profile manager opens.
2. From the Security tree, select Roles.
3. Highlight a role and click the Properties button on the toolbar.
The Role Properties dialog opens.
4. Select the Users tab.
The Select Users dialog opens.
5. Select a user and click and click Remove.
6. Click OK.
The user is deleted from the role.

Be aware of the following:

- You cannot remove a user that is assigned only one role.
- The caroot user cannot be removed from the user list.

Using the Audit Log

The Audit Log maintains a log of critical Arcserve Backup operations. For example, user login and logout information, adding a job, deleting a job, and so on.

This section contains the following topics:

- [Filter the Audit Log](#) (see page 99)
- [View the Audit Log](#) (see page 99)
- [View an Audit Log Record](#) (see page 100)
- [Copy Audit Log Records](#) (see page 100)
- [Export an Audit Log](#) (see page 100)
- [Print an Audit Log](#) (see page 101)
- [Delete Audit Log](#) (see page 101)
- [Configure System Event Log](#) (see page 102)


Filter the Audit Log

Arcserve Backup lets you use filter options to refine your Audit Log search.

To filter the Audit Log

1. Open the Job Status Manager.

Select the Audit Log tab.

Expand the header bar by clicking .

Choose from the following filter options:

- **Event Types**--Specify the type of event to sort by. The options are Success audit and Failure audit.
- **Source Machine**--Specify the machine to audit. The default is All.
- **Event**--Specify what event you want to view. You can choose a specific user task or All user tasks.

Note: The header bar will turn yellow if there is a change made to the type of event you want to view, indicating that the advanced filter was used.

- **User**--Specify the user whose audit log you want to view. The default is All.

Note: The header bar will turn yellow if a user is specified, indicating that the advanced filter was used.

- **Source Process**--Specify a specific Arcserve Backup process. The default is All.
- **From**--Specify the start day and time of an event. The option includes First Event and Events On.
- **To**--Specify the end day and time of an event. The option includes First Event and Events On.

2. Click Update.

The filtered results are displayed in the Properties panel.

View the Audit Log

Arcserve Backup lets you use the Audit Log to identify patterns or areas of repeated tasks.

To view the audit log

1. Open the Job Status Manager.
2. Select the Audit Log tab.

The Properties panel displays the audit log.

View an Audit Log Record

Arcserve Backup lets you view the details of a specific audit log.

To view an audit log record

1. Open the Job Status Manager.
2. Select the Audit Log tab.
3. Select the audit record you wish to view.
4. Right click and select Properties. Alternately, double click the record.
The Audit Record Properties dialog opens.
5. On the Audit Record Properties dialog the following options are available:
 - **Prev**--Navigate to the previous audit record.
 - **Next**--Navigate to the next audit record.
 - **Copy**--Copy all the audit record properties to the clipboard.
6. Click OK.

The Audit Record Properties dialog closes.

Copy Audit Log Records

Arcserve Backup lets you copy the audit log records in list format to the clipboard.

To copy audit log records

1. Open the Job Status Manager.
2. Select the Audit Log tab.
The Properties panel displays the audit log.
3. Select the audit log records you want to copy.
4. Press Ctrl+C to copy the records to the clipboard.
5. Open a text editing application and past the copied records.

Export an Audit Log

Arcserve Backup lets you export all audit records or selected audit records to a text file.

To export an audit log

1. Open the Job Status Manager.
2. From the domains directory tree, click the Arcserve Backup primary server containing the audit log that you want to export.

The Job Queue, Job History, Activity Log, and Audit Log tabs appear.

3. Click the Audit Log tab.
The Properties panel displays the audit log.
4. Select the audit log that you want to export.
Right-click the audit log that you want to export and click Export to a File on the pop-up menu.
The Save As dialog opens.
5. In the Save As dialog, specify a location and file name for the file.
Click Save.
The audit log is exported to a text file.

Print an Audit Log

Arcserve Backup lets you print the Audit log to a local printer.

To print an audit log

1. Open the Job Status Manager.
2. From the domains directory tree, click the Arcserve Backup primary server containing the audit log that you want to print.
The Job Queue, Job History, Activity Log, and Audit Log tabs appear.
3. Click the Audit Log tab.
The Properties panel displays the audit log.
4. Select the audit log that you want to print.
Right-click the audit log that you want to print and click Print on the pop-up menu.
The Print dialog opens.
5. Click OK.
The Audit Log is printed.

Delete Audit Log

Arcserve Backup lets you delete the audit log.

To delete the audit log

1. Open the Job Status Manager.
2. Select the Audit Log tab.
The Properties panel displays the audit log.
3. Click the Delete button on the toolbar.
The Delete Audit Log dialog opens.

4. Select delete options.

Choose from the following delete options:

- **Entire log table**--Delete all records in the Audit Log.
- **Partial**--Delete records in a specified time period.
- **Older than**--Enter a specific time based on the following criteria: 1 to 365 days, 1 to 54 weeks, 1 to 12 months, and 1 to 10 years.

5. Click OK.

The Audit Log records are deleted.

Configure System Event Log

Arcserve Backup lets you add the Audit Log information to the system event log.

To configure the system event log

1. From the Sever Admin, select the Configuration button on the toolbar.

The Configuration dialog opens.

2. Select the Log tab.

3. Select the Enable messages logging into Windows Event Log checkbox.

The Audit Log information is included in the Windows Event Log.

More information:

[Event Log Configuration \(Windows Servers\)](#) (see page 508)

Create an Audit Log Report

Arcserve Backup lets you generate an Audit Log report from the Report Manager.

More information:


[Generate Reports Using Report Manager](#) (see page 669)

How Arcserve Backup Processes Backup Data Using Multistreaming

Note: To process two or more streams of backup data using multistreaming, you must license the Arcserve Backup Enterprise Module.

Multistreaming is a process that divides your backup jobs into multiple sub-jobs (streams) that run simultaneously and sends data to the destination media (tape device or file system device). Multistreaming is used to maximize the effective use of the client machines during backup and recovery operations. Multistreaming is useful when performing large backup jobs, since it is more efficient to divide multiple jobs between multiple backup devices.

Multistreaming lets you use all of the available tape devices on the system by splitting your backup jobs into multiple jobs using all available tape devices. As a result, it will increase the overall backup throughput compared with the sequential method.



You can use all of the devices or you can specify a single group of devices. If the Arcserve Backup Tape Library Option is installed and the group with the library is selected, multistreaming uses all library devices. If the Arcserve Backup Tape Library Option is not installed, you can put devices into separate groups. For a changer, the total number of streams (child jobs) that are created depends on the number of tape devices. For a single tape drive device, the total number of streams depends on the number of device groups.

Multistreaming is performed at the volume level for regular files (two volumes can run simultaneously on two separate devices), and at the database level for local database servers. Multistreaming is performed at the node level for the Preferred Shares folder, remote database servers, and Windows Client Agents.

You can have only as many jobs running simultaneously as the number of devices or groups that are on the system. With multistreaming, one parent job is created that will trigger child jobs for as many volumes as you have. When a job is finished on one device, another job is executed until there are no more jobs to run.

Some characteristics and requirements of multistreaming are as follows:

- Each client machine can have multiple source streams, depending on the number of agents being backed up.
- Each agent can have a separate stream (one stream per agent).
- Multistreaming always requires a media pool selection to prevent the tapes from being overwritten.
- Separate tape devices should be configured in separate groups for regular drives, however for changers, they can be configured to be in the same group.
- Canceling the parent job cancels all of the child jobs. For Windows, canceling and monitoring is checked between jobs for performance considerations.
- If a job spawns child jobs, the number of child jobs spawned will not exceed the number of streams specified for the job. However, if a job spawns child jobs and you do not specify a number of streams to use, the child jobs will be created and backed up in one continuous stream.
- In the Job Status Manager, each child job has a default job description with this pattern:

```
JOB[ID][Servename](Multistream subjob [SID])[Status][Start time - End time][JOB No.]
```

Note: SID Represents the sub job (child) ID.
- The multistreaming option is ignored if the groups you choose have only one device, or if only one object (volume, database, or remote node) backup is submitted.

Be aware of the following:

- Backing up data using multistreaming to data deduplication devices can have an adverse affect on Tape Engine performance. For information about how you can remedy this problem, see [Increase Virtual Memory Allocation to Improve Tape Engine Performance](#) (see page 513).
- You should use the same types of tape devices for multistreaming jobs. In order to achieve the optimum performance with your multistreaming jobs, you should use a high-end server machine with multiple processors and at least 256 MB memory per processor.

Tasks Supported by Multistreaming

The following table describes tasks that are supported and not supported by multistreaming.

Supported	Not Supported
<ul style="list-style-type: none"> ■ Submitting rotation and GFS jobs using multistreaming. ■ Backing up data using pre and post operations and comments are supported at the parent job level. ■ Backing up Microsoft SQL Server data, Microsoft Exchange Server data, and Oracle RMAN data to a library. 	<ul style="list-style-type: none"> ■ Backing up data using the Image Option to libraries. ■ Backing up data using pre and post operations are not performed for child jobs. ■ Backing up Lotus Domino data to a library.
<p>Note: Local SQL Server data is backed up at the database level and remote SQL Server data is backed up at the instance level.</p>	

Multistreaming Support for Local Backup Jobs

The following considerations apply to multistreaming support for local backup jobs:

- If a local backup does *not* contain any VSS writers, the job can be a multistreaming job.
- If a local backup contains VSS writers, and *both* the following global options are unchecked, then the job can be a multistreaming job:
 - Writers and Components/Files included by Writer will be excluded from file system backups.
 - Writers and Components/Files excluded by Writer will be excluded from file system backups.
- If a local backup contains VSS writers, and *any* of the following global options are checked, then the job *cannot* be a multistreaming job:
 - Writers and Components/Files included by Writer will be excluded from file system backups.
 - Writers and Components/Files excluded by Writer will be excluded from file system backups, and following message is displayed in the activity log: The local node contains VSS writers, disk level multistreaming will be disabled.

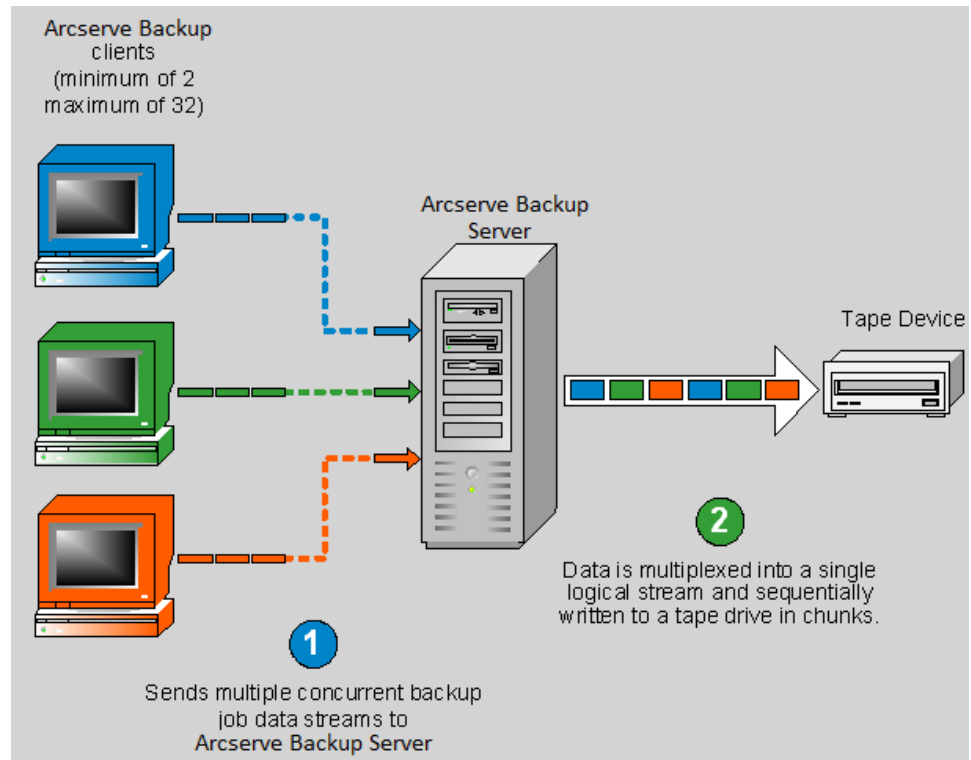
To enable multistreaming for local VSS backup, uncheck default global options for "Writers and Components".
- If a local backup contains VSS writers, and the registry key HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA Arcserve Backup\Base\Task\Backup\ForciblyMUSForLocalVSSBackup is set to 1, the global options for "Writers and Components" is ignored. The job can be a multistreaming job.

How Arcserve Backup Processes Backup Data Using Multiplexing

Multiplexing is a process in which data from multiple sources is written to the same media (tapes) simultaneously. Multiplexing is used to maximize the effective use of tape drives and libraries during backup and recovery operations and is useful when the tape drive is much faster than the backup source. Multiplexing keeps the backup hardware running at its maximum capability for the entire length of the backup process. A session included in a multiplexing backup should not be impacted by the speed of other sessions being multiplexed. The only factor that can limit the speed of a backup session is the speed of the hardware device.

The maximum number of jobs that you can multiplex is limited by the amount of available memory. The default number of jobs that you can multiplex is 4, and the minimum number is 2 while the maximum number is 32.

When a job that has multiple sources is submitted with the multiplexing option enabled, it is broken into child jobs with one for each source. These child jobs write data to the same media simultaneously. The number of child jobs spawned will, at most, be equal to the number of streams specified for multiplexing. However, if a job spawns multiple child jobs and the value specified for the Multiplexing Max # of Streams option is one, the child jobs will be created and backed up in one continuous stream (the default Max # Stream is 4).



Note: When using multiplexing, you can select the maximum number of streams that can write to a tape at the same time. For more information, see [Specify Multiplexing Options](#) (see page 190).

Multiplexing is useful when your tape drive throughput is faster than the rate at which data can be extracted from the source. Factors that can affect backup throughput are as follows:

- The kind of data being backed up. For example, backing up large number of small files reduces backup throughput because of the larger number of necessary file system operations (file open and close).
- Some databases may be inherently slow in providing data.
- The network throughput of the server being backed up.
- The disk performance on which the data resides.
- The server resources like CPU speed, memory size, page file size, network card, and amount of other activities on the server.
- Network backups that involve hundreds of servers.

When data is backed up over the network from multiple sources, most of the previous factors are involved, which reduces the throughput and increases the amount of time it takes to perform a backup. In addition, if the tape drive is not consistently streamed, the life of the tape drive is reduced drastically because of the "shoe shine" effect: when data is written intermittently, the drive has to stop, and then go back and forth on the media to adjust to the new position from where it has to write again. With multiplexing, data is continuously available and tape drives are constantly streaming. This behavior decreases the amount of time it takes to perform a backup while increasing the life of the hardware.

Multiplexing is performed at the volume level for regular files, two volumes can run simultaneously as two separate child jobs, and at the database level for local database servers. Multiplexing is performed at the node level for the Preferred Shares folder, remote database servers, and Windows Client Agents.

In the Job Status Manager, each child job has a default job description with this pattern:

```
JOB[ID][ServerName](Multiplexing subjob [SID])[Status][Start time - End time][JOB No.]
```

Note: SID Represents the sub job (child) ID.

Tasks Supported by Multiplexing

The following table describes tasks that are supported and not supported by multiplexing.

Supported	Not Supported
<ul style="list-style-type: none"> ■ Backing up Microsoft SQL Server data, Microsoft Exchange Server data, and Oracle RMAN data to a library. <p>Note: Local SQL Server data is backed up at the database level and remote SQL Server data is backed up at the instance level.</p> <ul style="list-style-type: none"> ■ Multiple jobs can write to the same tape drive. ■ Single session restore from multiplexing tapes. ■ QFA restore from multiplexing tapes. ■ Merge from multiplexing tapes. ■ Disaster recovery. ■ Session consolidation from a multiplexing tape to a non-multiplexing tape. ■ Scan and compare on multiplexing tapes. 	<ul style="list-style-type: none"> ■ Backing up Lotus Domino data to a library. ■ Backing up data using the Image Option to libraries. ■ Multiple restores simultaneously from a single multiplexing tape. ■ Multiple session consolidation simultaneously from a single multiplexing tape to multiple non-multiplexing tapes. ■ The Verify after Backup option. ■ Disk staging during multiplexing. ■ Multiplexing jobs cannot be submitted to NAS devices, File System Devices, RAID devices, and WORM media. ■ Multiplexing jobs cannot be submitted to a non-multiplexing media. ■ Multiplexing is not supported for NAS sources.

How Arcserve Backup Secures Data

Data security is the process of protecting sensitive information from unauthorized access or use. Data security helps you to ensure privacy and protect personal data. Arcserve Backup ensures that all sensitive data that is stored in a computer or on removable media cannot be read or compromised by any individuals without proper authorization.

Often these removable media can contain highly sensitive information which could be lost while in transit between company data centers and their off site storage vaulting service facilities. The data on these media needs to remain secure, even when in transit.

Encryption and Decryption

Most security measures involve data encryption and passwords. Data encryption is the translation of data into a form that is unintelligible without a deciphering mechanism. Decryption is the decoding or conversion of encrypted data into plain text and reversing the encryption process.

The Arcserve Backup data protection solution uses secure, industry standard encryption algorithms in various components to achieve the maximum security and privacy of customer data. Starting with Arcserve Backup r12, the Windows client agents will use a 256-bit AES algorithm provided in the RSA BSAFE cryptographic library for all encryption purposes. Any data collected by earlier versions of Arcserve Backup agents will use either a 168-bit 3DES or a proprietary Arcserve encryption algorithm for encryption purposes. In addition, the Windows base product also uses the same 256-bit AES algorithm to store any sensitive information on the Arcserve Backup server.

The AES (Advanced Encryption Standard) feature has been developed to replace the DES (Data Encryption Standard) and is designed to be more secure than DES. The AES is a FIPS-approved cryptographic algorithm that can be used to protect electronic data. The AES algorithm is a symmetric block cipher that can encrypt and decrypt information.

Note: Encryption and compression are not supported on deduplication devices.

Federal Information Processing Standards (FIPS)

The Federal Information Processing Standards (FIPS) are a set of standards that describe document processing, provide standard algorithms for searching, and provide other information processing standards for use within government agencies. The National Institute of Standards and Technology (NIST) issued the 140 Publication Series to coordinate the requirements and standards for cryptographic modules which include both hardware and software components for use by departments and agencies of the United States federal government.

Security Requirements for Cryptographic Module (FIPS 140-2) specifies the security requirements that will be satisfied by a cryptographic module utilized within a security system protecting sensitive but unclassified information within computer and telecommunication systems.

Arcserve Backup and FIPS Compliance

Arcserve Backup uses FIPS-compliant algorithms for backing up and restoring sensitive information such as username and password credentials.

- If you choose to encrypt your data during backup to a disk or tape, the algorithms used to encrypt this data will be FIPS-compliant.
- During backup time, the username and password will be sent to the Arcserve Backup server agent (running on the server to be protected). This username and password will be encrypted using FIPS-compliant algorithms and transferred to the agent.
- Arcserve Backup also supports tape drives (from external third-party vendors) that provide FIPS-compliant hardware encryption. This is in addition to FIPS-compliant tape or disk encryption provided by the Arcserve Backup software.
- Arcserve Backup provides additional agents and options that also use FIPS-compliant algorithms to support data encryption. These agents and options include: Agent for Microsoft Exchange Server, Agent for Microsoft SQL Server, Agent for Microsoft SharePoint Server, and Arcserve Replication.

Change the Current Encryption Algorithm

The current encryption algorithm (AES256) used for Arcserve Backup can be changed by modifying the `CryptoConfig.cfg` file. This file includes a list of all supported encryption algorithms for Arcserve Backup products installed on your machine. You can change the current encryption algorithm to any of the alternate candidate algorithm values that are listed. This change will affect all the Arcserve Backup products (base, agents, and options) that are installed on that machine.

To change the current encryption algorithm

1. Run the `cstop.bat` script to stop all services before making the change.
`ProgramFiles\CA\ARCserve Backup\cstop.bat`
2. Change the current encryption algorithm value to one of the candidate values.
`ProgramFiles\CA\SharedComponents\ARCserve Backup\CryptoConfig.cfg`
3. Run `Configencr.exe` to transfer the encrypted repositories to use the new encryption algorithm.
`ProgramFiles\CA\ARCserve Backup\Configencr.exe`
4. Run the `cstart.bat` script to start all services after making the change.
`ProgramFiles\CA\ARCserve Backup\cstart.bat`

Arcserve Backup Data Encryption

Arcserve Backup provides the flexibility to use encryption to protect sensitive data during various stages of the backup process. Generally, during the backup process, the sooner the data encryption occurs, the more secure your information will be. However, speed, performance, and scheduling restrictions are also factors to consider when choosing the best approach to securing your data.

The three different ways to encrypt data in a backup job are:

- Encryption at the agent server (or source) prior to the backup process
- Encryption at the Arcserve Backup server during the backup process
- Encryption at the Arcserve Backup server during the migration process (for a staging job)

These encryption options are accessible from the Encryption/Compression tab on the Global Options dialog for the Backup Manager. From this dialog you can choose to encrypt the data at the agent, at the backup server (during backup), or at the backup server (during migration).

You can also create a session encryption password that is saved to the Arcserve Backup database. This password is used to encrypt session data. For more information about passwords, see the topic [How Password Management Works](#) (see page 79).

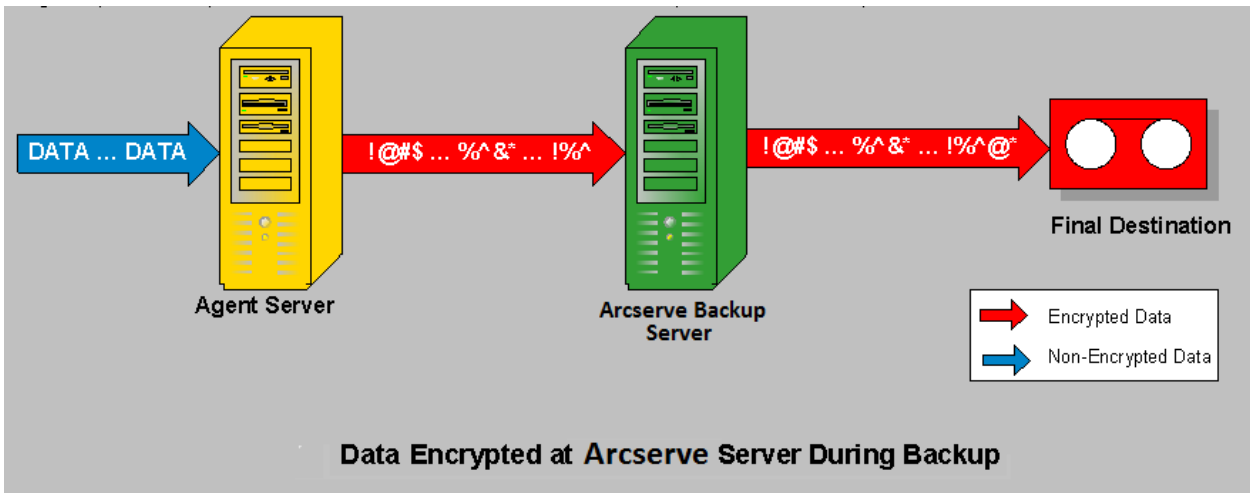
Note: Arcserve Backup will only encrypt data that is not already encrypted. If at any stage in the process Arcserve Backup detects that the data has already been encrypted, it will not attempt to encrypt it again. Since data deduplication is a form of encryption, you cannot encrypt data saved to a deduplication device.

In addition, there are also two basic methods for encrypting data; hardware encryption and software encryption. The advantages of hardware encryption are speed and improved CPU performance. Encryption using software is slower than encryption using hardware and can result in a larger backup window. By using hardware encryption, you can also avoid unnecessary CPU cycles on either the agent server or the backup server and the drive can compress the data before encrypting.

If you select to have your data encrypted during the backup or migration process, Arcserve Backup has the ability to detect if the final destination media (tape) is capable of hardware encryption and by default will automatically choose that hardware method if available.

How Arcserve Backup Encrypts Data at the Agent Server

Data can be encrypted at the Arcserve Backup agent server (agent server), prior to the actual backup process. The advantage of this method is that it does not transfer un-encrypted data from one location to another at all. However, this method puts added CPU cycles for encrypting the data on the agent server.



Not all Arcserve Backup agents have the capability to encrypt data prior to transferring it to the Arcserve Backup server.

The following Arcserve Backup agents support at the agent server data encryption:

- All Arcserve Backup file system agents
- Arcserve Backup Agent for Microsoft Exchange Server
- Arcserve Backup Agent for Microsoft SQL Server

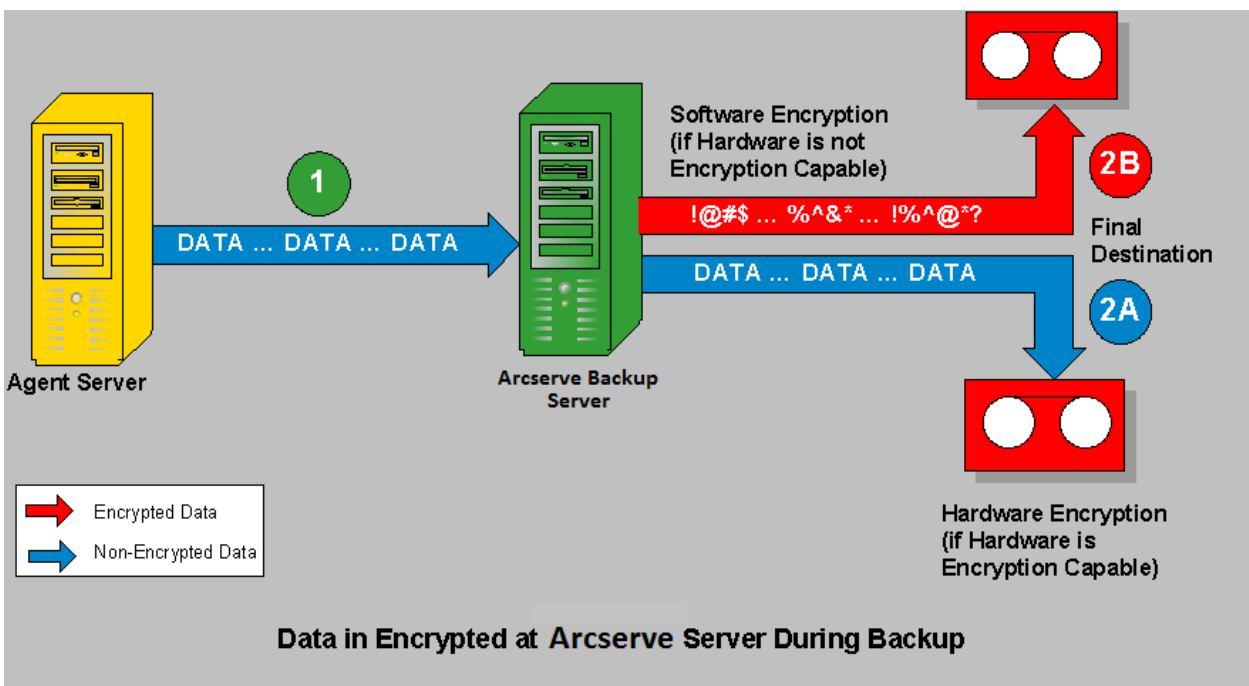
The following Arcserve Backup agents do not support at the agent server data encryption:

- Arcserve Backup Agent for IBM Informix
- Arcserve Backup Agent for Lotus Domino
- Arcserve Backup Agent for Microsoft SharePoint Server
- Arcserve Backup Agent for Oracle
- Arcserve Backup Agent for SAP R3 for Oracle

How Arcserve Backup Encrypts Data During Backups

Data can be encrypted at the Arcserve Backup server, during the backup process. Using this method, un-encrypted data is transferred from the agent server to the Arcserve Backup server. Arcserve Backup will then detect if the final destination media is capable of hardware encryption or not. If it is hardware encryption capable, then the un-encrypted data is transferred to the final destination media where it is then encrypted. This is the preferred and default method because it is faster and does not interfere with the backup window.

If Arcserve Backup detects that the final destination media is not capable of hardware encryption, Arcserve Backup uses software encryption to encrypt the data prior to transferring it to the final destination media.

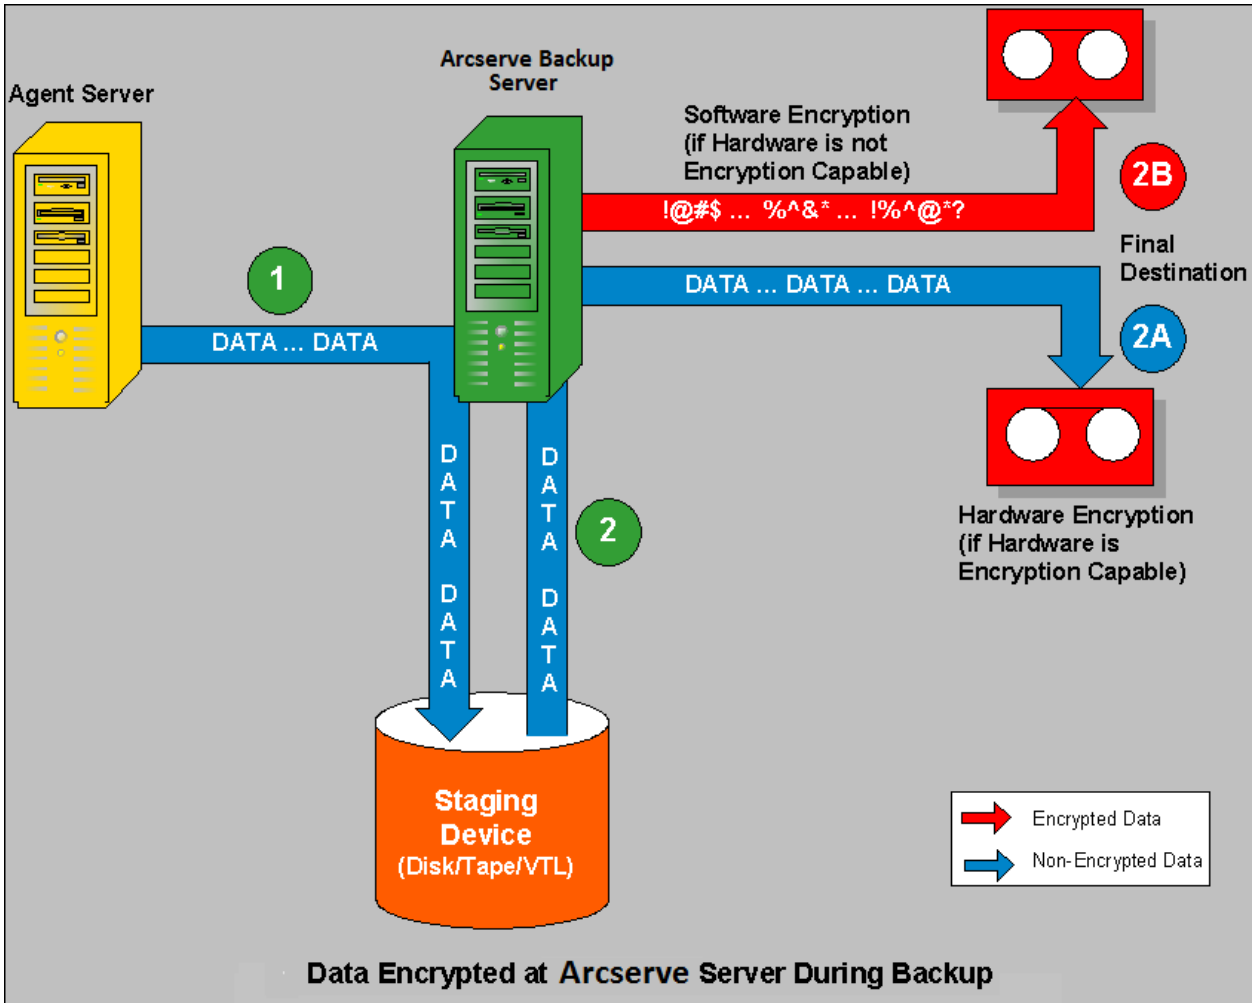


How Arcserve Backup Encrypts Data During Data Migration

Data can be encrypted at the Arcserve Backup server, during the migration process of a staging job.

Using this method, un-encrypted data is transferred during the backup process of a staging job from the agent server through the Arcserve Backup server to the staging device. The staging device can either be a disk, tape, or virtual tape library (VTL). When the data is ready for the migration process, Arcserve Backup will then detect if the final destination media is capable of hardware encryption or not. If it is hardware encryption capable, then the un-encrypted data is transferred from the staging device to the final destination media where it is then encrypted. This is the preferred and default method because it is faster and does not interfere with the migration window.

If Arcserve Backup detects that the final destination media is not capable of hardware encryption, it will then perform software encryption of the data prior to migration to the final destination media.



Effective Media Management

Effective media management provides valuable preparation for reliable backup and recovery performance. The type of media can be most types of SCSI or Fibre-attached removable storage.

Because functions such as tracking files to specific storage media are important requirements of your organization's daily production routine, effective media management requires that you know the contents and location of all removable media, such as magnetic tape. Arcserve Backup allows you to track your media through the Device Wizard and the Device Manager. Both the Device Manager and the Device Wizard allow you to manage and track your media easily.

Configure Devices Using the Device Wizard

You can start the Device Wizard from the Wizards menu. The Device Wizard helps you see all of the devices connected to your machine.

To configure devices using the Device Wizard

1. From the Administration menu in the Navigation Bar on the Home Page, click Device Wizard.

The Device Wizard Welcome screen appears.

2. Click Next.

The Login dialog appears.

3. Enter or select the server you want the device command to operate on, enter your user name and password, and click Next.
4. Select the device you want to target. Click More Information to view more information about the device.
5. Click OK, and click Next.
6. Select a device operation, and click Next.

Example: Select Format.


7. Enter a new media name and expiration date for the media Arcserve Backup is about to format, and click Next.
8. The schedule screen that appears lets you choose to run the device command immediately or schedule it for a later date and time. Select Run Now, and click Next to run the job immediately.

To schedule your job for a later time, select the Schedule option, and enter a date and time for the job to run.

9. Click Finish to execute the job.
10. You are prompted to confirm the action you are about to take. Click OK to start the device operation and display its status.
11. A message appears to notify you that Arcserve Backup has completed the device operation. Click Next to work with another device, or click Exit to close the Device Wizard.

Configure Device Groups

The Device Manager provides you with information about standalone Tape Drives, on the right side of the Device Manager window.



If you have more than one storage device connected to your network machine, Arcserve Backup lets you group the devices. This allows you to have one group perform a backup, while another group performs a restore operation, in a process known as parallel streaming.

If you have several devices in a group, and your job spans more than one media, the Device Manager can automatically span the media for you. You can then submit large backup jobs to Arcserve Backup and automatically span multiple media until the jobs are complete.

For example, if you have two media groups, GROUP1 (consisting of one storage device) and GROUP2 (consisting of two storage devices), and you have a large backup job that requires more than one media, you can insert blank (formatted) media into each GROUP2 drive and Arcserve Backup automates the media spanning for you. Without media spanning, you must change the media manually.

Note: For deduplication, device groups can contain only one deduplication device.

To configure device groups

1. In the Device Manager, click Configure Groups to open the Device Group Configuration dialog.
2. To assign a device to a new group, highlight it, and click Remove.
3. Click New to create a new group.
4. Enter a name for the new group, and click OK. The new group appears in the Groups field.
5. Highlight both the device and the new group, and click Assign to assign the device to the new group.
6. Click OK.

Back Up and Restore Data

Backing up and restoring your data is essential to the success of your organization. By efficiently and dependably protecting and retrieving files, Arcserve Backup helps you to ensure that your most valuable asset, your data, is protected.

This section introduces you to the essential Arcserve Backup functions of backing up and restoring data.

Backup Requirements Plan

Before you use Arcserve Backup for the first time, we recommend that you plan your backup requirements. You should consider the following:

- How much data do you need to back up?
 - What is the current disk capacity in your environment?
 - What is the server and data growth you anticipate over the next year?
- How do you want to manage the media you are using for backup?
- How do you plan to store your data? Are you using magnetic tape or does the stability of WORM media better suit your needs?

Add Computers to the Preferred Shares/Machines Tree

The Preferred Shares tree consists of a collection of your favorite backup shares. A share is a shared drive, directory, or entire system. You can manually add individual share points to the Preferred Shares tree; the share point is remembered and displayed regardless of the status of the network connection. This provides a quick access to commonly used shares on your machines. You can also set up preferred machines, which enables you to browse, backup, or restore all of the shared drives on a machine under a single machine.

When you set up a backup job, you must log in to and provide valid credentials on the preferred system to submit the job.

Note: Arcserve Backup does not support logging in to systems with passwords that are greater than 23 characters. If the password on the system you are attempting to log in to is greater than 23 characters, you must modify the password on the agent system such that it is 23 characters or less, and then you can log in to the agent system.

You must input the domain name as part of the user name. Otherwise, the preferred share job may fail because of invalid credentials with the following message:

W3301 Unable to find directory. (DIR=directory, EC=Logon failure: unknown user name or bad password)

To add computers to the Preferred Shares/Machines tree

1. From the Source tab on the Backup Manager window, right-click the Preferred Share/Machines object and select Add Object from the pop-up menu

The Add Preferred Shares dialog opens.

2. Select a Network Provider.

Enter a share name in Uniform Naming Convention (UNC) format.

Example: \\MACHINE\SHARE

Note: Arcserve Backup server names and Arcserve Backup domain names cannot exceed 15 bytes. A name totaling 15 bytes equates to approximately 7 to 15 characters.

Click Add.

The computer is added to the Preferred Shares tree.

3. To add more computers, repeat the previous step.
4. When you are finished adding computers, click Close.

Backup Media Rotations and Scheduling Options

Typically, the most convenient time for you to schedule backup operations is after business hours, when backup processing does not use valuable network bandwidth. Arcserve Backup provides you with the tools you need to automate your backup operations.

Arcserve Backup allows you to establish a schedule so that your backup automatically repeats at regular intervals, allowing you to regularly and reliably back up your data at any time. The Backup Manager provides you with scheduling options and rotation schemes to help you establish your automatic backup strategy.

Note: If you are using WORM media, you cannot use rotation schemes. By definition, WORM media cannot be overwritten, so you cannot recycle it in a rotation scheme or a media pool.

Types of Rotation Schemes

You can configure backup jobs using custom schedules, using the pre-defined rotation schemes provided by Arcserve Backup, or specifying your own rotation parameters. You can select a repeat method and choose from among the following three backup methods in your rotation scheme:

- **Full Backup**--Backs up all of your files. This backup method requires more time to process compared to incremental or differential backups. However, because all of your data is backed up, this strategy requires only the last backup media to restore your data completely.
- **Incremental Backup**--Backs up only those files that have changed since the last full or incremental backup was performed. Since this strategy backs up only new or newly changed files, incremental backups require less time to process. However, this strategy requires the full media set and every incremental set, including the latest set, to fully restore your data after a disaster.
- **Differential Backup**--Backs up only those files that have changed since the last full backup was performed. Since files that were backed up in the last differential job are backed up again, differential backup jobs require more time to process than incremental backup jobs. However, this strategy requires only two sets of media to restore a differential backup, the full media set, and the differential media set.

Note: For any rotation scheme that you use, you should include at least one full backup per week.

How Media Pools Work

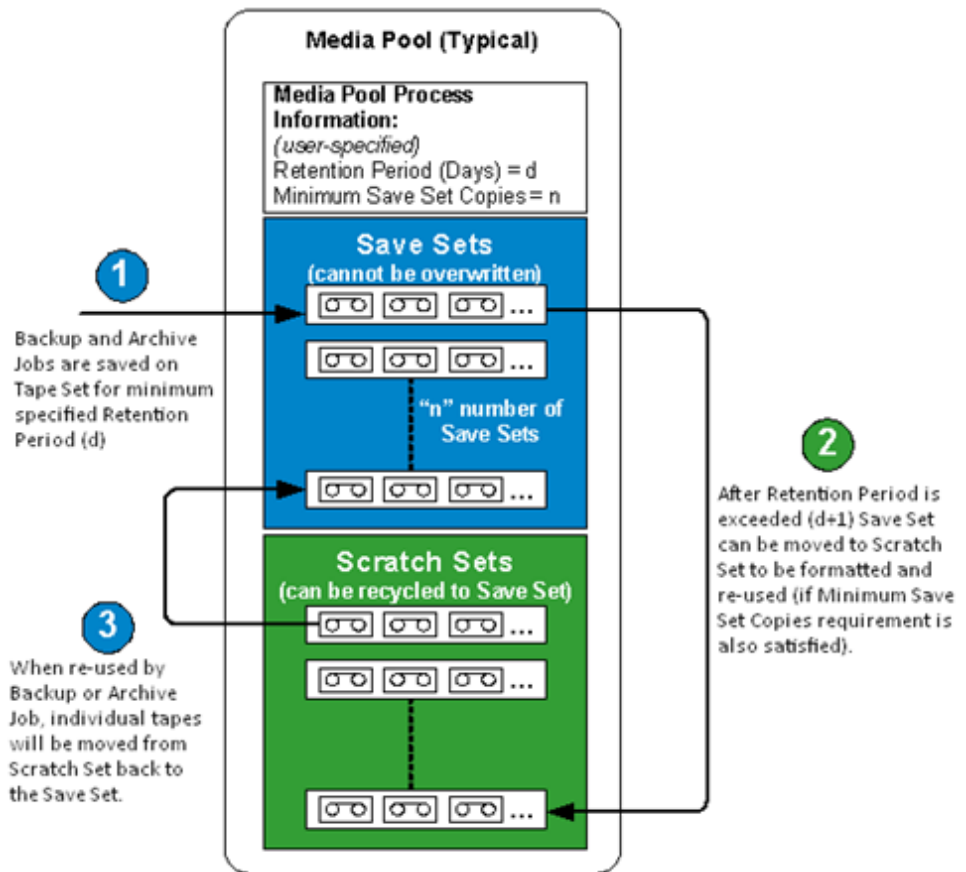
To prevent the accidental overwriting of needed data, Arcserve Backup manages media for rotation schemes in media pools. Media pools are logical collections of rewriteable, removable storage media managed as a single unit.

Important! Deduplication and Cloud devices cannot be assigned to media pools.

A media pool is a collection of backup and archive media (tapes) that is set aside for a specific job and managed as a unit. A media pool is a set of tapes that is logically grouped and used exclusively for a particular recurring backup and archive job. Within Arcserve Backup each media pool is automatically divided into a Scratch Set and a Save Set. Any media in a Save Set cannot be overwritten until certain user-specified criteria are met. This prevents the possibility of inadvertently overwriting a tape before adequate backups or archives are preserved. After the user-specified criteria is met, the Save Set becomes a Scratch Set and is recycled to be used again (overwritten).

Once the media has passed certain specified criteria, such as a minimum number of media in the Save Set and a minimum retention period, the media is moved to the Scratch Set. The retention period is the number of days media is kept in the Save Set of a media pool. When these criteria are met, the media is moved from the Save Set to the Scratch Set and is made available for use.

The Media Pool Manager lets you create and maintain the Arcserve Backup media pools. Each media pool is assigned a name, and is organized according to serial numbers. The serial numbers assigned are permanent. If you use a device with a bar code reader, the bar code labels are used as the serial number of the media. Media pools are organized by the range of serial numbers of the media they contain. Media pools apply to every media, regardless of which backup or archive type and method were selected.



How to Use GFS Rotations

The Grandfather-Father-Son (GFS) rotation strategy is a method of maintaining backups on a daily, weekly, and monthly basis. GFS backup schemes are based on a seven-day weekly schedule, beginning on any day of your choice. The primary purpose of the GFS scheme is to maintain a minimum standard and consistent interval at which to rotate and retire media. This scheme always uses the oldest media first.

You should perform a full backup at least once a week. On all other days, you can perform full or partial backups or no backup at all. The advantage of setting up a GFS rotation scheme is that once it is configured, you need only make sure the right media is in the drive for each day of the week.

From that time on, GFS tells you which media to use and manages the backups for you.

- Daily backups are the Son media
- A full backup is performed at least once a week. The last full backup of the week is the Father media
- The last full backup of the month (monthly backup) is the Grandfather media

Note: Monthly backups are saved throughout the year and the media on which they are stored should be taken off-site for safekeeping. You can track these media using the Media Management Admin.

Important! GFS rotations create three media pools--daily, weekly, and monthly pools. You cannot entirely customize this rotation and the media used for the rotation scheme must be named automatically. Custom rotation schemes allow you to configure the properties of the scheme, such as the pool or pools involved, the days to back up, and other properties. Deduplication devices are an exception: even though deduplication devices cannot be assigned to media pools, you may still set up GFS rotations. For more information, see [GFS Rotation Jobs on Deduplication Devices](#) (see page 754).

How GFS Rotations Work

The most commonly used media rotation schedule is the Grandfather-Father-Son (GFS) rotation. This schedule policy uses daily (Son), weekly (Father), and monthly (Grandfather) backup media sets (tapes). GFS rotation schedules allow you to back up your data for an entire year using a minimum of media (tapes). The number of tapes you use for GFS rotations is based on the number of workdays you specify for your backup policy.

The GFS rotation method works as follows:

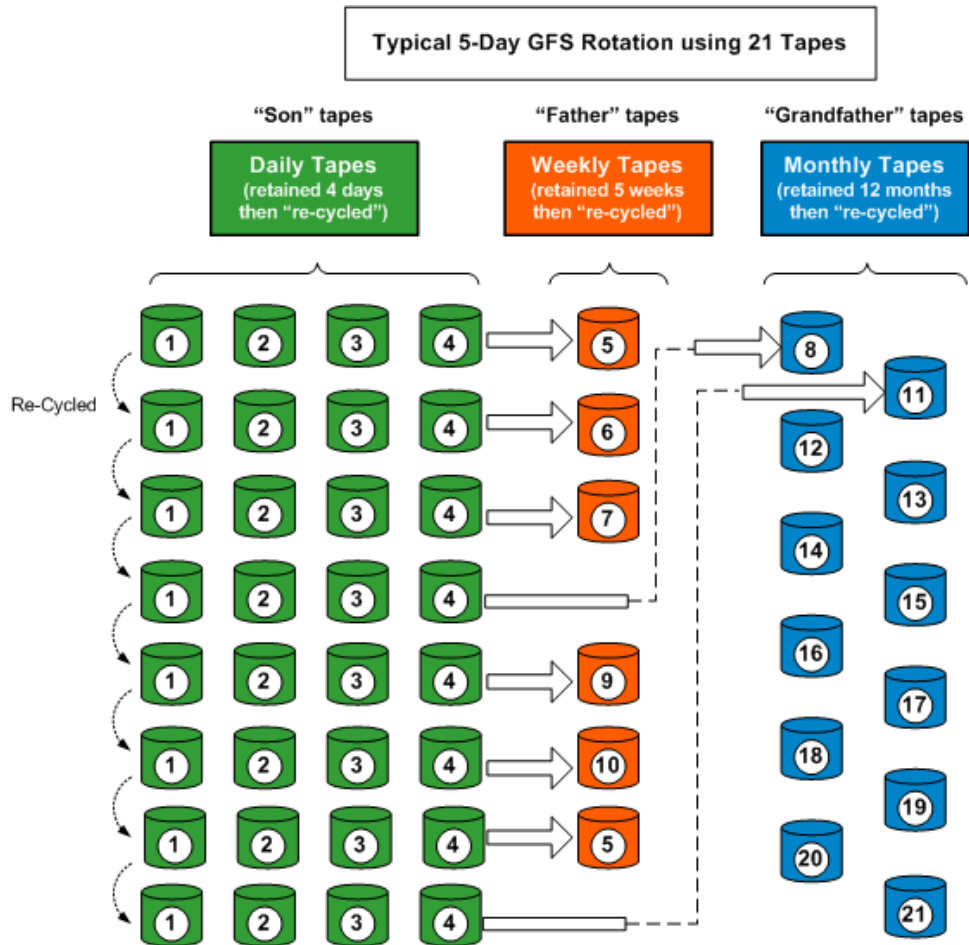
Note: To avoid confusion, it is important to clearly and properly label your tapes.

- You back up your data on a separate tape every working day. You should use a different tape for every daily backup. For example, if your backup cycle is based on a five-day workweek, you will need four "Daily" tapes before you use a weekly tape. (Maybe label the daily tapes Monday, Tuesday, Wednesday, and Thursday or Daily 1 through Daily 4, and so on.). You can perform Full, Incremental, or Differential backups for your daily backups. After the fourth day, the first daily tape used is then re-cycled and can be overwritten with the next scheduled daily backup.

Remember, because the daily tapes are used more frequently than the weekly and monthly tapes, you will need to replace them more often.

- On the fifth day, instead of using another daily tape, you will use a "Weekly" tape. You should always perform a Full backup for your weekly backups. You should also use five weekly tapes before you use a monthly tape. (Maybe label the weekly tapes Week 1 through Week 5). After the fifth week, the first weekly tape used is then re-cycled and can be overwritten with the next scheduled weekly backup.
- At the end of the third week, instead of using another weekly tape, you will use a "Monthly" tape. You should also perform a Full backup for your monthly backups. You should have 12 monthly tapes to safely backup a full year of data. (Maybe label the monthly tapes January through December or Month 1 through Month 12, etc.). After twelfth month, the first monthly tape used is then re-cycled and overwritten with the next monthly backup.

The following diagram shows an example of how a typical 5-day GFS rotation policy can be implemented to provide you with a safe and reliable method to perform data backups for an entire year while using a minimum amount of backup media:



Note: A five-day GFS rotation policy would require approximately 21 tapes per year, while a seven-day policy would require approximately 23 tapes per year (adding two additional daily tapes). For both of these schedules, the amount of media needed can vary depending upon your specified retention criteria and the quantity of data that you are backing up. Additionally, the amount of media needed in each schedule can also be affected by the use of multistreaming and if you are appending backup sessions to your media.

GFS Rotation Scheme Media Example

The following example illustrates how to determine the number of media you need for a GFS rotation scheme:

Your company's business hours are from Monday to Friday. You have specified daily incremental backups from Monday through Thursday, with a full backup on Friday. You have decided to retain monthly full backup data for six months before you recycle your media, and have specified that at least six monthly tapes are to be maintained in the Save Set of your media pool. In addition, you have specified that a minimum of four weekly tapes are to be retained in the Save Set.

Note: For more information about media pools, Save Sets, and Scratch Sets, see "Managing Devices and Media."

In the GFS rotation scheme you have selected, the incremental backups are the Son, the weekly full backups are the Father, and the monthly full backups are the Grandfather.

Your rotation scheme requires four daily incremental backups, requiring one tape for each day. Because the data these tapes contain is maintained on the weekly full backup, these tapes are recycled each week. Therefore, your scheme requires four daily (Son) tapes.

The backup performed each Friday, the weekly full backup, requires one tape for each week of the month. These tapes are retained for one month before they are recycled, and you have specified that a minimum of four tapes are to be maintained in the media pool Save Set. Therefore, you require a minimum of five weekly (Father) tapes.

The last full backup performed each month is the monthly backup. You specified that these tapes are to be retained for six months, and that six tapes are to be maintained in the media pool Save Set. The minimum number of monthly tapes required before the media recycles is six. Therefore, you need seven monthly (Grandfather) tapes.

The total media you need for this rotation scheme is 16.

Preflight Checks for Your Backups

The Preflight Check (PFC) utility enables you to run vital checks on the Arcserve Backup server and agents to detect conditions that may cause backup jobs to fail. The checks performed by PFC fall into the following categories:

- **System Checks**--These include checking system requirements for the server, available disk space for the database, and RPC service registration.
- **Arcserve Backup Checks**--These include checking the Arcserve Backup system account and its privileges, the status of the Arcserve Backup engines, SAN server connectivity (if the Arcserve Backup SAN Option is installed), and the health of the tape devices attached to the server.


- **Agent Checks**--These include checking the connection and credentials for any client and database agents needed for the job.

Note: The Preflight Check utility does not validate login credentials for the following database agents:

- Agent for Informix
 - Agent for Lotus Domino
 - Agent for Microsoft SharePoint Server
 - Agent for Microsoft Exchange Server
 - Agent for Microsoft SQL Server
 - Agent for Oracle
 - Agent for Sybase
 - Enterprise Option for SAP R/3 for Oracle
- **Media Checks**--These include checking the availability of media in the scratch set (if a media pool is specified for the job), checking the media expiration dates, and checking for source and destination conflicts for File System Devices.

The optimum time to run this command is a couple of hours before your jobs are scheduled to run so that you can have ample time to correct any problems that may appear in the PFC report. For more information on the PFC utility and its associated options, see the *Command Line Reference Guide*.

Before submitting a job, you can run a Preflight Check clicking the Preflight Check button on the Submit Job dialog.



Example: PFC Utility

You submit a job and run the PFC utility. If the PFC utility detects that a device is not assigned to the device group that you are using for the backup job, the PFC utility reports a failed job. To correct the problem, you must either use a device group with an assigned device or assign a device to the device group that you are using for the job. If you do not take corrective action, the job will eventually fail.

This capability is also supported when you run the PFC command line utility. For more information, see the *Command Line Reference Guide*.

Start Arcserve D2D

Arcserve D2D is a backup solution that lets you track changes to data at the block level, and back up only the changed blocks. Arcserve D2D lets you perform frequent incremental backups, which reduces the size of the backups and provides you with up-to-date backup data.

If Arcserve D2D is installed locally in your backup environment, you can start Arcserve D2D from the Arcserve Backup Manager Console.

If Arcserve D2D is not installed locally in your backup environment, you can specify the server name and port number to connect the remote Arcserve D2D server, or you can download and install Arcserve D2D.

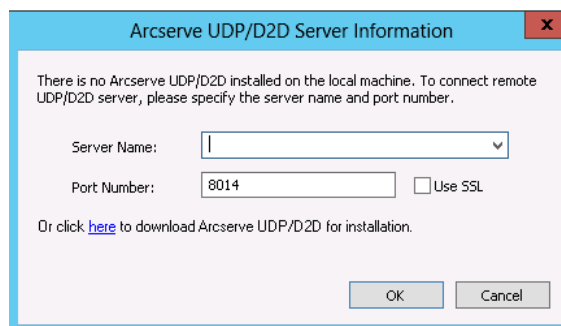
Follow these steps:

1. Open the Arcserve Backup Manager Console.

From the Quick Start menu on the Arcserve Backup Manager Console, select Protection & Recovery and click Arcserve D2D.

One of the following events occurs:

- If Arcserve D2D is not installed on the backup server, the Arcserve D2D Server Information dialog opens. From the Arcserve D2D Server Information dialog, you can log in to a remote Arcserve D2D server, or download and install Arcserve D2D.



- If Arcserve D2D is installed on the backup server, the Log in to Arcserve D2D screen opens.
2. On the Log in to Arcserve D2D screen, complete the following fields:
 - **Domain**--Specify the name of the Arcserve D2D domain
 - **User Name**--Specify the User name required to log in to the Arcserve D2D domain.
 - **Password**--Specify the password for the Arcserve D2D user name.

Click Log in.

Arcserve D2D opens.

Note: For more information about using Arcserve D2D, see the *Arcserve D2D Online Help* or the *Arcserve D2D User Guide*.

Start Arcserve Replication

Arcserve Replication is a data protection solution that uses asynchronous real-time replication to provide disaster recovery capabilities. This host-based software provides continuous data replication that transfers changes to application data as they occur to a standby replica server located locally or over the Wide Area Network (WAN).

If Arcserve Replication is installed locally in your backup environment, you can start Arcserve Replication from the Arcserve Backup Manager Console.

If Arcserve Replication is not installed locally in your backup environment, you can specify the server name and port number to connect the remote Arcserve Replication server, or you can download and install Arcserve Replication.

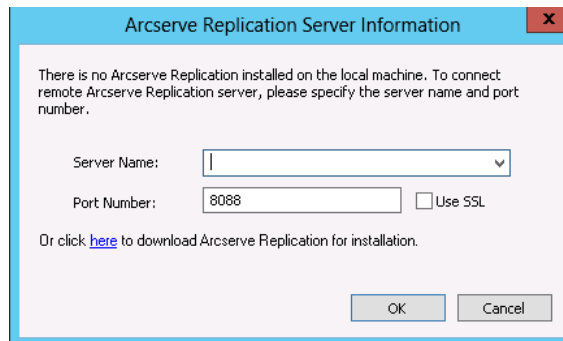
Follow these steps:

1. Open the Arcserve Backup Manager Console.

From the Quick Start menu on the Arcserve Backup Manager Console, select Protection & Recovery and Click Arcserve Replication.

One of the following events occurs:

- If Arcserve Replication is not installed on the backup server, the Arcserve Replication Server Information dialog opens. From the Arcserve Replication Server Information dialog, you can log in to a remote Arcserve Replication server, or download and install Arcserve Replication.



- If Arcserve Replication is installed on the backup server, the Log in to Arcserve Replication screen opens.

2. On the Log in to Arcserve Replication screen, complete the following fields:
 - **Domain**--Specify the name of the Arcserve Replication domain
 - **User Name**--Specify the User name required to log in to the Arcserve Replication domain.
 - **Password**--Specify the password for the Arcserve Replication user name.

Click Log in.

Arcserve Replication opens.

Note: For more information about using Arcserve Replication, see the Arcserve Replication documentation.

Chapter 3: Backing Up Data

This section contains the following topics:

[How Arcserve Backup Lets You Back Up Data](#) (see page 133)

[Submit a Backup Job](#) (see page 136)

[Backup Manager](#) (see page 137)

[Local Backup Options for UNIX and Linux Agents](#) (see page 152)

[Global Backup Options](#) (see page 154)

[Files and Objects that Arcserve Backup Does Not Back Up](#) (see page 184)

[Enable Arcserve Backup to Manage Open Files on Remote Computers](#) (see page 189)

[Multiplexing Job Options](#) (see page 189)

[Specify Multistreaming Options](#) (see page 192)

[Entire Node Backups](#) (see page 193)

[Create Repeating Backup Jobs](#) (see page 195)

[Back Up Remote Servers](#) (see page 197)

[Submit Static Backup Jobs](#) (see page 198)

[Backup Staging Methods](#) (see page 200)

[Archive Staging Methods](#) (see page 241)

[Backing up Multiple Data Mover Servers in a Single Job](#) (see page 246)

[Disaster Recovery](#) (see page 254)

[How NTFS Deduplication Works](#) (see page 255)

How Arcserve Backup Lets You Back Up Data

Arcserve Backup allows you to back up most machines attached to your Windows network using one of the following sources:

- Administrative shared drives
- User-shared files, directories, and drives

Because Arcserve Backup separates and lists Windows machines by the domain or workgroup to which they belong, you can easily back up all the machines belonging to a specific domain or workgroup, by selecting the name of the domain or workgroup.

The optional Arcserve Backup Client Agents allow you to communicate with remote workstations in various environments. This provides complete system backups, including system information from non-Windows systems, such as UNIX.

Similarly, the optional Backup Agents allow Arcserve Backup to back up and restore online databases such as Microsoft Exchange Server, Lotus Domino, Microsoft SQL Server, Oracle, and IBM Informix.

Specify Local Backup Options

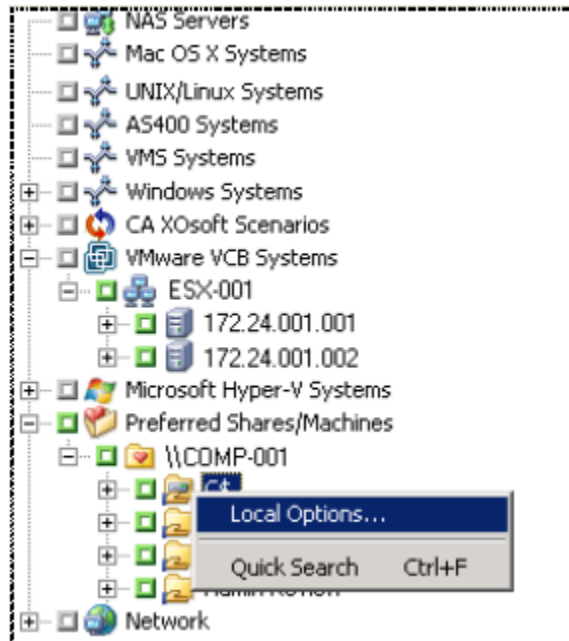
Arcserve Backup gives you the power and flexibility to customize local options for specific drives that you want to back up.

To specify local backup options

1. Open the Backup Manager and select the Source tab.

Click the check box next to the drive directory.

Right-click the drive and select Local Options options from the pop-up menu to open the Local Options dialog.



Important! When setting local options, you must select drives individually as your source even if you want to back up an entire server. You cannot click the green box next to the server name, and then customize local backup options for individual drives.

2. Specify the desired options:

Backup Verification Options

The Backup Verification options enable you to verify that your data was backed up correctly. The following options are available.

- **None**--If you select this, no verification will be performed on the data backed up.
- **Scan Backup Media Contents**--If you select this, Arcserve Backup scans the media and check that the header is readable for each file that is backed up.
- **Compare Backup Media to Disk**--Select this if you want Arcserve Backup to read blocks from the media and compare, byte for byte, the data on the media against the files.

Session/Encryption Password Option

Use this option to specify a password to protect the data.

- **Session/Encryption Password**--Enter a password for your backup job.

Important! It is important that you remember the Session/Encryption password to restore this session. There is no way to reset this password.

Compression and Encryption Options

Use these options to specify whether files should be compressed or encrypted before they are backed up. These options are not supported on deduplication devices. If you specify a deduplication device group as the backup destination or as the staging destination, compression and encryption are skipped if detected.

- **Compress Files Before Backup Using Software Compression**--Allows you to compress your files before running your backup job. Using this option directs Arcserve Backup to compress files before backing them up using a software compression algorithm. Since most tape devices are equipped with a hardware-based compression mechanism, using both software and hardware compression is unnecessary and can lead to a slow backup job and poor compression results. Therefore, you should select this option only if your tape drive is not equipped with a hardware compression mechanism.
- **Encrypt Files Before Backup**--Allows you to encrypt your files before running your backup job.

Important! Arcserve Backup performs local compression and encryption at the agent system. When you specify local compression and encryption and ARCserve server-based compression and encryption (global option), Arcserve Backup performs the compression and encryption at the agent system.

Note: For more information about specifying Arcserve server-based compression and encryption, see [Backup Manager Backup Media Options](#) (see page 164).

NTFS Data Deduplication

This option lets you perform optimized backups for full backups only on NTFS data deduplication enabled volumes. You can use this option on computers that are running the Windows Server 2012 operating system.

Default value: Disabled.

Note: For more information about NTFS Data Deduplication, see [How NTFS Deduplication Works](#) (see page 255) or the [Microsoft Windows Development Center website](#).

3. Click OK to apply the local settings to the specified volume.

More information:

[Local Backup Options for UNIX and Linux Agents](#) (see page 152)

Submit a Backup Job

This section summarizes how to submit a backup job.

For information about how to use disk staging (D2D2T) and tape staging (D2T2T) to manage your backup operations, see [How Backup to Disk to Tape Works](#).

To submit a backup job

1. From the Backup Manager, select the [Start](#) (see page 139), [Source](#) (see page 140), [Destination](#), and [Schedule](#) (see page 152) tabs to specify the options that you require for the job.

Click the Options toolbar button to specify global options that you require for the job. For more information, see [Global Backup Options](#) (see page 154).

Click the Submit toolbar button to submit your job.

The Security and Agent Information dialog opens.

2. On the Security and Agent Information dialog, edit or confirm the security and agent information for your job, and click OK.
3. When the Submit Job dialog opens, select Run Now to run the job immediately, or select Run On and select a date and time when you want the job to run.

Note: For more information about the Run Now option, see Job Queue Tab.

4. Enter a Job Name for your job.
5. If you selected multiple sources to back up and want to set the priority in which the job sessions initiate, click Source Priority. Use the Top, Up, Down, and Bottom buttons to change the order in which the jobs are processed. When you finish setting priorities, click OK.

6. To save the job as a Arcserve Backup job script, click the Save Job button.
7. To save the job template, click the Save Template button.
8. To preflight check the job, click the Preflight Check button. If the preflight check failed, click the Cancel button to modify the job settings.
9. On the Submit Job dialog, click OK.
The job is submitted.

More information:

[How to Manage Jobs Using the Job Queue Tab](#) (see page 323)

Backup Manager

The Backup Manager lets you customize your backup jobs using filters, options, and scheduling.

You can use the Backup Manager to:

- Create groups of backup sources.
- Back up to various media or create a customized backup scheme.
- Use filters to selectively exclude or include directories and files from backup jobs.
- Create an automated backup scheme using the Grandfather-Father-Son (GFS) rotation scheme.
- Apply filters to local source objects (such as volumes and nodes) or globally to the entire backup job, or to both at the same time.

Arcserve Backup allows you to back up the Windows registry as well as the system state for Windows systems. Each backup job requires a source and a destination (media). The Backup Manager screen provides tabs to customize your backup job:

- **Start**--Lets you specify the type of backup: Normal, Deduplication, or UNIX/Linux Data Mover. You may also enable staging for each backup type. For Normal or Deduplication backups, you can enable synthetic full backups.
- **Source**--Lets you specify the data that you want to back up.
- **Schedule**--Lets you specify a schedule, repeat method, or rotation scheme for the job.
- **Destination**--Lets you specify the location where you want to store your backup data.

The topics that follow provide full details about the options available on each tab.

This section contains the following topics:

[Options on the Backup Manager Start Tab](#) (see page 139)

[How to Specify Source Data Using the Classic View and the Group View](#) (see page 140)

[Backup Manager Destination Options](#) (see page 149)

[Backup Job Schedules and Rotations](#) (see page 152)

Options on the Backup Manager Start Tab

From the Backup Manager Start tab, you can select the backup type.

- **Normal backup**--Normal backup lets you backup a data source to a target destination, using a custom schedule, repeat method or rotation scheme.
Note: Use Normal backup when you want to submit a backup job to one data mover server.
- **Deduplication backup**--Deduplication backup lets you save only unique data chunks to disk, allowing you to fit more backup sessions on media, retain backups for longer periods of time and speed up data recovery. For more information about submitting deduplication backup jobs, see [Back Up Data with Deduplication](#). (see page 738)
- **UNIX/Linux Data Mover backup**--UNIX/Linux Data Mover backup lets you consolidate multiple Data Movers into a single backup job as long as they share a single library.

For each backup type, you must click the [Source](#) (see page 140), [Schedule](#) (see page 152), and Destination tabs to complete backup job configuration.

You may also [Enable Staging](#) (see page 200). Staging operations allow you to back up data to a staging device and then migrate the backed up data to a final destination (usually a tape). You can choose to Enable Staging on Normal, Deduplication, or Data Mover backup jobs.

You can also select [Enable Synthetic Full Backup](#) (see page 808). A synthetic full backup lets you synthesize a previous full session and its subsequent incremental sessions to a full session. You can choose to enable synthetic full backups on Normal backup with staging option, Deduplication backup with staging option, or Deduplication backup jobs. In addition to completing backup job information on the Source, Schedule, and Destination tabs, ensure you complete the information on the Staging Location tab and define a synthetic schedule on the Schedule tab.

How to Specify Source Data Using the Classic View and the Group View

The source is the path to the data that you want to back up. You can easily find the files you want to back up by browsing through the Backup Manager directory to select the user-shared drives and directories.

Arcserve Backup lets you browse and specify the source data using the following views:

- **Classic View**--This is traditional source view. Machines are listed first, allowing you to expand and then select specific data sources. With the Classic View, Arcserve Backup categorizes source computers based on the platform that is running on the computer. For example, Windows systems, UNIX/Linux systems, and Hyper-V systems.
- **Group View**--This view categorizes source computers based on the Arcserve Backup agent that is installed on the computer. The agents are listed as branches on the source tree. Within each branch, the computers that contain the specified agent are listed.

You may also create customized groups that allow you to group machines according to your own criteria. For example, using the Group view is effective approach to specifying source when you want to back up database files such as Microsoft SQL Server, Microsoft Exchange Server, and Microsoft SharePoint Server data that reside on a large quantity of machines, without having to expand each machine and then select the database node.

Note: The Agent for Microsoft Exchange Server 2010/2013 appears only in the Exchange Organization object. You cannot add Agent for Microsoft Exchange Server 2010/2013 systems to the Microsoft Exchange Server group.

When selecting a source, you can select to back up:

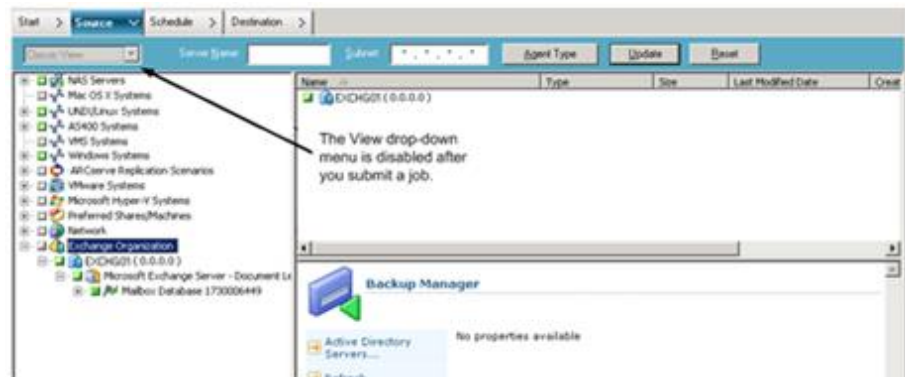
- an entire application
- a customized source group
- an entire server
- individual drives, directories, and files
- To select individual drives, directories, and files to back up, expand a server name and click the green boxes next to each drive, directory, and file.

To select an entire source group, click the green box next to the group name. When you do this, all the servers, nodes, volumes, drives, directories, and files included in the source group are automatically selected.

Be aware of the following behavior:

- The view that you specify when you submit a job cannot be modified.

For example, you submit a job using the Classic View. Subsequently, you want to modify the source selections for the job. When you modify the job and click the Backup Manager, Source tab, the view drop-down menu is disabled. The following screen illustrates this behavior.

**Custom Local Backup Options**




You can right-click individual drives to customize local backup options. If you have database agents installed, you can also right-click them to customize local backup agent options. If you want to customize local backup or local backup agent options, your job must be packaged explicitly, which means you must select drives, directories, files, or database agents individually as your source even if you want to back up an entire server. You cannot click the green box next to the server name, and then customize local backup options for individual drives, directories, files, or database agents. For more information, see [Dynamic Job Packaging](#) (see page 302) and [Static Job Packaging](#) (see page 305).

Backup Manager Markers

Each object displayed in the Backup Manager window has a green or gray box to its left called a marker.

- **Green marker**--Lets you control the extent of the backup for an object directly. Click a marker to exclude an object from a backup or to indicate that you want the backup 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 backup.
- **Gray marker**--These markers are associated with objects that are not real and that you cannot back up/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 back up.

The following table describes the different marker configurations and corresponding backup levels:

Marker	Description
	Full backup.
	Partial backup.
	Do not back up.

Note: Gray marker configurations follow the same pattern as green marker configurations, but reflect the proportion of files under them that are selected for backup.

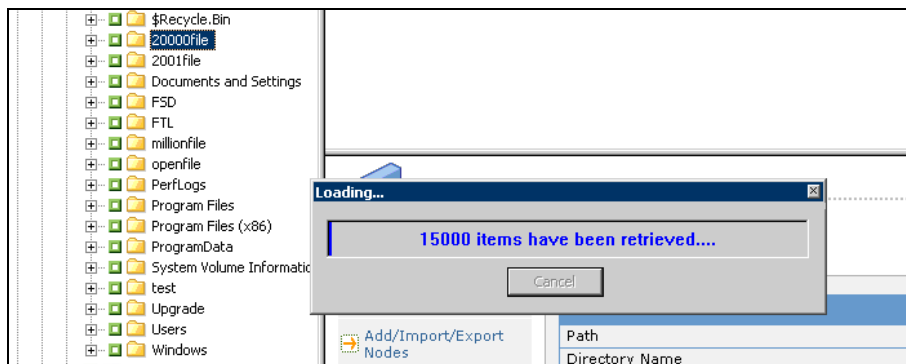
The fill proportion of a marker at a higher level of the directory tree depends on the fill proportions of the markers of the objects at the lower levels.

- If you click a marker at a higher, parent level so that it is completely filled, all the markers at the lower, child levels are automatically filled completely.
- If you click all the markers at the lower, child levels so that they are completely filled, then the marker at the higher, parent level is automatically partially filled.
- If the markers at the lower, child levels are a mix of completely filled and partially filled, the marker at the higher, parent level is automatically partially filled.

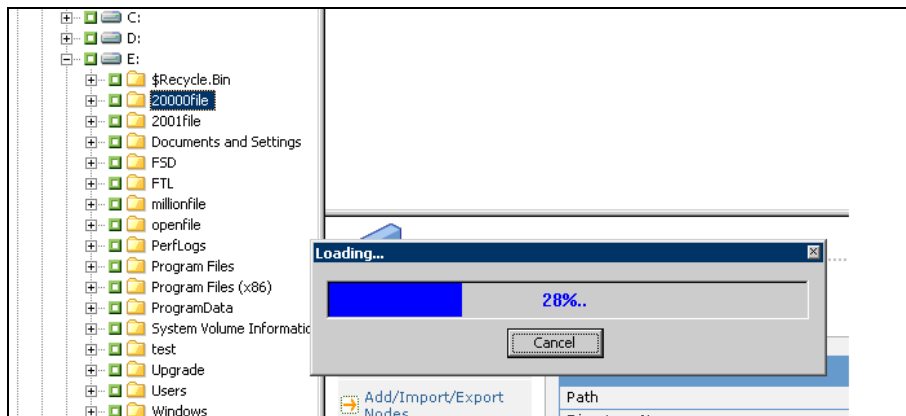
How Arcserve Backup Lets You Browse a Large Number of Items in the Backup Manager

Arcserve Backup lets you pause the process of loading items in the Backup Manager when you browse a large number of directories, files, and so on. The steps that follow describe how Arcserve Backup lets you browse a large number of items in the Backup Manager window.

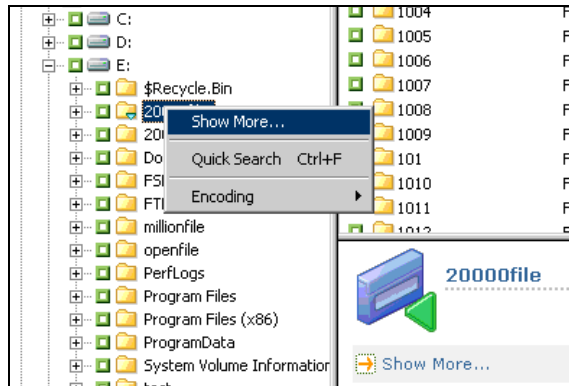
1. When you select a directory in the Backup Manager, Source tree, Arcserve Backup displays a Loading dialog to inform you that a large number of items need to be retrieved and loaded into the Backup Manager window. You cannot click Cancel while Arcserve Backup is retrieving the list of items to display in the Backup Manager window.



2. After Arcserve Backup retrieves the list of items to display in the Backup Manager window, the Loading dialog then displays the percentage of items that are loaded into the Backup Manager. If there are a large number of items to display, you can click Cancel to pause the loading process.



3. After you pause the Loading process, you can continue the Loading process by right-clicking target directory and selecting Show More from the pop-up menu.



4. If you pause the loading process, the icon for the target directory appears as follows:



5. You can pause and continue the loading process as often as necessary. To load more items, right-click the target directory and click Show More from the pop-up menu.

6. When the loading process is complete, the icon for the target directory displays as follows:



Browse a Large Number of Files in the Backup Manager

Use the following procedure when you need to browse a directory that contains a large number of items in the Backup Manager.

Note: The Backup Manager can display file name paths that contain up to 512 characters. This behavior includes the drive letter or network server and the share names. On NTFS and FAT32 file systems, the maximum length of file names is limited to 255 characters.

To browse a large number of items in the Backup Manager

1. Open the Backup Manager and specify a target directory from the Source tree.

The Loading message box appears, Arcserve Backup retrieves a list of items to display in the Backup Manager Window, and then Arcserve Backup loads the files into the Backup Manager window.

2. From the Loading message box, click Cancel to stop the loading process.

If Arcserve Backup did not load all items, the To show more objects, right-click the target directory and select Show More from the pop-up menu warning message appears.

Note: The message only appears the first time you click Cancel on the Loading message box.

3. From the Source tree, right-click the target directory and click Show More from the pop-up menu.

The Loading message box appears and Arcserve Backup continues loading the items.

4. You can pause and continue the loading process as often as necessary until Arcserve Backup loads all items in the target directory.

If you pause the loading process, the icon for the target directory displays as follows:



When the loading process is complete, the icon for the target directory displays as follows:



Browse Computers by Agent Type

By default, the Backup Manager view lists backup source in the Group View. The Group View lets you browse computers based on the Arcserve Backup agent that is installed on the computer. You may also set up your own groups, if desired.

When you close the Backup Manager, the view selected opens the next time that you open the Backup Manager. For example, if you select Classic View and then close the Backup Manager, the Classic View opens the next time that you open the Backup Manager.

Because a computer can belong to more than one source group, it is possible that you can specify the same backup source more than once. When Arcserve Backup detects the same source specified in multiple jobs, a warning message appears, allowing you to determine whether the same data should be backed up more than once.

To browse computers by agent type

1. Open the Backup Manager.

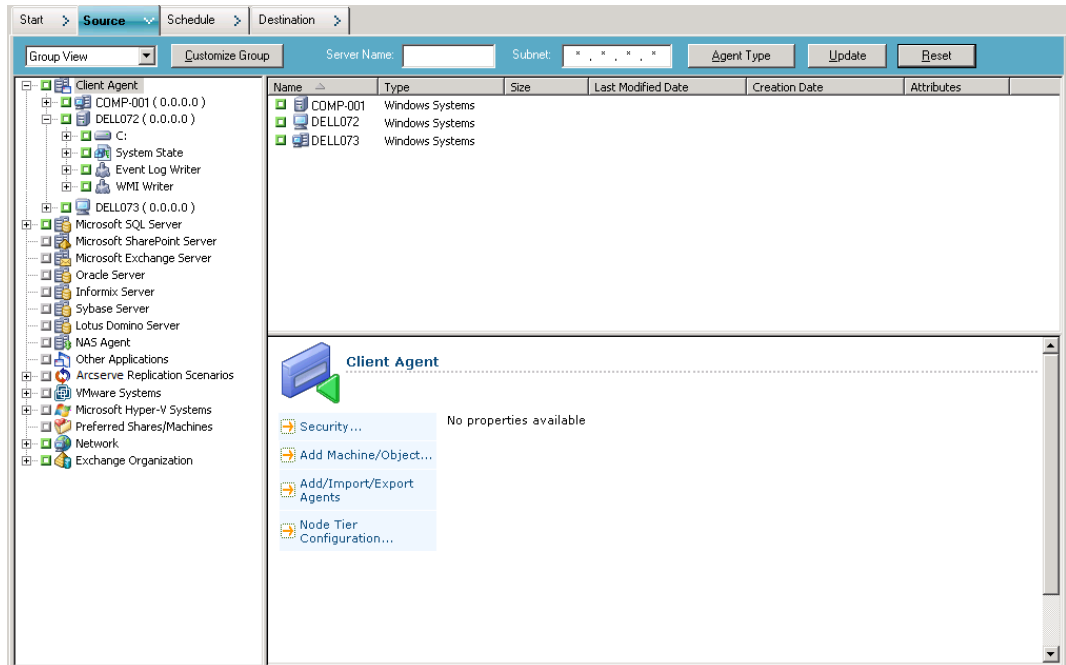
Click the Source tab.

The Source directory tree displays in the Group view.

Note: If there is more than one agent installed on a computer, (for example, Agent for Microsoft Exchange Server, Agent for Microsoft SharePoint Server) the computer can appear under more than one group).

2. Expand the computers in the source tree. If prompted, you must provide the required security information.

The following screen illustrates the computers available in Group View for the Client Agent for Windows.



Note: If you want to customize the groups that appear in the Source tree, click Customize Group. For more information, see [Configure Customized Groups in Group View](#) (see page 147).

3. In the source tree, search for the desired agent type and expand it to view a list of computers.

4. (Optional) Search for a computer using global filters that are saved when you exit Backup Manager and remain set until changed.
 - **Server Name**--Lets you filter source computers by the string you enter.
 - **Subnet**--Lets you filter computers by their IP addresses.
 - **Agent type**--Lets you filter computers by agent type. In Group and Classic Views, the Agent Type filter lets you view only the agent groups that correspond to the agent selected. For more information, see [Filter Nodes](#) (see page 339).

Job History also lets you view results by source group. For more information, see [How to Analyze Jobs Using Group View](#). (see page 328)

Configure Customized Groups for Group View

To help you manage large environments, you can create custom groups and then add computers to the groups based on criteria that you determine.

Example: Customized Groups for Group View

Suppose your Sales department contains SQL database files distributed across 100 computers. You can add all the machines that contain the sales data you want to back up to a customized group named Sales Data. Customized groups appear in the source tree as main branches, thus allowing you to quickly locate and select groups when you define your backup jobs.

To configure customized groups for group view

1. Open Backup Manager and click the Source tab.

The Backup Manager opens, showing the default Group View.

The source tree displays the Arcserve Backup agents as main branches. The Customize Group button appears next to the views drop-down list.

2. Click Customize Group.

The Customized Group Configuration dialog opens.

Customized Group Configuration displays existing groups, by name, on the left side of the dialog, and the servers that belong to each group on the right.

3. Click New to create a new group.

- a. In the Name field, enter a name for your group.
- b. If desired, enter a comment that describes your group.

Click OK.

The new group is added to the list of groups on the left.

4. Select the group that you created, if it is not already selected.

- From the list of servers on the right, click a server to add to the group and click Assign.

The server appears below the group.

Repeat this step, as required, to add more servers to the group.

- Click OK to save settings and exit Customized Group Configuration.

Manage Customized Groups

Arcserve Backup lets you change the name of a customized group, delete a customized group, and change the servers in the group, as needed.

To manage customized groups


- Open Backup Manager and click the Source tab.

Backup Manager opens with the default Group View displayed.

Note: If the source tree displays in the Classic View, click the drop-down list above the tree and select Group View.

- Click Customized Group.

The Customized Group Configuration dialog opens.



3. Do one of the following:
 - To delete a group, select the group and click Delete.
 - To rename a group, select the group and click Modify.
Enter a new name, and click OK.
 - To reassign servers, select a server that you want to remove from the left side of the dialog and click Remove.

The server is added to the list of available servers on the right. From the list of available servers, choose a server to add. Click Assign. The server is added to the group.
4. Click OK when done managing customized groups to save settings and exit Customized Group Configuration.

Backup Manager Destination Options

The destination is the backup media device or disk. You can use the Destination tab in the Backup Manager to browse to and select the groups and device.

Note: Cloud devices are not available as a backup destination.

The Backup Manager Destination tab includes the following backup options:

Multiplexing

The following options regulate how Arcserve Backup handles multiplexing.

- **Maximum Number of Streams**--Sets the maximum number of streams that can write to a tape at the same time. The default number of streams is 4 and the supported range is from 2 through 32.

Note: Multiplexing is not supported for UNIX/Linux Data Mover backup jobs.

Multistreaming

The Multistreaming option lets you split single backup jobs into multiple jobs and use all of the available tape devices in your system to complete the backup. For more information, see [Multistreaming](#) (see page 103).

Group and Media field

Use the Group and Media field to specify the device group that you want to use for the backup job.

- Place an asterisk in the Group or Media field to use the first available drive and media in the group.
- If you want to use any available group, click the Use Any Group option.

Media Pool

Select this option if you want to use a specific media pool for the backup job.

Be aware of the following:

- If you select a Media Pool, Arcserve Backup automatically checks that the other destination and backup options you selected to verify that no restrictions or conflicts occur when you run the job. If Arcserve Backup detects a conflict, a warning dialog opens.
- Cloud Device cannot be assigned to a Media Pool.

Server

This field displays the name or the primary server and member servers in your Arcserve Backup domain.

Note: If you did not install the Central Management Option, the name of the current server displays in this field.

With Enable Staging checked:

Deduplication Policy

Click this button to open the Deduplication Purge Policies dialog. This dialog lets you set the purge policy for full and differential/incremental backups.

Create Duplicate in-line

Select this check box to enable the Group, Duplicate Media prefix, and Duplicate Media pool prefix fields. These fields enable you to perform a Double Tape Copy, where you can migrate the data to two final destinations at the same time. To ensure that the migration job continues when one of the devices become unavailable, you can set a timeout value for additional backup media to continue the migration job in another device. If you do not set a timeout value, then the data migration job fails after the first backup media timeout completes. For more details about setting the timeout value, see [Backup Manager Backup Media Options](#) (see page 164).

Note: Ensure that a Disk Staging device is configured on your environment to perform a Double Tape Copy.

Group

Select a device group from the drop-down to determine where the duplicate is created.

Note: Tape/Changer groups and File System Device (FSD) groups are the only device groups available. When the backup job is submitted and the same device group is selected for both the final destination and the duplicate destination, the application will verify that there are two or more drives.

Duplicate Media prefix

This field lets you specify a prefix for the media name you want to duplicate. For example, "CPY_<media name>".

Duplicate Media pool prefix

This field lets you specify a prefix for the media pool name you want to duplicate. For example, "TEST_<media pool name>".

Note: To back data using disk staging, use Device Configuration and Device Group Configuration to configure the staging device. For more information, see [Backup Staging Methods](#) (see page 200).

More information:

[How Arcserve Backup Processes Backup Data Using Multistreaming](#) (see page 103)

How to Use Wildcards with Tape Library Groups

The wildcard characters' asterisk and question mark are supported in the Group field. When wildcard characters are used to specify a job's library group destination, the job is sent to a group whose name matches the criteria and has at least one available media, as long as there is at least one available drive associated with the library. A media is available when it is not being used by another job (Note: no special consideration is given to media suitability as determined by the job schema; for example, Media Pool). If more than one job uses wildcards and more than one group matches the selection criteria, all jobs go to the first group with an available media.

Typing a name in the media field forces the job to be directed to a group that matches the criteria and contains the specified media, even if the media is busy. If there is no media with the specified name in any of the matching groups, but there is a blank media in a matching group, it is used and renamed. If there is no blank media, the user is prompted to insert one.

Note: The media field does not support wildcard characters.

When a media pool is specified, a media from that pool is used if there is one available in the first matching group. If there is no such media in the group, but there is a blank media, it is renamed and added to the pool. If there is no blank media the user is prompted to insert one.

Backup Job Schedules and Rotations

You can configure your backup job to use a custom schedule, synthetic schedule, or a rotation scheme by using the Arcserve Backup template schemes or by specifying your own rotation parameters. You can also specify a repeat method and the following backup methods for each backup:

- **Full (Keep Archive Bit)**--Performed each time the job is repeated and keeps the archive bit.
- **Full (Clear Archive Bit)**--Performed each time the job is repeated and clears the archive bit.
- **Incremental backup**--Backs up only those files whose archive bits have been set since the last full or incremental backup was performed. After each backup, archive bits are reset so that they are not backed up during the next incremental backup job.
- **Differential backup**--Backs up only those files whose archive bits have been set since the last full backup was performed. Because differential backup jobs do not clear a file's archive bit, the files that were backed up in the last differential job are backed up again. It takes longer to process backup jobs using this method. However, this strategy requires only two sets of media to restore a differential backup; the full media set, and the differential media set. In the case of an incremental backup, you require the full media set and every incremental set until the latest set. For a synthetic full backup job, if the agent supports synthetic full backups, the differential backup job is converted to an incremental backup job.

Note: The above-described backup methods do not apply to the Linux Client Agent.

For a description of detailed job scheduling features, see [Customizing Jobs](#) (see page 301).

Local Backup Options for UNIX and Linux Agents

The following are the local options available when backing up a UNIX or Linux computer using the Client Agent for UNIX or the Client Agent for Linux.

Additional Options

- **Traverse Symbolic Link File**--Arcserve Backup follows symbolic links and backs up the linked files.
- **Traverse NFS--Backs up NFS**--mounted drives.
- **Traverse Across File System**--Arcserve Backup automatically includes locally mounted UNIX file systems in the backup.
- **Estimation Off**--Disables the estimation of the number of files and the amount of data to be backed up that takes place at the beginning of the backup job. Selecting this option decreases the time it takes to perform the backup.

- **Preserve File Access Time**--This option directs Arcserve Backup to preserve the last access time of files when a backup is performed.

Note: The Access Time of a file is automatically updated by the operating system whenever a file is accessed (read or write). However, after a compare is performed, the Access Times of all the backed up files are also updated. Therefore, if you want to track whether or not a file has actually been accessed (and not just compared), you need to preserve the original access time.

- If this option is selected (check in box), Arcserve Backup preserves the last file access time of any files that are backed as the original value that was present before the backup was performed (Change Time will be updated). This is the default setting.
- If this option is not selected (no check in box), the last file access time of any files that are backed up is updated to the new value that is present when the backup is completed (Change Time will not be updated).

Note: For Windows based agents, you must apply this option globally. For more information, see Global Backup Options.

Media format to use for backup

- **Arcserve Backup format**--This is a Arcserve Backup proprietary tape format. This format is designed to overcome the limitations of tar/cpio formats and leverage other features like compression/encryption provided by Arcserve Backup. For example, there are certain limitations with tar/cpio while backing up large files and huge data that may span across multiple tapes.
- **Posix tar format**--This is a Standard Posix Tar format. When you select this option, Arcserve Backup creates a backup image in Posix Tar format. Arcserve Backup or any tar utility can be used to restore data from an image created in this format. Using Arcserve Backup format is recommended.
- **Posix cpio format**--This is a Standard Posix CPIO format. When you select this option, Arcserve Backup creates a backup image in Posix CPIO format. Arcserve Backup or any CPIO utility can be used to restore data from an image created in this format. Using Arcserve Backup format is recommended.

More information:

[Global Backup Options](#) (see page 154)

[Specify Local Backup Options](#) (see page 134)

Global Backup Options

This section describes the global backup options you can select when submitting your backup job. To access the global options dialog, click the Options toolbar button in the Backup Manager.

Global options (also referred to as job-wide options) define the procedures and rules that Arcserve Backup uses to back up all nodes (servers, workstations, and applications) that you define for the particular backup job. The global options that you define for one backup job do not affect the global options that you define for any other backup job.

For a description of additional backup job options and filtering features, see [Customizing Jobs](#) (see page 301).

This section contains the following topics:

- [Backup Manager Alert Options](#) (see page 154)
- [Backup Manager Media Exporting Options](#) (see page 155)
- [Backup Manager Advanced Options](#) (see page 156)
- [Backup Manager Encryption/Compression Options](#) (see page 159)
- [Backup Manager Volume Shadow Copy Service Options](#) (see page 162)
- [Backup Manager Backup Media Options](#) (see page 164)
- [Backup Manager Verification Options](#) (see page 166)
- [Backup Manager Operation Options](#) (see page 167)
- [Backup Manager Pre/Post Options](#) (see page 171)
- [Backup Manager Agent Options](#) (see page 173)
- [Backup Manager Job Log Options](#) (see page 184)

Backup Manager Alert Options

You can use the Alert notification system to send messages about events that appear in the Activity Log during your backup operation. Choose one or more of the following events for which you want to be notified:

- **Job Completed Successfully**--All of the nodes and drives/shares were processed.
- **Job Incomplete**--Some nodes, drives, or shares were missed.
- **Job Canceled by User**--The user canceled the job.
- **Job Failed**--The job was started but could not be completed.
- **Virus Detected**--A virus was detected in one of the files to be backed up. See Virus options (Backup, Copy, Count).
- **Media not Available**--Media was not available during the execution of a job.

Note: The backup media must be tape media.

- **Format Blank Tape**--A tape was formatted during the execution of a job.
- **Customized Event**--A customized event occurred. To specify this type of event, enter an error, warning, or notification code in the space below the Event drop-box.

Choose one or more of the defined Alert configurations. The <default> configuration means that you will use whatever is configured in Alert Manager. Click Configure to define further configurations. Arcserve Backup provides the following defined Alert configurations:

- Broadcast
- Pager
 - Note:** Pager options are not supported on Japanese versions of Arcserve Backup.
- SMTP
- SNMP
- Event
- Printer
- E-Mail
- Lotus Notes
- Unicenter TNG

Specify miscellaneous options:

- **Attach Job Log**--Lets you include the job log information in the Alert message. (This option applies for Trouble Tickets and Mail only.)
 - Note:** The list you create using Alert Options is saved with the Job Script and the configuration defined using the Configuration button.
- **Send alert messages only for master jobs**--Lets Arcserve Backup send you alerts that reference only the master job number in the Alert message. The alert messages will not reference child and subjob numbers. You can specify this option on all jobs, including multiplexing and multistreaming jobs.

Backup Manager Media Exporting Options

At the end of a backup job, you can move media out of the library or to an off-site location for safe storage. Arcserve Backup provides the following media exporting options:

- **None**--No media exporting will take place at the end of a backup job.
- **Export RAID1 Duplicate Tape After Job**--If the job spanned to multiple media, all the duplicate media used in this job is exported.
 - Note:** This option is for RAID 1 support with libraries and mail slots only.

- **Export All Tapes After Job**--Arcserve Backup exports all the media for the related backup. If the job spanned to multiple media, all the media used in this job is exported. If there are not enough mail slots to export all the media, the media that could not be exported is moved back to the original home slot. In single mail slot libraries, Arcserve Backup retries a few times to check if the mail slot is empty to move the next media to the mail slot. If the operator does not move the media, Arcserve Backup writes this information in the activity log.

Note: This option is for RAID 1 support with libraries and mail slots only.

Media Exporting Limitations

Be aware of the following media exporting limitations:

- For staging backup jobs, media exporting options are only effective during the migration phase of the job.
- Media exporting options are functional only for regular and rotation jobs and are supported on media libraries and Tape RAID.
- Media exporting options are not supported when you are performing tape staging (B2T2T) backups and the staging device or the final destination device is a RAID device.
- If the job includes verification, the export is done at the end of the verification.

Backup Manager Advanced Options

The Advanced options determine how Arcserve Backup handles the file system extensions during a backup.

More information:

[Windows System Options](#) (see page 157)

[Disaster Recovery Options](#) (see page 158)

[Microsoft SQL Server Backup Options](#) (see page 158)

[NTFS Data Deduplication](#) (see page 159)

[Point in Time Restore Options](#) (see page 159)

Windows System Options

The Windows systems options are supported only on Windows XP and Windows Server 2003 operating systems.

The Windows system options are as follows:

- **Traverse Directory Junctions and Volume Mount Points**--Selecting this option causes the backup job to traverse the volume or the directory being specified to and take a backup of it. At the time of restore of this session, you can restore files and directories contained in the referred to volume or directory. When this option is not selected, the backup job does not back up the volume or the directory being referred to by the volume mount point or the directory junction respectively. Therefore, at the time of restore, you cannot restore a file or directory contained in the referred to volume or directory.
- **Backup Mount Points as Part of the volume that they are mounted on**--If you select this, the volumes referred to by the Volume Mount Points will be backed up as part of the same session as the Volume Mount Points. When this option is not selected, the volumes referred to by the Volume Mount Points are backed up as separate sessions. This option is available only when the previous option, Traverse Directory Junctions and Volume Mount Points, is selected.
- **Preserve File Hard Links**--If you enable this, Arcserve Backup preserves hard links during a restore.

Note: When you apply the *Traverse Directory Junctions and Volume Mount Points* and *Backup Mount Points as Part of the volume that they are mounted on* options to named, mounted volumes that contain virtual hard disks (VHDs), Arcserve Backup creates separate backup sessions for mounted volumes that contain VHDs.

Example: Mounted Volumes that Contain VHDs

A server contains physical disk (C:\) that contains VHDs D:\ and E:\. VHD files (D.vhd and E.vhd) that reside in C:\ are mounted as drive D:\ and drive E:\. Drive D:\ is mounted to C:\MountD, and drive E:\ is mounted to C:\MountE.

If you back up C:\MountD and specify the *Traverse Directory Junctions and Volume Mount Points* option, and the *Backup Mount Points as Part of the volume that they are mounted on* option is enabled or disabled, Arcserve Backup creates separate backup sessions for drive D:\ and C:\MountD.

Disaster Recovery Options

The Disaster Recovery options available are:

- **Generate DR information for partially selected nodes**--Disaster recovery information is normally generated when performing a full machine backup. However, there are special cases where you may need to keep the disaster recovery information updated but cannot perform full machine backups too often (like in a SAN shared disk environment). By enabling this option, you can generate or update a machine's disaster recovery information without having to back up everything on the machine.
- **Include filtered sessions when generating restore session information**--When generating disaster recovery information, the Arcserve Backup server keeps track of only the latest non-filtered backup sessions pertaining to the machine. By default, if you back up a machine using filters, the filtered backup sessions will not be used by disaster recovery when recovering the system. By enabling this option, you can alter the default behavior and have disaster recovery use the filtered backup sessions when recovering the system.

Important! Enabling this option is very risky, especially for system volumes. Missing system files may lead to incomplete recovery.

This option is disabled by default. When you enable this option, it works at the job level. If the job contains multiple machine backups, this option will apply to all machines.

Microsoft SQL Server Backup Options

For Microsoft SQL Server, Arcserve Backup supports the following global option:

- **Do not apply Scheduled Job Method or Rotation Phase to Microsoft SQL Server databases**--Lets you exclude the backup method specified on the Backup Manager, Schedule tab. With this option specified, Arcserve Backup behaves as follows:
 - Arcserve Backup ignores the custom schedule, rotation, and GFS rotation method that was specified for the job.
 - Arcserve Backup converts the backup method specified on the Schedule tab to Full backup, only if the logic for database level backups and global backup options requires the backup method specified on the Schedule tab.

Note: For more information about backing up and restoring Microsoft SQL Server databases, see the *Agent for Microsoft SQL Server Guide*.

NTFS Data Deduplication

This option lets you perform optimized backups for full backups only on NTFS data deduplication enabled volumes. You can use this option only on computers that are running Windows Server 2012.

Default value: Enabled.

Note: For more information about NTFS Data Deduplication, see [How NTFS Deduplication Works](#) (see page 255) or the [Microsoft Windows Development Center website](#).

Point in Time Restore Options

Arcserve Backup lets you perform point in time restores from data that is stored in synthetic backup sessions.

Note: To view this option, enable the Enable Synthetic Full Backup option from the Start tab in Backup Manager. For more details, see [Specify the Synthetic Schedule for a Synthetic Full Backup](#) (see page 813).

- **Enable Point in Time Restore (PIT Restore)**--With this option specified, Arcserve Backup executes point in time, incremental backups on all daily backup executions except for the days that full backups are scheduled to execute. You can specify the point in time restore option to all scheduled jobs that apply to Arcserve Backup file system agents.

For jobs with this option specified, Point in Time Incremental appears as the Backup Method in the Activity Log and as the Session Method in the Restore by Session window. However, in the Backup Manager window, the Backup Method appears as Incremental.

Important! The Enable Point in Time Restore option applies to only synthetic full backup jobs.

Backup Manager Encryption/Compression Options

Arcserve Backup lets you encrypt, compress, or encrypt and compress backup data.

Be aware of the following:

- Arcserve Backup does not support compressing and encrypting data on deduplication device groups.

Note: For more information, see [Compression and Encryption with Deduplication](#) (see page 742).

- If you specify encryption and compression options, and the backup destination is a drive that does not support compression, Arcserve Backup encrypts the backup data and does not compress the backup data.

The following options define how Arcserve Backup processes backup data during a backup job and during the migration phase of a staging backup job.

Session/Encryption Password

- **Session/Encryption password**--Specify a Session/Encryption password to restore this data from media.

If you specify a Session Encryption password, you must specify the password to perform the following operations:

- Restore operations where the encryption, compression, or both were processed at the agent or at the backup server.
- Compare operations where the encryption, compression, or both were processed at the agent or at the backup server.
- Merge and Scan operations where the encryption, compression, or both were processed at the backup server. (You do not need to specify the password to perform Merge and Scan operations where the encryption, compression, or both operations were processed at the agent.)

Note: The Session/Encryption password is not required when you Merge or Scan only the session headers.

- **Save Current Session/Encryption Password to the Arcserve Backup database**--Use this option to save the password to the Arcserve Backup database and enable password management. This option is selected by default. This option is available for both local and global option passwords.

Note: You can modify only the Global Option password from the Session/Encryption password dialog by right clicking on the job in the job queue.

- **Remind to change password n days after specifying a password**--Specify the number of days a password is valid. Seven days prior to the specified number of days, a message prompting you to change your password will be logged in the Activity Log.

Example:

On Jan. 1 you set n to 30 days. On Jan. 24 the message The backup job password will expire in 7 days, will appear in the Activity Log. On Jan. 31 the message The backup job password has been expired. Please change it now appears in the Activity Log.

Encryption/Compression methods

- **Encrypt data**--Use this option to encrypt the backup data. You can specify one of the following options:
 - **At agent**--Select this option to encrypt the backup data prior to the actual backup process. For more information about this option, see [Data Encryption at the Agent Server](#) (see page 113).
 - **At backup server during backup**--Select this option to encrypt the backup data at the backup server during the backup process. For more information, see [Data Encryption During Backup](#) (see page 114).
 - **At backup server during migration**--Select this option to encrypt the backup data during the migration phase of a staging backup job. For more information, see [Data Encryption During Migration](#) (see page 114).

Note: This option can be enabled when submitting migration jobs for cloud devices whether the compression is enabled or not.

If you encrypt data during the backup phase, Arcserve Backup will not encrypt the data again during the migration phase of the staging backup operation.

- **Compress data**--Use this option to compress the backup data. You can specify one of the following options:
 - **At agent**--Select this option to compress the backup data on the system where the agent is installed and running.

Note: Arcserve Backup does not support data compression at the agent system when the backup source consists of UNIX, Oracle RMAN data.
 - **At backup server**--Select this option to compress the backup data at the Arcserve Backup server during the backup process. This option lets you compress files before backing them up using a software compression algorithm.

Be aware of the following behavior:

- You must specify either Encrypt data at backup server during backup or Encrypt data at backup server during migration to enable compression at backup server.
- With the Compress data, At backup server options specified, and the Encrypt data at backup server during backup option, or the Encrypt data at backup server during migration option specified, Arcserve Backup uses software compression to compress the data at the backup server before the data is encrypted at the backup server.
- If the storage device associated with the job does not support hardware compression, Arcserve Backup ignores the setting Compress data, At backup server.

Specify Encryption and Compression Options by Migrating Jobs to Cloud Devices

To specify encryption and compression options by migrating jobs to cloud storage

1. Open the Backup Manager and click Options.
2. Click the Encryption/Compression tab.
 - a. Enter Session/Encryption password
 - b. Check Encrypt data and select at Backup Server during migration option
 - c. Check Compress data and select at Backup Server option.

Note: The Compress data option is disabled when Compression for the cloud-based device is disabled.

3. Click OK.

The encryption and compression options are applied.

Backup Manager Volume Shadow Copy Service Options

You can specify global options for using the Volume Shadow Copy Service (VSS). These options affect all Writers for VSS backups, but they do not apply to transportable VSS backups.

Note: For more information on VSS, see the *Microsoft Volume Shadow Copy Service Guide*.

On the Volume Shadow Copy Service tab, the File System Backup group box lets you specify how you want Arcserve Backup to handle open files during file system backups. These options do not affect Writers and Components.

- **Use VSS**--Directs Arcserve Backup to use VSS to handle the backup of open files.

If this check box is not selected, VSS support is not used and the Arcserve Backup Agent for Open Files (if available) is used to handle open files. If the Arcserve Backup Agent for Open Files is not available and Use VSS is not selected, a traditional backup is performed. However, the backup will be incomplete if there are any open files that cannot be backed up.

- **Revert to traditional backup if VSS fails**--Lets Arcserve Backup execute a traditional backup if an attempt to create a VSS backup fails.

Note: This option requires you to install and license the Arcserve Backup Agent for Open Files on the source computer.

Be aware of the following behavior:

- If the Agent for Open Files is installed and licensed on the source computer, Arcserve Backup uses the agent to handle open files when this option is specified and the VSS backup fails.
- If the Agent for Open Files is not installed or not licensed on the source computer, Arcserve Backup processes VSS backups as traditional backups, regardless of whether this option is specified.

The Writers and Components group box lets you specify how you want Arcserve Backup to treat Writers and Components. These global options affect all Writers, except for those with Writer-specific options in place. For more information about setting Writer-specific options, see the *Microsoft Volume Shadow Copy Service Guide*.

- **Files included by a writer will be excluded from file system backups**--Prevents files that belong to a Component from being backed up by a traditional file system backup. This option offers the following advantages:

- Avoids backing up files that have already been backed up by VSS.
- By excluding files from traditional backups, fewer files are processed, and traditional backups take less time to complete.
- Helps achieve successful backups by eliminating certain problems associated with files that must be processed as a group; for example, files associated with a Writer or database application. In a traditional backup, there is no mechanism to ensure that the files are processed together.

- **Files excluded by a writer will be excluded from file system backups**--Prevents files that have been excluded from being backed up by a Component from being backed up by a traditional file system backup.

There may be files associated with an application that should never be backed up (for example, the Windows page file). Each Writer is aware of whether its associated application maintains any such files. Selecting this option allows Arcserve Backup to use this information when performing traditional backups.

- **If a component file fails to backup the writer, the backup will terminate**--Cancels the backup of a Writer if the backup of any of the Components fail. The backup of a Component fails if one or more of its files cannot be successfully backed up.

Selecting this option ensures that any backup is consistent and that all of the files associated with a Writer are backed up before the backup is considered successful, regardless of how many components are associated with the Writer.

Backup Manager Backup Media Options

You can specify the overwrite/append rules for the media used in your backup job while you are configuring the job. This section describes the rules so that you can determine which method is best for your purposes.

Arcserve Backup allows up to 20000 sessions on a single tape and up to 101 sequences of a series of spanned tapes. Keep this in mind when planning your backups, because, if your sessions are small, you can reach 20000 sessions quickly. If you have a large amount of data to back up, you can quickly exceed 101 sequences, depending upon how much data each tape can hold. You can stop appending the data to the tape when the sequence number reaches the maximum limit and start a new tape set by selecting the Overwrite Same Media Name, Blank Media or Overwrite Same Media Name, or Blank Media First, then Any Media option.

For a file system device (FSD), the limitation is 4,294,967,295 sessions on a single FSD.

Jobs may fail when you back up older versions of Arcserve Backup database and application agents to FSDs that contain more than 65,535 sessions. Therefore, you must upgrade any older versions of Arcserve Backup database and application agents to this release to avoid job failure on these devices.

First Backup Media

The first backup media is the media you use when the backup job begins. The options for the first media determine the overwrite rules for the first media that is used for the backup job:

Note: If the "Use Rotation Scheme" backup option is selected on the Schedule tab, the Rotation Rules override these options.

- **Append to Media**--Append job sessions to the selected media.
- **Overwrite Same Media Name, or Blank Media**--Overwrite the media in the drive only if it is the one you specified for the job or if the media is blank. If neither of these conditions are met, Arcserve Backup prompts you to supply the specific media name.

- **Overwrite Same Media Name, or Blank Media First, then Any Media**--Overwrite any media found in the drive. If you select this media option, Arcserve Backup checks to see if the media in the drive is the one specified for the job. If it is not, Arcserve Backup checks to see if the media is blank. If the media is not blank either, Arcserve Backup reformats whatever media it finds in the device and starts backing up files at the beginning of the media.
- **Timeout for First Media**--Number of minutes Arcserve Backup attempts to write to media before canceling job or selecting a different media.

Be aware of the following:

- The deduplication device does not support **Overwrite Same Media Name**. The backup job will always append to the deduplication device even if **Overwrite Same Media Name, or Blank Media** or **Overwrite Same Media Name, or Blank Media First, then Any Media** is selected.
- If you want to format the deduplication device, format it manually in Arcserve Backup Manager.
- Arcserve Backup ignores the First Backup Media options that you specified when the Append days option is specified as a custom job option. For more information, see Repeat Method Options.

Additional Backup Media

These options apply to jobs that require more than one media to determine the overwrite rules for the additional media. You need to specify which media Arcserve Backup can use when the job spans media.

Note: If the "Use Rotation Scheme" backup option is selected on the Schedule tab, the Rotation Rules will override these options.

- **Overwrite Same Media Name, or Blank Media**--Write to the media in the device only if it has the same media name (but a different media ID) or if it is blank. Arcserve Backup remembers the name and ID of the job's first media. When the job requires additional media, Arcserve Backup checks if the new media has the same name (but different media ID) or if it is a blank media. As long as the ID is different, Arcserve Backup reformats the media, giving it the same name and ID as the first media. The sequence number changes.

Note: To overwrite media based on its name only, select the Distinguish Media by Name Only option.

- **Overwrite Same Media Name, or Blank Media First, then Any Media**--Overwrites any media found in the device (as long as it has a different ID from the first media's ID). If neither of these conditions are met, Arcserve Backup reformats whatever media it finds in the drive and starts backing up files at the beginning of the media. All subsequent media are reformatted with the same name and ID as the first media. Only the sequence number changes.

Note: To overwrite media based on its name only, select the Distinguish Media by Name Only option.

- **Timeout for Additional Media**--Number of minutes Arcserve Backup pauses before attempting write backup data to the same media, write backup data to different media, or cancel the job.

Distinguish Media by Name Only

Arcserve Backup writes to any media that has the name specified in the Media text box on the Destination tab, regardless of the media's ID or sequence number. This option is useful if you are running a repeating Overwrite job with a specific media and you want to ensure that the same media is used for the job each time.

When this option is not enabled, the second time the backup job is run, Arcserve Backup might not be able to locate the original tape because some of its identifying features will have changed. When this option is enabled, however, Arcserve Backup simply looks for a media that has the name specified in the Media text box and uses it, regardless of the media's other identifying features.

Note: If more than one media in the tape library have the same name, Arcserve Backup will use the first media in the device group that matches the specified name. Therefore, we do not recommend that you use this option to perform a one-time overwrite.

Backup Manager Verification Options

Arcserve Backup allows you to verify that your data was correctly backed up to media. You can verify data for the entire backup job or for a selected drive in your backup job. The global verification options (applied to the entire job) will be overridden by the options selected for a drive. Arcserve Backup provides the following options for verification:

- **None**--The backup will not be verified.

- **Scan Backup Media Contents**--Check the proprietary Arcserve Backup data area (the header) of each file on the backup media. If it is readable, Arcserve Backup assumes the data is reliable. If it is not readable, the Activity Log is updated with this information. This is the fastest verification method.

If you selected Calculate and Store CRC Value on Backup Media on the Operation tab, Arcserve Backup automatically performs CRC verification. This method assigns a value to the data that you copied to media and compares it to the value assigned to the data that you backed up. This enables you to identify the individual data packets that were backed up.

- **Compare Backup Media to Disk**--Data from the backup media is read and compared byte for byte against the source files. This option takes time, but ensures that all data on the backup media are exactly as on the disk. If Arcserve Backup finds a mismatch, the errors are recorded in the Activity Log.

Backup Manager Operation Options

The operation options for backup determine related actions that occur during or after the backup, and the level of information that is recorded in the database. Arcserve Backup provides the following options:

Append Backup of Arcserve Backup data at the end of job Options

The following options affect how the level of information that is recorded in the Arcserve Backup database for the Arcserve Backup underlying database.

- **Arcserve Backup database**--This option allows to explicitly select the Arcserve Backup database or instance from the Backup Manager, Source directory tree with all backup jobs.
- **Catalog files**--This option allows you to back up the related Arcserve Backup database catalog files when the backup job is complete.
- **Job scripts**--This options allows you to back up the related job scripts when the backup job is complete.
- **SQL Server Disaster Recovery Elements for the Arcserve Backup Database**--This option ensures that the elements required to recover a SQL Server database from a disaster are backed up after jobs are complete.

Operation Options

The following options affect only backup operations.

- **Disable File Estimate**--By default, file estimation is disabled. To enable file estimation, deselect this option so that before any file is backed up to media, Arcserve Backup performs an estimate of how long the job will take.

Be aware of considerations that follow:

- File estimation is no longer the default value.
- If you are in a Novell server environment and you select Disable File Estimate from the Operations tab of the Global Options dialog, when you look at the View Job Queue/Statistics window from the back-end, there is no status bar at the bottom of the window.
- **Calculate and Store CRC Value on Backup Media**--Calculating and storing the CRC value on the backup media will enable Arcserve Backup to perform CRC verification during the backup job. To instruct Arcserve Backup to use the CRC value stored on media, see the Backup Options, Verification tab.
- **Delete Source Files After Backup to Media (use with caution)**--This argument deletes the files from the hard disk after the file backup is completed. Select this option if you want to delete source files from the source machine after they have been backed up to media. This option deletes only the files from the specified unprotected folder. It does not delete the empty folder itself.

You can use this option to perform disk grooming. For example, if you set up a backup job with a filter to back up files that haven't been accessed for a certain period of time, you could then include this option to delete those files from the source disk.

Be aware of the considerations that follow:

- On Windows computers, protected system files and files that are excluded from the backup by other filters are not deleted. For a remote backup job or a 64-bit operating system local backup job or a Windows Server 2008 local backup, the Windows Client Agent backs up the files. After the backup, this option deletes only the files from the specified unprotected folder. It does not delete the empty folder itself. Boot files, however, are not protected and can be deleted.
- On Linux/UNIX and Mac computers, all files that are backed up are deleted, except for those in protected directories, such as /bin, /etc, and /lib. To designate additional directories as protected, add them to the groom.cntl file on the client agent machine.

Note: As a best practice, you should specify Verification options when using Delete Source Files After Backup to Media. With verification options, Arcserve Backup compares the source files to the backup data to ensure that backup data is identical to the source data. For more information, see [Backup Manager Verification Options](#) (see page 166).

- **Preserve File Access Time (Used for Windows file system only)**--This option directs Arcserve Backup to preserve the last access time of files when a backup is performed.

Note: The Access Time of a file is automatically updated by the operating system whenever a file is accessed (read or write). However, after a compare is performed, the Access Times of all the backed up files are also updated. Therefore, if you want to track whether or not a file has actually been accessed (and not just compared), you need to preserve the original access time.

- If this option is not selected (no check in box), the last file access time of any files that are backed up is updated to the new value that is present when the backup is completed. This is the default setting.
- If this option is selected (check in box), Arcserve Backup preserves the last file access time of any files that are backed as the original value that was present before the backup was performed.

Note: For UNIX based agents, you must apply this option locally. For more information, see [Local Backup Options for UNIX Agents](#) (see page 152).

- **Reset Archive bit for backup to deduplication device**--Select this option for custom backup jobs in which optimization is enabled to reset the archive bit on all files included in the job after the job completes. Optimization deduplicates only files that have changed since the last backup job, indicated by archive bits with a value of 1. Archive bits must be reset back to 0 so that subsequent backup jobs in which optimization is enabled can deduplicate only changed files. If you do not select this option with optimization-enabled jobs, files that have not changed since the previous backup could be included in subsequent jobs, resulting in significant performance reductions.
- **Back up deduplication device data**--Select this option if you want to forcibly include deduplication device files (index, reference and data files) in the backup job. These files are normally skipped in local backups. If you select this option, you should also enable the Use VSS option and disable the Revert to traditional backup if VSS fails option, both on the Volume Shadow Copy Service tab. If you forget to enable the Use VSS option and disable the Revert to traditional backup if VSS fails option, then the backup job automatically enables this option when the backup job runs.

- **Eject Backup Media upon Completion**--Select one of the following options:
 - **Use Default Device Setting**--Select this if you want to use the setting you selected during library configuration.
 - **Eject Media**--Select this if you want to eject media from the drive after the job finishes. This helps prevent any other job from overwriting information on this media. If you select this, it overrides the setting you selected during library configuration.
 - **Do not Eject Media**--Select this if you do not want to eject media from the drive after the job finishes. If you select this, it overrides the setting you selected during library configuration.
- Note:** For more information on library configuration, see Tape Cleaning and Changing Configuration Details.

Retry Missed Targets Options

- **Retry Missed Targets**--Reschedule a backup for any workstations, file systems, databases, and so on, that failed during the backup job.

You can specify one of the following reschedule options for a backup job:

After Job Finishes

Specifies the number of minutes that you want to elapse after the original job finished to start the makeup job.

Default: 5 minutes

Maximum: 1439 minutes

At

Specifies the time when the makeup is to run.

- **Max Times**--Specifies the maximum number of times to repeat the makeup jobs.

Default: 1 time

Maximum: 12 times

With respect to the Retry Missed Targets options, be aware of the following considerations:

- By default, Retry Missed Targets is enabled, After Job Finishes is selected, and the value of Max Times is 1.
- **File system backups**--If the backup job requiring a makeup job consists of file system backups, and the file system contains directories that reside in different volumes, the makeup job backs up only the failed volumes or directories. The makeup job does not back up the entire file system if it contained successful volume or directory backups.
- **Child jobs**--The child makeup jobs (makeup of makeup jobs) run based on the value specified for After Job Finishes.

- **Microsoft SQL Server backups**--If the backup job requiring a makeup job consists of Microsoft SQL Server instance backups, the makeup job backs up only the failed databases. The makeup job does not back up the entire instance if it contained successful database backups.
- **Microsoft Exchange Server backups**--If the backup job requiring a makeup job consists of Microsoft Exchange Server, database level backups, the makeup job backs up only the failed storage groups or mailbox database. The makeup job does not back up the entire database if it contained successful storage group backups. If the backup job consists of (Microsoft Exchange Server) document level backups, the makeup job backs up only the failed mailbox stores and databases. The makeup job does not back up the all of the items selected if it contained successful mailbox store and database backups.
- **Agent-based backups**--If the backup job requiring a makeup job consists of agent-based backups (for example, Sybase, Informix, Oracle, and so on), the makeup job attempts to back up all of the source that was specified (instances, databases, tables, and so on) for the backup. If the makeup job fails after one unsuccessful attempt, Arcserve Backup will create another makeup job that consists of all of the source selected for the original job, and submit the makeup job with a status of Hold.

More information:

[Local Backup Options for UNIX and Linux Agents](#) (see page 152)

Backup Manager Pre/Post Options

Pre and Post options let you run commands before and after jobs execute.

The list that follows describes commands that you can run using Pre and Post options.

- You can use the Pre option to stop the application that owns the data you are about to back up, and then use the Post option to restart the application after the backup is complete.
- You can use the Pre option to defragment a disk before a backup job starts.
- You can use the Post option to delete files from a disk after that backup is complete.

Be aware of the following behavior when using Pre and Post options:

- Arcserve Backup does not support running commands with executables that reside on remote systems.
- Using a Pre option and specifying an exit code prevents the backup operation from starting until after the Pre option process is complete.

- Using a Pre option and specifying an exit code and the Skip Operation option causes Arcserve Backup to skip the backup operation and, if specified, prevents the Post option process from starting.
- Post option processes start unless the following conditions are present:
 - An exit code is specified, the Skip Post Application option is specified, and the result exit code is equal to the exit code specified.
 - The result of the backup operation is equal to the value specified for the Do not run Command if option.
- Pre and Post options specified as global options run commands before a job starts or after a job finishes. Pre and Post options specified as node-level (local) options run commands before a node is backed up or after a node is backed up.

For example, a user submits a backup job consisting of nodes A and B. A Pre option is specified as a global option and a Pre option is specified for node B. Immediately before the job runs, the global Pre option executes. While the job is running, the Pre option specified for node B executes before node B is backed up.

Run Command Before Operation Options

Enter the path to and name of the file to be executed on the machine before the job takes off.

- **On Exit Code**--Arcserve Backup detects exit codes of other programs. You can specify the following options for a particular exit code:
 - **Run Job Immediately**--The job runs immediately if the selected exit code is returned.
 - **Skip Job**--The job does not run if the appropriate exit code is detected.
 - **Skip Post Application**--Skip any commands specified to run after the job if the appropriate code is detected.
- **Delay in Minutes**--Specify the delay in which Arcserve Backup waits before running a job when the appropriate exit code is detected.

Run Command After Operation Options

Enter the path and name of the file to be executed on the machine after the job is completed.

Do Not Run Command If Options

Specify for a command not to run if Arcserve Backup detects that a job fails, a job is incomplete, or a job is complete.

Note: This option is not available when you are using Arcserve Backup to manage a UNIX or Linux based server.

Run Before/After Command As Options

The User Name and Password corresponds to the system of the host server selected, and is required to check the system privileges on that server.

The user name and password entered into these fields should not be confused with the Arcserve Backup User Name and Password.

Example: Submitting a Job Using Pre and Post Commands

A user submits a job that backs up local volume C. The user wants to check and fix errors using chkdsk.exe on local volume C before local volume C is backed up. After the job is complete, the user wants to generate an error report using CAAdvReports.exe.

Pre Command for the node:

The command that follows checks and fixes errors on local volume C before the backup job starts.

```
chkdsk.exe C: /F", On Exit Code = 0, Run operation immediately
```

Post Command for the job:

The command that follows generates an error reports and saves it in a specified location.

```
CAAdvReports.exe -reporttype 5 -maxSize 5 -Server DUVDO01 -outfile "C:\Program Files  
(x86)\CA\ARCserve Backup\Reports\Backup Error Report_data.xml" -PastDays 1 -AutoName
```

Backup Manager Agent Options

You may select backup options on a per-database basis or define a set of default options for all databases in a backup job. The Agent then applies the options to each database as appropriate.

- **Agent Options/Database Level Options**--These are agent backup options and apply to only the selected database. They can either extend or override the Global Agent options. Access Database Level Options by right-clicking the database object and selecting Agent Option from the shortcut menu.
- **Global Options/Agent Options**--These options let you specify default job options for all selected objects in the Agent type. Global Agent Options are not supported by releases of the Agent prior to r12.5. Access Global Agent Options from the Agent Options tab of the Global Options dialog.

Global Agent Options applied at a global level let you specify default job options for all databases for the Agent selected. Options selected for a specific object at the database level can either extend or override the options specified as a global option. As a general rule, options applied at the global level will extend or override options that you specify on the Job Schedule tab.

Certain options are available from only one Agent Option dialog; they are noted.

Note: The agent combines options that you specify at the database level for a specific database with the appropriate global agent options.

You can specify Global Agent Options for the Arcserve Backup components that follow:

- [Agent for Microsoft SQL Server](#) (see page 174)--Includes the Agent for Arcserve Backup Database, which is supported by Arcserve Backup Agent for Microsoft SQL Server, r12.5 and later.
- Agent for Virtual Machines--Supported by the Arcserve Backup Agent for Virtual Machines, r12.5 and later.
- [Agent for Microsoft Exchange Server](#) (see page 183) - Includes Microsoft Exchange Server Database Level and Document Level options.

Consider the behavior that follows when packaging jobs using the above-described agents:

- Global agent options are not supported by releases of the above-described agents prior to Arcserve Backup r12.5, nor are they applied if you use any other agent to back up Microsoft SQL or Exchange Server databases.
- When you upgrade an older agent to Arcserve Backup r12.5 or later, the agent applies both any pre-existing local options and all global options that apply and do not conflict with the database level (local) options.
- For jobs packaged using older agents, local options are carried over as local options.

Agent for Microsoft SQL Server Options

The options described in the sections that follow affect all backups that include Microsoft SQL Server databases and the Arcserve Backup database at the job level.

Note: Database level agent options override settings made on the Global Agent Options tab.

Backup Method Options

The following backup methods are provided on both the Agent Options (database level) and Agent Options (Global Options) dialogs:

- **Use Global or Rotation Options**--Use Global or Rotation Options is the default setting.

Arcserve Backup can apply Incremental and Differential global backup methods from the Job Scheduler when backing up Microsoft SQL Server databases. This lets you use rotation schemes to perform differential and Transaction Log backups of Microsoft SQL Server databases, which are dynamically adjusted based on the limitations of each individual database.

In releases of the agent that pre-date Arcserve Backup r12, the Global Backup Method or Rotation Scheme option from the Job Scheduler overrides local database options. In this release, the Global Backup Method or Rotation Scheme is applied only if you selected Use Global or Rotation Options in the database level options for the database and in the Global Agent Options for SQL Server.

This option backs up the selected database using the Backup Method from the Job Schedule. The Job Methods are applied using the logic that follows:

- The Full job method will result in a Full backup of the database.
 - The Differential job method will result in a Differential backup of the database, unless this database has not yet had a Full backup.
 - The Incremental job method will result in a Transaction Log backup With Truncation for databases using the Full and Bulk-Logged Recovery Models, and a Differential backup of databases using the Simple Recovery Model, unless this database has not yet had a Full backup.
 - The three main System databases are exempt from the Job Method and from the Backup Method in the Global Agent Options; selecting this option for databases master, model, or msdb will always result in a Full backup.
- **Full**--A Full backup is performed. All files included in the Database Subset selected will be backed up in their entirety.
 - **Differential**--Backs up only data that has changed since the last Full backup. For example, if you ran a complete backup of your database on Sunday night, you can run a differential backup on Monday night to back up only the data that changed on Monday.

Note: When selected in the Global Agent Options, this option is ignored by system databases. Databases that have not received a Database Full Backup will revert to a Full Backup.

- **Back up Transaction Log After Database**--Backs up only the Transaction log. This option is only available for databases using the Full and Bulk-Logged Recovery Models. For databases using the Simple Recovery Model, Arcserve Backup performs a Differential backup when you select Transaction Log Only from the Global Agent Options tab.

Note: When selected in the Global Agents Options, this option is ignored by system databases. Databases that have not received a Database Full Backup will revert to a Full Backup.

The backup method selected on the Global Agent Options tab overrides the selection made in a job's global backup method or rotation phase in the Job Scheduler. If you select the backup method using the Global Agent Options tab, note the following:

- The three system databases (master, model, and msdb) are exempt from the Backup Method in the Global Agent Options tab.
- For databases that have not yet received a Database Full backup, Arcserve Backup ignores the backup method set in the Global Agent Options tab and performs a full backup by default.
- For databases using the Simple Recovery Model, Arcserve Backup performs a Differential backup when you choose Transaction Log Only on the Global Agent Options tab.

Because any selection other than Use Global or Rotation Method for a database overrides the selection in the Global Agent Options dialog, the Backup Method is not affected by the Override Global Options setting on the database's Agent Options (database level) dialog.

Database Subset

Database Subset options let you define the types of database components that you want to back up. You can use this option to choose between the entire database, or a selection of files and FileGroups contained within the database, when the size of the database and performance requirements do not allow you to back up the entire database.

Database Subset options are disabled if the selected Backup Method for a database is Transaction Log Only.

Important! Of the following Database Subset options, only the Back up Transaction Log After Database option is available on the Global Options/Agent Options dialog.

- **Entire Database**--Backs up the entire database.
- **Files and FileGroups**--Backs up selected files in a database. Use this option to back up a file or FileGroup when the database size and performance requirements make it impractical to perform a full database backup. This option is only available for databases using the Full and Bulk-Logged Recovery Models.

Note: For Microsoft SQL Server 7.0 databases, Arcserve Backup performs a Files and FileGroups Full backup if you set the database subset to Files and FileGroups for the database level and backup method to Differential on the Global Agent Options tab.

- **Partial Database**--Backs up the Primary FileGroup, and any other Read-Write FileGroups. For a Read-Only database, only the Primary FileGroup will be backed up. This option requires SQL Server 2005 or later.
- **Back up Transaction Log After Database**--Backs up the Transaction Log after the database, partial database, or selected set of data files is backed up. This allows you to perform a Full backup or Differential backup and a Transaction Log backup in the same job. This option is only available for databases using the Full and Bulk-Logged Recovery Models and is ignored for databases using the Simple Recovery Model if set in the Global Agent Options dialog.

Be aware of the following behavior:

- If you select this option on the Global Options/Agent Options tab and specify Incremental backup method using the Job Scheduler, Arcserve Backup performs only one transaction log backup on the database and uses the Transaction Log Truncation Options from the Global Agent Options tab instead of the Incremental backup's default behavior (truncate the transaction log).
- If you select this option using the Database Level Agent options, set the backup method at the database level to Use Global or Rotation, and set the backup method in the Global Agent Options to Transaction Log Only, Arcserve Backup performs only one transaction log backup on the database and uses the Transaction Log Truncation Options set from the database level.
- If you select this option in the Global Agent Options dialog and specify Transaction Log Only in Database Level Options, this option and the accompanying global Log Truncation Options setting is ignored for that database.

Log Truncation Options

Log Truncation Options are accessible from the Database Level/Agent Options and Global Options/Agent Options dialogs:

- **Remove inactive entries from transaction log, after backup**--(Truncation)
Truncates the Transaction Log files, removing entries included in the backup so that the space in the files can be reused. This is the default option.
- **Do not remove inactive entries from transaction log, after backup**--(No truncation)
Retains backed up log entries after backup. These entries will be included in the next Transaction log backup.
- **Back up only the log tail and leave the database in unrecovered mode**--(No recovery)
Backs up the log and leaves the database in a restoring state. This option is available for Microsoft SQL Server 2000 or later. Use this option to capture activity since the last backup and take the database offline prior to restoring or repairing it.

The Log Truncation Options are available only when the selected Backup Method is Transaction Log, or when the Backup Transaction Log After Database option is checked.

Important! Do not use the "Backup only the log tail and leave the database in unrecovered mode" log truncation option to back up the ARCserve Database. Performing a backup with this option causes the database to be placed in an offline status, and you can lose the ability to find the backups of the ARCserve Database in order to perform a restore and bring the database online. If you perform a backup of the ARCserve Database using this option, you can use ARCserve Database Recovery Wizard to recover the Arcserve Backup database and bring it back online.

Log Truncation Options are not affected by the Override Global Options setting on the Database Level Agent Options dialog. If Transaction Log or Backup Transaction Log After Database is selected for the database, the database Log Truncation Options are used.

Database Consistency Check (DBCC) Options

A database consistency check (DBCC) tests the physical and logical consistency of a database. DBCC provides the following options:

- **Before Backup**--Checks consistency before the backup of the database.
- **After Backup**--Checks consistency after the backup of the database.
- **Continue with backup, if DBCC fails**--Performs a database backup even if a consistency check before backup reports errors.
- **Do not check indexes**--Checks the database for consistency without checking indexes for user-defined tables.

Note: The system table indexes are checked regardless of whether you select this option.

- **Check only the physical consistency of the database**--Detects torn pages and common hardware failures, but does not check the data against the rules of the database schema. It still checks the integrity of the physical structure of the page and record headers, and the consistency between the page's object ID and index ID. This option is available for Microsoft SQL Server 2000 or later. If this option is selected from the Global Agent Options tab, it is ignored for SQL Server 7.0 databases.

All error messages that are generated during the DBCC are recorded in the Agent for Microsoft SQL Server log file called sqlpagw.log. The log is located in the Backup Agent directory.

Example: How DBCC Options Work

The following example illustrates how DBCC options work in conjunction with Override Global Options on the Agent Backup Options dialog.

- With Override Global Options specified, the DBCC options selected at the database level will be the only DBCC options specified.
- With Override Global Options not specified, all of the DBCC options specified for the database and all of the DBCC options selected in the Global options will be applied together.

On the Global Options/Agent Options tab, the Database Consistency Check options that follow are specified:

- After backup
- Do not check indexes

On the Agent Backup Options dialog, Override Global Options is not selected and the Database Consistency Check options that follow are specified:

- Before backup
- Continue with backup, if DBCC fails

Note: To open the Agent Backup Options dialog, open the Backup Manager, click the Source tab, browse to and expand the Arcserve Backup server, right-click the Arcserve Backup database, and then select Agent Option from the pop-up menu.

When you submit the backup job, Arcserve Backup applies the DBCC options specified in logical order: Perform the DBCC before the backup starts. If the DBCC fails, perform the backup. After the backup is complete, do not check the indexes.

Other Options

From the Global Agent Options tab, you can specify the following additional options:

- **Include Checksum Generated by SQL Server**--Includes error checking information from Microsoft SQL Server, which can be used to validate the integrity of the backed-up data during restore. This option requires SQL Server 2005 or later and is ignored when set in the Global Agent Options dialog for SQL Server 7.0 or 2000 databases.
- **SQL Native Backup Compression**--This option applies to only SQL Server 2008 (Enterprise) and later versions. If enabled, this option directs Arcserve Backup to use SQL Server database backup compression settings, resulting in faster backup times and smaller sessions.

From the local Agent Option dialog, you can choose to Override Global Options. This setting lets you choose a backup method and database consistency check that applies to only the database selected for this job.

Agent for Virtual Machines Options

The options that follow affect all VM backups in your environment at the job level.

Backup Mode Options

The options that follow determine the backup method used for the backup.

- **File Mode**--Lets you protect individual files and directories. File mode backup lets you perform the tasks that follow:
 - Back up files and directories at file level granularity contained in VM.
 - Perform full, incremental, and differential backups.
 - Restore data at file level granularity.
 - Process multiple streams of data simultaneously using the Multistreaming option.
 - Filter data using the Filter option.

Note: The elapsed time required to perform a file level backup of a full VM is greater than the elapsed time required to perform a raw (full VM) level backup of the same volume.

- **Raw Mode**--Lets you protect entire systems for disaster recovery. Raw mode backup lets you perform the tasks that follow:
 - Perform full backups of full VM images only.
 - Process multiple streams of data simultaneously using the multistreaming option.

Note: Raw mode does not let you restore data at file level granularity or filter raw (full VM) data. Filters applied to raw mode (full VM) backups are ignored at runtime.

- **Mixed Mode**--Mixed mode is the default backup mode. Mixed mode lets you perform the tasks that follow:
 - Perform GFS and rotation backup jobs that consist of weekly full backups in full VM (raw) mode and daily incremental and differential backups in file mode in a single backup job.

Note: Rotation and GFS rotation jobs are advantageous in that they contain backup data that provides you with daily protection (file level backups) and disaster recovery protection (raw, full VM backups) in a single backup job.

- **Allow file level restore**--Lets you back up data using Raw Mode efficiency and restore data with File level granularity. To perform granular file level restores from raw (full VM) backups, you must specify the name of the Arcserve Backup server on your VMs. For more information, see Specify the Name of the Arcserve Backup Server.

Allow file level restore lets you perform the tasks that follow:

- Restore data at file level granularity from Raw Mode (full VM) backups.
- Restore data at file level granularity from Mixed Mode backups.

With the Allow file level restore option, Arcserve Backup demonstrates the following behavior:

You can use the Allow file level restore option with all types of backups, including custom backups, rotation backups, and GFS rotations that consist of full, incremental, and differential backups. The full backups are captured in raw (full VM) mode and the incremental and differential backups are captured in file level backup mode. If you do not specify Allow file level restore, Arcserve Backup restores only the incremental and differential backups. The full backup, which is captured in Raw mode, is not included with the restore.

Transport Mode for VMware VM Options

The options that follow determine the transport method used for backup for VMware virtual machines.

- **Dynamic**--(*default*) This option lets VMware Virtual Disk Development Kit (VDDK) select the available transport mode.
- **SAN**--(Storage Area Network) This option lets you transfer backup data from proxy systems connected to the SAN to storage devices using Fibre Channel communication.
- **HOTADD**--This option lets you back up virtual machines configured with SCSI disks.
- **NBDSSL**--(Network Block Device Secure Sockets Layer) This option lets you use the Network File Copy (NFC) protocol to communicate. NBDSSL transfers encrypted data using TCP/IP communication networks.
- **NBD**--(Network Block Device, also referred to as LAN transport mode) This option lets you use 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.

Note: If the specified transport mode is not available, the transport mode defaults back to the Dynamic option.

Incremental / Differential Method for VMware VM Options

Lets you specify the communication method that Arcserve Backup will use to transfer incremental and differential backup data on VMware VMs to the backup proxy system.

- **Use VDDK**--Lets Arcserve Backup use VMware Virtual Disk Development Kit to transfer incremental and differential backup data to backup proxy systems. Use this option when you want to reduce the load on your network.
Note: Use VDDK is the default setting.
- **Use Client Agent**--Lets Arcserve Backup use Client Agent for Windows communication to transfer incremental and differential backup data to backup proxy systems. With this option specified, Arcserve Backup transfers data via your network.

Agent for Microsoft Exchange Server Options

Global Agent Options are available at the Database Level and at the Document Level.

Database Level Options

- **Use globally scheduled Custom or Rotation backup method**--This option is enabled by default. Clear this option to activate the Backup Method options. If you enable this option, you must specify a backup method on the Schedule tab of the Backup Manager when you configure the job, which is applied to all databases selected for backup.
- **Full backup**--This option backs up the entire database.
- **Copy backup**--This option also backs up the entire database but does not purge log files.
- **Incremental backup**--This option backs up only the changes that have occurred since the last backup, regardless of backup method.
- **Differential backup**--This option backs up only the changes that have occurred since the last full backup.
- **Exchange Server 2010/2013 Options:**

For Exchange Server 2010/2013 database backups, you can back up from a replica or from an active database. The replica is selected based on the Exchange Server database copy activation preference.

 - **Back up from replica**--This is the default setting and activates the Database Availability Group options. If you choose to back up from a replica, you may enable the option, Back up from active if there is no healthy replica available.
 - **Database Availability Group Options:**
 - First Preferred
 - Last preferred
 - Back up from active database

Document Level Options

- **Use globally scheduled Custom or Rotation backup method**--This option is enabled by default. Clear this option to activate the Backup Method options. If you enable this option, you must specify a backup method on the Schedule tab of the Backup Manager when you configure the job, which is applied to all databases selected for backup.
- **Full backup**--This option backs up the entire database.
- **Incremental backup**--This option backs up only the changes that have occurred since the last backup, regardless of backup method.

- **Differential backup**--This option backs up only the changes that have occurred since the last full backup.
- **Time-Based Backup**--Select this option to back up mailboxes newer than, or older than the date specified in the Date field, or in the field labeled, Days prior to when the job runs. You may also enable the Purge document after backup option to remove documents as they are backed up.

Backup Manager Job Log Options

The log options determine the level of detail that is included in the log report for the operation. The log options can be set in the following windows: Backup, Restore, Compare, Media Assure & Scan, Copy, Count, and Purge. Arcserve Backup provides the following log options:

- **Log all activity**--Record all of the activity that occurs while the job is running in the Job Log.

Note: When you specify Log all activity, Arcserve Backup creates a log file named JobLog_<Job ID>_<Job Name>.Log. With this log file, you can view detailed logging information about the job. Arcserve Backup stores the log files in the following directory:

C:\Program Files\CA\ARCserve Backup\LOG

- **Log summary only**--Record summary information on the job (including source, destination, session number, and totals) and errors.
- **Log disabled**--Do not record any information about this job in the Job Log.

Files and Objects that Arcserve Backup Does Not Back Up

Arcserve Backup does not back up the following files while processing backup jobs:

- DOS system files
- The following Windows files:
 - 386SPART.PAR
 - 386SPART.TMP
 - SPART.PAR
 - WIN386.SWP

- DoubleSpace files (DBLSPACE with any extension)
- DriveSpace files (DRVSPACE with any extension)
- Stacker files (STACVOL.DSK)
- BTrieve delta files
- The following Win32System files:
 - PAGEFILE.SYS
 - NETLOGON.CHG
 - NTUSER.DAT.LOG
 - CPL.CFG
 - EA DATA.SF
- The following Arcserve Backup files:
 - RDS.BAK
 - RDS.LOG
- The following registry keys:

\\HKEY_LOCAL_MACHINE\SYSTEM\CLONE

\\HKEY_LOCAL_MACHINE\HARDWARE

\\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\BackupRestore\FilesNotToBackup (and all files specified)

The above registry key is controlled by the Arcserve Universal Agent. The Arcserve Universal Agent service runs under the Local System account. As a result, Arcserve Backup will back up the files specified under FilesNotToBackup for the Windows Administrator account. If you do not want to back up the files specified under the Windows Administrator account, you must exclude the files explicitly under the above registry key. Alternately, you can modify the Arcserve Universal Agent service to run as the specific Windows Administrator account. To allow the Arcserve Universal Agent service to run as the specific Windows Administrator account, do the following:

1. From the Control Panel, choose Administrative Tools and then choose Services.
2. Click the Arcserve Universal Agent service from the Services list.
3. Click Action, Stop to stop the service from running.
4. Right-click the service and click Properties.
5. On the Log On tab of the Properties dialog, click This Account and provide the required credentials.
6. Click OK.
7. Restart the Arcserve Universal Agent service.

- Files with the extensions *.ALT and *.LOG that are located in the paths listed in the following registry key:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\hivelist

- Cluster installation directory files (typically, the WINNT\cluster folder), files with no extensions, and files with the extension *.LOG
- \RECYCLER folder
- \Document and Settings\Administrator\Local Settings\Temp folder
- %systemroot%\Temp (all files and subfolders inside)
- Folders for file system devices.
- If database agents are installed, all files backed up by the agents are skipped.
- Database folders in the Arcserve Backup home directory are skipped during regular file backup operations.

Skip or Include Database Files in Backups

Effective with Arcserve Backup r12, there are two registry keys used to include or skip certain database files during backup jobs. Use of these keys is determined by the type of database agent you are using.

SkipDSAFiles Registry Key

The following agents use the SkipDSAFiles registry key:

- Agent for Oracle, Agent for SAP R/3 (r12.1 and earlier versions)
 - *.dbf
 - Control*.*
 - Red*.log
 - Arc*.001
- Agent for Domino
 - *.nsf
 - *.ntf
 - Mail.box
- Agent for Sybase
 - Physical file of Master device
 - Physical file of non-Master device
 - Physical file of Mirror device
- Agent for Informix
 - *.000

To use the SkipDSAFiles registry key

1. When performing agent backups:
HKEY_LOCAL_MACHINE\SOFTWARE\Computer Associates\CA ARCserve Backup\ClientAgent\Parameters
2. Set the registry key to Value Name: SkipDSAFiles
Type: DWORD
Value: 0 to back up and 1 to skip

BackupDBFiles Registry Key

The following agents use the BackupDBFiles registry key:

- Agent for Microsoft SQL Server

The list of data and transaction log files that are part of online databases is retrieved from Microsoft SQL Server at the start of a file backup. This list typically includes, but not exclusively:

- *.ldf
- *.mdf
- *.ndf

Except distmdl.mdf, distmdl.ldf, mssqlsystemresource.mdf, mssqlsystemresource.ldf, which cannot be skipped. Also, if a SQL Server instance is shut down, the database files will not be skipped.

- Exchange Database Level Agent/Exchange Document Level Agent

- *.chk
- *.log
- Res1.log
- Res2.log
- *.edb
- *.stm

- Agent for Oracle (r12.5 or later)

- *.dbf
- Control*.*
- Red*.log
- Arc*.001

To use the BackupDBFiles registry key

1. When performing agent backups:

HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA
ARCserveBackup\ClientAgent\Parameters

2. Set the registry key to Value Name: BackupDBFiles

Type: DWORD

Value: 0 to skip and 1 to back up (0 is default)

Enable Arcserve Backup to Manage Open Files on Remote Computers

If the Arcserve Backup Agent for Open Files is installed on any of your computers, you can manage the BAOF Engine directly from the Backup Manager.

To enable Arcserve Backup to manage open files on remote computers

1. Open the Backup Manager and select the Source tab.
2. From the Source directory tree, select the system that you want to manage remotely.
3. If this server has the BAOF Engine installed, right-click the computer and select the following menu items or simply click these items in the Additional Information pane, on the bottom-right of the screen:
 - **Configure Open File Agent**--This displays the BAOF Configuration screen. From here, you can set various global settings for BAOF on the selected computer.
Note: For more information about the General, File/Group, and Clients options, see the *Agent for Open Files Guide*.
 - **View Open File Agent Status**--This displays the BAOF Status screen. This shows which files and groups BAOF is currently processing on the selected computer.
 - **View Open File Agent Log File**--This displays the Log File Viewer screen. This shows the log file for the selected computer.
4. Click OK.

You have successfully applied the open file settings.

Multiplexing Job Options

To submit a multiplexing job, you must enable the Multiplexing feature on the Destination tab in the Backup Manager. In addition, you can select any of the following:

- Multiplexing media (multiplexing media appear with a blue circle with an M next to them)
- Blank media
- Media pool

Note: You cannot submit a multiplexing job to a tape library that has WORM media unless you use the Virtual Library option to split the tape library into groups so that one has WORM media and the other does not. If you do this, you can submit a multiplexing job to the group that does not have WORM media. For more information about the Virtual Library option, see Virtual Library Configuration Option.

More information:

[Virtual Library Configuration Option](#) (see page 380)

Specify Multiplexing Options

Arcserve Backup lets you process backup data using multiplexing. For more information, see [How Arcserve Backup Processes Backup Data Using Multiplexing](#) (see page 106).

Be aware of the following behavior:

If a backup job with multiplexing spawns child jobs, the actual number of streams spawned will not exceed the number of streams specified for the job. However, if a job spawns multiple child jobs and the value specified for the Multiplexing Max Number of Streams option is one, the child jobs will be created and backed up in one continuous stream (the default Max # Streams is four).

To specify multiplexing options

1. Open the Backup Manager window and select the Destination tab.
Note: If the job is a staging backup, click the Staging Location tab.
2. Check the Multiplexing check box to enable multiplexing.
3. Specify a maximum number of streams.

The Maximum Number of Streams option defines the maximum number of streams that can write to media at the same time.

Default: 4

Range: 1 to 32

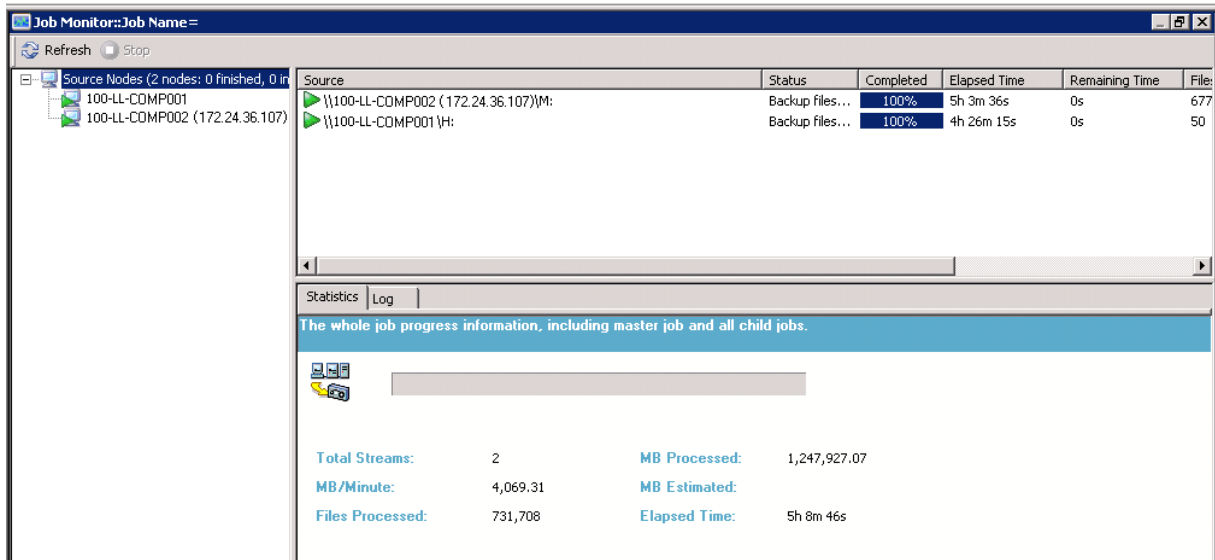
How the Job Status Manager Monitors Multiplexing Jobs

After you submit a multiplexing job, you can monitor of the job using the Job Status Manager. In the Job Queue, multiplexing jobs appear in levels so that you can view the status of child jobs related to the parent job.

To view the child jobs, open the Job Status Manager, select the Job Queue tab, and then select and double-click the parent job as illustrated by the following screen:

Job Name	Backup Se...	Job No.	Job ID	Status	Execution T...
[name] - Backup [Custom, MUx] Every 6 h...	100-LL-SE...	11	125	HOLD	3/12/2009 ...
[name] - Backup [Custom, MUx] Every 6 h...	100-LL-SE...	28	125	DONE	3/24/2009 ...

After you double-click the parent job, the child jobs appear in the Job Monitor screen as illustrated by the following screen:



In addition, the status of the parent job is the highest severity status of a child job. For example, if Child 1 is successful, Child 2 is incomplete, and Child 3 has failed, the parent job will denote a FAILED status.

Verify Multiplexing Data Integrity

If you want to verify the integrity of your data after your multiplexing job completes, use the Media Assure & Scan Utility to enable the Scan files global option with CRC verification and perform a scan media job.

For more information, see [Media Assure & Scan Utility](#) (see page 32).

Using Multiplexing with Microsoft Exchange Backup Jobs

Use the following registry key to control over how Arcserve Backup backs up Exchange data when using multiplexing. You can back up Exchange storage groups or mailbox database on the same server sequentially, with one sub-job, or simultaneously. You must set this registry value on the backup server running the multiplex job.

Note: In prior versions of Arcserve Backup, all storage groups on the same server were backed up simultaneously when multiplexing was enabled.

Registry key

SingleStreamExchangeAgent

Location

HKEY_LOCAL_MACHINE\Software\Computer Associates\Arcserve Backup\Base\task\backup

Value

- 0 (Default) Arcserve Backup executes Multiplexing jobs at the storage group level.
- 1 Arcserve Backup executes Multiplexing jobs at the Exchange server level, which means all storage groups are backed up sequentially by one sub-job.

Specify Multistreaming Options

Arcserve Backup lets you process backup jobs using multistreaming. For more information, see [How Multistreaming Processes Backup Data](#) (see page 103).

To specify multistreaming options

1. Open the Backup Manager window and select the Destination tab.
2. Check the Multistreaming check box.

Specify a **Max (Maximum) Number of Streams** to use. The default number of streams is 4. If you installed the Enterprise Module, the supported range is between 2 and 32.

Note: If a backup job with multistreaming spawns child jobs, the actual number of streams spawned will not exceed the number of streams specified for the job. However, if a job spawns child jobs and you do not specify a number of streams to use, the child jobs will be created and backed up in one continuous stream.

Entire Node Backups

If you want to back up an entire node, Arcserve Backup provides the capability to backup all file systems and databases on the specified node. The benefits of backing up an entire node are as follows:

- You can direct Arcserve Backup to back up a selected node and all of its contents with a single click in the Backup Manager directory tree. Arcserve Backup backs up all file systems, databases, and drives in the directory tree when you specify the node.
- You can create a single backup job for the entire node. Tracking several to many backup jobs on a single node can become a difficult and time consuming maintenance task.
- You can modify the node without having to modify preconfigured backup jobs. For example, if you add a drive to the node, Arcserve Backup detects the new drive automatically and backs up the entire node when you run the backup job.

Note: This feature supports Centralized Cross-platform Management.

Back Up an Entire Node that Contains Database Files

When backing up a node that includes database files, you must provide proper authentication to access all databases when creating the backup job. Proper authentication includes the User Name and Password for the corresponding databases. You do not need to provide this authentication when the backup job runs.

To facilitate database authentication, Arcserve Backup presents the Security and Agent Information dialog when you are creating a backup job on an entire node. The Security and Agent Information dialog opens as you click the Submit toolbar button, or if you select Save or Save As from the File menu on the Backup Manager window.

The Security and Agent Information dialog serves two purposes:

- Display a list of all database files on the node.
- Set or change the User Name and Password for the database item selected in the Security and Agent Information dialog.

To back up an entire node that contains database files

1. Open the Backup Manager and select the Source tab.

2. From the Source directory tree, select the node that you want to back up and click Submit on the toolbar to submit the job.

If the node contains database files, the Security and Agent Information dialog opens to display a list of all databases on the node, User Names, and Passwords.

Important! Client Agent Password Security is only supported for the Client Agent for Windows, including the Agent for Open Files and VSS Writers. If Password Security is enabled and any database, application, or messaging agent is installed on the same machine as the Client Agent, whole node backup is not supported. System Security is the default setting. If you have enabled Password Security on the primary or stand-alone server, the automatically-generated Database Protection Job will fail for the job queue and the Catalog Database. You must not enable Password Security in the Agent Configuration on any machine running a database, application, or messaging Agent before submitting the job.

3. Optionally, to set or change a User Name or Password, click the Security button.

Enter the appropriate User Name and Password and click OK.

In the Security dialog, you must specify User Name and Password with backup rights on that machine. For example, Administrator or root.

Note: Arcserve Backup does not support logging in to systems with passwords that are greater than 23 characters. If the password on the system you are attempting to log in to is greater than 23 characters, you must modify the password on the agent system such that it is 23 characters or less, and then you can log in to the agent system.

4. Click OK.

The Submit dialog opens.

5. Complete the fields are required for the job and click OK.

The backup job for the entire node is submitted.

More information:

[Submit a Backup Job](#) (see page 136)

Create Repeating Backup Jobs

Repeating backup jobs let you automate the process of protecting systems in your environment. Arcserve Backup lets you create repeating back up jobs using Arcserve Backup rotation schemes or create custom rotation schemes.

You can perform full, synthetic full, incremental, differential, and Grandfather-Father-Son (GFS) backups.

The following steps describe how to set up schedule for a Normal backup, which includes backups to one data mover server. For information about performing staging backups, see [Back Up Data Using Disk Staging](#) and [Back Up Data Using Tape Staging](#) (see page 236).

To create repeating backup jobs

1. Open the Backup Manager window by clicking Backup in the Quick Start menu.
The Backup Manager window opens and the Start, Source, Schedule, and Destination tabs appear.
2. Click the Start tab if it is not selected.
Select the Normal backup type of backup job.
3. Click the Source tab.
The backup source directory tree appears.
4. From the Source tab, browse to and select the files you want to back up.

5. Click the Schedule tab

The scheduling options appear.

6. Select Use Rotation Scheme to use one of the pre-designed backup schemes.
 - You can choose a five or seven day schedule using incremental, differential, or full backups.
 - To modify a rotation scheme, highlight the day you want to change and click the Modify button.

For example, you might want to change a routine to initiate a full backup on Saturday.

Make the necessary modification to your schedule, and click OK.

- Use the calendar to review the backup plan you selected. Click the Calendar View tab to see a calendar view of your rotation scheme.
- You can also make or view changes to your backup schedule by clicking the Exceptions tab. To make additional exceptions to your schedule, click the Add button to open the Exceptions dialog. Click the Date drop-down menu to open a calendar from which you can select the date you want to change.

Select the Start Date and Execution Time as required.

(Optional) Click Enable GFS to perform Grandfather-Father-Son (GFS) backups.

Note: For more information, see [How to Use GFS Rotations](#) (see page 122).

(Optional) Specify a Daily Backup Method. This option lets you specify to perform full, incremental, or differential backups for daily backups.

(Optional) Specify Use WORM Media to back up data to read once, write many (WORM) backup media.

Note: For more information, see [How Arcserve Backup Supports Write Once Read Many \(WORM\) Media](#) (see page 428).

7. Click the Destination tab.

The destination options appear in a directory tree.

8. From the Destination tab, specify the (media) group where you want to back up your data.

(Optional) In the Media Pool Used field, enter the name of the media pool that you want to use for the rotation.

9. When you are finished, click Submit on the toolbar to submit the job.

Cross-Job Duplicated Source Check

When you submit a repeat or rotation or GFS backup job, Arcserve Backup checks if any of the source is part of an existing repeat backup job. If a duplicate source is found, a message appears asking whether you want to proceed with the duplicate source. If you click Yes, the job is submitted and Arcserve Backup backs up the duplicated source multiple times. If you click No, the backup job is not submitted; you can remove the duplicate sources and submit the job again.

For example, suppose you create two customized source groups, one for the sales servers of all geographic locations and the other for all the servers of a specific geographic location. Suppose you create two repeating backup jobs, one to back up all the sales servers every Friday and the other to back up all the servers of the geographic location every Friday. Your sales servers will be part of both the backup jobs. When you submit the second job, Arcserve Backup displays a message that the sales servers are part of two backup jobs and asks if you want to proceed with creating the second job. If you click Yes, the sales servers are backed up twice; if you click No, the job is not submitted, and Source tab is displayed, so that you can exclude the sales servers from the second job and submit it again.

Back Up Remote Servers


Before you can back up remote servers, Arcserve Backup must be installed and running on your server, and you must install the appropriate agent (in this case, the Client Agent for Windows) on the remote server.

To back up remote servers

Note: The scenario that follows describes how to back up a server running Windows Server 2003.

1. From the Backup Manager, select the Windows Systems object. Right-click, and select Add Machine/Object from the pop-up menu.

The Add Agent dialog opens.



Note: Alternatively, you can add servers using the Add/Import/Export Nodes method. For more information, see [Add, Import, and Export Nodes Using the User Interface](#) (see page 334).

2. Enter the host name of the remote server in the Host Name field.

Note: Arcserve Backup server names and Arcserve Backup domain names cannot exceed 15 bytes. A name totaling 15 bytes equates to approximately 7 to 15 characters.

Check the Use Computer Name Resolution check box or specify the IP address of the computer. Click Add to include the remote server for backup, and click Close.

Note: Arcserve Backup lets you treat the backup of multiple servers as one job. If you choose this method, Arcserve Backup automatically submits the tasks in the job queue as one job and backs up the servers one after the other. Alternatively, you can select the Schedule tab and specify when each job should run to have Arcserve Backup back up each machine as a separate job.

3. Select the remote machine, and click the + to the left of it. The Security dialog appears, prompting you for security and agent information.

Enter your user name and password. You must supply this information to verify that you have sufficient rights to browse the machine and perform a backup. Click OK.

Note: Arcserve Backup does not support logging in to systems with passwords that are greater than 23 characters. If the password on the system you are attempting to log in to is greater than 23 characters, you must modify the password on the agent system such that it is 23 characters or less, and then you can log in to the agent system.

4. Repeat Steps 1, 2, and 3 for each machine you want to add.
5. Choose the sources and a destination for each machine you want to include in the backup.

Before you click Submit to submit the backup job, you can set up an Alert to send you notification after the job runs.

Submit Static Backup Jobs

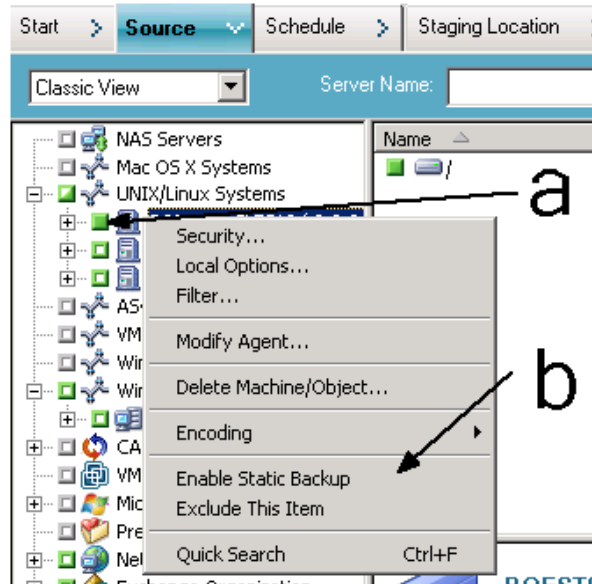
The Enable Static Backup option lets you submit scheduled backups of source groups and computers and maintain a static set of source volumes. With static backups, you can submit staging and deduplication backups that consist of full, incremental, and differential backups.

Static backups affect only the immediate subordinate objects of the source group or the computer specified when the job was submitted. Subsequent backups will not include objects and volumes that were added to the source group or the computer after you submitted the backup. Arcserve Backup backs up the original source volumes dynamically. Subsequent job executions will include changes to the files and folders contained in the original source volumes.


Note: For more information, see [Static Job Packaging](#) (see page 305) and [Dynamic Job Packaging](#) (see page 302).

To submit static backup jobs

1. Open the Backup Manager and click the Start tab.
The backup types appear.
2. Select Normal backup.
Click the Source tab.
The Source directory tree appears.
3. Select Classic View from the drop-down list.
 - a. Click the check box next to the computer name.
 - b. Right-click the computer and click Enable Static Backup on the pop-up menu.



The Static Backup option is applied to the computer, as illustrated by the following screen:



Note: The Static Backup option remains applied to the specified computer until you disable this option. You can repeat this step to disable the Enable Static Backup option.

4. Click the Schedule tab and specify the schedule that you want to use for the backup job.

Note: For more information, see [Rotation Schemes](#) (see page 307) and [Custom Schedules](#) (see page 318).

5. Click the Destination tab.

The Destination groups directory tree appears.

6. Specify the group that you want to use to store the backup data.

The storage group is applied.

7. Click Options on the toolbar.

The Options dialog opens.

8. Specify the options that you require for the job.

Note: For more information, see [Global Backup Options](#) (see page 154).

Click OK.

The backup options are applied.

9. (Optional) Expand the contents of the computer.

Select a drives or volume contained in the computer.

Repeat the previous two steps to apply options to the specified drive or volume.

(Optional) Repeat this step for all drives or volumes in the computer.

10. Click Submit on the toolbar.

The Submit Job dialog opens.

11. Complete the required fields on the Submit Job dialog and click OK.

The job is submitted.

Backup Staging Methods

Arcserve Backup provides two methods to backup to a staging area and then migrate (or copy) this data to a final destination (usually a tape).

- The disk staging method utilizes a disk as the staging area and is commonly referred to as Backup to Disk to Tape (D2D2T).
- The tape staging method utilizes a tape library or a virtual tape library as the staging area and is commonly referred to as Backup to Tape to Tape (D2T2T).


Each staging method contains specific options to control the behavior of Arcserve Backup during the backup process.

How Backup to Disk to Tape Works

Backup to disk to tape is a method of protecting data that lets you back up data to a temporary data storage location (staging area), and then based on selected policy options, migrate (copy) the data to the final destination (which could be a tape or disk), or automatically purge the data from the staging area after a specified duration time. When necessary, Arcserve Backup lets you restore data directly from the staging area.


Backup to disk to tape (B2D2T) is a two-part backup process.

- **Backup Process**--Arcserve Backup backs up data from the source to the staging location. The staging location is a file system device (FSD).
- **Copy Process**--Arcserve Backup copies or migrates the backup data from the staging location to the final destination media. The final destination is usually tape media, but can be another FSD.



Note: When you back up data from a disk staging location to the final destination, you are then able to perform a Double Tape Copy. Double Tape Copy lets you migrate (copy) the data to two destinations simultaneously. Using this feature writes to the destinations in parallel to save time.

To ensure that the migration job continues when one of the devices become unavailable, you can set a timeout value for additional backup media to continue the migration job in another device. If you do not set a timeout value, then the data migration job fails after the first backup media timeout completes. For more details about setting the timeout value, see [Backup Manager Backup Media Options](#) (see page 164).



Note: Arcserve Backup lets you transmit up to 32 streams of data using multistreaming. To back up data using multiple streams and transmit more than two streams of backup data, you must license the [Arcserve Backup Enterprise Module](#) (see page 40).

How to Use Disk Staging to Manage Backup Data

The following list describes how you can use staging to manage backup data.

- Using disk staging you can back up data to file system devices (FSD) that are used as a temporary staging area. A staging job can divide your backup job into several subjobs that run simultaneously. Disk staging lets you use simultaneous streaming to send multiple streams of data to the FSD. Since the data is split among several different streams, backup jobs with simultaneous streaming enabled can be completed significantly faster than regular backup jobs.

- You can then migrate (copy) the data from the FSD to a final storage media (or from disk to tape). As a result, the tape drive can be kept streaming, thereby minimizing the shoeshine effect (starting, stopping, and repositioning the tape), and increasing both the life expectancy and efficiency of the tape drive. While the backup image is still on the FSD, data can be restored directly from the FSD. The restore time is significantly reduced because restoring data from disk is generally faster than restoring from a tape (no delays due to tape load and seek latency).

Note: When you back up data from a disk staging location to the final destination, you are then able to perform a Double Tape Copy. Double Tape Copy lets you migrate (copy) the data to two destinations simultaneously which writes to the destinations in parallel to save time. To use this feature, it is required to configure a Disk Staging device to your environment. See [Device Configuration](#) (see page 383) for more details.

To ensure that the migration job continues when one of the devices become unavailable, you can set a timeout value for additional backup media to continue the migration job in another device. If you do not set a timeout value, then the data migration job fails after the first backup media timeout completes. For more details about setting the timeout value, see [Backup Manager Backup Media Options](#) (see page 164).

- During the backup-to-FSD process, if the FSD gets full or reaches the specified maximum threshold, Arcserve Backup lets you create makeup jobs which would then back up the data directly to the final destination after the staging backup job fails. This increases the success rate of backups. In addition, if there are any errors during the copy-to-final destination process, Arcserve Backup lets you create makeup jobs.

Note: Under disk full conditions, the makeup job created to back up the data to tape will always try to use a blank tape or a media from a scratch set. It will never try to append to an existing media.

- The backup images are kept on the FSD until the retention time expires (as determined by the specified purge policy). At that time, Arcserve Backup automatically purges the data from the FSD, and reclaims disk space so that backups can continue.
- For rotation jobs or GFS rotation jobs, Arcserve Backup lets you specify policies that disable staging for any particular day. This feature is helpful in situations where the FSD is full, is scheduled for maintenance, or has a problem.

More information:

[How to Use Tape Staging to Manage Backup Operations](#) (see page 232)

Disk Staging Capabilities

Using disk staging to store backup data provides the following capabilities:

- **File System Device Capacity Management**--Arcserve Backup lets you specify minimum capacity and maximum capacity thresholds of the file system device. The maximum threshold can be represented as either an absolute value or as a percentage of the capacity of the volume.
- **Ensures that Arcserve Backup does not use the full capacity of a disk**--A backup job will fail when writing to a file system device if the total disk space used exceeds the maximum threshold.

Important! File System Devices (FSD) that are part of a staging group cannot be erased or formatted using the corresponding utility from the Device Manager window. To prevent accidental erasing or formatting of an FSD prior to the staged data being migrated to a final destination media, the Erase and Format toolbar buttons on the Device Manager window are disabled. If you want to erase or format the FSD, you can either use the command line (`ca_devmgr`) or disable the staging option for the selected FSD.

- **Increases your overall backup success rate**--You can define staging policies that let you create makeup jobs that back up directly to tape if an exceeds maximum threshold condition occurs or to create a makeup job on hold if a data migration failure occurs.
- **Pause Data Migration**--Arcserve Backup lets you pause the migration of data from the FSD to the final destination (tape) by enabling the Pause Data Migration option. This feature allows you to continue backing up to the FSD, but pause the migration from the FSD to the final destination in case the tape library is scheduled for maintenance or has hardware problems.
- **Simultaneous Streaming**--Simultaneous streaming is a process that divides your backup jobs into several subjobs that run simultaneously. Disk staging allows you to utilize the simultaneous streaming feature to send multiple streams of data to the temporary staging device (FSD) at the same time. Since the work is split up among several different streams (for concurrent writing to the FSD), simultaneous streaming-enabled backup jobs can be completed significantly faster than regular backup jobs. Simultaneous streaming also provides the capability to restore data while backup jobs are running.

- **SnapLock Support**--SnapLock™ is technology from Network Appliance that provides non-erasable, non-rewritable, Write Once Read Many (WORM) data protection. Arcserve Backup lets you use SnapLock protection on the backup operation. When you back up data with SnapLock protection enabled, you cannot purge or over-write the backed up data until the specified retention time elapses. This ensures that the data on the FSD cannot be deleted by any user, thus providing WORM support on disk with a retention time out. The retention time for the enabled SnapLock protection is determined by the specified settings for the staging purge policies.

Note: The device must support SnapLock technology. If you enable SnapLock on a device that does not support SnapLock WORM protection, Arcserve Backup write-protects the data, however, the data can be deleted from the device.

- **Copy Image Tracking**--Arcserve Backup provides the capability to track copied images on different media. As a result, the merging of catalogs only has to be performed one time, and then all sessions which are copies of each other would point to the same catalogs.
- **Flexible Restore Options**--During the time period that the backed-up data is located both on the final destination media (tape) and on the FSD (prior to purging), Arcserve Backup provides you with a choice for selecting the source for restoring the data. If the backup image is located on both the FSD and the final destination, you can choose where to restore it from.
- **Smart Restore**--Arcserve Backup provides a transparent Smart Restore feature that lets you restore backup data from multiple locations. If a media or drive error occurs during the restore process, from either the FSD or from the final destination media, Arcserve Backup internally finds the alternate media and starts restoring the data from the alternate media. Smart Restore helps to increase the success rate of restores in the event hardware problems occur while the job is running.
- **Optimize Restore Option**--If Arcserve Backup detects duplicate backup sessions, where one session resides on tape media and another session resides on an FSD, the Optimize Restore option lets you restore the data from the session that resides on the FSD.
- **Command Line Support**--Arcserve Backup lets you create backups to FSDs using either the graphical user interface (GUI) or the command line utility. In the event that a copy-to-tape operation fails, you can use the Query tool to analyze the file and session contents on the FSD. If you need to purge sessions from the FSDs, you can use the Purge tool to remove data and free extra space on the FSDs.
- **Disk Staging Reports**--Arcserve Backup lets you generate reports that are dedicated to disk staging backups. Using these reports you can find the status of backup sessions, whether a session was copied, when the session was copied, where the session was copied, whether the session was SnapLocked, when the session will be purged from the FSD, and other valuable information.

How to Manage Backup Data Using Staging

The following sections provide information about how to protect data using disk staging (B2D2T) and tape staging (B2T2T) operations.

More information:

[How to Manage Backup Data Using Tape Staging](#) (see page 229)

Tasks You Can Perform Using Disk Staging

The operations and tasks associated with the staging include the following:

- Specify and configure file system devices, tape libraries, and virtual tape libraries.
- Configure devices as a staging group and configure staging group policies.
- Submit backup jobs to staging groups.
- Define policies for managing backup, data migration, data security, data purge, alert messages, and postscript operations.
- Perform simultaneous backup operations to devices in a staging group.
- Disable staging in rotation and GFS rotation backup jobs on any specified day of the week.
- View the status of master and child jobs in the Job Status Manager. The Job Status Manager displays a tree view of all master jobs and their corresponding child jobs for backup and migration operations.
- View the Activity Log (in Windows) displaying the logs of all the child jobs and migration jobs, and the purging activities of the master job in a tree format.
- Restore data from a staging device. If the data from a backup job resides in two locations (on the file system device and on the final destination media), you can restore the data from either location.
- Run command line tools that can analyze and purge data stored on a FSD in a staging group.
- Access reports using the Report Manager to capture information about purge and migration activities on FSDs.
- Perform a Double Tape Copy migration job from the disk staging location to two final destinations simultaneously.

To ensure that the migration job continues when one of the devices become unavailable, you can set a timeout value for additional backup media to continue the migration job in another device. If you do not set a timeout value, then the data migration job fails after the first backup media timeout completes. For more details about setting the timeout value, see [Backup Manager Backup Media Options](#) (see page 164).

How the Max Number of Streams Option Affects Backup and Restore Operations

Arcserve Backup provides you with the capability of streaming multiple jobs simultaneously to FSDs. Simultaneous streaming is a process that divides your backup jobs into several subjobs that run simultaneously. Arcserve Backup lets you use simultaneous streaming to send multiple streams of data to a device in a staging group. Since the data is split up among several different streams, simultaneous streaming-enabled backup jobs can be completed significantly faster than regular backup jobs.

When you back up data using disk staging, a backup job can spawn child jobs. Each child job employs one stream of data. The actual number of child jobs that the parent job can spawn varies based on whether the backup job is a node-level or a volume-level backup job. However, the number of child jobs will never exceed the number of streams specified for staging.

Note: If a job spawns child jobs and you do not specify a number of streams to use, the child jobs will be created and backed up in one continuous stream.

For a node-level backup job, the number of child jobs spawned depends upon the number of agents specified in the backup job. Similarly, for a volume-level backup job, the number of child jobs spawned depends upon the number of volumes specified in the backup job.

Example: Staging Backup Jobs with Multiple Streams

If a backup job consists of backing up four nodes and the backup level is at the node level, the parent job can spawn a minimum of four child jobs. In this example, if you specify three streams, the master job can stream three child jobs simultaneously and start the fourth child job as one of the three previous child jobs end. After all child jobs are complete, the parent job is considered finished.

Staging Location Tab

To access the information and options on the Staging Location tab, start the Backup Manager and select the Staging Location tab.

The Staging Location tab contains the following options and informational fields:

Group Field

Displays the name of the group selected for this job.

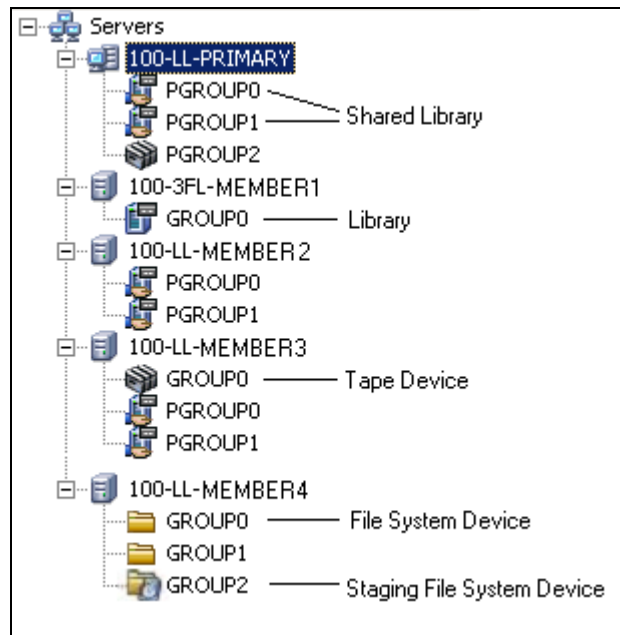
Note: A staging group must be selected in a staging job. Specifying a "*" group is not allowed for staging.

Max Number of Streams

Specifies the maximum number of simultaneous data streams that this job would be allowed to use while writing to the FSD in the staging group. For example, if the maximum number of streams is specified at 4, this means that at any point of time this staging job will have no more than 4 child jobs writing to the FSD simultaneously. To specify more than two streams, you must license the [Enterprise Module](#) (see page 40).

Staging Groups Directory Tree

Displays the names of the groups which were configured as staging groups.



Properties View

From the properties view in the Backup Manager, you can perform the following tasks:

- **Create Disk-based Devices**--Lets you open Disk-based Device Configuration so that you can configure Windows file system devices and deduplication devices.
- **Configure Groups**--Lets you open Device Group Configuration so that you can configure device groups.
- **Set Disk-based Device Group Properties**--Lets you open File System Device Group Configuration so that you can configure staging group properties.

How to Configure Arcserve Backup to Perform Disk Staging Backups

If you plan to back up your data to disk, the best practice is to use disk staging, rather than backing up data to a file system device (FSD). Disk staging lets you do the following:

- Back up to disk and then copy the data to two final destinations simultaneously (Double Tape Copy) and delete the data on the staging device by creating staging groups.

To ensure that the migration job continues when one of the devices become unavailable, you can set a timeout value for additional backup media to continue the migration job in another device. If you do not set a timeout value, then the data migration job fails after the first backup media timeout completes. For more details about setting the timeout value, see [Backup Manager Backup Media Options](#) (see page 164).

- Create flexible policies that determine when you want to copy and delete data.
- Reduce the backup window when a single job breaks into multiple simultaneous streams while writing to a disk staging area. The number of simultaneous streams can be controlled according to your disk network throughput capabilities.

Before you can back up data using disk staging, you must perform the following tasks:

- Create the staging devices. First, you must specify the devices in your environment that you will use for staging operations.

Important! Staging backup operations can quickly consume a large amount of free disk space on FSDs. Due to the maximum file size limitations of FAT 16 and FAT 32 file systems, you should not use these file systems on FSDs designated for staging operations.

- Configure the staging groups. After specifying the devices in your environment, you must configure the device group to function as a staging group.
- Configure the staging policies. To perform backup operations using staging, you must define the copy and purge policies that Arcserve Backup will use to manage data stored on staging devices.

The following sections provide you with information about how to configure Arcserve Backup to perform staging backups.

More information:

[Configure Device Groups to Use Staging](#) (see page 210)

[Specify Staging Groups Settings](#) (see page 212)

[Specify Copy and Purge Policies for Disk Staging Backups](#) (see page 213)

[Specify Miscellaneous Options for Disk Staging Backups](#) (see page 216)

[Specify Postscripts Options for Disk and Tape Staging Backups](#) (see page 220)

Configure Device Groups to Use Staging

This section describes how to configure device groups for staging operations.

Note: Before you can configure device groups, you must specify the devices that you will use for staging operations. For more information, see [Create File System Devices](#) (see page 385).

To configure device groups to use staging

1. From the Administration menu in the Navigation Bar on the Arcserve Backup Manager Console, click Device Group Configuration.
Device Group Configuration opens.
2. Click Next.
The Login Page dialog opens.
3. Complete the required fields on the Login Page dialog and click Next.
The Options dialog opens.
4. From the Options dialog, select the server that you want to configure, click Configure Disk-based Groups, and then click Next.
5. From the Groups list, select the group that you want to configure. To enable staging for the selected group, click the Enable Staging option and then modify the following options as needed:

- **Max Threshold**--Lets you specify the maximum amount of used space on a disk that Arcserve Backup will use for staging backups. When Arcserve Backup detects that the amount of used disk space exceeds the Max Threshold, Arcserve Backup pauses the backup job and purges the oldest migrated sessions from the FSD until the amount of used disk space is equal to or less than the Purge to Threshold.

Default value: If % is specified, 80%; if GB is specified, 8 GB; if MB is specified, 4000 MB.

The Max Threshold can be represented as either the total number of MB or GB used, or as a percentage of the total capacity used on the FSD. If the Max Threshold value is set as a percentage of the capacity of the FSD, the Max Threshold value must be equal to or less than 100% and the Purge to Threshold value must be greater than 0%.

Note: The Max Threshold must be greater than the Purge to Threshold. If you specify an absolute value (for example, MB or GB), the value must be equal to or greater than 1 MB.

- **Purge data when the used disk space exceeds the Max Threshold**--Lets Arcserve Backup purge migrated sessions from the FSD when the amount of used disk space exceeds the Max Threshold.

Note: To ensure that the purge mechanism starts in a timely manner, best practice is to specify a Max Threshold value that is at least 100 MB less than the total disk space.

- **Purge to Threshold**--Lets you specify the amount of used space on a disk when Arcserve Backup stops purging the oldest migrated sessions from the disk.

Default value: If % is specified, 60%; if GB is specified, 6 GB; if MB is specified, 3000 MB.

Arcserve Backup automatically specifies the units specified in the Max Threshold value (for example, %, MB, or GB). The Purge to Threshold value must be less than the Max Threshold value. If you specify an absolute value (for example, MB or GB), the value must be equal to or greater than 1 MB.

Example: The capacity of an FSD is 100 GB. The amount of used disk space is 75% (75 GB). The Max Threshold is 80% (80 GB) and the Purge to Threshold is 50% (50 GB). The administrator submits a job totaling 10 GB. Arcserve Backup detects that the job, when complete, will be greater than the Max Threshold. Arcserve Backup pauses the job and purges the oldest migrated sessions from the FSD until the amount of used disk space is equal to or less than the Purge to Threshold - in this example, 50% (or 50 GB). Arcserve Backup then continues the backup job.

Note: If Arcserve Backup purges all of the migrated sessions from the disk, but the amount of used continues to exceed the Purge to Threshold, Arcserve Backup restarts the job and attempts to complete the job using the available disk space.

- **Max # Streams**--Lets you specify the maximum number of simultaneous streams to the selected file system device group. If you have licensed the Enterprise Module, you may specify up to 32 streams, otherwise, the maximum number is 2.

Note: If a job spawns child jobs, the number of child jobs spawned will not exceed the number of streams specified for the job. However, if a job spawns child jobs and you do not specify a number of streams to use, the child jobs will be created and backed up in one continuous stream.

- **Enable SnapLock for this group**--Lets you enable SnapLock WORM protection on the file system device.

Note: This option is not available for libraries. To use this option, the file system device must support SnapLock technology. If you enable SnapLock on a device that does not support SnapLock WORM protection, Arcserve Backup write-protects the data, however, the data can be deleted from the device.

- **Pause data migration**--Lets you pause the data migration operation.

6. Repeat the previous step, as necessary, to configure other groups.
7. Click Next and then click Finish.
The options are applied to the job.

Specify Staging Groups Settings

Use the following procedure to modify staging group settings.

To specify staging groups settings

1. From the Backup Manager window, click the Staging Location tab.
2. From the groups list, right-click a group and select Set Disk-based Device Group Properties from the pop-up menu.
The Disk-based Group Property Configuration dialog opens.
3. Select the desired group and click Enable Staging.
4. In the Max Threshold field, specify the maximum file system device threshold. From the drop-down list, choose MB, GB, or %.
5. Click the Purge data when the used disk space exceeds the Max Threshold option if you want Arcserve Backup to automatically purge migrated sessions from the FSD when the amount of used disk space is greater than the Max Threshold
In the Purge to Threshold field, specify the amount of used space on a disk when Arcserve Backup stops purging the oldest migrated sessions from the disk.
6. When you enable disk staging, multiple streaming is enabled by default. If you:
 - Did not license the Enterprise Module, you can specify one or two streams.
 - Licensed the Enterprise Module, you can specify up to 32 simultaneous streams.
7. If the file system device supports SnapLock, and you want to WORM-protect the backed up data, click the Enable SnapLock option.
8. Click OK.

After you complete these steps, the newly created file system device group appears in the Groups directory tree on the Staging Location tab.

Note: Device groups identified as staging device groups do not appear in the Destination tab of the Backup Manager.

More information:

[Configure Device Groups to Use Staging](#) (see page 210)

Specify Copy and Purge Policies for Disk Staging Backups

Arcserve Backup lets you specify copy (migration) and purge policies for disk staging backups. Copy policies let you define when to migrate backup data to its final destination media after Arcserve Backup completes the backup to a disk staging device. Purge policies let you define when to delete backup sessions for the device to reclaim disk space.

The following information describes how to define policies for full and differential/incremental backups to a file system device or a deduplication device (B2D2T). For information about how to set policies for full and incremental/differential backup policies to a library or virtual library, see [How to Configure Arcserve Backup to Perform Tape Staging Backups](#) (see page 233).

Note: The copy and purge policies specified for disk staging backups apply to file system devices and deduplication devices.

To specify copy and purge policies for disk staging backups

1. Open the Backup Manager and select the Start tab.
From Start tab, click Normal Backup and Enable Staging.
The Staging Location and Policies tabs appear in the Backup Manager.
2. Click the Policies tab.
The copy and purge options appear.
3. Specify the following Copy Policies, as required, for the job:
 - Click **Full Backup** to specify policies for full backup jobs and select **Differential/Incremental Backup** to specify policies for differential and incremental backup jobs.
 - **Do not copy data**--Choose this option if you do not want to migrate the backup sessions to final destination media. For example, consider differential and incremental backup operations. Operations of this type tend to have short retention periods and are small with respect to overall size. If you do not copy the incremental and differential backups to final destination media, the need for tapes to store your backups diminishes.

- **After**--Lets you start the migration job after specified length of time elapses.

Note: Be aware that physical disks and volumes do not support differential and incremental backups. As a result, Arcserve Backup applies full backup policies to incremental and differential backups of physical disks and volumes. The Copy Time is the only exception to this behavior. With staging backups, Arcserve Backup copies incremental and differential backups of physical disks and volumes to final destination media based on the copy policies specified for incremental and differential backups.

Arcserve Backup starts the copy to final destination media operation based upon the occurrence of one of the following events:

- **After job starts**--Lets you start the migration operation at a fixed point in time after the backup to disk operation starts.
- **After job ends**--Lets you start the migration operation after the backup to disk operation ends.

Due to variations in the overall size of backup jobs and the length of time needed to complete backup to disk operations, simultaneous read and write operations to the disk staging device can occur. This option prevents simultaneous read and write operations to disk staging devices.

- **After each session is finished**--Choose this option if you want to start the copy to media operation immediately after the backup to disk operation for the session is complete.

Most backup jobs consist of several sessions. When you specify this option, you can direct Arcserve Backup to copy backup sessions to their final destination immediately after the backup job is finished. This option manifests simultaneous backup and copy operations. By performing backup and copy operations simultaneously, you can reduce the overall backup window and copy window.

Because this option induces simultaneous read and write operations on the FSD, you should only specify this option if you are using a high-speed device that can process many read and write operations simultaneously.

Note: For all Copy data after options, Arcserve Backup will not migrate sessions to their final destination media until after the backup job for the session is complete. This capability includes scenarios when the copy retention period expires before the backup operation is complete.

- **At**--Lets you start the migration operation at a specific time of day. When you use this option you can direct Arcserve Backup to start the migration process at a specific time on a daily basis.
 - Select the Or after the job is finished whichever happens later option if you suspect or anticipate the backup to disk operation to end after the specified start time for the copy to final destination operation. This option prevents Arcserve Backup from copying sessions from disk to tape while the backup operation is in progress.

- **Copy data for specified backups only**--Lets you migrate only monthly or weekly backups associated with rotation jobs.

Note: Copy data for specified backups only options do not apply to incremental and differential backups.

Default value: Disabled.

With this option enabled, you can specify one of the following migration options:

- **Copy data for monthly backups only**--Lets you migrate only the monthly full backup sessions, not the weekly full backup jobs, associated with rotation jobs.

Note: This option can be applied on only GFS rotation jobs.

- **Copy data for weekly backups only**--Lets you migrate only the weekly full backup sessions, not the daily backup sessions, associated with rotation jobs.

Seven day rotations--Lets you migrate data in the following scenarios: For 7-day weekly full backups, Arcserve Backup migrates the Saturday (full) backup sessions. For 7-day weekly incremental/differential backup, full backup on Sunday backups, Arcserve Backup migrates the Sunday (full) backup sessions.

Five day rotations--Arcserve Backup migrates only the Friday (full) backup sessions.

Note: This option can be applied on rotation jobs and GFS rotation jobs. For more information, see [Rotation Schemes](#) (see page 307).

4. Specify the following Purge Policies, as required, for the job:

- **After**--Lets you start the purge operation after specified length of time elapses. Arcserve Backup starts the purge operation based upon the occurrence of one of the following events:

- **After job starts**--Lets you start the purge operation at a specified time after the backup to staging device operation starts.
- **After job ends**--Lets you start the purge operation at a specified time after the backup to staging device operation ends.

- **At**--Lets you start the purge operation at a specific time of day. Use the spin box to specify the time of day that you want the operation to start.

5. To enable SnapLock protection, click the Enable SnapLock check box.

Note: The device must support SnapLock protection. If you specify Enable SnapLock on a device that does not support SnapLock WORM protection, Arcserve Backup write-protects the data; however, the data can be deleted from the device.

Specify Miscellaneous Options for Disk Staging Backups

To perform disk to disk to tape (B2D2T) backups, you can optionally specify policies that control how Arcserve Backup processes backup job data.

Note: The miscellaneous options specified for disk staging backups apply to file system devices and deduplication devices.

To specify miscellaneous options for disk staging backups

1. Open the Backup Manager and select the Policies tab.

Click Miscellaneous in the policies list.

Specify the options that follow that you require for the job:

- **Purge canceled sessions from disk**--Use this option to direct Arcserve Backup to delete sessions from the staging device immediately after a backup to staging device is canceled.

This option helps to reclaim free disk space on the staging device as quickly as possible.

- **Purge failed sessions from disk**--Use this option to direct Arcserve Backup to delete sessions from the staging device immediately after a backup to disk staging device fails.

This option helps to reclaim free disk space on the staging device as quickly as possible.

- **Create makeup jobs to back up data directly to final destination under disk full conditions**--Use this option to direct Arcserve Backup to back up data directly to its final destination media if there is insufficient free space on the file system device in a staging group.

A backup operation will fail if there is insufficient free disk space on the staging device. To remedy this situation, Arcserve Backup can divert the backup operation from the file system device in a staging group directly to the final destination media. A makeup job searches for blank media and media from a scratch. As such, specifying this option can increase the overall success rate of your backup operations when a "disk full" condition exists.


- **Create makeup jobs on hold if Data Migration jobs fail**--Use this option to direct Arcserve Backup to create makeup jobs on HOLD if data migration (copy to tape) jobs fail.

A data migration job can fail if a media or tape drive error occurs during the copy to tape operation. Use this option to create a makeup job with a HOLD status that you can change to a Ready status after correcting the tape drive or media errors. If an error condition exists, this option minimizes the need to create tapecopy jobs.

Consolidate Data Across Jobs While Copying Options

Lets you consolidate backup data during the migration operation.

- **Copy Method**--Specify a copy method (Append or Overwrite) that you want to use for the consolidation operation. The method that you specify must be the same for all jobs that you want to consolidate.
 - If you have a requirement to consolidate data across multiple jobs and ship the tapes on a daily basis, you should choose the Overwrite option.
 - If you have a requirement to consolidate data across multiple jobs (for daily backups) for the whole week to a single tape and ship the tapes on a weekly basis, you should choose the Append option.
- **Limitations and Considerations**
 - If the backup is a rotation or a GFS backup, you must specify the prefix of the target media and the prefix of the media pool that you want to use for consolidation on the Destination tab.



Note: For more information about consolidation options and examples, see [Consolidation During Migration](#) (see page 441).

Specify Alert Options for Disk and Tape Staging Backups

Arcserve Backup lets you use the alert notification system to send messages about migration events that occur during staging operations. For more information about setting up alerts, see Using the Alert Manager.

Note: The alert options specified for disk staging backups apply to file system devices and deduplication devices.

To specify alert options for disk and tape staging backups

1. Open the Backup Manager and click the Start tab.
From the Start tab, click Normal backup or Deduplication backup and then click Enable staging.
The Staging Location and Policies tabs appear in the Backup Manager.
2. Click the Policies tab.
The copy policy options appear.
3. Click Alert in the Policies list.
The Alert options display.

4. From the Event list, select one of the following migration job events for which you want to send an alert notification:

- **Job Completed Successfully**--A migration job completed successfully.
- **Job Incomplete**--A migration job did not complete successfully.
- **Job Canceled by User**--A migration job was canceled by the user.
- **Job Failed**--A migration job failed.
- **Media not Available**--Media was not available during the execution of a migration job.
Note: The migration media (final destination media) must be tape media.
- **Format Blank Tape**--A tape was formatted during the execution of a migration job.
- **Customized Event**--A user-defined event, such as an error message, warning message, and critical error message, occurred during the execution of a migration job, and the message appeared in the Activity Log.

Event Codes:

Note: You can specify event codes only when you select Customized Event.

- **E***--An error occurred and the error message appeared in the Activity Log.
- **W***--A warning occurred and the warning message appeared in the Activity Log.
- **N***--A notification message occurred and the notification message appeared in the Activity Log.
- **C***--A critical message occurred and the critical message appeared in the Activity Log.
- **AE***--An agent error message occurred and the agent error message appeared in the Activity Log.
- **AW***--An agent warning message occurred and the agent warning message appeared in the Activity Log.

Examples:

- AE* is specified in the Event Code field. An alert will be sent when any agent error message occurs and the agent error message appears in the Activity Log.
- AE0006 is specified in the Event Code field. An alert will be sent only when AE0006 occurs and the error message appears in the Activity Log.
- E*;AE0006 is specified in the Event Code field. An alert will be sent when any errors occur and the error messages appear in the Activity Log, when AE0006 occurs and the error message appears in the Activity Log, or both.

Note: You must separate event codes with a semi-colon ";".

5. From the Methods & Recipients field, you can accept the default options, or create a custom alert for the event. The <Default> configuration means that you will use the alert options configured using the Alert Manager.

To create custom alerts, click the Configure button.

The Methods & Recipients Configuration dialog opens. You can specify one or more of the defined alert configurations. Arcserve Backup provides the following defined alert configurations:

- Broadcast
- Pager
 - Note:** Pager options are not supported on Japanese versions of Arcserve Backup.
- SMTP
- SNMP
- Event
- Printer
- Email
- Lotus Notes
- Unicenter NSM

6. To add a new Methods & Recipients configuration, click the New button.

The Configuration Name dialog opens. Specify a name for the configuration and click OK.

A new configuration tree displays in the browser at the left of the dialog. The new configuration tree contains one branch for all available notification methods. You must now add recipients to the methods branches of your tree. For example, if you want to use the Printer notification method, you must add an available printer to the tree.

7. To add a recipient to a configuration, you must first select a method (for example, Broadcast) from the configuration tree and then click the Add button.

The appropriate Add Recipient dialog opens for the selected configuration. Configure the new recipient in this dialog. For more information about the different Recipient dialogs, click the Help button.

After you configure the new recipient, it is added to the tree.

Note: You cannot add recipients for the Unicenter TNG alerts. If you click Modify, the Unicenter TNG Event Map dialog opens. You can then have messages sent to the Unicenter console or the World View repository when an alert is generated.

8. To modify a Methods & Recipients configuration, select the configuration from the Configuration drop-down list.

The selected Configuration tree displays in the browser. You can add, modify, or delete recipients from the configuration tree by clicking the Add, Modify, or Delete button.

To delete a configuration, select the configuration from the Configuration drop-down list and click the Delete button.

To rename a configuration, select the configuration from the Configuration drop-down list, and click the Rename button.

Specify Postscripts Options for Disk and Tape Staging Backups

Arcserve Backup lets you specify postscripts that run based on particular migration events that occur during the course of staging operations.

A script is a set of instructions that are stored in user-defined files that can be created in any format, such as .bat and .exe. Scripts can be executed before or after an event occurs. A postscript is a set of instructions that you can run or execute after an event occurs, such a migration event. Postscripts are not limited to Arcserve Backup-based scripts.

Note: The postscript options specified for disk staging backups apply to file system devices and deduplication devices.

To specify postscripts options for disk and tape staging backups

1. Open the Backup Manager and click the Start tab.

From the Start tab, click Normal backup or Deduplication backup and then click Enable staging.

The Staging Location and Policies tabs appear in the Backup Manager.

2. Click the Policies tab.

The copy policy options appear.

3. Click Postscripts in the policies list.

Specify the options that follow that you require for the job:

- From the Event list, select one of the following migration events for which you want to run a postscript:

- **Migration job complete**--All sessions in the current migration job migrated successfully.

Example:


A backup job can consist of one or multiple migration jobs. This event will occur after each individual migration job is complete.

- **Migration job incomplete**--One or more sessions in a migration job did not complete successfully. For example, a session was skipped during the migration job.
- **Migration job canceled**--A migration job was canceled by a user while it was in an Active, Ready, or Hold status. A makeup job was not created.
- **Migration job failed**--One or more sessions in a migration job failed.
- **Makeup of migration job created**--A migration job failed and Arcserve Backup created a makeup job.
- **All sessions migrated**--All sessions corresponding with a staging job migrated successfully.

Example:

A backup job consists of multiple migration jobs. The migration jobs consist of several sessions. This event will occur when all sessions in all migration jobs for the backup job migrated successfully.

- Click in the Postscripts field adjacent to the selected event and do one of the following:

- Enter the path to script that you want to run after the event occurs.
- Click the ellipsis button  to browse to the script that you want to run after the event occurs.

Note: You can specify one postscript per migration event.

- From the Run As section, complete the following fields:

- User name
- Password
- Confirm password

Note: You must provide Windows credentials to run postscripts.

- Repeat Steps 3, 4, and 5 to specify postscripts for other migration events.

Back Up Data Using Disk Staging

Prior to performing a backup job using disk staging (B2D2T), you must have already configured the staging groups. If you did not configure Arcserve Backup to use disk staging, see [How to Configure Arcserve Backup to Perform Disk Staging Backups](#).

Arcserve Backup provides you with the capability to submit a backup job using either the Backup Manager or the command line utility. This information describes how to perform a disk staging backup job using the Backup Manager. For information about how to submit a staging backup job using the command line, see the *Command Line Reference Guide*.

Note: Before you can back up data using disk staging, ensure that all preconfiguration tasks are complete and licensing requirements are fulfilled. For more information, see [Licensing Requirements for Staging Backups](#).

Back up data using disk staging

1. Open the Backup Manager.

Click the Start tab, select Normal, Deduplication, or UNIX/Linux Data Mover, and then click the Enable Staging check box.

The Staging Location and Policies tabs appear.

2. Click the Source tab, browse to and select the source objects that you want to back up.
3. Click the Schedule tab and specify the schedule that you want to use for the backup job.

Note: For more information, see [Rotation Schemes](#) (see page 307) and [Custom Schedules](#) (see page 318).

4. Click the Staging Location tab and expand the Staging Servers object.

Browse to and select the staging group that you want to use for the backup job. If you selected a deduplication job on the Start tab, choose a deduplication device group from the Staging Location tab.

5. Click the Policies tab. Complete the following Staging Policies required for the job:

- **Full Backup**--Specify the Copy and Purge policies for full backups required for the job.

(optional)--Click Enable Snaplock.

- **Incremental/Differential Backup**--Specify the Copy and Purge policies for incremental and differential backups required for the job.

(optional)--Click Enable Snaplock.

Note: For more information, see [Specify Copy and Purge Policies for Disk Staging Backups](#) (see page 213).

- **Miscellaneous**--Specify the Miscellaneous policies required for the backup job.

Note: For more information, see [Specify Miscellaneous Options for Disk Staging Backups](#) (see page 216).

- **Alert**--Specify the Alert policies required for the backup job.

Note: For more information, see Specify Alert Options for Disk and Tape Staging Backups.

- **Postscripts**--Specify the Postscript policies required for the job.

Note: For more information, see [Specify Postscripts Options for Disk and Tape Staging Backups](#) (see page 220).

6. Click the Destination tab and expand the Servers object.

Browse to and select the destination device you want to use for the migration job. If desired, you may choose a deduplication device group as the final destination provided it is not the same deduplication device group selected as the staging group.

Note: You can also have the option to perform a Double Tape Copy. Double Tape Copy lets you migrate (copy) the data to two destinations simultaneously. Using this feature writes to the destinations in parallel to save time. To do this, select the Create Duplicate in-line check box from the Destination tab. For more details, see [Backup Manager Destination Options](#) (see page 149).

7. Click the Options on the toolbar.

The Options dialog opens.

8. Select the Encryption/Compression tab and complete the following fields for the backup job, provided you are not using deduplication devices:

- **Session/Encryption password**--Specify a Session/encryption password to restore this data from media.

Important! If you specify a Session/Encryption password, you must provide this password to restore the session.

- **Encrypt data**--Use this option to encrypt the backup data. You can specify one of the following options:
 - **At agent**--Select this option to encrypt the backup data prior to the actual backup process. For more information about this option, see [Data Encryption at the Agent Server](#) (see page 113).
 - **At backup server during backup**--Select this option to encrypt the backup data at the backup server during the backup process. For more information, see [Data Encryption During Backup](#) (see page 114).
 - **At backup server during migration**--Select this option to encrypt the backup data during the migration phase of a staging backup job. For more information, see [Data Encryption During Migration](#) (see page 114).
- **Compress data**--Use this option to compress the backup data. You can specify one of the following options:
 - **At agent**--Select this option to compress the backup data on the system where the agent is installed and running.
 - **At backup server**--Select this option to compress the backup data at the Arcserve Backup server during the backup process. Using this option directs Arcserve Backup to compress files before backing them up using a software compression algorithm.

Note: When you use data compression at the backup server before data encryption, the amount of space required to store the data on the staging device can be two times the size of the source files. Because of this limitation, we discourage the use of compression before encryption when backing up to disk.

Be aware of the following considerations that affect encryption and compression:

- If you want to apply other options that affect the migration job, you should do so at this time. For example, to eject the tape from a library after the migration job is complete, click the Operation tab on the Global Options dialog and select the Eject Media option.
- If you specify encryption and compression options, and the backup destination is a drive that does not support compression, or is a file system device (FSD), Arcserve Backup encrypts the backup data and does not compress the backup data.
- Encryption and compression are not supported on deduplication devices. However, if you select a regular FSD for either the staging or final destinations, you may enable encryption or compression, as needed. For more information, see the topic, [Compression and Encryption with Deduplication](#) (see page 742).

9. Click Submit on the toolbar to submit the backup job.

The Submit Job dialog opens.

10. Complete the required fields on the Submit Job dialog and click OK.

The job is submitted.

More information:

[Submit a Backup Job](#) (see page 136)

Modify a Staging Rotation Scheme

If you are using rotation or GFS rotation disk staging jobs, Arcserve Backup provides you with the flexibility to disable staging on any specified day of the week.

To modify staging when using a rotation scheme

1. Open the Backup Manager and select the Schedule tab.
2. Select the Use Rotation Scheme option, and then select the scheme name from the Scheme Name drop-down list.
3. Click the Rotation Rules tab.

The Staging column displays the current status of staging as it applies to your rotation scheme.

4. Select the Day Of Week for which you want to modify staging, and then click the Modify button.

The Configuration dialog opens.

5. From the Staging drop-down list, select Enable or Disabled.
6. Click OK.

Note: To disable staging for a staging group, see [Disable Staging](#) (see page 227).

Pause Data Migration

The Pause Data Migration option lets you temporarily stop the process of migrating data from the FSD to its final destination media.

Example: When You Should Pause Data Migration

You need to take a tape library offline to perform maintenance on the library. You can pause the data migration process, complete the maintenance tasks, bring the library back online, and then restart the migration process.

To pause data migration

1. From the Staging Groups tree on the Staging Location tab, select the group that you want to pause.
2. Right-click the group name and select Configure Disk-based Groups from the pop-up menu.

The Disk-based Group Property Configuration dialog opens.

3. Select the group in the groups list.

Click Pause Data Migration.

Click OK.

Arcserve Backup pauses the migration.

Note: To restart the data migration operation, repeat the above steps and clear the checkmark from Pause data migration.

Disabling Disk Staging Rotations

When you back up data using regular or GFS rotation rules, Arcserve Backup provides you with the capability to suspend or disable staging in the backup jobs on any specified day of the week, bypassing the FSD, and backing up your data directly to its final destination media.

Example: When You Should Disable a Staging Backup Job

If you discover that your FSD in a staging group is approaching or has exceeded its storage capacity threshold, backup jobs can fail. You can modify the staging job to disable staging on that day so that the data is backed directly to the final destination.

To verify whether staging for rotation and GFS rotations is disabled or enabled, open the Backup Manager, select the Schedule tab, and select the Rotation Rules tab. The Staging column in the Rotation Rules schedule displays the current status of all rotations and GFS rotations. To modify a rotation rule, click the Modify button below the schedule.

Disable Staging

Arcserve Backup provides you with the capability to disable (or bypass) backup to FSD operations. When you use this option, data is backed up directly to its final destination media, rather than being backed up to the FSD.

There are two methods that you can use to perform this task:

- From the Rotation Rules tab on the Schedule tab of the Backup Manager.
- Using File System Device Group Configuration.

Backup Manager - Schedule Tab

To disable backup to staging device operations from the Backup Manager, perform the following steps:

1. Open the Backup Manager window and click the Schedule tab.
2. Select the Scheme Name from the drop-down list.
3. Click the Rotation Rules tab and select the rotation that you want to disable.
4. Click the Modify button.

The Configuration dialog opens.

5. From the Staging drop-down list on the Configuration dialog, select Disabled.
6. Click OK.

File System Device Group Configuration Dialog

To disable back up to staging device group operations using Device Group Configuration:

1. Open the Backup Manager window and click the Staging Location tab.
2. Right-click the group that you want to disable and select Set Disk-based Device Group Properties.

The File System Device Group Configuration dialog opens, displaying all groups in your environment that are specified as file system device groups.

Note: The groups that are enabled for staging display with a corresponding dark blue flag. The groups that are not enabled for staging display with a corresponding light blue flag.

3. Select the group that you want to disable.
4. Clear the check mark from the Enable Staging check box.
5. Click OK.

How You Can Manage Staged Data When the Database Fails

When you use disk staging to back up data, the information about the backup jobs, sessions, staging policies, and so on is stored in the Arcserve Backup database. If the database fails, and you need to recover the Arcserve Backup database, the staging policies for the data residing on the staging device (for example, a file system device or a library) that specify when to copy the data to the final destination media and, if file system device (FSD), when to purge the data from the staging device are no longer available.

If this situation occurs:

- Arcserve Backup cannot copy (migrate) the data on staging device to its final destination media.
- Arcserve Backup cannot purge data from a file system device (FSD) to reclaim disk space. As a result, future backup jobs will probably fail due to an insufficient amount free disk space on the staging device.

To remedy this situation and retain all of the backup data that is stored on the staging device, you can use the `tapcopy` command line utility to copy all the backup data from the staging device to final destination media. (When you use this approach, media rotation rules, such as Friday tape or Monday tape may not be adhered to.) Then, you can use the `-purge` option from the Device Manager command line utility (`ca_devmgr`) to delete the data from the FSD and reclaim disk space.

How You Can Reclaim Disk Space

This section provides examples of how you can quickly reclaim disk space using the **Purge Data At** and **Purge Data After** options.

Example 1

You have a high-performance disk with a limited amount free disk space. You can quickly reclaim disk space by specifying a short length of time under the Purge data After option and select the After job starts option. This approach ensures that the purge operation starts shortly after the copy to final destination media operation starts, as opposed to the using the After job ends option, which starts the purge operation after the copy to final destination media operation ends.

Example 2

You have a backup job rotation or a GFS rotation scheme that starts at the same time daily and your high-performance disk maintains a limited amount of unused space. Using the Purge data At option to schedule the purge operation to start before the next backup operation starts. This approach helps to ensure that you have freed enough disk space to prevent the backup job from failing.

Important! If you specify that the data is to be copied to final destination media, Arcserve Backup does not start the purge operation until after the copy to final destination media operation is finished.

How to Manage Backup Data Using Tape Staging

The following sections provide information about how to protect data using backup to tape operations.

More information:

[How to Manage Backup Data Using Staging](#) (see page 206)

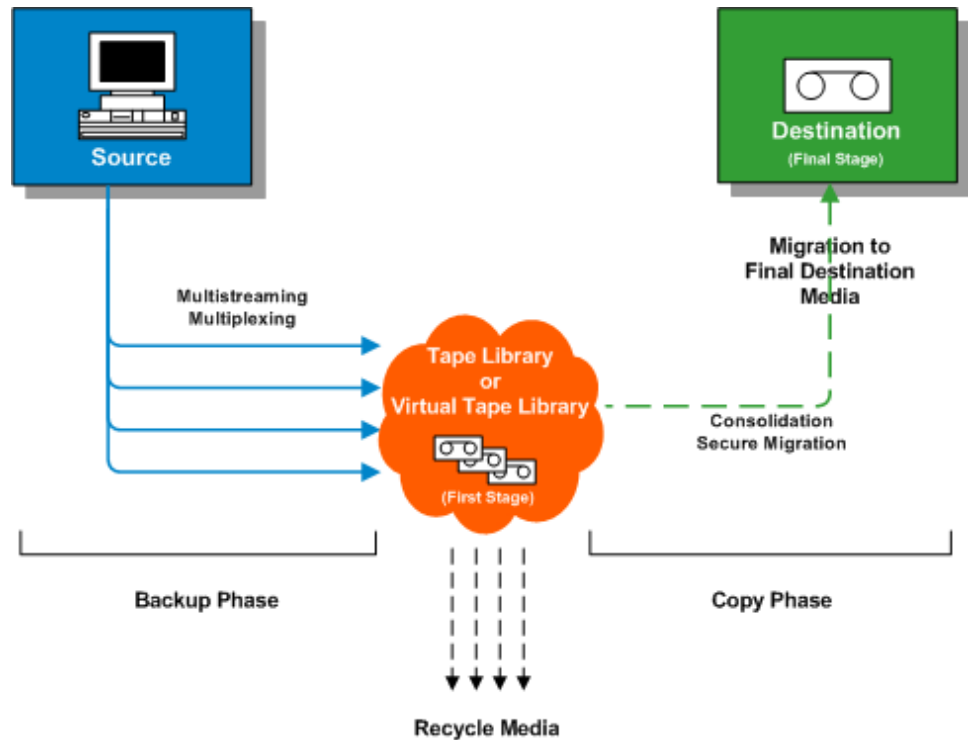
How Backup to Tape to Tape Works

Backup to tape to tape is a data protection solution that lets you to back up data to a tape library or a virtual tape library, and then copy the data to a different tape library or other type of device. Copy operations, also known as migration, are governed by user-specified copy policies.

Backup to tape to tape (B2T2T) is a two-part backup process.

- **Backup Process**--Arcserve Backup backs up data from the source to the staging location. The staging location is a tape library or a virtual tape library (VTL).
- **Copy Process**--Arcserve Backup copies or migrates the backup data from the staging location to the final destination media. The final destination is tape media.

The following diagram illustrates the flow of data from the source to the first stage tape library (or virtual tape library) and then on to the final destination.



Note: Arcserve Backup lets you transmit up to 32 streams of data using multistreaming. To back up data using multiple streams and transmit more than two streams of backup data, you must license the [Arcserve Backup Enterprise Module](#) (see page 40).

When you use backup to tape to tape (B2T2T) to protect data, the backup to tape to tape operation consists of two phases:

Backup Phase

Arcserve Backup backs up data from the source to the tapes in the first stage, based on user-specified policies.

- Backup jobs can consist of full, incremental, or differential backups.
- During the backup job, global options, media selection rules, media pool usage, rotation rules, GFS rotation rules, Alert messages, Export options, and so on are identical to that of backing up directly to tape.

Note: Various global backup options do not apply to backup phase operations.

- Multiplexing and multistreaming can be used to transmit and save data to first stage media.

Note: The Multiplexing option can be used for backup operations to all tape devices, with the exception of file system devices. The Multistreaming option can be used for backup operations to tape libraries that contain two or more drives.

Copy Phase

Arcserve Backup copies data from the first stage to the final destination based on user-specified policies.

- Arcserve Backup copies data from the first stage media to the final destination media one session at a time. Multiple sessions cannot be copied to a single tape simultaneously.
- If you need to copy data from more than one first stage media to one final destination media, Arcserve Backup copies each session in succession until all the sessions are copied to the final destination media.
- Arcserve Backup sessions associated with different jobs can be consolidated during migration. You can activate this capability using the consolidation option.
- If a hardware error occurs during the process of copying data to final destination, the job stops and Arcserve Backup creates a Makeup Job On Hold. After you correct the hardware error, you can change the job status to Ready, and then the job resumes.

More information:

[How to Use Tape Staging to Manage Backup Operations](#) (see page 232)

[How to Configure Arcserve Backup to Perform Tape Staging Backups](#) (see page 233)

[How to Submit a Tape Staging Backup Job](#) (see page 236)

How to Use Tape Staging to Manage Backup Operations

The following are common scenarios that describe how you can use Backup to Tape to Tape operations (B2T2T) to manage backup operations:

- If you need to store two copies of backup data, one copy on site and one copy at an off-site storage location, B2T2T lets you back up data directly to tape. After the backup job is complete, you can use Arcserve Backup copy utilities to automate and create copies of the backup tapes, and then ship the tapes to an off-site storage location.
- B2T2T lets you encrypt backup data when you are copying the data to its final destination media. This capability is beneficial when you are copying data from a virtual tape library or a library that does not support encryption to a library that supports encryption. This capability ensures that your backups are as fast as possible and the tapes that you need ship to an off-site storage location are encrypted.
- While backup operations are in progress, you may have many jobs that are backing up data to many different tapes. This can result in media that is not being used to its full capacity. B2T2T operations let you consolidate backups to ensure that media is being used to capacity when copying data to final destination. This capability helps to reduce the cost of media because you are using fewer tapes for final destination media, off-site storage, or both.
- If you need to reduce the length of time required to back up data and copy the data from a staging area to final destination, you can use virtual tape libraries (VTL) in your environment to manage backup operations.

A VTL is a temporary storage location, such as a disk drive, that is configured to behave like a library. Since most backup data is transmitted across a network, Arcserve Backup lets you use multiplexing to reduce the backup window. When you use a VTL to store backup data, you can quickly read data from a multiplexing formatted data in a VTL because your operations do not encounter tape positioning overhead. As a result, the processes of backing up data to a VTL, reading from a VTL (disk), and copying the data to final destination media is fast. Arcserve Backup lets you automate the process of copying to final destination media when you use a VTL to stage your backup data.

How to Configure Arcserve Backup to Perform Tape Staging Backups

Before you can back up data using tape staging, you must perform the following tasks:

- Create the staging devices.

If you plan to use a virtual library, open the Device Manager to confirm the library is properly configured. Arcserve Backup automatically configures libraries when you stop and restart the Tape Engine. If a library is not properly configured, you can run Device Configuration to manually set up libraries and virtual libraries for staging operations.

Note: For more information about using Device Configuration to set up libraries and virtual libraries, see [Tape Library Configuration](#) (see page 373).

- Specify device groups as staging groups.
- Configure staging policies.

Note: To perform backup operations using staging, you must define the migration policies that Arcserve Backup will use to manage data stored on staging devices. For more information, see [Back Up Data Using Tape Staging](#) (see page 236).

More information:

[Specify Postscripts Options for Disk and Tape Staging Backups](#) (see page 220)

[Specify Migration Policies for Tape Staging Backups](#) (see page 233)

[Specify Miscellaneous Options for Tape Staging Backups](#) (see page 235)

Specify Migration Policies for Tape Staging Backups

Arcserve Backup lets you specify migration (copy) policies for tape staging backups. Migration policies let you define when to migrate backup data to its final destination media after Arcserve Backup completes the backup to a tape staging device (for example, a library, a virtual library, a tape drive, an FSD, and so on).

To specify migration policies for tape staging backups

1. Open the Backup Manager and select the Start tab.

From Start tab, click Normal Backup and Enable Staging.

The Staging Location and Policies tabs appear in the Backup Manager.

2. Click the Policies tab.

The copy policy options appear.

3. Specify the following Copy Policies, as required, for the job:

- Click **Full Backup** to specify policies for full backup jobs and click **Differential/Incremental Backup** to specify policies for differential and incremental backup jobs.
- **Do not copy data**--Choose this option if you do not want to migrate the backup sessions to final destination media. For example, consider differential and incremental backup operations. Operations of this type tend to have short retention periods and are small with respect to overall size. If you do not copy the incremental and differential backups to final destination media, the need for tapes to store your backups diminishes.

Be aware of the following behaviors:

- Be aware that physical disks and volumes do not support differential and incremental backups. As a result, Arcserve Backup applies full backup policies to incremental and differential backups of physical disks and volumes. The copy time is the only exception to this behavior. With staging backups, Arcserve Backup copies incremental and differential backups of physical disks and volumes to final destination media based on the copy policies specified for incremental and differential backups.

- **Copy data for specified backups only**--Lets you migrate only monthly or weekly backups associated with rotation jobs.

Note: Copy data for specified backups only options do not apply to incremental and differential backups.

Default value: Disabled.

With this option enabled, you can specify one of the following migration options:

- **Copy data for monthly backups only**--Lets you migrate only the monthly full backup sessions, not the weekly full backup jobs, associated with rotation jobs.

Note: This option can be applied on only GFS rotation jobs.

- **Copy data for weekly backups only**--Lets you migrate only the weekly full backup sessions, not the daily backup sessions, associated with rotation jobs.

Seven day rotations--Lets you migrate data in the following scenarios: For 7-day weekly full backups, Arcserve Backup migrates the Saturday (full) backup sessions. For 7-day weekly incremental/differential backup, full backup on Sunday backups, Arcserve Backup migrates the Sunday (full) backup sessions.

Five day rotations--Arcserve Backup migrates only the Friday (full) backup sessions.

Note: This option can be applied on rotation jobs and GFS rotation jobs. For more information, see [Rotation Schemes](#) (see page 307).

Specify Miscellaneous Options for Tape Staging Backups

To perform tape staging backup operations, you can optionally specify policies that control how Arcserve Backup processes backup job data.

To specify miscellaneous options for tape staging backups

1. Open the Backup Manager and select the Policies tab.

Click Miscellaneous in the policies list.

Specify the options that follow that you require for the job:

- **Create a makeup job on hold if a data migration job fails**--Use this option to direct Arcserve Backup to create makeup jobs on HOLD if data migration (copy to tape) jobs fail.

A data migration job can fail if a media or tape drive error occurs during the copy to tape operation. Use this option to create a makeup job with a HOLD status that you can change to a READY status after correcting the tape drive or media errors. If an error condition exists, this option minimizes the needs to create tapecopy jobs.

- **Schedule a makeup job for a data migration job if it cannot proceed because the source group or tape is not available**--Use this option to direct Arcserve Backup to schedule a makeup job when the source group or tape is not available.

The source may not be available do to a variety of reasons. For example, the backup phase for the job is not complete, or a hardware problem exists in the tape library or virtual tape library.

- **Reschedule after**--Specify how many minutes must elapse before the makeup will be rescheduled.

2. To consolidate the backup data during the migration operation, click the Consolidate data across jobs while copying option and complete the following fields.

Note: If you want to consolidate data across multiple jobs to the same tape, you should run the backup jobs on the same machine.

- **Copy Method**--Specify a copy method (Append or Overwrite) that you want to use for the consolidation operation.

The method that you specify must be the same for all jobs that you want to consolidate.

- If you have a requirement to consolidate data across multiple jobs and ship the tapes on a daily basis, you should choose the “Overwrite” option.
- If you have a requirement to consolidate data across multiple jobs (for daily backups) for the whole week to a single tape and ship the tapes on a weekly basis, you should choose the “Append” option.

Note: For more information about consolidation options and examples, see [Consolidation During Migration](#) (see page 441).

How to Submit a Tape Staging Backup Job

The following sections provide you with information about how to submit a tape staging backup job.

Licensing Requirements for Tape Staging Backups

To perform successful backup to tape to tape backup (B2T2T) operations, verify that the following licensing requirements are fulfilled.

- You must license the [Arcserve Backup Enterprise Module](#) (see page 40) to back up data using multiple streams and transmit more than two streams of backup data.
- You must license the Arcserve Backup Tape Library Option to back up data to a tape library that contains more than one tape drive.

Back Up Data Using Tape Staging

Prior to performing a backup job using tape staging (B2T2T), you must have already configured the staging groups. For more information, see [How to Configure Arcserve Backup to Perform Tape Staging Backups](#) (see page 233).

Arcserve Backup provides you with the capability to submit a backup job using either the Backup Manager or the command line utility. This information describes how to perform a tape staging backup job using the Backup Manager. For information about how to submit a staging backup job using the command line, see the *Command Line Reference Guide*.

Note: Before you can back up data using tape staging, ensure that all preconfiguration tasks are complete and all licensing requirements are fulfilled. For more information, see Licensing Requirements for Staging Backups.

To back up data using tape staging

1. Open the Backup Manager.
Click the Start tab and click the Enable Staging check box.
The Staging Location and Policies tabs appear.
2. Click the Source tab, browse to and select the source objects that you want to back up.
3. Click the Schedule tab and specify the schedule that you want to use for the backup job.

Note: For more information, see [Rotation Schemes](#) (see page 307) and [Custom Schedules](#) (see page 318).

4. Click the Staging Location tab and expand the Staging Servers object.
Browse to and select the staging group that you want to use for the backup job. If you wish to enable deduplication, choose a deduplication device group from the Staging Location tab.
5. Click the Policies tab. Complete the following migration policies required for the job:

- **Full Backup and Incremental/Differential Backup**--Specify the migration policies that you require for full backups and for incremental/differential backups.

Note: For more information, see [Specify Migration Policies for Tape Staging Backups](#) (see page 233).

- **Miscellaneous**--Specify the Miscellaneous policies required for the backup job.

Note: For more information, see [Specify Miscellaneous Options for Tape Staging Backups](#) (see page 235).

- **Alert**--Specify the Alert policies required for the backup job.

Note: For more information, see Specify Alert Options for Disk and Tape Staging Backups.

- **Postscripts**--Specify the Postscript policies required for the job.

Note: For more information, see [Specify Postscripts Options for Disk and Tape Staging Backups](#) (see page 220).

6. Click the Destination tab and expand the Servers object.

Specify the final destination device group for the job by doing one of the following:

- If the staging device contains two or more drives, you can select any device group on the Destination tab.

Example: You are required to stage your backup data to a virtual tape library and ship the final media to an off-site storage facility. To manage a backup of this type, you can specify a group that corresponds with a virtual library on the Staging Location tab, and then specify a group that corresponds with a tape library on the Destination tab.

- If the final destination device contains one drive (for example, an FSD or a single-drive library), you must select a device group on the Destination tab that is different from the device group specified on the Staging Location tab.

7. Click Options on the toolbar.

The Options dialog opens.

8. Select the Encryption/Compression tab and complete the following fields, as required, for the backup job:

- **Session/Encryption password**--Specify a Session/encryption password to restore this data from media.

Important! If you specify a Session/Encryption password, you must provide this password to restore the session.

- **Encrypt data**--Use this option to encrypt the backup data. You can specify one of the following options:
 - **At agent**--Select this option to encrypt the backup data prior to the actual backup process. For more information about this option, see [Data Encryption at the Agent Server](#) (see page 113).
 - **At backup server during backup**--Select this option to encrypt the backup data at the backup server during the backup process. For more information, see [Data Encryption During Backup](#) (see page 114).
 - **At backup server during migration**--Select this option to encrypt the backup data during the migration phase of a staging backup job. For more information, see [Data Encryption During Migration](#) (see page 114).

- **Compress data**--Use this option to compress the backup data. You can specify one of the following options:
 - **At agent**--Select this option to compress the backup data on the system where the agent is installed and running.
 - **At backup server**--Select this option to compress the backup data at the Arcserve Backup server during the backup process. Using this option directs Arcserve Backup to compress files before backing them up using a software compression algorithm.

Note: The Compress data options do not apply to backups to UNIX and Linux data mover servers.

Click OK.

The Options dialog closes and the specified encryption and compression options are applied.

Be aware of the following:

- If you want to apply other options that affect the migration job, you should do so at this time. For example, to eject the tape from a library after the migration job is complete, click the Operation tab on the Global Options dialog and select the Eject Media option.
- If you specify encryption and compression options, and the backup destination is a drive that does not support compression, or is a file system device (FSD), Arcserve Backup encrypts the backup data and does not compress the backup data.

9. Click Submit on the toolbar to submit the backup job.

The Submit Job dialog opens.

10. Complete the required fields on the Submit Job dialog and click OK.

The job is submitted.

More information:

[Submit a Backup Job](#) (see page 136)

Submit a Tape Staging Backup Job Using a Rotation Scheme

This section topic describes how to set up tape staging (D2T2T) backup jobs using a rotation or GFS rotation scheme. A rotation scheme lets you determine the type of backup (full, differential, and incremental), when to run a backup job, and where to save the backup data (media).

To submit a tape staging backup job using a rotation scheme

1. Open the Backup Manager.
Click the Start tab and click the Enable Staging check box.
The Staging Location and Policies tabs appear.
2. Click the Source tab, browse to and select the source objects that you want to back up.
3. Click the Staging Location tab and the Destination tab to configure the media pool and group.
Specify the values that you require to submit the backup job in the following fields:
 - Media Pool or Media pool (prefix)
Note: The name of the media pool or the media pool prefix that you specify in this field is the name (or prefix) of the media pool that you will use for the tape staging job.
 - Group
4. Select the Schedule tab, and select the Use Rotation Scheme option.
A list of available schemes display in the Scheme Name drop-down list.
5. From the Scheme Name drop-down list, select the scheme that you require for your backups.
6. Click Submit on the toolbar.
The Submit Job dialog opens.
7. Complete the required fields on the Submit Job dialog and click OK.

Archive Staging Methods

Arcserve Backup provides two methods to archive to a staging area and then migrate (or copy) this data to a final destination (usually a tape).

- The disk staging method utilizes a disk as the staging area and is commonly referred to as Archive to Disk to Tape.
- The tape staging method utilizes a tape library or a virtual tape library as the staging area and is commonly referred to as Archive to Tape to Tape.

Each staging method contains specific options to control the behavior of Arcserve Backup during the archive process.

Note: The Archive Staging Methods have the same behavior as the Backup Staging Methods. Refer to the [Backup Staging Methods](#) (see page 200) for details.

Archive Disk Staging

The Archive Disk Staging method utilizes a disk as the staging area and is commonly referred to as Archive to Disk to Tape.

This section describes the following topics:


- [How Archive to Disk to Tape Works](#) (see page 241)
- [How to Use Disk Staging to Manage Archive Data](#) (see page 242)

How Archive to Disk to Tape Works

Archive to disk to tape is a method of protecting data that lets you archive data to a temporary data storage location (staging area), and then based on selected policy options, migrate (copy) the data to the final destination (which could be a tape or disk), or automatically purge the data from the staging area after a specified duration time. When necessary, Arcserve Backup lets you restore data directly from the staging area.

Archive to disk to tape is a two-part archive process.

- **Archive Process**--Arcserve Backup archives data from the source to the staging location. The staging location is a Disk Staging Device.
- **Copy Process**--Arcserve Backup copies or migrates the archive data from the staging location to the final destination media. The final destination is usually tape media, but can be another Disk Staging Device.



Note: Arcserve Backup lets you transmit up to 32 streams of data using multistreaming. To archive data using multiple streams and transmit more than two streams of archive data, you must license the [Arcserve Backup Enterprise Module](#) (see page 40).

How to Use Disk Staging to Manage Archive Data

The following list describes how you can use staging to manage archive data.

- Using disk staging you can archive data to disk staging devices that are used as a temporary staging area. A staging job can divide your archive job into several subjobs that run simultaneously. Disk staging lets you use simultaneous streaming to send multiple streams of data to the disk staging device. Since the data is split among several different streams, archive jobs with simultaneous streaming enabled can be completed significantly faster than regular archive jobs.

- You can then migrate (copy) the data from the disk staging device to a final storage media (or from disk to tape). As a result, the tape drive can be kept streaming, thereby minimizing the shoeshine effect (starting, stopping, and repositioning the tape), and increasing both the life expectancy and efficiency of the tape drive. While the archive image is still on the disk staging device, data can be restored directly from the disk staging device. The restore time is significantly reduced because restoring data from disk is generally faster than restoring from a tape (no delays due to tape load and seek latency).
- During the archive-to-disk staging device process, if the disk staging device gets full or reaches the specified maximum threshold, Arcserve Backup lets you create makeup jobs which would then archive the data directly to the final destination after the staging archive job fails. This increases the success rate of archives. In addition, if there are any errors during the copy-to-final destination process, Arcserve Backup lets you create makeup jobs.
Note: Under disk full conditions, the makeup job created to archive the data to tape will always try to use a blank tape or a media from a scratch set. It will never try to append to an existing media.
- The archive images are kept on the disk staging device until the retention time expires (as determined by the specified purge policy). At that time, Arcserve Backup automatically purges the data from the disk staging device, and reclaims disk space so that archives can continue.

Archive Tape Staging

The Archive Tape Staging method utilizes a tape library or a virtual tape library as the staging area and is commonly referred to as Archive to Tape to Tape.

This section describes the following topics:

- [How Archive to Tape to Tape Works](#) (see page 243)
- [How to Use Tape Staging to Manage Archive Operations](#) (see page 245)

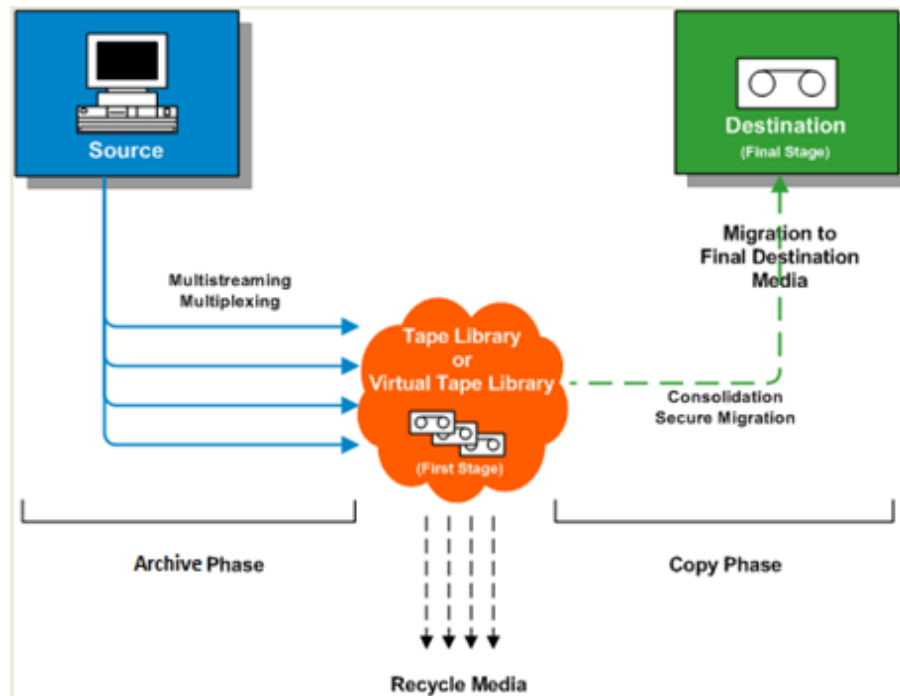
How Archive to Tape to Tape Works

Archive to tape to tape is a data protection solution that lets you to archive data to a tape library or a virtual tape library, and then copy the data to a different tape library or other type of device. Copy operations, also known as migration, are governed by user-specified copy policies.

Archive to tape to tape is a two-part archive process.

- **Archive Process**--Arcserve Backup archive data from the source to the staging location. The staging location is a tape library or a virtual tape library (VTL).
- **Copy Process**--Arcserve Backup copies or migrates the archive data from the staging location to the final destination media. The final destination is tape media.

The following diagram illustrates the flow of data from the source to the first stage tape library (or virtual tape library) and then on to the final destination.



Note: Arcserve Backup lets you transmit up to 32 streams of data using multistreaming. To archive data using multiple streams and transmit more than two streams of archive data, you must license the [Arcserve Backup Enterprise Module](#) (see page 40).

When you use archive to tape to tape to protect data, the archive to tape to tape operation consists of two phases:

Archive Phase

Arcserve Backup archive data from the source to the tapes in the first stage, based on user-specified policies.

- Archive jobs contain only a full backup method.
- During the archive job, global options, media selection rules, media pool usage, rotation rules, GFS rotation rules, Alert messages, Export options, and so on are identical to that of archiving directly to tape.

Note: Various global archive options do not apply to archive phase operations.

- Multiplexing and multistreaming can be used to transmit and save data to first stage media.

Note: The Multiplexing option can be used for archive operations to all tape devices, with the exception of file system devices. The Multistreaming option can be used for archive operations to tape libraries that contain two or more drives.

Copy Phase

Arcserve Backup copies data from the first stage to the final destination based on user-specified policies.

- Arcserve Backup copies data from the first stage media to the final destination media one session at a time. Multiple sessions cannot be copied to a single tape simultaneously.
- If you need to copy data from more than one first stage media to one final destination media, Arcserve Backup copies each session in succession until all the sessions are copied to the final destination media.
- Arcserve Backup sessions associated with different jobs can be consolidated during migration. You can activate this capability using the consolidation option.
- If a hardware error occurs during the process of copying data to final destination, the job stops and Arcserve Backup creates a **Makeup Job On Hold**. After you correct the hardware error, you can change the job status to **Ready**, and then the job resumes.

How to Use Tape Staging to Manage Archive Operations

The following are common scenarios that describe how you can use Archive to Tape to Tape operations to manage archive operations:

- If you need to store two copies of archive data, one copy on site and one copy at an off-site storage location, Archive to Tape to Tape lets you back up data directly to tape. After the archive job is complete, you can use Arcserve Backup copy utilities to automate and create copies of the archive tapes, and then ship the tapes to an off-site storage location.
- Archive to Tape to Tape lets you encrypt archive data when you are copying the data to its final destination media. This capability is beneficial when you are copying data from a virtual tape library or a library that does not support encryption to a library that supports encryption. This capability ensures that your archives are as fast as possible and the tapes that you need ship to an off-site storage location are encrypted.
- While archive operations are in progress, you may have many jobs that are archiving data to many different tapes. This can result in media that is not being used to its full capacity. Archive to Tape to Tape operations let you consolidate archives to ensure that media is being used to capacity when copying data to final destination. This capability helps to reduce the cost of media because you are using fewer tapes for final destination media, off-site storage, or both.

- If you need to reduce the length of time required to archive data and copy the data from a staging area to final destination, you can use virtual tape libraries (VTL) in your environment to manage archive operations.

A VTL is a temporary storage location, such as a disk drive, that is configured to behave like a library. Since most archive data is transmitted across a network, Arcserve Backup lets you use multiplexing to reduce the archive window. When you use a VTL to store archive data, you can quickly read data from a multiplexing formatted data in a VTL because your operations do not encounter tape positioning overhead. As a result, the processes of archiving data to a VTL, reading from a VTL, and copying the data to final destination media is fast. Arcserve Backup lets you automate the process of copying to final destination media when you use a VTL to stage your archive data.

Backing up Multiple Data Mover Servers in a Single Job

Arcserve Backup lets you submit backups that consist of multiple data mover servers to shared tape libraries in a single job. This capability lets you simplify the process of managing backups, and helps to minimize the amount of media used to store the backup data.

- **Backup types supported**--Arcserve Backup lets you submit normal UNIX/Linux Data Mover backups and tape staging UNIX/Linux Data Mover backups.
- **Storage devices supported**--Arcserve Backup lets you submit multiple UNIX/Linux Data Mover server backups in a single job to shared tape libraries.
- **Licensing requirements**--Arcserve Backup lets you submit UNIX/Linux Data Mover backups with the following licenses:

Arcserve Backup UNIX and Linux Data Mover must be installed on the data mover servers. The licenses for UNIX and Linux Data Mover must be registered with the primary server.

Note: For more information about licensing requirements for UNIX and Linux Data Mover, see the *UNIX and Linux Data Mover Guide*.

- **Backup considerations**--UNIX/Linux Data Mover backup does not support the following Arcserve Backup functionality:
 - Backing up data to Arcserve Backup data deduplication devices, file system devices, and Arcserve Backup Tape RAID devices.
 - Backing up data using multiplexing, using Arcserve Backup server-side encryption, Arcserve Backup server-side compression, and LTO encryption.
 - Protecting Oracle database data at Oracle object level granularity.

Note: For more information about installing and using UNIX and Linux Data Mover, see the *UNIX and Linux Data Mover Guide*.

This section contains the following topics:


[Back up Multiple Data Mover Servers in a Single Job](#) (see page 247)

[Back up Multiple Data Mover Servers in a Single Job Using Staging](#) (see page 250)

Back up Multiple Data Mover Servers in a Single Job

Arcserve Backup lets you submit backups consisting of multiple data mover servers to a shared tape library in a single job.

The following diagram illustrates how Arcserve Backup processes backups of data mover servers in a single job. Note that this type of job does not use tape staging processes.



Prerequisite Tasks

- Verify that you configure at least one library. For information about how to configure libraries, see [Configure Libraries](#) (see page 374).

To back up multiple data mover servers in a single job

1. Open the Arcserve Backup Manager Console.

From the Quick Start menu, select Backup.

The Backup Manager opens and the Start tab displays.

2. From the Start tab, click UNIX/Linux Data Mover backup.

The backup type is applied to the job.

Important! You cannot specify file system devices as staging devices for UNIX/Linux Data Mover backups. To use a file system device as a staging device for data mover backups, you must specify Normal backup.


3. Click the Source tab.

The backup source directory tree appears.

4. Expand the UNIX/Linux Data Movers object.

The data mover servers appear.

5. Specify the source that you want to back up, as illustrated by the following screen:



Click the Schedule tab.

The schedule options appear.

6. Specify the schedule options that you require for the job.

Note: For more information about scheduling jobs, see [Job Customization Methods](#) (see page 301).


Click the Destination tab.

The destination group directory tree appears.

- Expand the SAN object and expand the Shared Groups object.

Specify the Device Group where you want to store the backup data.

Note: Click the Device Group to identify the library associated with the group as illustrated by the following screen:



Important! Arcserve Backup prevents you from submitting backup jobs when the data mover server specified on the Source tab does not share the device group specified on the Destination tab.

- (Optional) On the Destination tab, click the Multistreaming check box to back up your data using multistreaming. Without multistreaming, Arcserve Backup processes backup sessions sequentially. With multistreaming, Arcserve Backup lets you distribute the backup sessions across multiple streams, which helps to decrease the overall length of time required to complete the backup.

Note: For more information, see [How Arcserve Backup Process Backup Data Using Multistreaming](#) (see page 103).

- (Optional) Click Options on the toolbar to specify additional options that you require for the backup.

Note: For more information about backup job options, see [Global Backup Options](#) (see page 154).

- Click Submit on the toolbar to submit the job.

The Security and Agent Information dialog opens.

11. On the Security and Agent Information dialog, edit or confirm the security and agent information for your job, and click OK.

The Submit Job dialog opens.


12. Complete the required fields on the Submit Job dialog and click OK.

The job is submitted.

Back up Multiple Data Mover Servers in a Single Job Using Staging

Arcserve Backup lets you submit backups consisting of multiple data mover servers to a shared tape library, using staging (D2T2T), in a single job.

The following diagram illustrates how Arcserve Backup processes backups of data mover servers in a single job using staging.



Prerequisite Tasks

- Verify that you configure at least one library. For information about how to configure libraries, see [Configure Libraries](#) (see page 374).


To back up multiple data mover servers in a single job using staging

1. Open the Arcserve Backup Manager Console.
From the Quick Start menu, select Backup.
2. The Backup Manager opens and the Start tab displays.
From the Start tab, click the following:
 - UNIX/Linux Data Mover backup
 - Enable staging

The backup type is applied to the job.

Important! You cannot specify file system devices as staging devices for UNIX/Linux Data Mover backups. To use a file system device as a staging device for data mover backups, you must specify Normal backup.

3. Click the Source tab.
The backup source directory tree appears.
4. Expand the UNIX/Linux Data Movers object.
The data mover servers appear.
5. Specify the source that you want to back up, as illustrated by the following screen:



- Click the Schedule tab.
The scheduling options appear.

- Specify the schedule options that you require for the job.

Note: For more information about scheduling jobs, see [Job Customization Methods](#) (see page 301).

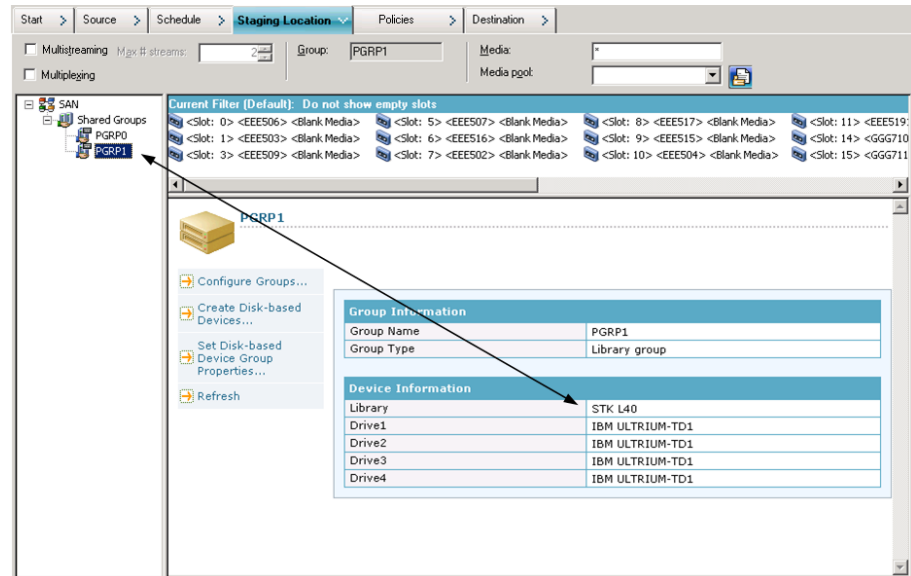
Click the Staging Location tab.

The staging location and group directory tree appears.

- Expand the SAN object and expand the Shared Groups object.

Specify the Device Group where you want to stage the backup data.

Note: Click the Device Group to identify the library associated with the group as illustrated by the following screen:



Important! Arcserve Backup prevents you from submitting backup jobs when the data mover server specified on the Source tab does not share the device group specified on the Staging Location tab.

- (Optional) On the Staging Location, click the Multistreaming check box to back up your data using multistreaming. Without multistreaming, Arcserve Backup processes backup sessions sequentially. With multistreaming, Arcserve Backup lets you distribute the backup sessions across multiple streams, which helps to decrease the overall length of time required to complete the backup.

Note: For more information, see [How Arcserve Backup Process Backup Data Using Multistreaming](#) (see page 103).

- Click the Policies tab.

The migration policy options appear.

10. Complete the following migration policies required for the job:

- **Full Backup and Incremental/Differential Backup**--Specify the migration policies that you require for full backups and for incremental/differential backups.

Note: For more information, see [Specify Migration Policies for Tape Staging Backups](#) (see page 233).

- **Miscellaneous**--Specify the Miscellaneous policies required for the backup job.

Note: For more information, see [Specify Miscellaneous Options for Tape Staging Backups](#) (see page 235).

- **Alert**--Specify the Alert policies required for the backup job.

Note: For more information, see [Specify Alert Options for Disk and Tape Staging Backups](#).

- **Postscripts**--Specify the Postscript policies required for the job.

Note: For more information, see [Specify Postscripts Options for Disk and Tape Staging Backups](#) (see page 220).


Click the Destination tab.

The Destination directory tree appears.

11. Expand the Servers object on the Destination tab.

Specify the device group containing the device where you want to store the data.

Note: Arcserve Backup lets you execute migration jobs via the primary server or the data mover server. Migration jobs execute from the primary server or the data mover server based on the method used to specify the final destination media, as illustrated by the following screen:



12. (Optional) Click Options on the toolbar and specify the options that you require for the job.

As a best practice, you should apply other options that affect the migration job at this time. For example, to eject the tape from a library after the migration job is complete, click the Operation tab on the Global Options dialog and select the Eject Media option.

Note: For more information about backup options, see [Global Backup Options](#) (see page 154).

13. Click Submit on the toolbar to submit the job.

The Security and Agent Information dialog opens.

14. On the Security and Agent Information dialog, edit or confirm the security and agent information for your job, and click OK.

The Submit Job dialog opens.

15. Complete the required fields on the Submit Job dialog and click OK.

The job is submitted.

Disaster Recovery

To ensure against data loss, maintain current backups of all your servers and workstations. If you do not have these backups, Arcserve Backup is limited in its ability to recover data. Be sure to create a media rotation scheme and a schedule to maintain current backups.

By default, the Arcserve Backup server always generates or updates disaster recovery information for all full backup systems, even when the Arcserve Backup Disaster Recovery Option is not installed. This ensures that the latest backup information is always available if the Arcserve Backup Disaster Recovery Option is installed a later time.

Note: To disable the Arcserve Backup server from generating or updating the disaster recovery information, create and set the following registry key value to 1 on the Arcserve Backup server machine:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Computer Associates\CA ARCserve  
Backup\Base\Task\backup\SkipDRSession
```

For more information, see the *Disaster Recovery Option Guide*.

How NTFS Deduplication Works

NTFS data deduplication is functionality that is included with Windows Server 2012 operating systems. NTFS data deduplication removes duplicate data from a volume, which in turn lets you store more data on the volume.

NTFS data deduplication should not be confused with Arcserve Backup data deduplication. NTFS data deduplication is Windows implementation that affects the data that resides on a volume. Arcserve Backup lets you back up and restore data that resides in NTFS data deduplication enabled volumes. Arcserve Backup data deduplication is an implementation that affects how backup data is stored on storage media. For more information about the Arcserve Backup implementation of data deduplication, see the section Using Deduplication. For more information about NTFS data deduplication, see the [Microsoft Windows Development Center website](#).

You define NTFS data deduplication as a [local backup option](#) (see page 134) or as a [global backup option](#) (see page 159).

The process of backing up NTFS data deduplication data is the same as the process for performing any other type of standard backup to disk or to tape media. This approach also applies to restoring data that was optimized by NTFS data deduplication. However, when you restore data that was backed up as optimized NTFS data deduplication, there must be a sufficient amount of free disk space on the destination volume that is greater than or equal to the logical size of the files that you want to restore.

Example: File A is 100 MB and file B is 100 MB. Both files contain the exact same data. With an optimized NTFS data deduplication backup, the storage space that you require is 100 MB. However, to recover both files, you require 200 MB of free disk space on the destination volume.

Note: This example does not apply to recovering data using Arcserve Backup disaster recovery operations.

With NTFS data deduplication you can perform optimized backups and non-optimized backups, given that NTFS data deduplication is enabled on the volume that you are backing up. With optimized backups, Arcserve Backup backs up the optimized files, the non-optimized files, and the chunk store. When you perform optimized NTFS data deduplication backups but NTFS data deduplication is not enabled on the volume, Arcserve Backup performs optimized backups.

The following limitations apply to NTFS data deduplication backups:

- By default, the optimize NTFS data deduplication option is enabled as a global backup option and disabled as a local backup option. With this configuration, Arcserve Backup performs optimized backups on all NTFS deduplication enabled volumes.

To enable or disable this option for backup jobs for specific volumes, disable this option as a [global backup option](#) (see page 159) and then enable or disable (as appropriate) as a [local backup option](#) (see page 134).

- Arcserve Backup can recover data from optimized backup sessions to only computers running Windows Server 2012.
- With NTFS data deduplication backups, Arcserve Backup backs up only the data deduplication settings. For example, the Enable data deduplication, Deduplicate files older than (in days), and File extensions to exclude options. For all other options, (for example, the Set Deduplication Schedule settings) Arcserve does not include the settings with full volume backups because it cannot restore the settings with full volume restores. When you want to recover all of the deduplication settings, you must recover the entire computer or perform a disaster recovery of the computer.
- Arcserve Backup cannot perform Compare jobs on data that was backed up using the NTFS data deduplication (optimize) option. When you submit Compare jobs of this type, the jobs fail and an error indicating as such (Compare job not supported) appears in the Activity Log.
- When you perform optimized NTFS data deduplication backups and specify the Compare Backup Media to Disk option, Arcserve Backup compares only the reparsed files (not the entire volume).

Chapter 4: Restoring Data

Arcserve Backup provides you with various tools and options that you can use to restore data. This section includes information about how you can safely and efficiently restore data.

This section contains the following topics:

[Restore Manager](#) (see page 257)

[How to Find Files That You Want to Restore](#) (see page 258)

[Restore Manager Markers](#) (see page 268)

[Restore Manager Location Options](#) (see page 269)

[Restore Job Schedules](#) (see page 269)

[Specify Run as Administrator on Windows Server 2008 and 2012 Systems](#) (see page 270)

[Global Restore Options](#) (see page 271)

[System State Restore Options](#) (see page 278)

[Restoring Data Scenarios](#) (see page 280)

[How to Restore Data from Cloud Storage](#) (see page 299)

Restore Manager

The aim of running a successful restore job is to identify the data you need quickly and to retrieve it from the appropriate backup and archive media.

Arcserve Backup allows you to restore data to most computers attached to your Windows network. Each restore job requires a source and destination. The files selected as your source must originate from backup or archive media created by Arcserve Backup, and the destination must be a hard drive. The Restore Manager provides three tabs to customize your restore job:

- Source
- Destination
- Schedule

The optional Arcserve Backup Client Agents allow you to communicate with remote workstations in various environments to restore data to non-Windows systems, such as UNIX.

Similarly, the optional Backup Agents allow Arcserve Backup to restore online databases and applications such as Microsoft Exchange Server, Microsoft SharePoint Server, Microsoft SQL Server, Lotus Domino, Oracle, and IBM Informix.

How to Find Files That You Want to Restore

Arcserve Backup makes it easy to find the files you want to restore. Because your requirements and circumstances can vary, Arcserve Backup provides you with the following methods for selecting the data (the source) you want to restore:

- **Restore by Tree**--Lets you restore a specific directory or drive from a display of files and directories that were backed up or archived with Arcserve Backup. Use this method when you do not know which media contains the data you need, but you know the computer from which the backup or archive originated.

The Restore by Tree view displays only the last instance of a backup or archive. To view and access all other instances, select the object that you want to restore and click the Recovery Point button. If there are multiple partial backups/archives of the same drive, the Restore by Tree view displays only the last backup/archive. However, if there is a full-volume backup/archive of the drive available, the last full backup/archive is displayed, instead of the last partial backup/archive.

The Computer Name field allows you to filter based on partial name searching. You can enter any part of the name and a list of relevant items are returned. For example, if there are some computers whose computer name contains 'BB', you can enter 'BB' in the Computer Name field and click the Update button. Relevant computers are found. The Computer Name field also supports full name searching and wildcard searching.

Note: The Restore Manager cannot display file paths that exceed 512 bytes. This limitation includes the drive letter or network server and the share names. For single-byte languages, this equates to approximately 500 characters. For multibyte languages with a combination of single, mixed, and multibyte characters, 512 bytes equates to 250 to 500 characters. For multibyte languages with all multibyte characters, 512 bytes equates to approximately 250 characters. If a file path exceeds 512 bytes, truncation occurs. To restore data from a truncated directory, you must submit the restore job from the last directory in the path whose name was not truncated.

Note: Arcserve Backup lets you export the results from the Restore Manager tree to a text file. Right-click the session, media, directory, or server node and click Export Data. Save the file to your local drive.

- **Search** --Click the Search button to search your backups or archives for a specific file or group of files with a similar file name. Arcserve Backup lets you specify file names up to 255 characters, including the file extension, in the Search for field. 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 Search for field.

Note: The Search restore method does not work if the Database Engine is stopped.

Examples:

1. Drive D:\ contains two directories that are backed up/archived on a weekly basis—D:\Temp and D:\Documents. D:\Temp and D:\Documents were both backed up/archived on April 21 and April 28. A full backup/archive of drive D:\ was performed on April 1.
 2. The Restore Manager displays instances relating to the full backup/archive of Drive D:\ performed on April 1.
 3. To restore the April 28 instance of D:/Documents, select the D:/Documents directory in the Restore by Tree view and click the Recovery Point button. From the Recovery Point dialog, select the April 28 instance and then click the Select button.
- **Restore by Session**--Lets you select the session, the files, and directories you want to restore. Use this method when you know the media name, but are not certain about the session you want to restore.

This restore method does not work if the Database Engine is stopped.

Deduplication devices support the Restore by Session, but will likely contain thousands of sessions. You will be prompted to specify a display option to manage the volume.

The Media Name field allows you to filter based on partial name searching. You can enter any part of the name and a list of relevant items are returned. For example, if there are some sessions whose media name contains 'BB', you can enter 'BB' in the Media Name field and click the Update button. Relevant sessions are found. The Media Name field also supports full name searching and wildcard searching.

Note: The Restore Manager cannot display file paths that exceed 512 bytes. For single-byte languages, this equates to approximately 500 characters. For multibyte languages with a combination of single, mixed, and multibyte characters, 512 bytes equates to 250 to 500 characters. For multibyte languages with all multibyte characters, 512 bytes equates to approximately 250 characters. If a file path exceeds 512 bytes, truncation occurs. To restore data from a truncated directory, you must submit the restore job from the last directory in the path whose name was not truncated.

Note: Arcserve Backup lets you export the results from the Restore Manager tree to a text file. Right-click the session, media, directory, or server node and click Export Data. Save the file to your local drive.

- **Restore by Query**--Lets you restore files based on the search pattern used to locate the names of the files or directories. Lets you specify the type of query to restore, for example, All (default), Backup Session, or Archive. Use this method when you know the name of the file or directory you want to restore, but do not know the computer that was backed up or archived from or the media it was backed up or /archived to.

Restore by query is not a case-sensitive operation.

Arcserve Backup lets you specify file names up to 255 characters, including the file extension, in the File Name field. If you do not know the complete file name, you can simplify the results of the query by specifying the wildcard characters "*" and "?" in the File Name field.

Note: This restore method will not work if the Database Engine is stopped.

- **Restore by Backup Media**--Lets you restore a complete backup session from a specified media in a storage device. All files in the session are restored to the destination, unless filters are added to the restore job. Use this method when media was created by a different version of Arcserve Backup or if the database does not recognize it.

Important! If you cannot see the items that you would like to restore, the corresponding records may have been pruned from your database. You can repopulate your restore source selection by running the Merge utility. For more information about the Merge utility, see the section Merge Utility.

- **Recover Virtual Machine**--There are two options recover a virtual machine:
 - **Search for Virtual Machine using the VM host name or VM type**--Lets you recover a virtual machine by VM name or VM type. If you search by VM type, you must have VMware Converter 3.0.2 installed on the proxy or recovery jobs fail. Using this restore method enables additional options on the Global Options Operation tab.
 - **Transport Modes for VMware VM**--Lets you recover a virtual machine by selecting a specific transport method.
 - **Dynamic**--(*default*) This option lets VMware Virtual Disk Development Kit (VDDK) select the available transport mode.
 - **SAN**--(Storage Area Network) This option lets you transfer backup data from proxy systems connected to the SAN to storage devices using Fibre Channel communication.
 - **HOTADD**--This option lets you back up virtual machines configured with SCSI disks.
 - **NBDSSL**--(Network Block Device Secure Sockets Layer) This option lets you use the Network File Copy (NFC) protocol to communicate. NBDSSL transfers encrypted data using TCP/IP communication networks.
 - **NBD**--(Network Block Device, also referred to as LAN transport mode) This option lets you use 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.

Note: If the specified transport mode is not available, the transport mode defaults back to the Dynamic option.

Recover Virtual Machine makes the specified VM unavailable while the recovery job is in progress.


If you installed Arcserve Backup Enterprise Module, you will also have the Restore by Image method available. Use this method when you must read and restore blocks of data quickly by bypassing the file system.

How Arcserve Backup Lets You Browse a Large Number of Items in the Restore Manager


Arcserve Backup lets you pause the process of loading items in the Restore Manager when you browse a large number of directories, files, and so on. You can pause the loading process when you click the Destination tab, clear the check mark from the Restore files to their original location(s) option, and browse items on the Destination tab.

The steps that follow describe how Arcserve Backup lets you browse a large number of items in the Restore Manager window.


1. When you select an item in from the directory tree in the Restore Manager, Destination tab, Arcserve Backup displays a Loading dialog to inform you that a large number of items need to be retrieved and loaded into the Restore Manager window. You cannot click Cancel while Arcserve Backup is retrieving the list of items to display in the Restore Manager window.



2. After Arcserve Backup retrieves the list of items to display in the Restore Manager window, the Loading dialog then displays the percentage of items that are loaded into the Restore Manager. You can click Cancel to pause the operation.



- After you pause the Loading operation, you can continue the Loading operation by right-clicking the target directory and selecting Show More from the pop-up menu.



- If you pause the loading process, the icon for the target directory displays as follows:



- You can pause and continue the loading process as often as necessary. To load more items, right-click the target directory and click Show More from the pop-up menu.

- When the loading process is complete, the icon for the target directory displays as follows:



Browse a Large Number of Files in the Restore Manager

Use the following procedure when you need to browse a directory that contains a large number of items in the Restore Manager.

Note: The Restore Manager can display file name paths that contain up to 512 characters. This behavior includes the drive letter or network server and the share names. On NTFS and FAT32 file systems, the maximum length of file names is limited to 255 characters.

To browse a large number of files in the Restore Manager

1. Open the Restore Manager and click the Destination tab.

On the Destination tab, clear the checkmark from Restore files to their original location, and then browse to and specify a target directory from the directory tree.

The Loading message box appears.

2. From the Loading message box, click Cancel to stop the loading process.

If Arcserve Backup did not load all items, the To show more objects, right-click the target directory and select Show More from the pop-up menu warning message appears.

Note: The message only appears the first-time you click Cancel on the Loading message box.

3. From the directory tree, right-click the target directory and click Show More from the pop-up menu.

The Loading message box appears and Arcserve Backup continues loading the items.

4. You can pause and continue the loading process as often as necessary until Arcserve Backup loads all items in the target directory.

If you pause the loading process, the icon for the target directory displays as follows:



When the loading process is complete, the icon for the target directory displays as follows:



Duplicate Backup Sessions

When you back up data using disk staging, or copy media using the `tapemedia` command line utility, duplicates of backup sessions can exist in multiple locations. For example, you can define your staging copy and purge policies such that backup sessions remain on the file system device used for staging for a period of time after the copy to final destination media operation occurs. If the backup session was not purged from the file system device, data will reside on the file system device and the final destination media. If this situation presents itself, you can restore the session by using data that resides on the file system device.

When you copy media, duplicate backup sessions exist on multiple media. If one media remains on site and the other media was vaulted, Arcserve Backup lets you use the media that is on site to facilitate the restore operation.

Smart Restore

Arcserve Backup provides a transparent Smart Restore feature that can increase the overall success rate of your restore operations. If a media read error or a hardware error occurs during a restore job, Arcserve Backup searches for an alternate media to use to complete the restore job.

Example: Smart Restore

During a restore job, the restore source media jams and disables the library. Arcserve Backup then searches for duplicates of the backup session. If a duplicate of the session exists, regardless of whether it exists on a file system device or another media, the restore operation continues without user intervention.

Note: If a second media error occurs during the restore job, the job will fail.

Export the Restore by Query Results and View the Results in a Spreadsheet

Arcserve Backup lets you query the Arcserve Backup database and export the results of the query to a text file. Arcserve Backup exports the values in tab-separated format. With a tab-separated format, you can import the data into a spreadsheet application (for example, Microsoft Excel) to analyze the results.

Example: Export the Restore by Query Results and View the Results in a Spreadsheet

Users asked you to restore several files that reside on different computers in your environment. The users do not know the precise file names. There are other files with similar file names on the computers. You can query the Arcserve Backup database using wildcard characters to obtain the host names, file paths, file names, and file modification dates. Using an Excel spreadsheet, you can sort the results and then ask your users to inform you which files to restore.

To export Restore by Query results and view the results in a spreadsheet

1. From the Quick Start menu on the Arcserve Backup Home Page, click Restore.

The Restore Manager window opens.

2. From the drop down list, click Restore by Query.

The query options fields appear.

3. Specify the values that you require and click Query.

The results of the query appear below the query fields.

4. Click Export Query Result.

Arcserve Backup gathers the query results and the Save as dialog opens.

5. Specify a location and a file name, and then click Save.

The query results are saved to a text file.

6. Open your spreadsheet application.

Import the text file that you just created.

Note: For information about how to import text files, see the documentation for your spreadsheet application.

The results of the query appear in the spreadsheet.

Restore Data by Query on UNIX and Linux Platforms

The Restore by Query method of restoring data lets you search for and restore files based on the search criteria used to locate the names of the files or directories stored in your backup data. On UNIX and Linux platforms, the syntax that you use to query the Arcserve database based on the Look in Directory (file location) option is different from that of Windows platforms. The following procedure describes the syntax that you will use to restore data by query on UNIX and Linux platforms.

To restore data by query on UNIX and Linux platforms

1. From the Restore Manager, select Restore by Query from the Source view drop-down list.

The Restore by Query fields appear.

2. To specify your search criteria, complete the following fields:

- **Computer Name**--Lets you specify the name of the computer that you want to search. You can specify a specific computer name or select <<ANY>> from the drop-down list to search all computers in your Arcserve environment.
- **File Name**--Lets you specify a wild card or specific file name search. On UNIX and Linux platforms, Arcserve Backup uses the standard 8.3 file naming convention. For example, if you specify *.txt, all files with a .txt file extension appear in the query results.

Note: Do not specify leading or trailing spaces in this field.

- **Look in Directory**--Lets you specify the directory that you want to search. You must specify an exact string match, starting with the drive letter, in this field.

UNIX and Linux platforms regard the backward slash character "\" as a separator. For example, \root\dir1\text.txt.

Examples: Look in Directory

If the mount point is "/", use the following search string:

```
^root\dir1\text1
```

If the mount point is "/root", use the following search string:

```
Vroot\dir1\text1
```

Note: Do not specify leading or trailing spaces in this field.

- **Include Subdirectories**--Lets you search the subdirectories of the directory specified in the Look in Directory field.

3. Click Query.

Arcserve Backup queries the database and returns the files that meet your search criteria.

4. Select the files and directories you want to restore by double-clicking the name of the file or directory.




A green light appears when a file or directory is selected.

Restore Manager Markers

Each object displayed in the Restore Manager window has a green or gray box to its left called a marker.

- **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 you click a marker at a higher, parent level so that it is completely filled, all the markers at the lower, child levels are automatically filled completely.
- If you click all the markers at the lower, child levels so that they are completely filled, then the marker at the higher, parent level is automatically partially filled.
- If the markers at the lower, child levels are a mix of completely filled and partially filled, the marker at the higher, parent level is automatically partially filled.

Restore Manager Location Options

Arcserve Backup provides you with two methods for selecting the location that you want to restore the data to:

- Restore files to their original location
- Restore to user-shared directories and drives

Note: The default method is to restore files to their original location. If you clear the Restore files to their original location check box, Arcserve Backup presents you with a list of machines, directories, and files that you can specify for the location to restore the data.



Restore Job Schedules

Jobs can be submitted so that they repeat as follows:

- **Once**--Do not repeat this job.
- **Every n frequency**--Repeat this job every specified number of Minutes, Hours, Days, Weeks, or Months.
- **Day(s) of the Week**--Repeat this job on the days that are checked off.
- **Week(s) of the Month**--Repeat this job on the weeks that are checked off.
- **Day of the Month**--Repeat this job on the specified day.
- **Custom**--Repeat this job on the month, day, hour, or minute specified.

Note: If you select the Run Job Now option when your storage device is busy, Arcserve Backup reports that the storage device is busy and the job is not submitted to the job queue. You should schedule your job, keeping the current date and time. This way, when Arcserve Backup discovers that the storage device is busy, it automatically retries the job until the drive becomes free.

For a description of detailed job scheduling features, see “Customizing Jobs.”

Specify Run as Administrator on Windows Server 2008 and 2012 Systems

On Windows Vista and Windows Server 2008 and 2012 operating systems, a security feature prompts you to provide or confirm administrator credentials (user name and password) each time you attempt to launch an executable or application. To bypass the continuous prompting you can specify to run each executable or application as an administrator.

For example, if you want to run the Windows Command Prompt, locate the Command Prompt icon (from the Start menu), right-click the icon, and select the Run as administrator from the pop-up menu. After your administrator privileges have been established for the Command Prompt console, all subsequent Command Prompt invocations can be launched without any further prompts until you close the Command Prompt console.

Note: This task should be performed on all Arcserve Backup executables and applications. For example, ca_auth, ca_backup, ca_restore, cabatch, and so on.

Follow these steps:

1. From Windows Explorer, locate the executable or application that you want to specify as Run as administrator.

Right-click the executable or application and select Run as administrator from the pop-up menu.

Windows prompts you to provide administrator credentials (a user name and a password).

2. Do one of the following when you are prompted to provide administrator credentials:
 - If you are not logged in as an administrator, enter the administrator user name and password.
 - If you are logged in as an administrator, click Continue.
3. Follow the prompts and complete the required fields to complete this task.

More information:

[Authentication Levels for Arcserve Backup Services, Components, and Applications](#) (see page 527)

Global Restore Options

This section describes the global restore options you can select when submitting your restore job. To access the global options dialog, click the Options button in the Restore Manager.

Global options (also referred to as job-wide options) define the procedures and rules that Arcserve Backup uses to restore all nodes (servers, workstations, and applications) that you define for the particular restore job. The global options that you define for one restore job do not affect the global options that you define for any other restore job.

The available options are as follows:

- [Backup Media options](#) (see page 271).
- [Destination options](#) (see page 272).
- [Operation options](#) (see page 274).
- [Pre/Post options](#) (see page 276).
- [Job Log options](#) (see page 277).
- [Alert options](#) (see page 277).

Restore Manager Backup Media Options

The Restore Manager supports the following backup media options:

- **Timeout Options**--You can specify a timeout period that Arcserve Backup will wait to provide the media you need to restore your data. Available media options are:
 - **Timeout for First Backup Media**--Period of time that Arcserve Backup waits for the first media required for your restore job. If the time expires, the job fails.
 - **Timeout for Additional Backup Media**--Period of time that Arcserve Backup waits for any additional media to become available.

- **Optimize Restore**--If, during a restore operation, Arcserve Backup discovers duplicate backup sessions, where one session resides on tape media and another session resides on a file system device, the Optimize Restore option directs Arcserve Backup to restore the data from the session that resides on the file system device.

The Optimize Restore option is a global setting that is applied to all restore operations, and is enabled by default.

Under most circumstances, restoring data from a file system device is faster than restoring from tape media. However, you may wish to consider disabling the Optimize Restore option if you are using tape media or a library with high-speed reading capabilities, or there is a known problem with your file system device.

To disable the Optimize Restore option, clear the check mark from the Optimize Restore check box.

Restore Manager Destination Options

The Destination options determine how the directory structure is created on the destination when files are copied or restored. They also determine which files (if any) can be overwritten.

Directory Structure Options

Select one of the following methods Arcserve Backup should use to create directories on your destination.

- **Do Not Create the Base Directories**--(default) Do not create the base directory on the destination path, but create all subdirectories below the source base directory. A base directory is considered the first directory selected in the source path.
- **Create Directories from the Base**--Create the destination path beginning from the base directory.
- **Create Entire Path from the Root**--Create the entire source path (except the root drive or volume name) on the destination. No files from any parent directories are restored. Only the directory path to the base directory is created on the destination.

File Conflict Resolution Options

Select the method Arcserve Backup should use when there are files on the destination disk that have the same name as files being copied from the source. The default is Overwrite All Files.

- **Overwrite All Files**--Restore all source files to the destination regardless of conflicting file names. The files from the source overwrite existing files on the destination.
- **Rename Files**--Copy the source file to the destination with the same file name but a different extension. The format of the renamed extension will vary based upon the file system that is present on the target partition.
 - If the length of file name is more than 251 characters, Arcserve Backup truncates the file name at 251 characters and appends '.__0' to the file name, after the first restore. For all subsequent restores, Arcserve Backup appends '.__1', '.__2', and so on to the truncated file name.
 - If the length of the file name is less than or equal to 251 characters and has a file extension, Arcserve Backup replaces the last character of the file extension with the character 1 (for example, filename.tx1). For subsequent restores, Arcserve Backup replaces the last character of the file extension with the character 2, 3, and so on. After the 10th restore, Arcserve Backup replaces the last two characters of the file extension with 10, 11, 12, and so on (for example, filename.t10). After the 100th restore, Arcserve Backup replaces the last three characters of the file extension with 100, 101, 102, and so on (for example, filename.100). After the 999th restore, Arcserve Backup cannot rename the file extension, which causes the restore to fail. If the length of the file name is less than or equal to 251 characters, and it does not have a file extension, Arcserve Backup appends '.__0' to the end of the file name. If Arcserve Backup appends '.__0' to the file name after the first restore, the renaming process appends two characters after the 10th restore (for example, filename.__10), and after the 100th restore, the renaming process appends three characters to the file name (for example, filename.__100). After the 999th restore, Arcserve Backup cannot rename the file name, which causes the restore to fail.
- **Skip Existing Files**--Do not restore a source file if a file with the same name already exists on the destination.
- **Overwrite with Newer Files Only**--Only restore source files whose modification date is later than the modification date of the file with the same name on the destination. Source files whose modification date is earlier are not copied to the destination.

VMS File Version Options

The following options indicate how Arcserve Backup should act when restoring VMS files that have the same names and version numbers as the files in the target restore directory.

- **Create New File Version**--Arcserve Backup will restore all files as new versions of the original. The files in the target directory will not be affected.
- **Replace Current File Version**--If a file in the target directory has the same name and version number as a file in the restore data, Arcserve Backup will overwrite the file.
- **Restore File Version**--If a file in the target directory has the same name and version number as a file in the restore data, Arcserve Backup will not restore the file. All other files will be restored with their original names and version numbers.

Restore Manager Operation Options

Operation options let you determine the actions or related action that you want to perform while a job is in progress or after a job completes, and the level of detail that is recorded in the Arcserve Backup database.

The options that follow affect the Arcserve Backup database:

- **Record Job Information Only**--Record job information.
- **Disable Database Recording**--Do not record job information.
- **Restore and Preserve Directory Attributes and Security Information**--Restore the existing directory attributes (such as Read Only, Archive, and Hidden) and security data to the machine.
- **Restore and Preserve File Attributes and Security Information**--Restore the existing file attributes (such as Read Only, Archive, and Hidden) and security data to the machine.
- **Restore Registry Files and Event Logs**--Restore registry files and event logs to the restore target machine if the sessions selected for restore have the registry files and event log files.

The following options apply to only virtual machine restores. These options appear on the Operation tab only when Recover VM is the restore method specified on the Restore Manager.

- **Power on VMware or Hyper-V VM after restore**--Powers on the VM after the restore job is complete.

Default value: Enabled.

- **Overwrite VMware VM, if it exists**--Lets you overwrite the VM, if the VM exists.

When you restore a VMware VM, Arcserve Backup detects the VMs that reside in the host system. If the VM exists in the host system, this option lets you overwrite the VM using the existing UUID of the VM.

Default value: Disabled.

Note: For Hyper-V VMs, the agent always overwrites the VM, if the VM exists in the Hyper-V host.

The option that follows affects Arcserve Replication scenarios:

- **Continue the restore job even when the scenario cannot be stopped**--Lets you restore a Arcserve Replication scenario while you are backing up the scenario.

When you attempt to restore a Arcserve Replication scenario while you are backing up the scenario, by default, the restore job will fail. With this option specified, Arcserve Backup will complete the restore job while a backup is in progress.

Note: This option appears on the Global Options dialog only when you integrate Arcserve Backup with Arcserve Replication.

Restore Manager Pre/Post Options

The Pre/Post options let you run commands on your system before or after jobs execute.

For example, you can use the Pre option to stop the application that owns the data you are backing up, and then use the Post option to start it the application after the backup is complete.

Note: Commands with executables on remote systems are not supported.

- **Run Command Before Job**--Select the following options to run a command on your machine before the job is executed:
 - Enter the path to, and name of, the file to be executed on the machine before the job starts.
 - **On Exit Code**--Arcserve Backup detects exit codes of other programs. For a specified exit code, you can choose to run the job immediately, skip the job, or skip post application.
 - **Delay in Minutes**--Specify the delay that Arcserve Backup waits before running a job when the specified exit code is detected.
- **Run Command After Job**--Enter the path to, and name of, the file to be executed on the machine after the job is completed.
- **Do Not Run Command If**--Specify that a command will not run if Arcserve Backup detects the following events:
 - **Job Fails**--If a job fails, then the command will not run.
 - **Job is Incomplete**--If a job is not completed, then the command will not run.
 - **Job is Complete**--If a job is completed, then the command will run.
- **Run Before/After Command As**--Specify the User Name and Password that corresponds to that of the Local Host server selected, and is required to check the system privileges on that server. The user name and password entered into these fields should not be confused with the Arcserve Backup User Name and Password.

Restore Manager Job Log Options

Using this option, you can determine the level of detail that is included in the log report for the restore job. You can view the log report in the Job Queue or Database Manager window (Job View). The log options are:

- **Log all activity**--Record all of the activity that occurs while the job is running.
Note: When you specify Log all activity, Arcserve Backup creates a log file named JobLog_<Job ID>_<Job Name>.Log. With this log file, you can view detailed logging information about the job. Arcserve Backup stores the log files in the following directory:
`C:\Program Files\CA\ARCserve Backup\LOG`
- **Log summary only**--Record summary information of the job (including source, destination, session number, and totals) and errors.
- **Log disabled**--Do not record any information about this job in the job log.

Restore Manager Alert Options

You can use the Alert notification system to send messages about events that appear in the Activity Log during your restore operation. Choose one or more of the following events for which you want to be notified:

- **Job Completed Successfully**--All of the nodes and drives/shares were processed.
- **Job Incomplete**--Some nodes, drives, or shares were missed.
- **Job Canceled by User**--The user canceled the job.
- **Job Failed**--The job was started but could not be completed.
- **Customized Event**--A customized event occurred. To specify this type of event, enter an error, warning, or notification code in the space below the Event drop-box.

Choose one or more of the defined Alert configurations. The <default> configuration means that you will use whatever is configured in Alert Manager. Click Configure to define further configurations. Arcserve Backup provides the following defined Alert configurations:

- Broadcast
- Pager
- SMTP
- SNMP
- Event

- Printer
- E-Mail
- Lotus Notes

Select **Attach Job Log** to include the job log information in the Alert message. (This option applies for Trouble Tickets and Mail only.)

Note: The list you create using Alert Options is saved with the Job Script and the configuration defined using the Configuration button.

System State Restore Options

Right-click the system state session to access the restore option context menu. The following options are available:

- **Make the Restored Copy of the Active Directory Authoritative**--Forces the restored copy to become the "authoritative" version of Active Directory on the system. This means that, even if the restored replica set is older than the current replicas, the older data is replicated to all of its replication partners. Authoritative restore is typically used to restore a system to a previously known state.

Note: Servers running Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012 do not support placing the Active Directory in an authoritative mode.

- **When Restoring replicated data sets, mark the data as primary for all replicas**--Forces the restored File Replication service data to be replicated to other servers. If this option is not enabled, the replicated data sets may not be replicated to other servers because the restored data will appear to be older than the data on other servers.
- **Stop the Cluster if necessary to Restore the Cluster Database**--Gives permission to stop a cluster service to restore the cluster database. This applies only to cluster machines. If this option is not enabled and the cluster service is running, Arcserve Backup dumps the cluster database files into the%SYSTEMROOT%\clubkup folder, but does not load them. Arcserve Backup provides a utility program (caclurst.exe) that lets you load the cluster database files at a convenient time.
- **Enable Quorum Drive Selection when Quorum Location Changes (Prior to Windows Server 2008 Cluster)**--Lets you to set the drive of the quorum resource that a cluster currently uses. If a cluster system was reconfigured to use a different quorum drive since the last system state backup, use this option to provide the new quorum drive. Otherwise, the backup copy of the quorum drive will be used which will cause the cluster database restore to fail.
 - **Select the drive letter in the case the quorum location changed since this backup**--Lets you specify a drive letter to restore data to when the location of the quorum changed since the backup was performed.

- **Authoritative Restore Cluster Database (Windows Server 2008 Cluster or later)**--Lets you perform an authoritative restore on Windows Server 2008 or later clusters. An authoritative restore lets you to restore the cluster database across all nodes. You should enable this option when you want to roll back the cluster configuration to the previous version.

Use the following guidelines to determine when to process an authoritative restore or a non-authoritative restore:

- **Authoritative restore**--An authoritative restore lets you use the cluster configuration that is stored in the backup data, not the current cluster node configuration, to recover the node. An authoritative restore lets you allow the cluster to use the restored configuration as the most recent configuration. If you recover the node using an authoritative restore, the current cluster configuration replicates to all of the nodes in the cluster.
- **Non-authoritative restore**--A non-authoritative restore lets you use the backup data to recover disabled nodes. With a non-authoritative restore, the latest cluster configuration information replicates to the recovered node after it becomes functional and joins the cluster.

Be aware of the following behaviors and considerations:

- The Authoritative Restore Cluster Database (Windows Server 2008 Cluster or later) option can be applied at the node level.
- If the node that you want to restore is corrupt or disabled, you must perform a node restore before you perform an authoritative restore. To perform a node restore, do not enable this option.
- You must restart the node after you perform an authoritative or a non-authoritative restore.

Note: For information about recovering clusters from a disaster, see "Recovering Clusters" in the *Disaster Recovery Option Guide*.

- **Do not Stop World Wide Web Service**--Lets you continue the www service while the certificate server is being restored. The IIS Publishing Service may be using the certificate service dynamic files at the time of certificate server restore. For this reason, by default, WWW service will be stopped during certificate server restore. If you do not want it to stop, use this option.

Restoring Data Scenarios

The following sections describe how to restore data in specific scenarios.

- [Restore Data Backed Up Using Staging](#) (see page 280).
- [Restore a Remote Agent on a System without the Disaster Recovery Option](#) (see page 282).
- [Restore Arcserve Backup Member Servers without the Using the Disaster Recovery Option](#) (see page 283).
- [Best Practices - How to Recover a Stand-alone Server from a Disaster Using the Disaster Recovery Option](#) (see page 285).
- [Best Practices - How to Recover a Arcserve Backup Server from a Disaster Without Using the Disaster Recovery Option](#) (see page 287).

Restore Data that was Backed Up Using Staging

The process for restoring data that was backed up using staging is identical to the process of restoring data that was backed up to any other type of storage media. However, staging provides you with the option to restore data from the location that is most suitable to your needs.

When you perform backup operations using staging, and the backed up data has been copied to its final destination media, the data can reside in two locations (the staging device and its final destination media). If you need to perform a restore operation and the data resides in two locations, you can restore the data directly from the staging device. Restore operations from staging devices are usually faster than tape-based restores.

Follow these steps:

1. Open the Restore Manager and select the Restore by Tree method.
2. In the left pane of the Restore Manager, select the volume, drive, directory, or file you want to restore.

Arcserve Backup searches the database for all backed up versions of this file, directory, drive, or volume.

Note: When using disk or tape staging, ensure that the staging tape is not offlined without formatting or erasing the staging tape. This will allow you to view the session details from the destination (migration) tape.

3. Click the Recovery Points drop-down list to specify the date of the recovery point that you want to use for the restore job.

Next, specify the precise recovery point that you want to use for the restore job, as illustrated by the following example:

Backup Time	Method	Session St...	Media Name
7/24/13 4:43 PM	Full	Finished	STAGE.FSD
7/24/13 4:46 PM	Full	Finished	DEST.FSD
7/24/13 4:46 PM	Full	Finished	7/23/13 3:43 PM
7/24/13 4:39 PM	Full (Clear Archive Bit)	Finished	DDD
7/24/13 4:42 PM	Full (Clear Archive Bit)	Finished	DEST.FSD
7/24/13 4:42 PM	Full (Clear Archive Bit)	Finished	7/23/13 3:43 PM
7/24/13 4:27 PM	Full (Clear Archive Bit)	Finished	DDD
7/24/13 4:30 PM	Full (Clear Archive Bit)	Finished	DEST.FSD
7/24/13 4:30 PM	Full (Clear Archive Bit)	Finished	7/23/13 3:43 PM
7/24/13 4:22 PM	Full (Clear Archive Bit)	Finished	DDD
7/24/13 4:24 PM	Full (Clear Archive Bit)	Finished	DEST.FSD
7/24/13 4:24 PM	Full (Clear Archive Bit)	Finished	7/23/13 3:43 PM
7/24/13 4:12 PM	Full (Clear Archive Bit)	Finished	DDD
7/24/13 3:58 PM	Full	Finished	STAGE.FSD
7/24/13 3:28 PM	Full	Finished	STAGE.FSD
7/24/13 3:28 PM	Full	Finished	STAGE.FSD

4. From this list, select the version you want to restore.

Note: Duplicate backup sessions can exist when clones of the session reside on the multiple media which might have happened because of staging backup jobs or tape copies. When the Media Name is an file system device (FSD) or data deduplication device (DDD), restoring data from a disk is generally faster than restoring from a tape. When you restore data from a disk, there are no delays caused by tape load and seek latency. When you need to restore data that exists in two locations (disk and tape), you can reduce the restore time by restoring directly from the disk rather than retrieving it from a tape.

5. Click the Destination tab and specify the location where you want to restore the data.
6. Click the Schedule tab and specify a schedule for the job.
7. (Optional) Click Options on the toolbar to specify any other options that you want to apply to the job.

8. Click Submit on the toolbar to submit the job.

Complete the required fields on the Submit dialog and click OK to start the restore process.

Restore a Remote Agent on a System without the Disaster Recovery Option

This section describes how to restore a remote agent on a system without the Arcserve Backup Disaster Recovery Option.

Before proceeding, ensure that the following prerequisite tasks are complete:

- Verify that there is one full backup of the remote agent machine, and verify that the backup media is available.
- Record the disk partition/volume configuration, including all volume drive letters and volume mount points, when the system is up and running.
- Record the network configuration when the system is up and running.
- Ensure the operating system CD, the device drivers, and the Arcserve Backup installation media are available.

To restore a remote agent on a system without the Disaster Recovery Option

1. Start the computer you want to recover, using the Windows operating system CD.
2. Create partitions which are necessary for installing the operating system. Other disk partitions/volumes can be restored manually after the operating system is installed. For dynamic disk configuration, it must be restored after the operating system is installed.
3. Install the operating system and verify that the host name is the same as the original system.
4. Restore the remaining disk/volume configuration, disk partition layout, dynamic disk volumes, etc.
Note: The volume drive letter should be the same as the original system.
5. Install the device drivers which are not included on the operating system CD. This includes SCSI/RAID/FC drivers and network adapter drivers.
6. Configure the network and verify that all configurations are the same as the original system.
7. Apply the operating system patch. This is necessary when the system is going to be connected to the network.
8. Install the same antivirus software as backup time, and update to the latest patch. This is necessary when the system is going to be connected to the network.
9. Install Arcserve Backup Client Agent.

10. Add this machine to the source node list of the Arcserve backup server if it is not on the existing node list.
11. Select restore by tree in the Arcserve Backup Restore Manager and submit the restore job.

Restore Arcserve Backup Member Servers without the Using the Disaster Recovery Option

This section describes how to restore Arcserve Backup member servers without using the Disaster Recovery Option.

Important! This procedure does not apply to restoring Arcserve Backup primary servers and stand-alone servers.

Prerequisite Tasks:

Before proceeding, ensure that the following prerequisite tasks are complete:

- Ensure that there is at least one full backup of the system and the backup media is available.
- Record the disk partition/volume configuration, including all volume drive letters and volume mount points, when the system is up and running.
- Record the network configuration when the system is up and running.
- Ensure that the operating system CD, the device drivers, and the Arcserve Backup installation media are available.

To restore Arcserve Backup member server without using the Disaster Recovery Option

1. Start the computer you want to recover, using the Windows operating system CD.
2. Create the partitions that are necessary for installing the operating system. Other disk partitions/volumes can be restored manually after the operating system is installed. For dynamic disk configuration, it must be restored after the operating system is installed.
3. Install the operating system and verify that the host name is the same as the original system.
4. Restore the remaining disk/volume configuration, disk partition layout, dynamic disk volumes, and so on.

Note: The volume drive letter should be the same as the original system.

5. Install the device drivers which are not included on the operating system CD. This includes SCSI/RAID/FC drivers and network adapter drivers.
6. Configure the network and verify that all configurations are the same as the original system.

7. Apply the operating system patch.

Note: This step is necessary when the system is going to be connected to the network.

8. Install the same antivirus software that was running when the last backup completed and update to the latest patch.

Note: This step is necessary when the system is going to be connected to the network.

9. Install all applications same as original system.

10. Install Arcserve Backup, agents, and options into the same directories as the original installation.

11. Open the Restore Manager and click the Options toolbar button.

The Options dialog opens.

Click the Operation tab, select the Restore Registry File and Event Log option and click OK.

The restore options are applied.

12. From the Restore Manager, specify the Restore by tree restore method and submit the restore job to restore the system.

Restart the system after the restore job is complete.

If the system is not a domain controller, go to Step 17.

13. When restarting the system press the F8 key to start the Windows Server 2003 Expansion Option Menu.

14. When prompted, select Directory Service Restore Mode to start the system in the restore mode.

15. Restore the System State using the following options:

- Specify the System State as the source.
- From the Arcserve server, specify the Global Options for restoring data (see Step 10).

16. Restore the system.

17. After the job is complete, restart the system.
18. Based on one of the following scenarios, confirm that the system has been restored successfully after the system restarts.
 - If the operating system detects that the backed up registry information does not reflect the currently-used hard disk device, you may need to change the drive letter assignment. If this occurs, re-assign the proper drive letter.
 - If a file is required for the system, in addition to the system drive, you may need to restart the system after the drive letter is re-assigned. If other drives are data only, restarting the system is not necessary. If you are not sure of the types of data that reside on the other drives, you should restart the system after re-assigning the drive letter.

Best Practices - How to Recover a Stand-alone Server from a Disaster Using the Disaster Recovery Option

The following scenario describes how to leverage the Disaster Recovery Option to protect a Arcserve Backup server that is hosting SQL Server databases.

Server Specifications

The Arcserve Backup server is configured as follows:

- The Arcserve Stand-alone Server installation option is installed on the server.
- The Arcserve database is hosted using Microsoft SQL Server 2008 Express Edition.
- The stand-alone server is hosting various Microsoft SQL Server databases.

Software Specifications

The following applications are installed on the Arcserve Backup server:

- Microsoft Windows Server
- Microsoft SQL Server
- Microsoft SQL Server 2008 Express Edition as the Arcserve Backup database
- Arcserve Backup for Windows
- Arcserve Backup Client Agent for Windows
- Arcserve Backup Agent for Microsoft SQL Server
- Arcserve Backup Disaster Recovery Option

Note: Microsoft SQL Server and Microsoft SQL Server 2008 Express reside on the same node. The Arcserve Backup installation routine installed the Microsoft SQL Server 2008 Express application.

Use the following guidelines to recover a Arcserve Backup server that is hosting SQL Server databases from a disaster:

1. During the backup operation, ensure that you do not status the SQL Server and SQL Server 2008 Express instances as offline.
2. Perform a full backup of the machine. The backup should be successful.
3. Create the Disaster Recovery Boot Kit.

Note: For information about creating a Disaster Recovery Boot Kit, *Disaster Recovery Option Guide*.

4. Perform Disaster Recovery restore. The restore should be successful.

During the restore operation, the Disaster Recovery Option recovers the master and model databases in the Arcserve Backup database (Microsoft SQL 2008 Express).

5. When prompted, restart the Arcserve Backup server.

After you restart the Arcserve Backup server, the disaster recovery process continues recovering the Arcserve Backup database. After the Arcserve Backup database is recovered, you can start Arcserve Backup normally.

6. Restore the disaster recovery element sessions for the SQL Server database instances.

Note: For more information, see the *Agent for Microsoft SQL Server Guide*.

7. Restart the SQL instances.
8. Restore the data to each of the SQL instances.

Best Practices - How to Recover a Arcserve Backup Server from a Disaster without Using the Disaster Recovery Option

Arcserve Backup lets you perform a full disaster recovery of a Arcserve Backup server without installing the Disaster Recovery Option. To enable this capability, you must complete the steps described in Recover the Arcserve Backup server described below. If the Arcserve Backup server that you are recovering is a primary or stand-alone server, you must also complete the steps in Recover the Arcserve Backup Database below.

The procedure consists of the following tasks:

1. Perform a full backup and restore of the Arcserve Backup server.
2. Recover the Arcserve Backup server.
3. Recover the Arcserve Backup database.
4. Reactivate the existing Arcserve Backup database (optional).
5. Recover the Job Queue Session.

Important! You must restore the Job Queue to the Arcserve primary or stand-alone server. Do not restore the Job Queue to a member server of any Arcserve domain.

6. Recover the Active Directory.

Note: This task applies only to Arcserve Backup servers that function as a domain controller.

7. Confirm Arcserve Backup licenses, if needed.

Be aware of the following behavior:

- During and after the recovery process, you will encounter error messages in the system log and the Arcserve Backup log. These messages are normal under recovery circumstances and will not result in a loss of data or functionality problems.
- When you restore the Arcserve Backup server without using the disaster recovery option on to the same server, verify that there is enough space available on the System Reserved Volume. The System Reserved Volume is allocated when installing Windows Server 2012 with the default size of 350 MB. It is recommended to change the size to about 490 MB (be careful to keep it less than 500 MB). If you cannot increase the System Reserved Volume, recover the Arcserve Backup server using the Disaster Recovery Option or create a new System Reserved Volume.

When creating new System Reserved Volumes, keep in mind that you can create the System Reserved Volume only on the primary partition of an MBR (Master Boot Record) disk. When the system volume and the boot volume are combined, the BitLocker feature cannot be used to encrypt the volumes on your computer.

Follow these steps:

1. Create a new volume. For example: F:, on the same disk that contains the boot partition of about 490 MB (verify that the size of the volume is less than 500 MB).
2. With Windows Server 2012 installed on the C:\ drive, open the DOS command prompt (run as administrator) and type the following command:

```
bcdboot.exe C:\Windows /s F:.
```
3. Type DISKPART
4. From the DISKPART command prompt, type the following:

```
DISKPART> select volume F  
DISKPART > active
```
5. Reboot your machine and the new System Reserved Volume is created on the F:.

Note: To revert to your previous configuration, repeat these same steps using the original system volume by assigning it the drive letter, for example, D:.

To recover the Arcserve Backup server

Important! You must have performed at least one full backup of the Arcserve Backup server before you can recover the Arcserve Backup server.

1. Reinstall the operating system on the Arcserve Backup server.
Verify that the hard disk partitions, hardware, and the operating system (version, edition, and service pack) configurations are identical to the configurations that were backed up.
2. Reinstall Arcserve Backup, agents, and options into the same directories as the original installation.

3. After you install Arcserve Backup, open the Merge utility and merge the media used for the last full backup.
4. After the merge is complete, open the Restore Manager, and verify that the Restore files to their original location(s) option is selected.

Locate the full backup sessions.

Select the backup sessions for the machine, excluding all the following Arcserve Backup-specific sessions:

- Disaster Recovery session
- Arcserve Job Queue session
- Arcserve Database session
- SQL Server Disaster Recovery Elements session

Note: If the Arcserve Backup Catalog database session was selected during the restore, you must close the Arcserve Backup Manager Console after you submit the restore job (Arcserve Backup enables the catalog database by default). This approach lets the restore process overwrite the catalog database. You can reopen the Job Status Manager or Job Monitor to monitor the status of the job, however, you must not open the Restore Manager or Database Manager until the job is complete.

5. Click Options on the toolbar.

The Restore Manager, Options dialog opens.

Click the Operations tab, click Restore Registry Files and Event logs, and click OK.

The Options dialog closes.

6. Click Submit on the toolbar to submit the restore job.

The Session User Name and Password dialog opens.

7. On the Session User Name and Password dialog, complete the fields that follow, as required, and then click OK.

- **User Name**--Specifies the user name for the target Arcserve Backup server.

Note: You must complete this field on Windows Server 2003 64-bit, Windows Server 2008, and Windows Server 2012 systems.

- **Password**--Specifies the password for the target Arcserve Backup server.

Note: You must complete this field on Windows Server 2003 64-bit, Windows Server 2008, and Windows Server 2012 systems.

- **Session Password**--Specifies the password for encrypted backup sessions.

- **IP Address**--Specifies the IP address of the target Arcserve Backup server.

On the Session User Name and Password dialog, click Edit to modify the User Name, Password, and IP address for the selected session.

The Enter User Name and Password dialog opens.

8. On the Enter User Name and Password dialog, specify the User Name and Password for the Arcserve Backup server, click the Apply [User Name and Password] to All Rows check box to apply the user name and passwords specified to all sessions.

Note: When you are editing IP addresses and passwords, you must edit the individual IP address and session password for each individual session.

Click OK.

The Enter User Name and Password dialog closes.

9. Click OK to close the Session User Name and Password dialog.

Note: After you click OK, a message box labeled Arcserve Backup may open and prompt you to specify the IP addresses for the sessions that require authentication to submit the restore job. If the Arcserve Backup dialog opens, you must specify all the IP address for all sessions to submit the job, and then click OK.

The restore job is submitted.

10. When the restore job is complete, restart the computer.

Be aware of the following behavior:

- After you restart Windows Server 2012 systems, the Startup menu may not appear when you roll the cursor over the Startup menu. When this behavior occurs, log off and then log in to the computer to display the Startup menu.
- When you log in to the operating system, you may receive an error message asking why the computer was unexpectedly shut down. This is expected behavior caused by the System State Recovery. If needed, you should select the appropriate response from the drop-down list and continue.
- You may also see the message, "At least one service or driver failed during system startup. Use Event Viewer to examine the event log for details." In addition, the service SQL Server (ARCSERVE_DB) cannot start, thus resulting in the error. This is normal behavior because the SQL Server service does not recover during this step. This error will be corrected when you complete the steps in the section Recover the Arcserve Backup database.

11. After the computer restarts, do one of the following:
 - If the SQL Server service started and the Arcserve Backup database *is not* Microsoft SQL Server 2008, continue to the next task: **To recover the Arcserve Backup database.**
 - If the SQL Server service started and the Arcserve Backup database *is* Microsoft SQL Server 2008, do the following:
 - a. Open the Windows Services Manager.
Stop the Arcserve Database Engine service and the SQL Server service.
 - b. Start the SQL Server service in single user mode using `sqlservr.exe -m`.
 - c. Using SQL Server Management Studio, log in to the local SQL Server.
From the Object Explorer pane, drill down to <hostname>, [Security], and [Logins].
Delete the original Windows account.
Example: <hostname>\Administrator
Note: If a message appears that warns you about deleting this account, you can safely ignore the warning message.
Add a Windows account that you want to use to log in SQL Server.
Example: <hostname>\Administrator
Specify a default language for the database.
Specify [public] and [sysadmin] privileges to this account.
 - d. Stop the single user mode for the SQL Server service.
 - e. Open the Windows Services Manager and start the SQL Server service.
 - f. Continue to the next task: **To recover the Arcserve Backup database.**
 - If the SQL Server service did not start, continue to the next step.
12. Right-click the Data folder and select Properties from the pop-up menu.
The Properties dialog opens.
13. Click the Security tab and then click Advanced.
The Advanced Security Settings for Data opens.
Note: On Windows Server 2008 systems, click Edit on the Data Security Detailed Settings screen.

14. Click the Permissions tab and select the Replace permission entries on all child objects shown here that apply to child objects option and click OK.

Note: On Windows Server 2008 systems, click the Replace the existing inheritable permissions on all descendants with inheritable permissions from this object option.

If the Arcserve Backup Database is configured such that the database data files are stored in a different directory, repeat Steps 12, 13, and 14 on this folder to change its security attributes.

15. Open the Windows Services Manager and start the SQL Server service (ARCSERVE_DB).
16. Do one of the following:
 - If the SQL Server service is started, continue to the next task, **To recover the Arcserve Backup database.**
 - If the SQL Server service is not started and you cannot start the SQL Server service, continue to the next step.
17. Open the Windows Computer Management Console, click Local Users, and then click Groups.

The following group name should appear:

```
SQLServer2008MSSQLUser$MACHINENAME$ARCSERVE_DB
```

Note: The value of MACHINENAME should be the name of your computer.

18. Record this Group name.

Return to the Data folder (see Step 12).

Right-click the Data folder and select Properties from the pop-up menu.

The Properties dialog opens.
19. Click the Security tab and then click Add.

The Select Users, Computer, and Groups dialog opens.
20. Click Locations and then click Local Machine.

Add the Group that you recorded in step 14 and click OK.

The Select Users, Computer, and Groups dialog closes.
21. Click the Advanced tab and then click the Permissions tab.

Select the Replace permission entries on all child objects shown here that apply to child objects option and click OK.

Note: If the Arcserve Backup database is configured such that the database data files are stored in a different directory, repeat Steps 12 to 21 on this folder to change its security attributes.

22. Open the Windows Services Manager and start the SQL Server service (ARCSERVE_DB).

To recover the Arcserve Backup database

Important! Arcserve Backup will not be available until you recover the database. Error messages may appear in the Arcserve Backup Activity Log that you can ignore.

1. Open the Windows Services Manager and start the Arcserve Database Engine service.
2. Open the Restore Manager.

From the restore methods drop-down list, select Restore by Session.

Locate and select the Arcserve Backup Database session as the restore source.

Click the Destination tab and verify that the Restore files to their original location(s) option is selected.

Note: If the Arcserve Backup database you are recovering is stored in an independent local SQL Server instance, you should select the "master" database to restore before restoring the Arcserve Backup "asdb."

3. Open the Restore Options by clicking Options on the toolbar.

Click the Operation tab, select the Disable Database Recording option, and click OK.

The Options dialog closes.

4. Right-click the Arcserve Database session and select Agent Option from the pop-up menu.
5. Click the Restore Options tab, select the Force Restore over existing option, and click OK.

Note: If you do not choose this option, the restore job may fail, and the Database Engine will not start. For troubleshooting assistance, see the steps under the section To Reactivate the existing Arcserve Backup database.

6. Click Submit on the toolbar to submit the restore job.

Note: After you click Submit on the toolbar to submit the restore job, you must specify a user name and password on the DBAgent tab on the Session User Name and Password dialog.

During the restore process, the Database Engine service may pause or stop and the Manager Console may respond slowly. Because the Database Engine is unavailable during the recovery process, clients cannot connect to it. Therefore, Error E1516 [Staging] may be recorded in the Activity Log: "Cannot inquiry the database (Error=4294967293)." This behavior is normal during the database recovery process.

After the restore job completes successfully, the Database Engine will automatically resume and Arcserve Backup will return to normal operations.

To reactivate the existing Arcserve Backup database

This is an optional task. If the restore job fails because the correct options were not selected in the previous steps, the database may have been in an offline state while the restore job was in progress. As a result, the Database Engine could not access the Arcserve Backup database during the restore. The following steps describe how to reactivate the Arcserve Backup database.

1. Browse to the Arcserve Backup Home directory and locate asdbe_start.bat.
2. Execute asdbe_start.bat.

Note: The script uses a Microsoft SQL CLI utility "sqlcmd" to run a series of commands that will bring the Arcserve Backup database online.

After you execute the script, the Database Engine service resumes.

3. Repeat the steps in the section To recover the Arcserve Backup database.

Note: Verify that you specify the Disable Database Recording and Force Restore over existing options before you start the job.

To Recover the Job Queue Session

1. After the Arcserve Backup server starts, open the Restore Manager, locate, and select the Job Queue session.

Note: When you select this session, Arcserve Backup requires a merge job of the Job Queue session.

Click Yes to continue the recovery of the Job Queue session.

2. From the Destination tab, specify an alternate location to restore the Job Queue session.
3. Click Submit on the toolbar to submit the job to restore the Job Queue session to an alternate location.

Note: Verify that the alternate location is an empty directory.

4. After the Job Queue session is restored to the alternate location, open the Server Admin and do the following:
 - a. Locate the Arcserve Backup primary server or stand-alone server.
 - b. Right-click the Arcserve Backup server and select Stop all services from the pop-up menu.

All Arcserve Backup services stop.

5. Access the alternate location and copy all Job Queue files under the folder that you restored to the following directory: ARCSERVE_HOME\00000001.qsd

6. From the Server Admin, restart all Arcserve Backup services by doing the following:
 - a. Locate the Arcserve Backup primary server or stand-alone server.
 - b. Right-click the Arcserve Backup server and select Start all services from the pop-up menu.

All Arcserve Backup services start.

Note: The status of the backup job that you used for restoration is in a "crashed" state. When the Job Queue was backed up, that job was in an active state, but the corresponding process was not running. Therefore, that job is now in a crashed state and error message E1311 is written to the Activity Log: Job has crashed. This behavior is normal while the Job Queue is being restored.

7. If the Arcserve Backup is not a domain controller, go to the section, **Confirm Arcserve Backup product licenses**. Otherwise, continue to **Recover the Active Directory**.

Recover the Active Directory

1. Restart the system.

After the system restarts, press F8.

The Advanced Option Menu appears.

2. Select Directory Service Restore Mode and start the system in the Restore mode.

Start Arcserve Backup.

Error message E3073 occurs:

Unable to logon as user, user =Administrator,EC=Logon Failure or W3073 Unable to logon as user, user =Administrator,EC=Logon Failure

3. Open the Restore Manager and select the Source tab.

From the Restore methods drop-down list, select Restore by Session.

Locate and select the System State session.

Perform one of the following options:

For Windows Server 2003 systems:

- To perform a **Non-authoritative Restore**, go to Step 5.
- To perform an **Authoritative Restore**, right-click the System State session and select Local Options from the pop-up menu.

The System State Restore Options dialog opens.

Continue to Step 4.

For Windows Server 2008 or later systems:

- Go to Step 5.

4. On the System State Restore Options dialog, click Make the Restored Copy of the Active Directory Authoritative and click OK.

5. Click Options on the Toolbar.

The Global Options dialog opens.

6. Click the Operation tab.

Click Restore Registry Files and Event logs and click OK.

The Global Options are applied.

7. Perform one of the following options:

- **Windows Server 2003 systems**--Click Submit on the toolbar to submit the restore job. After the restore job is complete, restart the system.

Note: To restore the active directory data in Authoritative mode, Arcserve Backup executes NTDSUTIL.exe on the Arcserve Backup server. However, NTDSUTIL.exe executes asynchronously with the restore job and may not complete at the same time as the restore job. If this behavior occurs, restart the system after NTDSUTIL.exe completes. To ensure that NTDSUTIL.exe is complete, open Windows Task Manager, click Processes, and search for NTDSUTIL.exe. If NTDSUTIL.exe does not appear in the Windows Task Manager, NTDSUTIL.exe is complete and you can restart the system.

- **Windows Server 2008 systems or later**--Complete the steps that are described in [Restore Active Directory Objects](#) (see page 583).

To recover Authoritative Active Directory in Windows Server 2008 or later systems, perform the following steps:

- a. After the System State restore job completes, reboot the Active Directory server in Directory Service Restore mode.
- b. Open a Command prompt window and run the following command:
ntdsutil
- c. Run "activate instance ntds"
- d. Run "authoritative restore"
- e. Run "restore subtree dc=Domain_Name, dc=xxx"

The Active Directory is restored.

After the restore job is complete, restart the system.

Confirm Arcserve Backup product licenses

You should confirm the product license after full restore is complete. The current Arcserve Backup license is restored back to the original state when you perform a full backup. If you have applied new licenses after a full backup, or the licenses were dynamically assigned to other servers, it is possible to encounter license errors. You should register or adjust product licenses accordingly.

Error Messages

After you recover the Arcserve Backup server, you may discover errors, warnings, and failure audits in the system Event Log similar to the error messages listed below, depending on how your system is configured. Such messages are caused by the intermediate state of recovery, or are related to the startup order in which Arcserve Backup and the SQL Server services started.

Arcserve Backup may report the following errors during the intermediate state of recovery.

Error 8355

This error message is reported when a "service broker" disabled setting is detected in the MSDB that was recovered. You can safely ignore this error because this is the default behavior of a system database recovery that is limited to SQL Server 2008 Express Edition. SQL Server 2000 and SQL Server 2005 are not affected by this behavior.

You can suppress this error by doing the following:

1. Open a Windows Command Line window.
2. Connect to the Arcserve Backup database (ARCSERVE_DB) by executing the following sqlcmd:

```
SQLcmd -S <machine name>\<Instance name>
```

For example:

```
C:\Users\Administrator>sqlcmd -S localhost\ARCSERVE_DB
```

3. Confirm the value of service_broker on msdb is 0:

```
select name,is_broker_enabled from sys.databases  
go
```

4. Execute the command using the following arguments:

```
alter database msdb set enable_broker  
go
```

5. Confirm the value of service_broker on msdb is 1:

```
select name,is_broker_enabled from sys.databases  
go  
Quit
```

Close the Command Line window.

Arcserve Backup corrects the following error conditions after the recovery process is complete:

Error 615

This error message is reported when the 'master' database is restored during recovery without using the Arcserve Backup Disaster Recovery Option and SQL Server is hosting a single Arcserve Backup database. This error may not occur when SQL Server 2008 Express Edition hosts the Arcserve Backup database.

The Arcserve Backup Agent for Microsoft SQL Server will not back up the tempdb database even when a full instance is selected. Tempdb is also excluded by the Arcserve Backup server and file system agent when the normal file system is backed up. However, tempdb is recorded as an existing database in the SQL Server master database, so when the master is restored, the SQL Server service reports that it cannot find tempdb.

Error 15466

This error message is reported when the system state is restored during recovery without using Arcserve Backup Disaster Recovery Option. Either SQL Server 2008 Express Edition or SQL Server can host the Arcserve Backup database.

When you recover the Arcserve Backup server without using the Disaster Recovery Option, the Windows operating system and SQL Server are reinstalled, and the Service Master Key (SMK) of SQL Server is created. The Service Master Key (SMK) is used to encrypt all database master keys and all server-level secrets such as credential secrets or linked server login passwords.

The key is a 128-bit 3DES key. The SMK is encrypted using DPAPI and the service account credentials. When the system state is restored but SQL Server sessions have not yet been restored, the system state is overwritten by the restore operation. However, the SQL Server instance is not yet overwritten. The SMK is in the system state, so it has been recovered to the old one, which is therefore inconsistent with the SQL Server instance. At this time, the recovery procedure requires that you restart the operating system.

During the restart, SQL Server reads the SMK and checks it against the SQL database. Because the SMK and SQL Server database are inconsistent, the error occurs.

Error 17113

This error message is reported when the user permission settings for either the file or its containing folders are incorrect. Using the procedures described in this topic, you will adjust these permissions. The error will be corrected after you adjust the permissions.

Errors not related to the recovery process

If the Arcserve Backup services and SQL Server services do not start in the proper sequences, SQL Server error messages may appear in the system log. This behavior is a known issue. For more information, see the readme file.

How to Restore Data from Cloud Storage

Arcserve Backup lets you restore data from cloud devices to most computers attached to your Windows network.

You can restore the data from cloud devices using the following restore methods:

- Restore by Session

Within this view, you can manually delete a selected session.

- Restore by Tree

- Restore by Media

Note: For more information about restoring data, see "Restoring Data."

Chapter 5: Customizing Jobs

This section contains the following topics:

[Job Customization Methods](#) (see page 301)

[Rotation Schemes](#) (see page 307)

[How Job Filters Work](#) (see page 312)

[Schedule Custom Jobs](#) (see page 317)

[Custom Schedules](#) (see page 318)

[Tasks You Can Perform Using the Job Status Manager](#) (see page 319)

[How Save Agent and Save Node Information Works](#) (see page 333)

[How to Use the Job Scheduler Wizard to Schedule Jobs](#) (see page 341)

[Job Scripts](#) (see page 342)

[Job Templates](#) (see page 343)

[Windows-Powered NAS and Storage Server 2003 Device Configuration](#) (see page 346)

[Scan Jobs for Cloud Storage](#) (see page 349)

[Merge Jobs to Cloud Storage](#) (see page 349)

[How Arcserve Backup Purges Data from Cloud Storage](#) (see page 350)

Job Customization Methods

Arcserve Backup provides a number of methods to customize your jobs to suit your needs. This section discusses the following customization methods in further detail.

- **Rotation schedules** let you to define standard and consistent intervals at which to rotate and retire backup media.
- **Filters** let you select the files and directories to be included in (or excluded from) your backup and restore jobs based on a wide variety of criteria.
- **Scheduling options** provide you with the ability to schedule your jobs to run immediately, later, or on a regular basis.
- The **Job Scheduler Wizard** is a powerful tool that lets you quickly and easily schedule and submit any job that can be entered at the command line.
- The **Job Status Manager** is a graphical tool that helps you centrally manage Arcserve Backup servers enterprise-wide.
- **Job scripts** allow you to save the options, filters, and scheduling information you define for your job as a file, so you can re-use, copy, or efficiently resubmit jobs with these settings.
- **Job templates** let you use preconfigured settings to submit jobs on any machine running Arcserve Backup without having to repeat the set up detail tasks for each job. The job template copies your configured backup schedule settings to be used again in the future on any machine.

Dynamic Job Packaging

If you click the box next to an item and the box turns completely green, this item is packaged dynamically. Dynamic job packaging means that the content of your selection is determined when the job runs. For example, if you choose to back up a source group or a server, and the nodes or volumes listed on that source group or server change between the time you scheduled the job and the time the job runs, the nodes and volumes changed from the time the job actually runs are backed up.

When you dynamically select a parent item, all of its associated (or child) items are automatically backed up dynamically as well.


Exclude Objects from Dynamically Packaged Jobs

When you package job to back up data dynamically, you can exclude nodes or disks that you do not want to back up. For example, if you choose to dynamically back up a customized source group, where a server is part of another customized source group, you can exclude the server from being backed up from one of the customized source groups.

With the Exclude This Item option, Arcserve Backup lets you exclude only machines and disks; you cannot select individual folders to be excluded from a dynamically packaged job.

When you exclude a machine or disk from a dynamically packaged job, Arcserve Backup does not back up the data that resides on the excluded object. However, Arcserve Backup captures the information that it needs to recover the excluded object from a disaster.


You can apply the Exclude This Item option to all of the server objects that appear in the Source tree. For example, you can exclude objects from servers that appear under the VMware Systems object, the Arcserve D2D Servers object, and so on, as illustrated by the following screen:



Note: For more information, see [Dynamic Job Packaging](#) (see page 302).

To exclude items from dynamically packaged jobs

1. Open the Backup Manager and dynamically package a backup job. For example,



2. Right-click the machine or disk that you do not want to back up, and then click Exclude This Item on the pop-up menu.

(Excluded) appears next to the machine or disk name.

3. Go to the other tabs and finish creating the backup job.

Arcserve Backup will not back up the marked items when you execute the backup job.

Exclude the Arcserve Backup Database from Backup Jobs

Arcserve Backup backs up the Arcserve Backup database when you submit a full node backup of the primary server or stand-alone server. This is default behavior that is designed to help ensure that you protect the most current data in the Arcserve Backup database as you back up the primary server or stand-alone server. Optionally, you can exclude the Arcserve Backup database from the backup job using the following steps.

To exclude the Arcserve Backup database from backup jobs

1. Open the Backup Manager.
2. Click the Start tab and specify the backup type that you require for the job.
3. Click the Source tab.

From the view drop-down list, select Group View.

The Backup Manager groups the agents installed in your Arcserve Backup environment.

4. Expand the Microsoft SQL Server object.

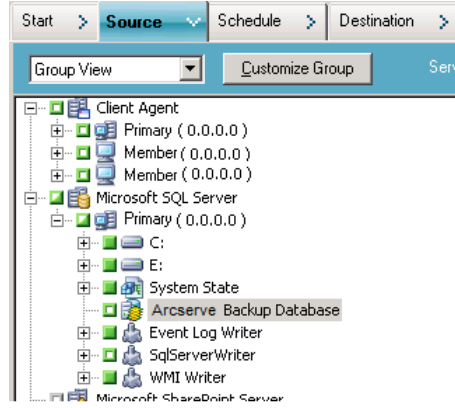
Locate the primary server or stand-alone server.

Click the check box next to the primary server or stand-alone server.

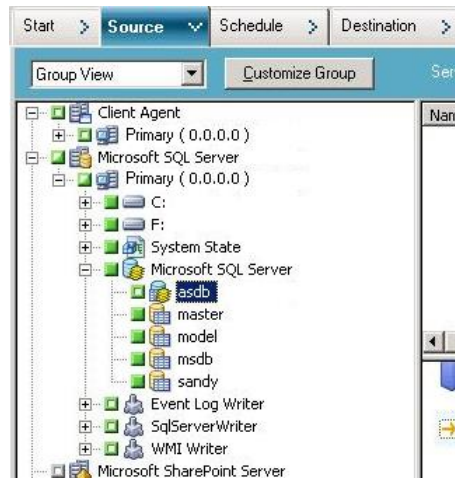
The Backup Manager selects all of the content in the backup server.

5. Do one of the following:

- **Microsoft SQL Server Express Edition databases**--Clear the check box next to Arcserve Backup Database.



- **Microsoft SQL Server databases**--Clear the check box next to ASDB.



The Arcserve Backup database is excluded from source for the backup job.

6. Click the Schedule tab and specify the schedule that you require for the job.
7. Click the Destination tab and specify the destination that you require for the backup data.
8. Click Options on the toolbar and specify additional options that you require for the job.
9. Click Submit on the toolbar to submit the job.

Arcserve Backup backs up the entire backup server and does not back up the Arcserve Backup database.

Static Job Packaging

If you click the box next to a child item and its parent box appears half green, the parent item is packaged statically. Static job packaging means that, in a parent object, you select only certain child items to include in your job. As a result, the content of what is packaged from the parent item is determined when you schedule the job rather than when the job runs.

When a parent item is explicitly selected, it applies only to its child items (the level that immediately follows). It does not apply to any selection you may make in the child items (for example, if you choose to back up only certain files within your child items).

Examples: How Static Backup Works

A computer contains drive c:\ and drive d:\ when you submit the job.

- Drive e:\ was appended to the computer after the job completed. The next time the job runs, Arcserve Backup backs up drive c:\ and drive d:\. Drive e:\ is not backed up.
- Directory c:\documents was appended to drive c:\ after the job completed. The next time the job runs, Arcserve Backup backs up all of drive c:\, including c:\documents, and drive d:\.
- Drive d:\ was deleted from the computer after the job completed. The next time the job runs, Arcserve Backup backs up drive c:\ and reports a failed backup of drive d:\.

A source group contains computers A, B, C, and D when you submit the job. Computer A contains drive c:\.

- Computer E was added to the source group after the job completed. The next time the job runs, Arcserve Backup backs up computers A, B, C, and D. Arcserve Backup does not back up computer E because it was not included in the original source group.
- Drive d:\ was appended to computer A after the job completed. The next time the job runs, Arcserve Backup backs up computers A, B, C, and D and drive d:\ in computer A. Arcserve Backup behaves in this manner because computer A was included in the original backup source group, and Arcserve Backup backs up the volumes in the source group dynamically.

Note: For more information, see [Submit Static Backup Jobs](#) (see page 198).

Create Static Backup Jobs

A static backup job backs up only those servers, nodes, and volumes in a source group or a server that you selected at the time of creating the job. If you add a server to a source group, or a node or volume to a server after creating a static backup job, it is not backed up when you run the job.

To create static backup jobs

1. Open the Backup Manager and go to the Source tab.
2. Right-click the source group or server that you want back up statically, and click Enable Static Backup on the pop-up menu.

(Static Backup) appears next to the group or machine name.

3. Go to the other tabs and finish creating the backup job.

When you run the job, only the nodes and volumes that you had selected at the time of creating the job are backed up.

Convert Jobs Submitted Using the Classic View to the Group View

Arcserve Backup lets you browse source data and submit jobs in two view formats:

- **Classic view**--Lets you browse source data and submit jobs based on the operating system that is running on the source computers. For example, Windows, UNIX/Linux, and so on.
- **Group view**--Lets you browse source data and submit jobs based on the Arcserve Backup agents that are running on the source computers. For example, Agent for Microsoft Exchange, Agent for Microsoft SQL Server, and so on.

Submitting jobs using the Group View is a convenient method of submitting jobs. The Group View lets you submit backup jobs that include specific agents. However, you cannot change the view that you specified when you submitted the job (Classic View to Group View and vice versa) after you submitted the job.

If you upgraded to this release, all of the jobs that you submitted using the previous release are packaged in the Classic View. The following steps describe how to convert jobs that were submitted using the Classic View to jobs that were submitted using the Group View.

To convert jobs submitted using the Classic View to Group View

1. Open the Job Status Manager and click the Job Queue tab.

The jobs in the Job Queue appear.

2. Locate the job that you want to convert.

Right-click the job and click Convert Job for Group View on the pop-up menu.

The Convert Job for Group View dialog opens.

3. Do one of the following:

- (Best practice) In the Group Name field, accept the default name provided or specify a new name for the group.

Click OK.

- From the Group Name drop-down list, select a name for the group and click OK.

The Duplicated Source Notification dialog opens only when both of the following conditions occur:

- One or more nodes in the job that you want to convert are included with full backup jobs that are configured as scheduled, repeating, or GFS rotation jobs.
- The group name specified, which is an existing group, is included with full backup jobs that are configured as scheduled, repeating, or GFS rotation jobs.

If you are sure that you want to use the group name specified, click Yes.

The job is converted.

Rotation Schemes

This section describes how to configure a rotation scheme for a backup job by using the Arcserve Backup default scheme or by specifying your own rotation parameters. To access the parameters for configuring a rotation scheme, select the Schedule tab in the Arcserve Backup Manager. The parameters that you can use are described below.

- **Scheme Name**--Select the type of rotation scheme you want, based on 5 or 7 days, and incremental, differential, or full backups. For more information on these standard schemes, see [Calendar View Tab](#) (see page 311) to modify your rotation scheme.
- **Start Date**--The date the backup will start.
- **Execution Time**--The time the backup will start.

- **Enable GFS**--Arcserve Backup allows you to select from pre-defined Grandfather-Father-Son (GFS) rotation schemes consisting of full weekly backup jobs combined with daily incremental and differential jobs. The GFS strategy is a method of maintaining backups on a daily, weekly, and monthly basis.

Accessible from the Backup Manager, the primary purpose of the GFS scheme is to suggest a minimum standard and consistent interval to rotate and retire the media. The daily backups are the Son. The last full backup in the week (the weekly backup) is the Father. The last full backup of the month (the monthly backup) is the Grandfather. GFS rotation schemes allow you to back up your servers for an entire year using a minimum of media.

GFS backup schemes are based on a five or seven-day weekly schedule beginning any day. A full backup is performed at least once a week. On all other days, full, partial, or no backups are performed. Using GFS rotation, you can restore data reliably for any day of the week by using the weekly full backup in conjunction with the daily incremental or differential backup jobs.

Note: A five-day GFS rotation scheme requires 21 media-per-year, while a seven-day scheme requires 23 media-per-year.

Although GFS rotation schemes are predefined, you can modify these schemes to suit your individual needs. You can deviate from your standard rotation scheme (for instance, if a holiday falls on Wednesday, your usual backup day).

- **Append Media**--If you specify the Enable GFS option, you can direct Arcserve Backup allow data from GFS rotation to append to existing jobs on the media.
- **Daily Backup Method**--The Daily Backup Method option lets you specify one of the following options for your daily backup jobs:
 - **Full**--All source files are backed up. This backup method clears the archive bit.
 - **Incremental**--Files that have changed since the last backup are backed up. This backup method clears the archive bit.
 - **Differential - Archive Bit**--Files that have changed since the last full backup job are backed up. This backup method does not change the archive bit.
- **Use WORM Media**--The Use WORM Media option lets you direct Arcserve Backup to use WORM media for all rotation rules. With this option enabled, you have the capability to use WORM media for daily, weekly, and monthly GFS backup jobs.

Important! Arcserve Backup does not support the use of WORM media for multiplexing and multistreaming backup jobs. As a result, when you enable the Multiplexing option or the Multistreaming option on the Destination tab of the Backup Manager, the Use WORM Media option is disabled.

More information:

[Calendar View Tab](#) (see page 311)

How You Can Manage GFS Rotation Jobs on File System Devices

Arcserve Backup supports using a GFS rotation scheme on File System Devices. A retention period for the media being used in the GFS rotation scheme can be determined using the following default retention cycle for a seven-day weekly rotation:

Frequency	Number of Media
Daily	6
Weekly	5
Monthly	12
Total	23

To run a rotation job beyond a year, a GFS rotation scheme requires 23 File System Devices to be created. These settings can be modified to meet your specific needs. Modifying the default values of the GFS rotation may change the number of FSDs required.

Note: Previously, only local disks were considered FSDs. You can now create FSDs that are accessible through a network share using a Universal Naming Convention (UNC) path.

Because a GFS rotation job may be using local disk drives and drive arrays, users must first make sure that there is enough space on the particular file system to store all the data being backed up for the entire retention period. Creating file system devices on a boot partition is not recommended because a boot disk that becomes full can cause the operating system to function abnormally.

Note: All file system devices need to be assigned to the same device group.

A configured GFS rotation job can run on a daily basis at a specified time. Arcserve Backup utilizes file system devices similar to a physical tape. As needed on a daily basis, Arcserve Backup moves tapes between the save sets and the scratch sets in the media pools, formats blank media, overwrites expired media, and tracks all operations in the database.

You can choose to duplicate backup data stored on the file system devices to physical tape media. The Job Scheduler Wizard and the TapeCopy utility provide the ability to automate the creation of the duplicate images.

The following sections describe the tabs available to customize your rotation job.

Set Up GFS Rotation Schemes

This section describes how to set up GFS rotation schemes using the Backup Manager. For detailed information about the functions, capabilities, options, and tasks that you can perform using the Backup Manager, see "Backing Up Data."

To set up GFS rotation schemes

1. From the Backup Manager, select a source and destination, and click the Schedule tab.
2. Enable the Use Rotation Scheme option. From the Scheme Name drop-down menu, choose one of the backup schemes.

Note: The Enable GFS option is automatically checked when a GFS scheme is selected.

3. If you want to add the data from one incremental or differential backup session to the same media as the previous backup session, enable the Append Media option.
4. In the Media Pool Name Prefix field, enter the prefix for all of your media names.

Arcserve Backup automatically creates and names your backup media using the name you designate.

Note: Arcserve Backup prevents you from using the underscore character (_) and the hyphen character (-) when specifying Media Pool names.

5. Specify the Start Date and the Execute Time as required.
6. Click Submit to submit the job.

The backup job runs precisely as you specified, and your media is recycled as you determined.

Rotation Rules Tab

You can modify the backup method or execution time for each day of the week.

Day of Week	Media Name	Method	Exec. Time	Staging
Sunday	<Auto Naming>	Off		
Monday	<Auto Naming>	Incremental	<Default>	Enabled
Tuesday	<Auto Naming>	Incremental	<Default>	Enabled
Wednesday	<Auto Naming>	Incremental	<Default>	Enabled
Thursday	<Auto Naming>	Incremental	<Default>	Enabled
Friday	<Auto Naming>	Full	<Default>	Enabled
Saturday	<Auto Naming>	Off		

Calendar View Tab

You can customize individual days. With GFS rotation either enabled or disabled, you can use the Calendar View feature to customize your rotation scheme according to the types of backups you want for particular days of the week or month, based on the calendar.

Rotation Rules Calendar View Exceptions Media							
Print...	Preview...	January 2010					< >
Sun	Mon	Tue	Wed	Thu	Fri	Sat	
					1	2	
3		4 Full F-<POOL>-TUE-01/05/10 <Default>	5 Incremental F-<POOL>-WED-01/06/10 <Default>	6 Incremental F-<POOL>-THU-01/07/10 <Default>	7 Full W-<POOL>-FRI-01/08/10 <Default>	8 Off	9
Off	10 Incremental F-<POOL>-MON-01/11/10 <Default>	11 Incremental F-<POOL>-TUE-01/12/10 <Default>	12 Incremental F-<POOL>-WED-01/13/10 <Default>	13 Incremental F-<POOL>-THU-01/14/10 <Default>	14 Full W-<POOL>-FRI-01/15/10 <Default>	15 Off	16
Off	17 Incremental F-<POOL>-MON-01/18/10 <Default>	18 Incremental F-<POOL>-TUE-01/19/10 <Default>	19 Incremental F-<POOL>-WED-01/20/10 <Default>	20 Incremental F-<POOL>-THU-01/21/10 <Default>	21 Full W-<POOL>-FRI-01/22/10 <Default>	22 Off	23
Off	24 Incremental F-<POOL>-MON-01/25/10 <Default>	25 Incremental F-<POOL>-TUE-01/26/10 <Default>	26 Incremental F-<POOL>-WED-01/27/10 <Default>	27 Incremental F-<POOL>-THU-01/28/10 <Default>	28 Full M-<POOL>-FRI-01/29/10 <Default>	29 Off	30
Off	31						


Note: This feature enables you to specify exceptions to the standard rotation scheme you are using.

Exceptions Tab

Define particular days on which the backup method and the execution time or date differs from the pre-existing schemes.

Media Tab

View information about the media pool you selected, including name, base serial number, next serial number, serial number range, minimum number of media, retention time, and prune retention time. You can also click the Daily, Weekly, or Monthly boxes to change the number of required media per year.



The screenshot shows a dialog box with four tabs: "Rotation Rules", "Calendar View", "Exceptions", and "Media". The "Media" tab is selected. Below the tabs is a section titled "Media Pool Information" containing several input fields:

- Pool Name: A dropdown menu.
- Base Serial Number: A text input field.
- Next Serial Number: A text input field.
- Serial Number Range: A text input field.
- Min. # of Media: A text input field.
- Retention Time (days): A text input field.
- Prune Retention Time (days): A text input field.

Media Pool Specification

Specify a (non-shared) media pool to the rotation scheme. If necessary, you can append data to media and change the media name.

Backup Method Options

A combination of three different backup methods is available: full, differential, and incremental. For detailed information about each of these methods, see [Custom Schedules](#) (see page 318).

How Job Filters Work

Filters allow you to include or exclude files and directories from your backup and restore jobs, as well as from the utilities, such as Copy, Count, and Purge.

For backup jobs, filtering can be performed on a per node basis. This means you can include a directory from one node and exclude the same directory from another node. A backup job can have node-level (local) and job-level (global) filters for the same job. Node-level filters apply to one specific node, not the entire job. If you want to add a filter that applies to the entire job, use a job-level, or global, filter. If you specify local (node-level) filters and global (job-level) filters for a backup job, Arcserve Backup applies the local filters and disregards the global filters.

You can include or exclude files based on the following criteria:

- Specific file names, patterns, attributes, and size.
- Specific directory names or patterns.
- Files accessed, modified, and created before, after, between, or within a specific date range.

Arcserve Backup uses wildcards or substitute characters, except when it detects that an absolute path is specified. If a valid absolute path is specified, Arcserve Backup will only exclude (or include) the absolute path specified, rather than excluding (or including) more directories, as it would for regular expression.

The wildcard characters supported for job filters based on file name or directory name are as follows:

- "*" --Use the asterisk to substitute zero or more characters in a file or directory name.
- "?" --Use the question mark to substitute a single character in a file or directory name.

Important! Exercise caution when specifying filters for your backup or restore operation. Incorrectly applied filters may not back up or restore the data you need, and can result in lost data and wasted time.

Examples: Back Up Data Using Wildcards

The following table describes examples of how you can use wildcards in conjunction with filters to back up data.

Note: The following examples assume that the source data resides in drive C:\.

Filter	Type and Criteria	Results
File	Include *.doc	Arcserve Backup backs up all files residing in drive C:\ that contain a .doc file extension.
File	Exclude *.doc	Arcserve Backup backs up all files residing in drive C:\ that do not contain a .doc file extension.
File	Include ?.doc	Arcserve Backup backs up files that contain a single character file name and a .doc file extension. For example, a.doc, b.doc, 1.doc, 2.doc, and so on.

Filter	Type and Criteria	Results
File	Include C:\myFolder\CA*.exe and Include C:\test\ms*.dll	Arcserve Backup backs up all of the following files: <ul style="list-style-type: none"> Files residing in C:\myFolder that start with CA and contain an .exe file extension. Files residing in C:\test that start with ms and contain a .dll file extension.
File	Exclude/Include C:\DOC\C*	Arcserve Backup restores all the files backed up in the 'C:\DOC\' folder except the files starting with 'C.'
Directory	Include m*t	Arcserve Backup backs up all directories residing in drive C:\ with directory names that start with m and end with t.
Directory	Exclude win*	Arcserve Backup backs up all directories residing in drive C:\ except directories that start with win.
Directory	Exclude C:\test\m* and Include C:\test\media*.gif	Arcserve Backup does not back up data. In this example, the exclude filter criteria directs Arcserve Backup to exclude all directories that reside in C:\test that start with m. As such, C:\test\media is excluded from the backup. Although the include filter directs Arcserve Backup to back up all files that reside in C:\test\media and contain a .gif file extension, Arcserve Backup will not back up any files because C:\test\media is excluded from the backup. Note: When you combine include filters with exclude filters, Arcserve Backup filters data based on exclude criteria first, and then by include criteria.

Filter Options

You can access the filter options from the Backup Manager, Restore Manager, Copy, Count, Scan, Compare, and Purge Utility windows.

- Exclude filters**--Exclusions always take precedence over inclusions. For example, if you add a filter to include files that have an .exe extension, and you add another filter to exclude your \SYSTEM directory, all .exe files in the \SYSTEM directory are excluded.

- **Include filters**--Results contain only those files that satisfy the filter specifications. For example, suppose you selected to back up your entire local hard drive, and you then set up a filter to include files in the \SYSTEM directory. The result would be that Arcserve Backup would only back up files from your \SYSTEM directory. No other files would be backed up.

Types of Filters

Filters are available which enable you to include and exclude files to suit your needs. The following list describes the types of filters that you can use to filter data.

- **File Pattern Filter**--Use the File Pattern filter to include or exclude files from a job. You can specify a particular file name or you can use wildcards to specify a file pattern.
Note: Wildcards "*" (asterisk) and "?" (question mark) can be used in the file pattern filter.
- **Directory Pattern Filter**--Use the Directory filter to include or exclude specific directories from a job. You can enter an entire directory name or you can use wildcards to specify a directory pattern.
Note: Wildcards "*" (asterisk) and "?" (question mark) can be used in the directory pattern filter
- **File Attributes Filter**--Use the File Attributes filter to include or exclude specific types of files from a job. Select as many of the following types of file attributes as you want:
 - **Hidden**--Files not shown in a directory listing. For example, IO.SYS is a hidden file.
 - **System**--Files that are unique to the machine you are using.
 - **Archive**--Files whose archive bit is set.
 - **Read Only**--Files that cannot be modified.
- **File Modified Filters**--Use the files last modified attribute to include or exclude files, based on the time they were last changed. There are four options from which to choose:
 - **Before**--Files whose date matches, or whose date is earlier than this date, are included or excluded.
 - **After**--Files whose date matches, or whose date is later than this date, are included or excluded.
 - **Between**--Files whose date falls between the two dates are included or excluded from the job. You must specify two dates for this selection.
 - **Within**--Files whose date falls within the specified time are included or excluded from the job. You must specify the number of days, months, or years.

- **File Created Filters**--Use the files last created attribute to include or exclude files based on when they were created. There are four options from which to choose:
 - **Before**--Files whose date matches, or whose date is earlier than, this date is included or excluded.
 - **After**--Files whose date matches, or whose date is later than, this date is included or excluded.
 - **Between**--Files whose date falls between the two dates are included or excluded from the job. You must specify two dates for this selection.
 - **Within**--Files whose date falls within the specified time are included or excluded from the job. You must specify the number of days, months, or years.
- **File Accessed Filters**--Use the file last accessed attribute to include or exclude files based on when they were last accessed. There are four options from which to choose:
 - **Before**--Files whose date matches, or whose date is earlier than, this date is included or excluded.
 - **After**--Files whose date matches, or whose date is later than, this date is included or excluded.
 - **Between**--Files whose date falls between the two dates are included or excluded from the job. You must specify two dates for this selection.
 - **Within**--Files whose date falls within the specified time are included or excluded from the job. You must specify the number of days, months, or years.
- **File Size Filters**--Use the file size attribute to include or exclude files based on the specific size of the file. You can specify a size range from 0 to 9999999999 bytes, KB, MB, or GB. There are four options from which to choose:
 - **Equal to**--Files whose size matches the size range are included or excluded from the job.
 - **Greater than**--Files whose size matches or whose size is greater than the size range are included or excluded from the job.
 - **Less than**--Files whose size matches or whose size is less than the size range are included or excluded from the job.
 - **Between**--Files whose size falls between the two sizes are included or excluded from the job.

Effective with this release, two new types of filters have been added:

- **Node name pattern**-- Only nodes whose name meets the pattern are included for backup. You can specify all or part of a search string using wildcards.
- **Node subnet pattern**--Only the node whose IP address is in the subnet is backed up.

Schedule Custom Jobs

All jobs can be scheduled using the Schedule options available in each Manager. Jobs can be submitted with a repeat method. For information about repeat methods, see [Rotation Schemes](#) (see page 307) or [Custom Schedules](#) (see page 318).

If you select the Run Job Now option when your storage device is busy, Arcserve Backup reports that the storage device is busy and the job is not submitted to the Job Queue. You should schedule your job, keeping the current date and time. This way, when Arcserve Backup discovers that the storage device is busy, it automatically retries the job until the drive becomes free.

You should select the Run Job Now option when:

- The job you are submitting is a one time only job that you want executed immediately.
- You want to monitor the job as it runs.

You should schedule your job when:

- You are submitting a single occurrence job but and you want it to run at a specific time.
- You are submitting a single occurrence job, but you do not want to run it now. You want to submit the job on Hold, and start it manually at a later time.
- You are submitting a job that should run regularly. This is especially useful for setting up a media rotation scheme for your network.
- You are submitting relatively low volume, daily, full backups. With custom schedules, Arcserve Backup lets you specify the Append days option. The Append days option helps to maximize the amount of data stored on tape and minimize the number tapes that need to be stored, relative to a job.
- Your storage device is busy and you want to run a backup job as soon as the drive is free. To do this, schedule your backup job with the current date and time.

Important! All scheduled times for Arcserve Backup jobs are based upon the time zone where the Arcserve Backup server is located. If your agent machine is located in a different time zone than the Arcserve Backup server, you will need to calculate the equivalent local time that you want the job to start.

Custom Schedules

You can select a custom schedule on the Schedule tab in the Backup Manager. A custom schedule enables you to run a backup job either once or on a repeating basis. You can specify the following parameters for backup and restore jobs:

- **Repeat Method**--All jobs can be scheduled using the Schedule options available in each Manager. Jobs can be submitted with a repeat method of:
 - **Once**--Do not repeat this job.
Note: This option is not available for synthetic full backup jobs.
 - **Every n frequency**--Repeat this job every specified number of Minutes, Hours, Days, Weeks, or Months.
Note: This option is not available for synthetic full backup jobs.
 - **Day(s) of the Week**--Repeat this job on the days that are checked off.
 - **Week(s) of the Month**--Repeat this job on the weeks that are checked off.
Note: This method is not available for synthetic full backup jobs.
 - **Day of the Month**--Repeat this job on the specified day.
Note: This method is not available for synthetic full backup jobs.
 - **Custom**--Repeat this at the specified interval, but exclude the days that are checked.
Note: This method is not available for synthetic full backup jobs.
- **Append days option**--The Append days option lets you append backups to one tape for the number of days that you specify. This option helps to maximize the amount of data stored on tape, and minimize the number tapes that need to be stored relative to a job. As a best practice, you should use this option when you perform relatively small volume, daily full backups. However, you can apply this option to all custom backup schedules except for run Once backup jobs and grandfather-father-son rotation jobs.
 - **Default value:** 6 days
 - **Range:** 1 to 3650 days**Note:** The Append days option cannot be applied to jobs where the destination device is data deduplication devices or cloud devices. In addition, with the Append days option applied to the backup, Arcserve Backup ignores the First Backup Media (global options) applied to the job. For more information about First Backup Media option, see [Backup Manager Backup Media Options](#) (see page 164).

You can specify the following parameters for a backup job:

- **Backup Method**--This specifies what data will be backed up. Jobs can be submitted with a backup method of:
 - **Full (Keep Archive Bit)**--Performed each time the job is repeated and *keeps* the archive bit.
 - **Full (Clear Archive Bit)**--Performed each time the job is repeated and *clears* the archive bit.
 - **Incremental backup**--Backs up only those files whose archive bit have been set since the last full or incremental backup was performed. After each backup, archive bits are reset so that they are not backed up during the next incremental backup job.
 - **Differential backup**--Backs up only those files whose archive bits have been set since the last full backup was performed. Because differential backup jobs do not clear a file's archive bit, the files that were backed up in the last differential job are backed up again. Using this backup method, the backup jobs require more time to process than incremental backup jobs. However, this strategy requires less effort to restore servers and workstations because you will probably require fewer media to restore your machines.

Note: For a synthetic full backup job, if the agent supports synthetic full backups, the differential backup job is converted to an incremental backup job.
- **Use WORM Media**--Directs Arcserve Backup to use WORM media when the backup job runs.

Tasks You Can Perform Using the Job Status Manager

The Job Status Manager is a graphical tool that helps you centrally manage Arcserve Backup servers across the enterprise.

You can use the Job Status Manager to:

- View all available Arcserve Backup servers, job history, job queues, and activity logs.
- Manage jobs--stop, add, run, delete, modify, reschedule jobs, and place jobs in a hold status.
- Monitor the progress of active jobs that are running on Arcserve Backup primary servers and member servers. You can view the real-time status of active jobs in the queue.
- View job detail and job log information about all the jobs that have been executed.

- View activity and media logs.
- Stop an active job.
- Modify user names and passwords associated with a job.
- Preflight Check the job.

Pop-up menus enable you to perform various operations with the Job Status Manager. These menus appear in both the left (server browser) and right (Job Queue) panels. To access a menu, right-click a selected item. When accessing a pop-up menu in the browser, the pointer must be on a selected group, server, or object. When accessing a pop-up menu in the Job Queue, a job must be selected.

Note: When you submit a job that spawns child jobs, the Job Queue tab displays details about the master job only. The Activity Log tab displays details about the master and child jobs. The Activity Log presents you with a description for the job.

Important! When you are executing a multistreaming, multiplexing, or disk staging job, the number of child jobs associated with a master job will never exceed the number streams specified for the job. However, if a job spawns multiple child jobs and the value specified for the Multiplexing Max # of Streams option is one, the child jobs will be created and backed up in one continuous stream (the default Max # Stream is four).

More information:

[Preflight Checks for Your Backups](#) (see page 126)

Modify Pending Data Migration Jobs

Migration is the process of moving backup data from a temporary staging location (device or media) to final destination media.

Note: With Disk Staging device as the staging location, you can perform a Double Tape Copy where the data can migrate to two final destinations simultaneously. See [Backup Manager Destination Options](#) (see page 149) for more details.

A migration job is the Arcserve Backup task associated with migrating data, or copying data, from the staging location to the final destination media. The parameters for the migration job, such as the schedule, copy policies, and so on, are defined by the staging policies that you specified when you submitted the job.

To help you manage pending migration jobs, Arcserve Backup provides you with a tool called the Migration Job Status dialog. The Migration Job Status dialog displays a list of all backup sessions for a job in the Job Queue that are pending data migration. With the Migration Job Status tool you can reschedule migration and purge jobs, cancel migration jobs, and change the device group that contains the final destination media.

To modify pending data migration jobs

1. Open the Job Status Manager and select the Job Queue tab.

Job Name	Backup Se...	Job No.	Job ID	Status	Execution T...	Job Type	Last Result	MB Processed	Elapsed Time	MB/Minute
123	LIY002...	5	3	DONE	<Run Now>	Backup	Finished			
Data Migration Job	LIY002...	14	14	DONE	1/22/2009 ...	Migration	Finished			
Data Migration Status	LIY002...									
221	LIY002...	7	15	DONE	1/22/2009 ...	Backup	Finished			
Data Migration Job	LIY002...	12	11	DONE	1/22/2009 ...	Migration	Finished			
Data Migration Job	LIY002...	16	17	DONE	1/22/2009 ...	Migration	Finished			
Data Migration Status	LIY002...									
333	LIY002...	9	9	DONE	<Run Now>	Backup	Finished			
Data Migration Job	LIY002...	13	12	DONE	1/22/2009 ...	Migration	Finished			
Data Migration Status	LIY002...									
Backup [Custom, Staging]	LIY002...	20	21	ACTIVE	Backup files...	Backup		0.00	3s	1:25
Data Migration Status	LIY002...									
Backup [Custom]	LIY002...	3	1	DONE	<Run Now>	Backup	Finished			
Database protection job	LIY002...	2		HOLD	1/22/2009 ...	Backup (Ro...				
Database pruning job	LIY002...	1	13	READY	1/23/2009 ...	DB Pruning	Finished			
eee	LIY002...	17	18	DONE	<Run Now>	Backup	Finished			

Locate the jobs with pending data migration jobs and click Data Migration Status.



The Migration Job Status <Backup Server Name> dialog opens as illustrated by the following screen.

No.	Backup Time	Copy Time	Purge Time	Num. of Sessions	Data Size(MB)	Status	Group Name	Tape Name
1	11/08/12 10:50 AM	11/08/12 10:59 AM	11/22/12 10:54 AM	1	0.61	Failed	DISKDEST	THURSDA
2	11/09/12 10:50 AM	11/09/12 10:59 AM	11/23/12 10:54 AM	1	0.61	Failed	DISKDEST	FRIDAY
3	11/10/12 10:50 AM	11/10/12 10:59 AM	11/24/12 10:54 AM	1	0.61	Pending	DISKDEST	SATURDA
4	11/11/12 10:50 AM	11/11/12 10:59 AM	12/09/12 10:54 AM	1	163.28	Pending	DISKDEST	SUNDAY
5	11/12/12 10:50 AM	11/12/12 10:59 AM	11/26/12 10:54 AM	1	0.61	Scheduled	DISKDEST	MONDAY

Note: The list doesn't display those sessions which have already been copied or were not supposed to be copied.
 ** A Scheduled status indicates that the sessions have not been copied because the scheduled start times has not passed.
 *** A Pending status indicates that the sessions have not been copied although the scheduled start time has passed.
 **** A Failed status indicates that the sessions are failed to be copied, you can modify the copy time to run the migration job again.

2. Locate and click the sessions that you want to modify and click Modify.

The Migration Job Configuration dialog opens.

3. Specify the options that follow:

- **Copy Time**--Lets you specify the date and time that you want to start the migration job.

Note: If you do not want to migrate the backup data to final destination media, clear the check box next to Do not copy data.

- **Purge Time**--Lets you specify the date and time that you want to purge the backup data from the staging device.
- **Target**--Lets you specify the group containing the final destination media.
- **Apply to the selected rows**--If you selected more than one session on the Migration Job Status dialog, this option lets you apply the Copy Time, Purge Time, and Target options that you specified to all of the sessions selected on the Migration Job Status dialog.

4. Click OK.

Note: You can also view the pending migration jobs associated with a staging location by completing the following steps:

1. Open the Backup Manager and select Enable Staging on the Start tab.
2. Select the Staging Location tab.
3. Right click a Deduplication group or device or a Disk Staging group or device.
4. Click the Data Migration Status on the pop-up menu to open the Migration Job Status dialog.

The dialog lists all pending migration jobs associated with the selected staging group or device. For example, the dialog lists pending migration jobs that will migrate data from the selected staging group or device.

More information:

[Back Up Data Using Tape Staging](#) (see page 236)

[How You Can Reclaim Disk Space](#) (see page 229)

Update Multiple Jobs

In your Arcserve Backup environment, you can have several to many jobs listed in the Job Queue. If a situation arises where you need to change the READY or HOLD status on more than one job, you can update multiple jobs simultaneously. The updates that you can perform simultaneously include changing the job status from HOLD to READY, from READY to HOLD, and delete the job.

Note: When you select a job that contains child jobs, Arcserve Backup applies the update to the parent job and all of its child jobs.

To update multiple jobs

1. From the Job Status Manager, select the Job Queue tab.
2. Click to select the job that you want to update.
 - To select multiple adjacent jobs, press and hold the Shift key as you select the jobs.
 - To select multiple non-adjacent jobs, press and hold the Ctrl key as you select the jobs.
3. Right-click the selected jobs.
4. From the pop-up menu, select HOLD, READY, or Delete Job as warranted by the situation.

How to Manage Jobs Using the Job Queue Tab

The Job Queue tab on the right panel displays information about all jobs. Every time you run or schedule a job with the Arcserve Backup Manager, you submit it to the Job Queue. Arcserve Backup continuously scans the Job Queue for jobs that are waiting to execute. Select a job and right-click for the following options:

- **READY/HOLD**--Changes the job's status to HOLD or to READY (if it is currently on hold). HOLD signifies that the job is not scheduled to be executed, while READY means that the job can be executed.
- **Add Job**--You can quickly submit a job to the queue by using a previously saved script. (A script is a job that you saved to a file. It contains the original source, destination, option, and schedule information for the job.)
- **Modify Job**--Modifies a job. Allows you to add options or additional sources to an existing job, without having to create a new job.
- **Reschedule Job**--Quickly change a job's execution date, time, or status. It also allows you to resubmit a Done job that is still in the Job Queue. You may want to do this if your job was not successful when it first ran.

- **Run Now**--Available for jobs that have a Ready or Done status. This option is useful in the following scenarios:
 - You want to run a job earlier than the time it is scheduled to run
 - A scheduled job did not run because of a hardware problem and you want to run it immediately after the problem is fixed

If a device group is available, this option runs the job immediately. If you select Run Now and a device group is not available, the job stays in the queue and waits for a group to become available.

If you select the Run Now option for a repeating, rotation, or GFS rotation job, the following conditions apply:

- The job runs immediately and the existing schedule is not affected unless the time it takes to run the job overlaps with the next scheduled run. In this scenario, the scheduled run is skipped for that day. For example, if you have a job scheduled to run Monday through Friday at 9:00 p.m., you select Run Now at 6:00 p.m. and it does not finish till 10:00 p.m., the 9:00 p.m. scheduled run for that day is skipped.
- The backup method used for the job is the same backup method that will be used for the scheduled run that day. For example, if you have an incremental backup job scheduled for 9:00 p.m. and select Run Now at 6:00 p.m., the job that runs at 6:00 p.m. will be an incremental backup. If you select Run Now on a day that does **not** have a scheduled run, the backup method of the next scheduled job will be used. For example, if you have an incremental job scheduled to run Monday and you select Run Now on Saturday, the job that runs on Saturday will be an incremental backup.

- **Stop Job**--Cancels an active job from the Arcserve Backup queue and reschedules it for its next regular interval.

Note: If you stop a job, the Last Result field displays “Canceled.”

- **Delete Job**--Cancels the job and deletes it from the Arcserve Backup queue completely.

You cannot use the Delete Job option on an active job. Use the Stop Job option if you want to delete an active job that repeats at intervals (determined when you create the job). Selecting the Delete Job button will interrupt and remove the job completely from the queue, and it will not be rescheduled. You will have to recreate the job if you did not save it to a script file.

Note: Arcserve Backup lets you re-create the Arcserve Backup database protection job and the database pruning job in the event they are deleted intentionally or unintentionally. For more information, see [Re-create the Arcserve Backup Database Protection Job](#) (see page 625) and [Re-create the Arcserve Backup Database Pruning Job](#) (see page 328).

- **Modify User Name**--Modify the user name and password for server and source nodes.

- **Modify Encryption Password**--Modify the encryption password that was previously specified for a job.
- **Preflight Check**--Runs vital checks on the Arcserve Backup server and agents to detect conditions that may cause backup jobs to fail.
- **Sort By**--Jobs in the queue are listed in order of execution time. This option changes the order in which jobs are listed in the queue. Sorting the Job Queue is for informational purposes only. It does not affect the order in which jobs are processed. To sort jobs, click any of the following fields: Status, Execution Time, Job Type, Server, Last Result, Owner, Total Files, and Description.
Note: You can resize these columns by using the “drag and drop” method with the mouse. Place the cursor on the divider between columns, click and hold down the left mouse button, and then move the divider in either direction until the column is the size you want.
- **Properties**--Double-click to call up the Job Properties dialog when the job is processing.

Job Status Types


When a job is in the Arcserve Backup queue, it is listed with a status. The status can be one of the following:

- **Done**--A non-repeating job that was successfully executed and completed.
- **Ready**--A new one-time or repeating job (a backup job that runs every Friday, for example) waiting to be executed.
- **Active**--A job that is currently being executed.
- **Hold**--A job that is in the queue and was placed in a hold status.

Note: A job with a hold status will not execute until you remove the hold status.

- Waiting for Target**--A job that is ready to execute and is waiting for the target device, media, or both to become available. A device or media may not be available, for example, because it is busy with another job. To determine the specific reason why a job is waiting for a device, click the Waiting for Target hyperlink to open the Target Group Status dialog.

The Target Group Status dialog represents stream-based backups or device-based backups.



The Target Group Status dialog describes the information that follows:

- Title bar**--Displays the name of the backup server, the job number, and the job ID of the job that you clicked in the Job Status Manager.
- The reason the job is waiting (for the device or media)
- Backup Server**--The name of the backup server using the listed device.
- Job No.**--The job number that is using the listed device, if available.
- Job ID**--The job ID that is using the listed device, if available.
- Media Name**--The name of the media in the listed device, if available.
- Drive Serial No.**--The serial number of the listed device.
- Status**--The status of the job that is using the listed device.
- Waiting for source group**--A migration job is waiting for a source group to be available.
- Waiting for source tape**--A migration job is waiting for the source tape to be available.
- Waiting for target tape**--A job that should be active, but it is not because it is waiting for the target device or media.
- Positioning source tape**--A migration job is waiting for the source tape to be positioned in the drive.

- **Positioning target tape**--A migration job is waiting for the target tape to be positioned in the drive.
- **Copying**--A migration job (copy to final destination media) is in progress.

Note: Completed jobs remain listed in the Job Queue for a specified number of hours. This period of time is set up through the Arcserve Backup Server Admin. For more information, see [Job Engine Configuration](#) (see page 500).

How to Analyze Jobs Using the Last Result Field

The Last Result field on the Job History tab indicates whether your executed job was successful. If it was not successful, the information in this field helps you determine why the job may have failed. The Last Result field may contain the following one of the following statuses:

- **Finished**--All of the nodes and drives and shares were processed.
- **Incomplete**--The job was partially successful. Review the Activity log information to check the exact nature of what occurred to prevent job completion.
- **Canceled**--The job was intentionally canceled. The following actions may have occurred:
 - A user canceled the job from the Job Queue.
 - Someone answered NO or CANCEL to a console prompt.
 - The job required either a confirmation of OK, or media to be inserted before the time out was reached. (Time out is set in the media options in the Backup Manager window.)
- **Failed**--The job failed to perform its designated task. This usually occurs if Arcserve Backup cannot back up any source nodes of a job (for example, if the agent is not loaded or an invalid password was entered) or if a hardware error occurs. If the job was started, but the Manager could not complete the job, you will receive "Run Failed" status. Review the Activity log information to check the exact nature of what occurred to prevent the job from completing.
- **Run Failed**--The job was started, but the program that runs the job failed, because either there was not enough memory to run the job or a DLL file was not found.
- **Crashed**--The job was started and a system error occurred which prevented Arcserve Backup from completing its task, such as a memory violation that caused Arcserve Backup or the operating system to be shut down. If a job has a status of Crashed, it can be retried after the Job Engine restarts. This can be set up through the Arcserve Backup Server Admin in the Job Engine Configuration tab.


How to Analyze Jobs Using Group View

The Job History tab lets you group results by Source Group. Source Group view lets you view backup results for jobs using the Group View feature.

Default view: Job view

To view jobs by source group, open the Job Manager, click the Job History tab and select Source Group in the Group By dropdown list. Click the Update button to refresh the display.

Results are displayed by group, allowing you to expand results for the specific machine, device and session detail within a group.



Re-create the Arcserve Backup Database Pruning Job

Arcserve Backup lets you re-create the Arcserve Backup Database Pruning Job in the event it was deleted intentionally or unintentionally.

To re-create the Arcserve Backup Database Pruning Job

1. Start the Arcserve Backup Server Admin and click the Configuration toolbar button.
The Configuration dialog opens.
2. Select the Database Engine tab.
3. Check the Submit prune job option.

Note: The Submit prune job option is active only if the Database Pruning Job was deleted.

4. Click OK.

The database pruning job is submitted to the job queue and will run at the specified time.

View Job Details Using the Activity Log


The Activity Log tab on the right panel contains comprehensive information about all the operations performed by Arcserve Backup.

The log provides an audit trail of every job that is run. For each job, the log includes the following information:

- Time the job started and ended
- Type of job
- Average throughput of the data
- Number of directories and files processed (backed up, restored, or copied)
- Job session number and job ID
- Result of the job
- Errors and warnings that occurred

When you install the Central Management Option, you can view Activity Log data as it relates to the domain primary server, a domain member server, or both.

The following diagram illustrates that the Central Management Option is installed, domain member server MEMBER01 is selected, and the activity log details for the MEMBER01 display.



Group by Week (if checked) is always the first level group. The date comes from the operating system's setting.

The format for the week node is as follows:

Week[start date - end date]

The format for the job node is as follows:

JobID [Server Name](Job Name)[Job Status][Start time - End time][Job No.]

The Generic Log appears at the end of the master job list. It contains the logs that do not belong to any job.

Note: If you do not install the Central Management Option, the Activity Log displays data relating to the Arcserve Backup server that you are currently logged in to.

You can scan this log every day to see if any errors occurred. You can also use it to find a session number in case you need to restore a specific session. You can organize the Activity log view or print it to a file.

Delete Activity Log Files

To conserve file space, you can delete the entire Activity log file or unnecessary log records older than a specific time period.

To delete files in the Activity log

1. Open the Job Status Manager and select the Activity Log tab. Click the Delete toolbar button.

The Delete dialog opens.

2. Select the criteria you want to apply in the Delete dialog. Chose one of the following options:
 - **Entire Log**—Deletes all log file records.
 - **Partial**—Lets you select specific logs based on a time period. You can choose from the following criteria:
 - **days:** range 1-365
 - **weeks:** range 1-54
 - **months:** range 1-12
 - **years:** range 1-10

3. Click OK.
A caution dialog appears.
4. Click OK.
The Activity log files are deleted.

You can also use the command line interface to purge job logs (or any other log file) from the Activity log. Use the `ca_log -purge` command to delete logs older than a specified period of time from any log file. You can also use the `ca_log -clear` command if you want to delete all log data from log files with no specified time period.

Note: For more information about the `ca_log` command, see the *Command Line Reference Guide*.

Activity Log Pruning

To conserve file space, you can schedule log pruning.

More information:

[Database Engine Configuration](#) (see page 514)

Set Activity Log Queries

Arcserve Backup provides you with the capability to customize the type of information and how the information displays in the Activity Log.

The default activity log query values are as follows:

- View: Group by Week and Group by Job
- Job Status: All Message
- Type: All
- Date: All Time
- Job ID: blank
- Session: blank
- Keywords: Message

Note: To return to the default setting at any time click Reset.

To set Activity Log Queries

1. Open the Job Status Manager and select the Activity Log tab.
2. Expand the Log Query Bar. By default the Log Query Bar is collapsed.
The Log Query Bar opens.

3. Specify the desired options.
 - **View**--Specify how you want to group the activity log messages. You can group by week, type and job.
 - The Group by Week lets you group the activity log messages by the week (default).
 - The Group by Jobs option, lets you group the Activity Log with the parent job together with all of its child jobs. For each parent job and its corresponding child jobs, the Activity Log presents you with a description for the job (default).
 - The Group by Type option lets you group error messages, warning messages, and information messages.
 - **Job status**--Specify the types of jobs that you want to view in the Activity Log. You can view All, Finished, All unsuccessful, Canceled, Failed, Incomplete, Crashed and Unknown.
 - **Message type**--Specify the types of messages that you want to view in the Activity Log.

You can view All, Error, Warnings, Errors and Warnings, Informations, Errors and Informations, and Warnings and Informations.
 - **Date**--Specify a date or range of dates of messages that you want to view in the Activity Log.

You can show all messages, filter messages such that only messages before or after a specified date display, or display a specific range of dates.
 - **Job ID**--Specify a know job ID.
 - **Session**--Specify a know session.
 - **Keywords**--Sort the activity log by keywords. You can specify Job Name or Message.
4. Click Update.

The Activity Log displays the results according to the specified query.

Note: To get the latest jobs, with the existing filters, click Refresh on the toolbar or F5.

Tape Log Tab

The Tape Log tab in the Job Status Manager displays if you enabled the option Show Tape Log on Job Status Manager while configuring the Tape Engine. For more information, see the section Tape Engine Configuration.

Note: After you enable the option "Show Tape Log on Job Status Manager," you must click Refresh in the Job Status Manager for the changes to take effect.

More information:

[Tape Engine Configuration](#) (see page 503)

Job Detail Tab

The Job Detail tab in the bottom panel displays details about any job in the queue, including the source and destination targets and the job's schedule. If you have selected customization options such as Pre/Post backup requirements, they will be displayed here. After a job has started, you can view its sequence and session number.

Job Log Tab

The Job Log tab in the bottom panel displays information about specific jobs that have been executed. It is generated for each job that Arcserve Backup runs. You can specify the level of detail in the log by choosing the Log options before you submit the job.

Arcserve Backup provides the following log options:

- **Log all activity**--Record all of the activity that occurs while the job is running.
Note: When you specify Log all activity, Arcserve Backup creates a log file named JobLog_<Job ID>_<Job Name>.Log. With this log file, you can view detailed logging information about the job. Arcserve Backup stores the log files in the following directory:
C:\Program Files\CA\ARCserve Backup\LOG
- **Log summary only (default)**--Record summary information on the job (including source, destination, session number, and totals) and errors.
- **Log disabled**--Do not record any information about this job.

How Save Agent and Save Node Information Works

Saving an agent or node to the Arcserve Backup database makes it accessible to all users on the same domain. Arcserve Backup views the Primary Server, Member Servers, and all agents in a domain as nodes.

The Save Agent/Node Information feature lets you perform the following tasks:

- Save new nodes to the Arcserve Backup database
- Save the user account information to the Arcserve Backup database
- Filter nodes by agent type
- Group agents by type of agent

Add, Import, and Export Computers Using the Source Tree

The Add, Import, and Export feature lets you add multiple computers (for example, nodes, agents, and Arcserve D2D servers) into the Backup Manager Source tree, whether computers will or will not be backed up.

This function is beneficial in backup environments that contain a large number of computers. You can add, import, and export computers from the Classic View and Group View in the Backup Manager.

To add, import, and export computers using the Source tree

1. Open the Backup Manager and select the Source tab.

From the view drop-down list, specify one of the following views:

- Group View

Note: The default view is Group View.

- Classic View

The view option is applied.

2. Based on the view option that you specified, do one of the following:

- **Group View**--Right-click one of the group objects in the browser (for example, the Microsoft SQL Server, Microsoft Exchange Server, and Oracle Server objects) and select Add/Import/Export Agents from the pop-up menu.

Note: If the group specified is Arcserve D2D Servers, click Add/Import/Export D2D Servers on the pop-up menu.

- **Classic View**--Right-click one of the classic objects in the browser (for example, the Windows Systems object) and select Add/Import/Export Nodes from the pop-up menu.

Note: If the classic object specified is Arcserve D2D Servers, click Add/Import/Export D2D Servers on the pop-up menu.

The Add/Import/Export dialog opens. Any any existing computers are populated in list of computers that will be added to the Source tree as viewed in the right-pane.

3. Add the computers to the list in the right-pane, which will be added to the Source tree. This can be done in the following ways:

- Specify the host name or host name (IP address) of the computers that you want to add in the text box and click Add.

The best practice is to specify the host name and the IP address of the target system. This approach helps ensure that Arcserve Backup can accurately detect the target system, based on its IP address, and display the system under the Windows Systems object.

Note: If you specify only the host name, Arcserve Backup sets the IP address value to 0.0.0.0.

- Select the computers from the list of computers detected by auto-discovery in the left-pane and click Add or Add All.

You can select multiple computers and agents using the CTRL + click and Shift + click keyboard shortcuts.

Note: The computers are removed from the left-pane list after they have been added to the right-pane list.

- Click Import to add a list of computers using a .csv or .txt file.

Example: .txt file:

```
Hoatname1(IP)
Hostname2(IP)
Hostname3(IP)
Hostname4(IP)
```

Example: .csv file:

```
Hostname1(IP), Hostname2(IP), Hostname3(IP), Hostname4(IP), ...
```

Note: For more information, see [Add Multiple Computers Using .csv and .txt Files](#) (see page 337).

The computers that will be added to the Backup Manager Source tree are displayed in the right-pane list.

4. (Optional) Click Delete or Delete All if necessary to remove items from the right-pane list.

The Delete and Delete All buttons are only enabled if you select a node or multiple nodes in the right-pane list. If the computer was originally entered in the text box or imported from a .csv or .txt file, and you click Delete, the computers will be removed from the right-pane list. If the computers were detected by auto-discovery, and you click Delete, the computers display in the left-pane list of computers detected by auto-discovery.

5. Select the computers in the right-pane list for which you want to enter a user name and password, and then click Security.

(Optional) From the list of computers that will be added to the Source directory tree, double-click the Host name or Address value of the target system.

The Security dialog opens where you can add the user name and password for multiple nodes and agents at one time. The nodes and agents displayed on the Security dialog are provided from the right-pane list on the Add/Import/Export dialog.

6. Enter the user name and password and click OK.

You are returned to the Add/Import/Export dialog and the user name and password are added to the right-pane list.

7. (Optional) Select a node or agent in the left-pane list and click Properties.

The Server Properties dialog opens and displays the Domain name, Server name, IP address, Last response time, and Products installed. These properties are detected by the auto-discovery service, so the Properties button will only be enabled when you select a node or agent in the left-pane list and click Properties.

8. Click OK.

If Arcserve Backup cannot access the newly added computers, the Add Agents Result dialog opens.

Be aware of the following:

- The Add Nodes Result or the Add D2D Servers Result dialog opens when adding Nodes or Arcserve D2D Servers.
- The result dialog provides you with a list of computers that Arcserve Backup cannot access and the corresponding reason (Status) for each computer. If the reason for the failure relates to security credentials, the result dialog lets you modify the user name and password that Arcserve Backup requires to log in to the computer. Follow the on-screen instructions on the result dialog to add the computers and click OK.

The computers are added to the Backup Manager Source tree. If an existing computer was deleted, it will be removed from the Backup Manager Source tree. If a computer name is duplicated, you will see a warning message indicating that this is a duplicate name and computer will not be added to the Source tree. In addition, Arcserve Backup lets you add multiple host names with the same IP address.

Add Multiple Agents and Nodes Using .csv and .txt Files

The Import function lets you add multiple computers (for example, nodes, agents, and Arcserve D2D Servers) using the Arcserve Backup user interface by importing them from a .csv or a .txt file.

Note: A .csv file is a file that uses a comma-separated value format.

To add computers using .csv and .txt files

1. Open the Backup Manager and select the Source tab.

From the view drop-down list, specify one of the following views:

- Group View

Note: The default view is Group View.

- Classic View

The view option is applied.

2. Based on the view option specified, do one of the following:

- **Group View**--Right-click one of the group objects in the browser (for example, the Microsoft SQL Server, Microsoft Exchange Server, and Oracle Server objects) and select Add/Import/Export Agents from the pop-up menu.

Note: If the group object specified is Arcserve D2D Server, click Add/Import/Export D2D Servers on the pop-up menu.

- **Classic View**--Right-click one of the classic objects in the browser (for example, the Windows Systems object) and select Add/Import/Export Nodes from the pop-up menu.

Note: If the classic object specified is Arcserve D2D Server, click Add/Import/Export D2D Servers on the pop-up menu.

The Add/Import/Export dialog opens.

3. Click Import.

The Windows Open dialog opens.

4. Browse to the file containing the list of computers that you want to import and click Open.

The computers added to the right-pane list on the Add/Import/Export dialog.

5. Select the computers in the right-pane list for which you want to specify user name and password and then click Security.

The Security dialog opens. The Security dialog lets you add the user name and password for computers at one time. The computers displayed on the Security dialog are provided from the right-pane list on the Add/Import/Export dialog.

6. Click OK.

The computers are added to the Backup Manager Source tree.

Export Multiple Agents and Nodes to a Text File

Nodes and agents that are already entered Backup Manager Source tab can be exported to a .txt file to make it easy to import the list of nodes and agents to another Arcserve Backup server.

To export multiple agents and nodes to a text file

1. Open the Backup Manager and select the Source tab.

From the view drop-down list, specify one of the following views:

- Group View

Note: The default view is Group View.

- Classic View

The view option is applied.

2. Based on the view option specified, do one of the following:

- **Group View**--Right-click one of the group objects in the browser (for example, the Microsoft SQL Server, Microsoft Exchange Server, and Oracle Server objects) and select Add/Import/Export Agents from the pop-up menu.

The Add/Import/Export Agents dialog opens.

- **Classic View**--Right-click one of the classic objects in the browser (for example, the Windows Systems object) and select Add/Import/Export Nodes from the pop-up menu.

The Add/Import/Export Nodes dialog opens.

3. Click Export

The Export dialog opens.

4. Select the agents or nodes that you want to export to a text file.

Note: By default, all agents or nodes are selected for you.

5. (Optional) Click Select All or Clear All to select or clear the agents and nodes in the list that you want to export.

6. Click OK.

The Windows Save As dialog opens.

7. Select a path where the file should be created and saved.

The selected agents or nodes are exported to a .txt file.

Note: The user name and password are not exported.

Filter Nodes by Agent Type

You can determine which computers in the domain have the same agents installed. Arcserve Backup lets you filter nodes from the Backup Manager Source tree and the Restore Manager Source tree and Destination tree.

Note: From the Classic View, you can filter only the Windows System object and the UNIX/Linux Systems object.

To filter nodes by agent type

1. Open the Backup Manager and select the Source tab.

From the view drop-down list, specify one of the following views:

- Group View

Note: The default view is Group View.

- Classic View

The view option is applied.

2. Click Agent Type on the Source toolbar.

The Filter by Agent Type dialog opens. The default value is Show all nodes.

3. Click Filter by Agent Type.

From the list of Arcserve Backup agents, specify the agents that you want the Source tree to display.

Click OK.

The nodes are filtered and only the nodes containing the specified agents display.

Note: To clear the filters, so that you can view all nodes, click Reset on the Source toolbar.

Modify the IP Address or Host Name of Agents and Nodes

Arcserve Backup lets you modify the IP address, host name, or both for agents and nodes.

To modify the IP address or host name of agents and nodes

1. Open the Backup Manager and select the Source tab.
From the view drop-down list, specify Classic View.
The view option is applied.
2. Locate the agent or node that you want to modify.
Select and right-click the agent that you want to modify.
From the pop-up menu, click Modify Agent.
The Agent Option dialog opens.
3. From the Agent Option dialog, clear the Use computer name resolution check box.
Enter a Host name and an IP address.
Click OK.
The new IP address, host name, or both are applied to the agent or node.

Delete Agents and Nodes from the Source Tree

Arcserve Backup lets you delete agents and nodes from the Backup Manager source tree. The agent and node information, all accounts associated with the agents and are deleted from the Arcserve Backup database.

To delete agents and nodes from the Source tree

1. Open the Backup Manager and select the Source tab.
From the view drop-down list, specify one of the following views:
 - Group View
Note: The default view is Group View.
 - Classic ViewThe view option is applied.

2. Locate the agent or node that you want to delete.

Based on the view that you specified, do one of the following:

- **Group View**--Select and right-click the agent that you want to delete.
- **Classic View**--Select and right-click the node that you want to delete.

From the pop-up menu, click Delete Machine/Agent.

The delete confirmation dialog opens.

3. Click Yes.

The agent or node is deleted.

How to Use the Job Scheduler Wizard to Schedule Jobs

The Arcserve Backup command line enables direct control over all operations that can be performed by a Arcserve Backup server. The Job Scheduler Wizard provides an alternative to entering job scheduling commands in the Command Prompt window.

The benefits of using this wizard rather than the command line include:

- Jobs can be scheduled and repeated.
- Jobs appear in the Job Queue and Activity log.
- Jobs can be stopped in the Job Queue.
- The commands you can enter are not limited to Arcserve Backup. You can use this wizard for virtually any executable, such as Notepad.exe.
- It provides an easy way to quickly package and submit jobs.

Important! All scheduled times for Arcserve Backup jobs are based upon the time zone where the Arcserve Backup server is located. If your agent machine is located in a different time zone than the Arcserve Backup server, you will need to calculate the equivalent local time that you want the job to be run.

When you submit a job using the Job Scheduler Wizard, it is labeled as a generic job in the Job Queue and Activity log. Although you can modify a generic job using the Job Queue; you can only reschedule and stop it.

Note: You must have Administrator rights on the local Windows machine to submit jobs using the Job Scheduler Wizard.

Job Scripts

A script is a job that you saved to a file. It contains the original source, destination, options, and schedule information for the job. It will also contain any filters you created to include and exclude files and directories.

Creating a script has the following advantages:

- You can re-use the same settings later.
- You can copy your settings to a different Windows machine running Arcserve Backup.
- You can quickly resubmit regularly executed jobs after a job has been accidentally deleted.

Create a Job Script

You can save almost any type of job as a script. A script is a set of Arcserve Backup instructions that let you execute jobs.

To create a job script

1. After you create the job, click the Submit toolbar button.
The Submit Job dialog opens.
2. Click Save Job to save the job criteria in a script.
The Save Job Script dialog opens.
3. Enter a name for the script and click Save.
The job script is saved.
4. Click OK to submit the job to the queue.
The job is submitted and a job script is created.

Execute a Job Using a Script

You can execute almost any type of job using a script. A script is a set of Arcserve Backup instructions that let you execute jobs.

To execute a job using a script

1. Open the Job Status Manager and select the Job Queue tab.
Click Load on the toolbar.
The Add Job dialog opens.

2. Browse to and select the script for the job that you want to execute.
Click Open.
The Select a Server dialog opens.
3. From the drop-down list, select the server from which you want the job to execute.
Click the Submit on Hold option to submit the job in a Hold status.
Note: The Submit on Hold option is selected by default.
Click OK.
The job information for the previously saved script will be displayed in the Job Queue as a new job.

Job Templates

A job template contains a series of settings such as the destination, options, and schedule information for the job. Similar to job scripts, a template can also contain any filters you created to include and exclude files and directories.

Job templates are different, however, from job scripts because they provide the flexibility to repeat custom backup schedule settings on other Arcserve Backup machines. Because the job template does not retain the backup source information as the job script does, the template files can be copied and applied to any new server source running Arcserve Backup. But, job scripts cannot be modified to accommodate new server sources.

You can choose from seven default job templates or you can create a custom template to meet your individual backup needs. The default job templates are designed to meet specific backup tasks such as rotation scheme, backup method, and GFS options. The default job templates can be accessed from the File menu when you choose the Open Job Template option.

Create Custom Job Templates

You can create a custom job template that you can save for future jobs on any Arcserve Backup system.

To create a job template

1. From the Arcserve Backup Manager window, select Backup from the Quick Start menu.
The Backup Manager Window opens.

2. Make selections for your backup job by accessing the Start, Source, Destination, and Schedule tabs.

Click Submit in the toolbar to submit the job.

The Submit Job dialog opens.

3. Click Save Template.

The Save Job Template dialog opens.

4. In the File Name field, specify a name for the job template and click Save.

The job is saved as a job template with an .ast file name extension.

Note: While default job templates are stored in the Templates/Jobs folder in the Arcserve Backup directory, you can save your template in any directory you want. To open your custom job template on a local machine or from a remote server, access the File menu and choose the Open from Template option. After the job template is open, specify the source data, and then you can submit the job.

Save Job Template

You can create custom job templates that you can save for future jobs on any Arcserve Backup machine. The Save Job Template dialog allows you to save your custom job template with an .ast filename extension.

To save job templates

1. Open the Backup Manager by selecting Backup Manager from the Quick Start menu.

The Backup Manager window opens.

2. Make selections for your backup job by accessing the Source, Staging, Destination, and Schedule tabs.

3. Submit your job by clicking the Submit button.

The Security and Agent Information dialog opens.

4. Edit or confirm the security and agent information for your job and click OK.

The Submit Job dialog opens.

5. Click Save Template in the Submit Job dialog box to save your new job template with an .ast filename extension.

The Save Job Template dialog opens.

6. Complete the following fields and click Save.
 - **Save In**--Specify the path for the job template file. While default job templates are stored in the Templates/Jobs folder in the Arcserve Backup directory, you can save your template in any directory you want.
 - **File Name**--Enter a name for the job template. By default, it will be saved with an .ast extension.
 - **Save as type**--To save the job template in another format, select the file format you want.
 - **Job Information**--Enter a brief description of the job.

The job template is saved.

Open Job Template

The Open Job Template dialog allows you to open an existing job template you previously saved or choose from one of the default job templates.

To open a job template

1. Select File from the main menu and then select Open from Template.

The Open Job Template dialog opens.

2. Complete the following fields and click Open.
 - **Look In**--Indicates the path for the job template file. While default job templates are stored in the Templates/Jobs folder in the Arcserve Backup directory, you can save your template in any directory you want.
 - **File Name**--Enter a name for the job template you want to open.
 - **Files of type**--If you saved the job template in another format, select the file format you want to open. By default, the job template will be saved with an .ast extension.

Once the job template is open, you can submit your job.

Note: You can choose from seven default job templates or you can create a custom template to meet your individual backup needs. The default job templates are designed to meet specific backup tasks such as rotation scheme, backup method, and GFS options. The default job templates can be accessed from the File menu when you choose the Open from Template option.

Windows-Powered NAS and Storage Server 2003 Device Configuration

Arcserve Backup provides support for backup and restore of Windows-powered NAS and Storage Server 2003 devices (referred to as Windows-powered NAS).

When you install Arcserve Backup on Windows-powered NAS, a new Arcserve Backup tab is available on the Web administration user interface for the device. By accessing the tab, you can connect directly with the Arcserve Backup components.

Access Arcserve Backup Through the Windows-powered NAS Device

You can administer backup and restore jobs as well as perform agent maintenance for Windows-powered NAS devices through the Web interface. A seamless integration of the Arcserve Backup Home Page is easily accessible from the Windows-powered NAS Web administration interface.

Links to Arcserve Backup Manager, Device Configuration, or Client Agent Admin are displayed from the Windows-powered NAS menu option. The options available are dependent upon the options installed on the Windows-powered NAS Device.

Use the following table to determine what options are available in the Windows-powered NAS Web administration interface based on a specific Arcserve Backup component install.

Installed Arcserve Backup Component	Options Available in Windows-powered NAS Interface
Arcserve Backup Manager	Arcserve Backup Manager
Arcserve Backup Server	Device Configuration
Arcserve Backup Windows Client Agent	Agent Admin

Arcserve Backup and Windows-powered NAS Device Configuration

The following describes basic Windows-powered NAS configurations supported by Arcserve Backup.

This section contains the following topics:

- [Backup Devices Connected Directly to Windows-powered NAS Devices](#) (see page 347).
- [How You Can Back Up Devices Connected to Arcserve Backup Server](#) (see page 347).
- [How You Can Back Up Devices Shared Between Arcserve Backup and Windows-powered NAS](#) (see page 348).

Backup Devices Connected Directly to Windows-powered NAS Devices

You can configure and deploy Arcserve Backup directly on a Windows-powered NAS as displayed in the following illustration:

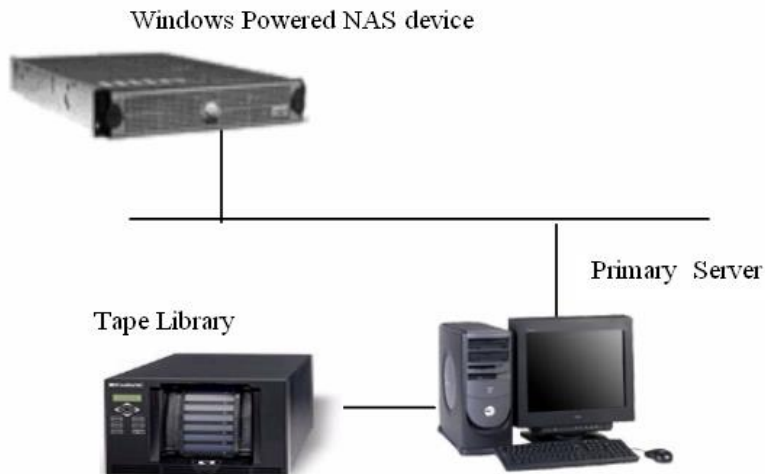


You can use the Web interface integration provided by Arcserve Backup on a remote server and perform backup and restore tasks as well as monitor scheduled jobs configured for the installation.

How You Can Back Up Devices Connected to Arcserve Backup Server


You can configure Arcserve Backup Windows Client Agents on the Windows-powered NAS device. Agents can be administered using the integrated Web administration interface provided by Arcserve Backup.

The Agents can be backed up from the remote Arcserve Backup Server which may be running on another Windows-powered NAS device as shown in the following illustration.



How You Can Back Up Devices Shared Between Arcserve Backup and Windows-powered NAS

You can configure Arcserve Backup Server, Manager, and the SAN Option on a Windows-powered NAS device and create a secondary remote Arcserve Backup server with a SAN Option. Both machines can connect to a shared backup device such as a tape library through a fiber switch as displayed in the following illustration.



Scan Jobs for Cloud Storage

The process of running a Scan job for cloud storage is the same as a regular Scan job. If desired, click the Media Assure button to access Media Assure options, then click the Enable Media Assure option to select it.

The Media Assure option works with all types of media. With cloud storage, Media Assure must randomly sample large quantities of sessions that meet your criteria to help ensure that the data can be recovered. In addition, cloud storage access speeds can be slow. As a result, scan jobs associated with cloud storage may require additional time to complete.

You can scan specific sessions or the entire cloud storage, as well as set filter options. You can specify one node or multiple nodes, separating names with commas in the field provided. You cannot specify * groups for cloud storage Media Assure scan jobs, but you can specify a * tape in a specific group, or you can select a specific tape. The job scans the sessions that satisfy your selection criteria and repeats periodically until deleted.

Note: Media Assure supports scanning only one group and is suspended by other jobs targeting the same media. If suspended, an error is produced, "E3708 Unable to reserve group <group name> in <minutes> minutes."

More information:

[Media Assure & Scan Utility](#) (see page 32)

Merge Jobs to Cloud Storage

The process of performing a merge job associated with cloud storage is identical to that of a regular merge job. For example, you can merge specific session details or the entire cloud device to the database. For more information, see [Merge Utility](#) (see page 30).

Note: Cloud storage access speeds can be slow. As a result, merge jobs associated with cloud storage may require additional time to complete.

How Arcserve Backup Purges Data from Cloud Storage

There are several ways to purge cloud sessions:

- From the Backup Manager console, Destination tab, set the retention time through Cloud Purge Policy when submitting a migration job. When the purge policy expires, the cloud session is deleted to free up space and keep costs down.
- From the Device Manager console, click Browse Cloud Folder from the toolbar, select a folder and right-click to delete.
- From the Restore Manager console:
 1. Select from the Source tab, the Restore by Session method.
 2. Select the cloud sessions you want to delete from the Sessions directory tree.
 3. Right-click the sessions and click Delete selected sessions.
The sessions are removed immediately from the cloud storage.

Chapter 6: Archiving Data

Arcserve Backup provides the ability to archive your files by setting schedules, policies, and selecting your source and destinations on Windows, UNIX/Linux and MAC file systems.

This section contains the following topics:

[Overview of How Arcserve Backup Archives Data](#) (see page 351)

[Submit Archive Jobs](#) (see page 351)

[Archive Manager](#) (see page 353)

[Global Options for Archiving](#) (see page 359)

Overview of How Arcserve Backup Archives Data

Arcserve Backup lets you archive data that resides on Windows, UNIX/Linux and MAC file systems to tape media or cloud storage devices. Arcserve Backup provides the ability to archive data using the [Archive Manager](#) (see page 353). You open the Archive Manager from the Arcserve Backup home page. You can archive old files to media that can be retained for a predetermined period of time, and you retrieve the archived files using the [Restore Manager](#) (see page 257) at any time.

Arcserve Backup also lets you set schedules and policies and select your source and destination for your files to be archived in order to [Submit Archive Jobs](#) (see page 351).

Submit Archive Jobs

This section includes steps on how to submit an archive job.

To submit archive jobs

1. From the Archive Manager, select the [Source](#) (see page 355), [Schedule](#) (see page 356), [Policies](#) (see page 356), and [Destination](#) (see page 358) tabs to specify the options that you require for the job.
2. Click the Options toolbar button to specify global options that you require for the job. For more information, see [Global Options for Archiving](#) (see page 359).
3. Click the Submit toolbar button to submit your job.

The Security and Agent Information dialog opens.

4. On the Security and Agent Information dialog, edit or confirm the security and agent information for your job, and click OK.
5. When the Submit Job dialog opens, select Run Now to run the job immediately, or select Run On and select a date and time when you want the job to run.
Note: For more information about the Run Now option, see [Job Queue Tab](#) (see page 74).
6. Enter a description for your job.
7. If you selected multiple sources to archive and want to set the priority in which the job sessions initiate, click Source Priority. Use the Top, Up, Down, and Bottom buttons to change the order in which the jobs are processed. When you finish setting priorities, click OK.
8. To save the job as a Arcserve Backup job script, click the Save Job button.
9. To save the job template, click the Save Template button.
10. To preflight, check the job, click the Preflight Check button. If the preflight check failed, click the Cancel button to modify the job settings.
11. On the Submit Job dialog, click OK.

The job is submitted.

If you would like to view the status of the archive job, go to the [Job Status Manager](#) (see page 319).

More information:

[How to Manage Jobs Using the Job Queue Tab](#) (see page 323)

Archive Manager

The Archive Manager lets you customize your archive jobs using filters, specifying options, and setting schedules.

You can use the Archive Manager to:

- Set archive schedules and policies for selected Windows, UNIX/Linux, and MAC file system data.
- Submit archive jobs.
- Find data that is eligible for archiving and backing up, and then delete the archived data from its original location immediately.
- Use filters to selectively exclude or include directories and files for an archive job.
- Specify the File Size and File Age for an archive job.
- Set retention time to the archived data for how long to keep data in the archive media.
- Retrieve archive data from archive media with restore job.
- View archive status through Dashboard reports. For more information on reports, see the *Dashboard Guide*.
- View list files that were archived. Arcserve Backup records the list of files in a CSV file, and stores the CSV files in the following directory:

BAB_HOME\Archived Files

- The CSV file is backed up in the Arcserve Backup Database protection job.
- The CSV file is backed up in a backup job that includes the Catalog database, as specified in the Global options for the job.
- The CSV file is compressed on NTFS file systems.

The Archive Manager provides tabs that let you customize your archive job:

- **Start tab**--Lets you specify the type of archive: Normal, Deduplication, or Cloud. You may also enable staging for each archive type. For the Cloud Archive option, the Enable Staging option is selected automatically.
- **Source tab**--Lets you specify the data that you want to archive.
- **Schedule tab**--Lets you specify a repeat method for the job.
- **Policies tab**--Lets you specify a file type, file size, when the file was last accessed, modified or created, and data verification.
- **Destination tab**--Lets you specify the location where you want to store your archive files.

The topics that follow provide full details about the options available on each tab.

This section contains the following topics:

[Options on the Archive Manager Start Tab](#) (see page 354)

[How to Specify Source Data for Archiving Using the Classic View and the Group View](#) (see page 355)

[Archive Job Schedules](#) (see page 356)

[Options on the Archive Manager Policies Tab](#) (see page 356)

[Options on the Archive Manager Destination Tab](#) (see page 358)

Options on the Archive Manager Start Tab

From the Archive Manager Start tab, you can select one of the archive types below:

- **Normal Archive**--Normal Archive lets you archive data source to a target destination, using a repeat method schedule.
- **Deduplication Archive**--Deduplication Archive lets you save only unique data chunks to disk, allowing you to fit more archive data on media, retain files for longer periods of time.
- **Cloud Archive**--Cloud Archive lets you archive data to a cloud device destination only.

Note: [Enable staging](#) (see page 241) is automatically selected if you choose this option.

For each archive type, you must click the [Source](#) (see page 355), [Schedule](#) (see page 356), [Policies](#) (see page 356), and [Destination](#) (see page 358) tabs to complete archive job configuration.

You may also [Enable Staging](#) (see page 241). Staging operations allow you to archive data to a staging device and then migrate the archive data to a final destination (usually a tape). You can choose to Enable Staging on Normal Archive or Deduplication Archive jobs.

How to Specify Source Data for Archiving Using the Classic View and the Group View

The source is the path to the data that you want to archive. You can easily find the files you want to archive by browsing through the Archive Manager directory to select the user-shared drives and directories.

Arcserve Backup lets you browse and specify the source data using the following views:

- **Classic View**--This is traditional source view. Machines are listed first, allowing you to expand and then select specific data sources. With the Classic View, Arcserve Backup categorizes source computers based on the platform that is running on the computer. For example, Windows systems and UNIX/Linux systems.
- **Group View**--This view categorizes source computers based on the Arcserve Backup agent that is installed on the computer. The agents are listed as branches on the source tree. Within each branch, the computers that contain the specified agent are listed.

You may also create customized groups that allow you to group machines according to your own criteria. For example, using the Group view is an effective approach to specifying source when you want to create a group for a specific file server.

When selecting a source, you can select to archive individual drives, directories, and files

To select individual drives, directories, and files to archive, expand a server name and click the green boxes next to each drive, directory, and file.

To select an entire source group, click the green box next to the group name. When you do this, all the servers, nodes, volumes, drives, directories, and files included in the source group are automatically selected.

Be aware of the following behavior:

The view that you specify when you submit a job cannot be modified.

For example, you submit a job using the Classic View. Subsequently, you want to modify the source selections for the job. When you modify the job and click the Archive Manager, Source tab, the view drop-down menu is disabled.

Archive Job Schedules

You can configure your archive job by using the Repeat Method on the Schedule tab in the Archive Manager. The schedule lets you run an archive job on a repeating basis. You can specify the following parameters for an archive job:

- **Repeat Method**--All jobs can be scheduled using the Schedule options available in each Manager. Jobs can be submitted with the following repeat methods:
 - **Day(s) of the Week**--Repeat this job on the days that are specified. (Default) Friday is the default day of the week.
 - **Week(s) of the Month**--Repeat this job on the weeks that are specified.
 - **Day of the Month**--Repeat this job on the specified day.
 - **Custom**--Repeat this at the specified interval, but exclude the days that are checked.
- **Append**--Select this option to add additional days to your scheduled archive job.
- **Use WORM Media**--Select this option to back up data to write once, read many (WORM) backup media when the backup job runs.

Options on the Archive Manager Policies Tab

You can use the Policies tab in the Archive Manager to set archive policies for the selected Windows, UNIX/Linux and MAC file system data. The Archive Manager Policies tab includes the following filters:

- **File Filter**--The File Filter options let you include or exclude files, directories, or both from your archive job.

Types of filters offered in this section are:

- **File Pattern**--Use the File Pattern filter to include or exclude files from a job. You can specify a particular file name or you can use wildcards to specify a file pattern.
Note: Wildcards "*" (asterisk) and "?" (question mark) can be used in the file pattern filter.
- **Directory Pattern**--Use the Directory filter to include or exclude specific directories from a job. You can enter an entire directory name or you can use wildcards to specify a directory pattern.
Note: Wildcards "*" (asterisk) and "?" (question mark) can be used in the directory pattern filter

- **File Attributes**--Use the File Attributes filter to include or exclude specific types of files from a job. Select as many of the following types of file attributes as you want:
 - **Archive**--Files whose archive bit is set.
 - **Read Only**--Files that cannot be modified.
- **Host Name Pattern**--Use the Host Name Pattern filter to include or exclude specific computer systems from a job.
- **File Size**--The File Size options let you specify the file size attribute to include files. You can specify a size range from 0 to 999999 bytes, KB, MB, or GB. There are three file size options that you can specify:
 - **Greater than**--Files with a file size that is equal to or greater than the range are included or excluded from the job.
 - **Less than**--Files with a file size that is equal to or less than the range are included or excluded from the job.
 - **Between**--Files with a file size that falls between the ranges are included or excluded from the job.
- **File Age**--The File Age options let you use the file age attribute to exclude files based on when they were last accessed, modified, and created in days, months, and years.

Note: Arcserve Backup preserves the last access time for the file.
- **Verification**--The Verification options let you verify that your data was archived correctly. The following options are available.
 - **None**--No verification will be performed on the data backed up.
 - **Scan Archive Media Contents**--Lets Arcserve Backup scan the media and check that the header is readable for each file that is archived.
 - (Default) **Compare Archive Media to Disk**--Lets Arcserve Backup read blocks from the media and compare the data (byte for byte) on the media against the files.

Note: With this option specified, the throughput for the job may be slow.

Options on the Archive Manager Destination Tab

The destination for a Normal Archive job can be File System Device (FSD) and Tape (includes standalone, tape library, and virtual tape library).

The destination for a Deduplication Archive job can only be Deduplication Device.

You can use the Destination tab in the Archive Manager to browse to and select the groups and devices. The Archive Manager Destination tab includes the following archiving options:

Multiplexing

The following options regulate how Arcserve Backup handles multiplexing.

- **Maximum Number of Streams**--Sets the maximum number of streams that can write to a tape at the same time. The default number of streams is 4 and the supported range is between 2 and 32.

Multistreaming

The Multistreaming options let you split single archive jobs into multiple jobs and use all of the available tape devices in your system to complete the archive.

Group and Media field

Lets you specify the device group that you want to use for the archive job.

- Place an asterisk in the Group or Media field to use the first available drive and media in the group.
- If you want to use any available group, click the Use Any Group option.

Media Pool

You must specify a media pool for the archive job. A media pool is required so that you can specify tape media or group media.

Be aware of the following:

- You cannot use inactive media pools that were created using the Media Pool Manager for archive jobs. If you want to use this type of media pool, delete the media pool using the Media Pool Manager and then submit the archive job using the name of the media pool that was deleted.
- When you specify a Media Pool, Arcserve Backup automatically checks the other destination and archive options that you selected to verify that restrictions and conflicts do not occur when you run the job. If Arcserve Backup detects a conflict, a warning dialog opens. For more information on Media Pools, see [How Media Pools Work](#) (see page 121).
- You do not need to specify a media pool only if the final destination is a deduplication device or deduplication device group.

Retention Time

Click this button to define the length of time that archive media is not used before it is moved to the scratch set (tape media) or purged (deduplication devices and backup to cloud devices).

Examples:

The device is a library and the retention time is 30 days: The media will remain in the Save Set as long as it has been used within 30 days. If the media was not used for 30 days, Arcserve Backup moves the media from the Save Set to the Scratch Set.

The device is a deduplication device and the retention time is 30 days: Arcserve Backup purges the data from the deduplication device after the 30-day retention period expires.

Servers

This field displays the name or the primary server and member servers in your Arcserve Backup domain.

Note: If you did not install the Central Management Option, the name of the current backup server displays in this field.

More information:

[How Arcserve Backup Processes Backup Data Using Multistreaming](#) (see page 103)

Global Options for Archiving

This section describes the global options for archiving when submitting your archive job. For a description of additional job options and filtering features, see [Customizing Jobs](#) (see page 301).

To access the global options dialog, click the Options toolbar button in the Archive Manager.

This section contains the following topics:

[Archive Manager Media Options](#) (see page 360)

[Archive Manager Operation Options](#) (see page 362)

[Archive Manager Pre/Post Options](#) (see page 364)

[Archive Manager Job Log Options](#) (see page 366)

[Archive Manager Alert Options](#) (see page 366)

[Archive Manager Advanced Options](#) (see page 369)

[Archive Manager Encryption/Compression Options](#) (see page 370)

Archive Manager Media Options

You can specify the overwrite/append rules for the media used in your archive job while you are configuring the job. This section describes the rules so that you can determine which method is best for your purposes.

Arcserve Backup allows up to 20,000 sessions on a single tape and up to 101 sequences of a series of spanned tapes. Keep this in mind when planning your archives, because, if your sessions are small, you can reach 20,000 sessions quickly. If you have a large amount of data to archive, you can quickly exceed 101 sequences, depending upon how much data each tape can hold. You can stop appending the data to the tape when the sequence number reaches the maximum limit and start a new tape set by selecting the Overwrite Same Media Name, Blank Media or Overwrite Same Media Name, or Blank Media First, then Any Media option.

First Archive Media

The first archive media is the media you use when the archive job begins. The options for the first media determine the overwrite rules for the first media that is used for the archive job:

- **Append to Media**--Append job sessions to the selected media.
- **Overwrite Same Media Name, or Blank Media**--Overwrite the media in the drive only if it is the one you specified for the job or if the media is blank. If neither of these conditions are met, Arcserve Backup prompts you to supply the specific media name.
- **Overwrite Same Media Name, or Blank Media First, then Any Media**--Overwrite any media found in the drive. If you select this media option, Arcserve Backup checks to see if the media in the drive is the one specified for the job. If it is not, Arcserve Backup checks to see if the media is blank. If the media is not blank either, Arcserve Backup reformats whatever media it finds in the device and starts to archive files at the beginning of the media.
- **Timeout for First Media**--Number of minutes Arcserve Backup attempts to write to media before canceling job or selecting a different media.

Be aware of the following:

- Deduplication devices do not support the **Overwrite Same Media Name** option. The archive job will always append to the deduplication device even if the **Overwrite Same Media Name, or Blank Media** or the **Overwrite Same Media Name, or Blank Media First, then Any Media** options are selected.
- If you want to format the deduplication device, format it manually in the Arcserve Backup Manager.

Additional Archive Media

The Additional Archive Media options apply to jobs that require more than one media to determine the overwrite rules for the additional media. You need to specify which media Arcserve Backup can use when the job spans media.

- **Overwrite Same Media Name, or Blank Media**--Write to the media in the device only if it has the same media name (but a different media ID) or if it is blank. Arcserve Backup remembers the name and ID of the job's first media. When the job requires additional media, Arcserve Backup checks if the new media has the same name (but different media ID) or if it is a blank media. As long as the ID is different, Arcserve Backup reformats the media, giving it the same name and ID as the first media. The sequence number changes.

Note: To overwrite media based on its name only, select the Distinguish Media by Name Only option.

- **Overwrite Same Media Name, or Blank Media First, then Any Media**--Overwrites any media found in the device (as long as it has a different ID from the first media's ID). If neither of these conditions are met, Arcserve Backup reformats whatever media it finds in the drive and starts to archive files at the beginning of the media. All subsequent media are reformatted with the same name and ID as the first media. Only the sequence number changes.

Note: To overwrite media based on its name only, select the Distinguish Media by Name Only option.

- **Timeout for Additional Media**--Number of minutes Arcserve Backup pauses before attempting write archive data to the same media, write archive data to different media, or cancel the job.

Distinguish Media by Name Only

Arcserve Backup writes to any media that has the name specified in the Media text box on the Destination tab, regardless of the media's ID or sequence number. This option is useful if you are running a repeating Overwrite job with a specific media and you want to ensure that the same media is used for the job each time.

When this option is not enabled, the second time the archive job is run, Arcserve Backup might not be able to locate the original tape because some of its identifying features will have changed. When this option is enabled, however, Arcserve Backup simply looks for a media that has the name specified in the Media text box and uses it, regardless of the media's other identifying features.

Note: If more than one media in the tape library have the same name, Arcserve Backup will use the first media in the device group that matches the specified name. Therefore, we do not recommend that you use this option to perform a one-time overwrite.

Archive Manager Operation Options

The operation options for archive determine related actions that occur during or after the archive, and the level of information that is recorded in the database. Arcserve Backup provides the following options:

Operation Options

The following options affect only archive operations.

- **Disable File Estimate**--By default, file estimation is disabled. To enable file estimation, deselect this option so that before any file is archived to media, Arcserve Backup performs an estimate of how long the job will take.

Be aware of the following consideration:
 - File estimation is no longer the default value.
- **Calculate and Store CRC Value on Archive Media**--Calculating and storing the CRC value on the archive media will enable Arcserve Backup to perform CRC verification during the archive job. To instruct Arcserve Backup to use the CRC value stored on media, see the [Verification](#) (see page 356) section in the Policies tab.

- **Eject Archive Media upon Completion**--Select one of the following options:
 - **Use Default Device Setting**--Select this if you want to use the setting you selected during library configuration.
 - **Eject Media**--Select this if you want to eject media from the drive after the job finishes. This helps prevent any other job from overwriting information on this media. If you select this, it overrides the setting you selected during library configuration.
 - **Do not Eject Media**--Select this if you do not want to eject media from the drive after the job finishes. If you select this, it overrides the setting you selected during library configuration.

Note: For more information on library configuration, see the chapter Managing Devices and Media.

- **Retry Missed Targets**--Reschedule an archive for any workstations, file systems, databases, and so on that failed during the archive job.

You can specify one of the following reschedule options for an archive job:

After Job Finishes

Specifies the number of minutes that you want to elapse after the original job finished to start the makeup job.

Default: 5 minutes

Maximum: 1439 minutes (less than 24*60 minutes)

At

Specifies the time when the makeup is to run.

- **Max Times**--Specifies the maximum number of times to repeat the makeup jobs.

Default: 1 time

Maximum: 12 times

Be aware of the considerations that follow:

- By default, Retry Missed Targets is enabled, After Job Finishes is selected, and the value of Max Times is 1.
- **File system archives**--If the archive job requiring a makeup job consists of file system archives, and the file system contains directories that reside in different volumes, the makeup job archives only the failed volumes or directories. The makeup job does not archive the entire file system if it contained successful volume or directory archives.
- **Child jobs**--The child makeup jobs (makeup of makeup jobs) run based on the value specified for After Job Finishes.
- **Agent-based archives**--If the archive job requiring a makeup job consists of agent-based archives, the makeup job will attempt to archive all of the source selected for the archive. If the makeup job fails after one unsuccessful attempt, Arcserve Backup will create another makeup job that consists of all of the source selected for the original job, and submit the makeup job with a status of Hold.

More information:

[Local Backup Options for UNIX and Linux Agents](#) (see page 152)

Archive Manager Pre/Post Options

Pre and Post options let you run commands before and after jobs execute.

The following describes a command that you can run using Pre and Post options.

- You can use the Pre option to defragment a disk before an archive job starts.

Be aware of the following behavior when using Pre and Post options:

- Arcserve Backup does not support running commands with executables that reside on remote systems.
- Using a Pre option and specifying an exit code prevents the archive operation from starting until after the Pre option process is complete.
- Using a Pre option and specifying an exit code and the Skip Operation option causes Arcserve Backup to skip the archive operation and, if specified, prevents the Post option process from starting.
- Post option processes start unless the following conditions are present:
 - An exit code is specified, the Skip Post Application option is specified, and the result exit code is equal to the exit code specified.
 - The result of the archive operation is equal to the value specified for the Do not run Command if option.
- Pre and Post options specified as global options run commands before a job starts or after a job finishes. Pre and Post options specified as node-level (local) options run commands before a node or after a node is archived.

For example, a user submits an archive job consisting of nodes A and B. A Pre option is specified as a global option and a Pre option is specified for node B. Immediately before the job runs, the global Pre option executes. While the job is running, the Pre option specified for node B executes before node B is archived.

Run Command Before Job Options

Enter the name of the file to be executed on the machine before the job starts.

- **On Exit Code**--Arcserve Backup detects exit codes of other programs. You can specify the following options for a particular exit code:
 - **Run Job Immediately**--The job runs immediately if the selected exit code is returned.
 - **Skip Job**--The job does not run if the appropriate exit code is detected.
 - **Skip Post Application**--Skip any commands specified to run after the job if the appropriate code is detected.
- **Delay in Minutes**--Specify the delay in which Arcserve Backup waits before running a job when the appropriate exit code is detected.

Run Command After Job Options

Enter the name of the file to execute after the job completes.

Do Not Run Command If Options

Specify for a command not to run if Arcserve Backup detects that a job fails, a job is incomplete, or a job is complete.

Note: This option is not available when you are using Arcserve Backup to manage a UNIX or Linux based server.

Run Before/After Command As Options

The User Name and Password corresponds to the system of the host server selected, and is required to check the system privileges on that server.

The user name and password entered into these fields should not be confused with the Arcserve Backup User Name and Password.

Example: Submitting a Job Using Pre and Post Commands

A user submits a job that archives local volume D. The user wants to check and fix errors using `chkdsk.exe` on local volume D before local volume D is archived. After the job is complete, the user wants to generate an error report using `CAAdvReports.exe`.

Pre Command for the node:

The command that follows checks and fixes errors on local volume D before the archive job starts.

```
chkdsk.exe D: /F", On Exit Code = 0, Run operation immediately
```

Post Command for the job:

The command that follows generates an error report and saves it in a specified location.

```
CAAdvReports.exe -reporttype 5 -maxSize 5 -Server DUVDO01 -outfile "C:\Program Files  
(x86)\CA\ARCserve Backup\Reports\Backup Error Report_data.xml" -PastDays 1 -AutoName
```

Archive Manager Job Log Options

The log options determine the level of detail that is included in the log report for the operation. The log options can be set in the following Arcserve Backup managers:

- Backup
- Archive
- Restore
- Compare
- Media Assure & Scan
- Copy
- Count
- Purge

Arcserve Backup provides the following log options:

- **Log all activity**--Record all of the activity that occurs while the job is running in the Job Log.

Note: When you specify Log all activity, Arcserve Backup creates a log file named JobLog_<Job ID>_<Job Name>.Log. With this log file, you can view detailed logging information about the job. Arcserve Backup stores the log files in the following directory:

C:\Program Files\CA\ARCserve Backup\LOG

- **Log summary only**--Record summary information on the job (including source, destination, session number, and totals) and errors.
- **Log disabled**--Do not record any information about this job in the Job Log.

Archive Manager Alert Options

You can use the Alert notification system to send messages about events that appear in the Activity Log during your archive operation. You can specify one or more of the following events for which you want to be notified:

- **Job Completed Successfully**--All of the nodes and drives/shares were processed.
- **Job Incomplete**--Some nodes, drives, or shares were missed.
- **Job Canceled by User**--The user canceled the job.
- **Job Failed**--The job was started but could not be completed.

- **Media not Available**--Media was not available during the execution of a job.
Note: The archive media must be tape media.
- **Format Blank Tape**--A tape was formatted during the execution of a job.
- **Customized Event**--A customized event occurred. To specify this type of event, enter an error, warning, or notification code in the space below the Event drop-box.

Specify one or more of the defined Alert configurations. The default configuration means that you will use whatever is configured in Alert Manager. Click Configure to define further configurations. Arcserve Backup provides the following defined Alert configurations:

- Broadcast
- Pager
Note: Pager options are not supported on Japanese versions of Arcserve Backup.
- SMTP
- SNMP
- Event
- Printer
- E-Mail
- Lotus Notes
- Unicenter TNG

Specify miscellaneous options:

- **Attach Job Log**--Lets you include the job log information in the Alert message. (This option applies for Trouble Tickets and Mail only.)

Note: The list you create using Alert Options is saved with the Job Script and the configuration defined using the Configuration button.

- **Send CSV as attachment**--Lets you save archived file names to a CSV file to be sent in the Alert message.

The CSV file includes the following fields:

- ArchiveDate (Archive Date)
- HostName (Host Name)
- TapeName (Tape Name)
- TapeId (Tape ID)
- TapeSerialNo (Tape Serial Number)
- SessionNumber (Session Number)
- PathName (Path Name)
- FileName (File Name)
- FileModifiedTime (File Modified Time)
- FileSize (File Size)

- **Send alert messages only for master jobs**--Lets Arcserve Backup send you alerts that reference only the master job number in the Alert message. The alert messages do not reference child and subjob numbers. You can specify this option on all jobs, including multiplexing and multistreaming jobs.

Archive Manager Advanced Options

The Advanced option for archiving determines how Arcserve Backup handles the file system extensions when an archive job is running.

The Windows system options are as follows:

Note: The following options are not supported on Windows operating systems that are older than Windows XP.

- **Traverse Directory Junctions and Volume Mount Points**--Selecting this option causes the archive job to traverse the volume or the directory being specified to and gets archived. At the time of restore of this session, you can restore files and directories contained in the referred to volume or directory. When this option is not selected, the archive job does not archive the volume or the directory being referred to by the volume mount point or the directory junction respectively. Therefore, at the time of restore, you cannot restore a file or directory contained in the referred to volume or directory.
- **Archive Mount Points as Part of the volume that they are mounted**--If you select this, the volumes referred to by the Volume Mount Points will be archived as part of the same session as the Volume Mount Points. When this option is not selected, the volumes referred to by the Volume Mount Points are archived as separate sessions. This option is available only when the previous option, Traverse Directory Junctions and Volume Mount Points, is selected.

Note: When you apply the *Traverse Directory Junctions and Volume Mount Points* and *Archive Mount Points as Part of the volume that they are mounted* options to named, mounted volumes that contain virtual hard disks (VHDs), Arcserve Backup creates separate archive sessions for mounted volumes that contain VHDs.

Example: Mounted Volumes that Contain VHDs

A server contains physical disk (C:\) that contains VHDs D:\ and E:\. VHD files (D.vhd and E.vhd) that reside in C:\ are mounted as drive D:\ and drive E:\. Drive D:\ is mounted to C:\MountD, and drive E:\ is mounted to C:\MountE.

If you archive C:\MountD and specify the *Traverse Directory Junctions and Volume Mount Points* option, and the *Archive Mount Points as Part of the volume that they are mounted* option is enabled or disabled, Arcserve Backup creates separate archive sessions for drive D:\ and C:\MountD.

Archive Manager Encryption/Compression Options

Arcserve Backup lets you encrypt, compress, or both encrypt and compress archive data.

Be aware of the following:

- Arcserve Backup does not support compressing and encrypting data on deduplication device groups.

Note: For more information, see [Compression and Encryption with Deduplication](#) (see page 742).

- If you specify encryption and compression options, and the archive destination is a drive that does not support compression, Arcserve Backup encrypts the archive data and does not compress the archive data.

The following options define how Arcserve Backup processes archive data during an archive job.

Session/Encryption Password

- **Session/Encryption password**--Specify a Session/Encryption password to restore this data from media.

If you specify a Session Encryption password, you must specify the password to perform the following operations:

- Restore operations where the encryption, compression, or both were processed at the agent or at the archive server.
- Compare operations where the encryption, compression, or both were processed at the agent or at the archive server.
- Merge and Scan operations where the encryption, compression, or both were processed at the archive server. (You do not need to specify the password to perform Merge and Scan operations where the encryption, compression, or both operations were processed at the agent.)

Note: The Session/Encryption password is not required when you Merge or Scan only the session headers.

- **Save Current Session/Encryption Password to the Arcserve Backup database**--Use this option to save the password to the Arcserve Backup database and enable password management. This option is selected by default. This option is available for both local and global option passwords.

Note: You can modify only the Global Option password from the Session/Encryption password dialog by right clicking on the job in the job queue.

- **Remind to change password n days after specifying a password**--Specify the number of days a password is valid. Seven days prior to the specified number of days, a message prompting you to change your password will be logged in the Activity Log.

Example:

On Jan. 1 you set n to 30 days. On Jan. 24 the message, "The archive job password will expire in 7 days.", will appear in the Activity Log. On Jan. 31 the message, "The archive job password has been expired." Please change it now appears in the Activity Log.

Encryption/Compression methods

- **Encrypt data**--Use this option to encrypt the archive data. You can specify one of the following options:
 - **At agent**--Select this option to encrypt the archive data prior to the actual archive process. For more information about this option, see [Data Encryption at the Agent Server](#) (see page 113).
 - **At backup server during archive**--Select this option to encrypt the archive data at the backup server during the archive process.
- **Compress data**--Use this option to compress the archive data. You can specify one of the following options:
 - **At agent**--Select this option to compress the archive data on the system where the agent is installed and running.

Note: Arcserve Backup does not support data compression at the agent system when the archive source consists of UNIX.

- **At backup server**--Select this option to compress the archive data at the Arcserve Backup server during the archive process. This option lets you compress files before archiving by using a software compression algorithm.

Be aware of the following behavior:

- You must specify Encrypt data at backup server during archive to enable compression at backup server.
- With the 'at Backup Server' option specified in the Compress data section and the 'at Backup Server during archive' option specified in the Encrypt data section, Arcserve Backup uses software compression to compress the data at the backup server before the data is encrypted.
- If the storage device associated with the job does not support hardware compression, Arcserve Backup ignores the 'at Backup Server' option selected in the Compress data section.

Chapter 7: Managing Devices and Media

This section contains the following topics:

[Device Management Tools](#) (see page 373)

[Device Manager](#) (see page 395)

[How Media Pools Work](#) (see page 445)

[Media Management Administrator \(MM Admin\)](#) (see page 457)

[MM Admin Interface](#) (see page 458)

[How the Media Management Process Works](#) (see page 465)

[How to Manage Cloud Storage](#) (see page 477)

Device Management Tools

Arcserve Backup provides a number of ways to help you manage, monitor, and maintain your devices and media:

- The Device Manager gives you information about storage devices connected to your system, the media in these devices, and the status of these devices. It is the starting point for all media and device monitoring and maintenance operations.
- The Media Pool Manager lets you create, modify, delete, and manage media pools, collections of media managed as a unit to help you organize and protect your media.
- The Media Management Administrator (MMO) provides the tools you need to control, manage, and protect media resources.

Tape Library Configuration

The Tape Library configuration option lets you configure a single-drive tape libraries in Windows environments.

The following sections describe the tasks that you can perform to fully configure your library.

Note: For information about working with and configuring multiple-drive tape libraries and Tape RAID libraries, see the *Tape Library Option Guide*.

More information:

[Configure Devices Using the Device Wizard](#) (see page 116)

Device Assignment

Assigning a drive to a library allows Arcserve Backup to recognize the drive's existence within the library.

Usually the manufacturer configures a library in such a way that the first library drive has the lowest SCSI ID number and the last library drive has the highest SCSI ID number.

Note: This is not always the case. See the documentation that came with your library for information on how its drives are configured.

To manually assign a drive to a library, highlight the drive you want to assign from the Available Devices list and the library in which the drive should reside from the Library Devices list then use the Assign button to move the drive to the library. To un-assign a drive from a library, highlight the drive in the Library Devices list and click Remove.

Note: All drives must be empty for Arcserve Backup to complete the drive configuration. The process may take a few minutes, depending on the number of drives in your library.

Configure Libraries

Arcserve Backup automatically detects and configures your libraries as the Tape Engine starts. You do not need to run a wizard or other external application to enable Arcserve Backup to detect your libraries.

Note: If Arcserve Backup does not automatically configure your libraries, use Device Configuration to manually configure your libraries.

To configure a library, verify that the following prerequisite tasks are complete:

1. Install the Arcserve Backup base product.
2. Install the license for the Arcserve Backup Tape Library Option as required for your environment.
3. Start the Tape Engine.

Arcserve Backup automatically detects and configures your libraries.

4. If you want Arcserve Backup to read the data on the tapes, do the following:

- a. Open Arcserve Backup Device Manager
- b. Browse to and select the library.
- c. Click Inventory on the toolbar.

Arcserve Backup reads the tapes.

To configure libraries

1. Open the Device Manager window and browse to the library.

Right-click the library and select Library Properties from the pop-up menu.

The Library Properties dialog opens.

2. Click the General tab.

Modify the following General options as required for your library:

- **Bar code reader installed**--If your library contains a bar code reader, this option lets you use the bar code reader in the device to inventory the tapes in the library.

- **Set unknown bar code media to not inventoried during initialization**--To enable this option, you must select the Bar code reader installed option.

This option lets Arcserve Backup initialize faster by designating media with a bar code that is not recorded in the Arcserve Backup database as "Not Inventoried." This option prevents Arcserve Backup from inventorying the "not inventoried" slots as the Tape Engine starts. Media that is designated as not inventoried can remain in its slot until you need it. To use media that is designated as "Not Inventoried," you must inventory the media using the Manual Inventory option from the Device Manager window.

- **Bar code reader not installed**--Specify this option if your library does not contain a bar code reader.

- **Library Quick Initialization**--To enable this option, you must select the Bar code reader not installed option.

This option is designed for libraries that cannot read bar codes. With this option enabled, Arcserve Backup retains information about the library's slots in the Arcserve Backup database. As a result, Arcserve Backup does not repeat the inventory process when the Tape Engine is restarted. Arcserve Backup ignores this option on libraries that contain a bar code reader.

Note: If the library does not support bar codes and this option is disabled, Arcserve Backup inventories the entire library when Arcserve Backup starts.

This option lets Arcserve Backup initialize faster by bypassing the inventory slots process when the Tape Engine starts. When you use this option, Arcserve Backup assumes that the media in the slot have not been added, removed, moved, or swapped since the last shutdown. If you added, removed, moved, or swapped media, you should manually inventory the entire library or inventory the slots that changed.

Note: Arcserve Backup must inventory the library after you configure the library. The quick initialization option takes affect after you complete the first full inventory of the library.

- **Eject media upon backup job completion**--Lets you direct Arcserve Backup to move the tapes back to their original slots after the backup job is complete rather than allow them to remain in the drives.

Note: You can override this option on a job-by-job basis by enabling the global option for jobs called Do not Eject Media. In addition, if you do not enable the ejection of media after a backup job completes and later decide that you want to eject media after a particular job, you can enable the global option for jobs called Eject Media.

- **Library is a VTL**--This option lets you set up a library to function as a virtual tape library (VTL).

Be aware of the following behavior:

- Arcserve Backup ignores media expiration dates when you select this option.
 - Read performance improves when you identify a library as a VTL. This capability lets Arcserve Backup maximize drive efficiency and overall VTL backup and data migration performance.
 - You should not identify a physical library as a VTL. When you identify a physical library as a VTL, the library's backup and data migration performance can be adversely affected.
- **Monitor blank tapes**--Lets you log a warning message in the Activity Log when the number of blank tapes available to the library is less than the specified value. This option can be applied to single and multiple drive libraries.

Default value--1

Range--1 to 65535

3. Click the Cleaning tab.

Modify the following Cleaning options as required for your library:

- **Clean by slot**--Lets you designate specific slots as cleaning slots. You can specify one or more cleaning slots and they do not need to be in a successive order.

- **Clean by Barcode**--Lets you specify cleaning slots for your library based on a specific bar code or a range of bar codes using a prefix and a wildcard character. In the Clean Bar Code Prefixes field, enter the prefixes of your bar coded cleaning tapes.

Specify the bar code prefixes into the Clean Barcode Prefix(s) field.

Note: The asterisk is a wildcard character.

Click OK.

The cleaning slots are set based on their bar code prefix.

Examples:

- The bar code on your cleaning tape is CLN123. In the Clean Barcode Prefix(s) field, specify CLN123.
- There are several cleaning tapes in your library. The bar code prefix for the cleaning tapes is ABC. In the Clean Barcode Prefix(s) field, specify ABC*.
- There are several cleaning tapes in your library. The cleaning tapes' bar code prefixes are ABC, CLN1, and MX. In the Clean Barcode Prefix(s) field, specify ABC*; CLN1*; MX*.
- **Automatic tape cleaning**--Lets you direct Arcserve Backup to manage your tape cleaning tasks automatically. When you enable this option you must specify the number of hours that must elapse between cleaning tasks.

4. Click OK.

The library is configured.

Specify the Block Size for Backups to Tape Drives

Arcserve Backup lets you specify the block size for tape backups. This capability lets Arcserve Backup transfer more data per block to the device. Before you specify a larger tape block size, consider the following circumstances:

- The largest default tape block that Arcserve Backup can transfer is 64 KB. You can increase the tape block size to 128 KB, 256 KB, or 512 KB when using tape drives in your backup environment that support larger tape blocks.

Note: See the device and host bus adapter (HBA) manufacturer documentation to verify that the device supports larger data transfers.

- Large block backups require HBAs that support large transfer sizes. When Arcserve Backup detects the HBA cannot support the data transfer size required for the larger block size, Arcserve Backup disables the larger block size and records an error message in the Tape engine log file.
- When Arcserve Backup overwrites or uses blank tapes, Arcserve Backup formats the media using the newly defined tape block size.
- When Arcserve Backup appends data to media, Arcserve Backup writes the appended data to the media using the block size originally applied to the media.
- For SAN attached libraries, verify that the HBA for all SAN member servers can support the transfer size that the tape block size requires. Backups to SAN member servers that do not support the larger transfer size usually fail.

To specify the tape block size for tape drives

1. Stop the Arcserve Backup Tape Engine service on the backup server that you want to configure.
2. Open Windows Registry Editor and browse to the following:
 - **x86 platforms:**
HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA ARCServe Backup\Base\TapeEngine
 - **x64 platforms:**
HKEY_LOCAL_MACHINE\SOFTWAREWow6432Node\ComputerAssociates\CA ARCServe Backup\Base\TapeEngineLocate the DEVICE# key for the tape drive that you want to configure.
3. Under the DEVICE# key, create a DWORD registry key named ForceBlockSize. Specify one of the following DWORD values:
 - 128 KB: 0x80
 - 256 KB: 0x100
 - 512 KB: 0x200Close the Edit DWORD Value dialog and close Windows Registry Editor.
4. Restart the Arcserve Backup Tape Engine service.

The tape block size is applied to the tape drive.

RAID Device Configuration Option

The RAID Device configuration option lets configure a RAID device in the Windows environment.

To configure a RAID Device, the Tape Engine must be stopped. If your Tape Engine is running, a pop-up window is displayed to allow you to stop the engine.

The following sections describe the steps required to fully configure your RAID device.

Note: For information about configuring Tape RAID device, see the *Tape Library Option Guide*.

More information:

[Configure Devices Using the Device Wizard](#) (see page 116)

RAID Level Configuration

Choose the RAID device from the Device Configuration dialog.

When you click Next, the RAID Option Setup dialog appears, enabling you to:

- Create a new RAID device
- Assign a RAID level
- Delete an existing RAID
- Change the RAID level

To review the attributes of each RAID level, instructions on selecting a RAID level, and instructions on assigning drives to the RAID device, see the *Tape Library Option Guide*.

RAID Group Configuration

The RAID device must be added to a group in the Device Manager to perform backup, restore, and copy operations using that RAID device. When running the job, Arcserve Backup automatically assigns a RAID device to a group, if it is not already assigned.

For information on how to manually assign a RAID device to a RAID group, see the *Tape Library Option Guide*.

Virtual Library Configuration Option

The Virtual Library configuration option lets you configure or modify the configuration of a virtual library in the Windows environment.

To configure a virtual library, the Tape Engine must be stopped. If your Tape Engine is running, a pop-up window is displayed to allow you to stop the engine.

The Virtual Library option is used to set up all virtual libraries. The option lets you define virtual libraries and their parameters, including the number of slots and drives required for the library. At minimum, a virtual library must have at least one slot and one drive associated with it.

Because the Virtual Library feature works on existing configured libraries, you must install the Arcserve Backup Tape Library Option and configure your physical libraries before configuring virtual libraries.

Note: You must separate WORM (Write Once Read Many) and non-WORM media in the same library using the Virtual Library configuration option. If WORM and regular media are not separated, the Job Manager treats all media as WORM media. However, the Device Manager can manage these media correctly.

For more information see [Configure Libraries to Function as VTLs](#) (see page 418).

More information:

[Configure Devices Using the Device Wizard](#) (see page 116)


Control Devices Using Removable Storage Management

The Enable/Disable Devices (for RSM) configuration option lets you enable or disable devices for Removable Storage Management (RSM) in a Windows Server 2003 environment.

Be aware of the following considerations and expected behavior:

- To enable or disable devices for RSM, the Tape Engine must be stopped. If your Tape Engine is running, a pop-up window displays that lets you stop the Tape Engine.
- Servers Windows Server 2003 take control of all the devices attached to the server when the Removable Storage service is enabled. This service manages removable media, drives and libraries. To control these devices Arcserve Backup must have exclusive access to them.

- By default, RSM maintains exclusive control of all devices. When the Tape Engine starts, Arcserve Backup detects all devices under the control of RSM and attempts to obtain exclusive control of the devices by disabling the devices in RSM. However, Arcserve Backup can obtain exclusive control from RSM only if the devices are not being used by other applications. If RSM is not running when the Tape Engine starts, Arcserve Backup detects the devices but Arcserve Backup cannot disable RSM's control of the devices. As a result, RSM obtains exclusive control of the devices the next time RSM starts. To help ensure that Arcserve Backup can obtain exclusive control of RSM devices, you must specify the Enable/Disable Devices (for RSM) option and disable the devices. You can access the Enable/Disable Devices (for RSM) from Device Configuration as illustrated by the following screen:



- When RSM has exclusive control of device, Arcserve Backup cannot send SCSI commands directly to the device. However, when Arcserve Backup has exclusive control of a device, it can communicate (input and output commands) directly to the device.
- When you choose the Enable/Disable Devices (for RSM) option, you are provided with a list of all the devices that are currently available in the system. Arcserve Backup manages the devices currently selected. If you want another application to manage any device, clear the selected device.
- You do not need to disable a device in the RSM if the device driver is not installed on the Arcserve Backup server. RSM functions in this manner because the lack of a device driver prevents RSM from detecting the device. Arcserve Backup does not require the presence of a device driver to be able to detect a device.

More information:

[Configure Devices Using the Device Wizard](#) (see page 116)

Configure Devices Using Enterprise Module Configuration

Enterprise Module Configuration is a wizard-like application that lets you configure the following devices:

- **StorageTek ACSLS**--The StorageTek ACSLS configuration option lets you configure or modify the configuration of a StorageTek ACSLS Library. With this option, the Arcserve Backup server can interface with the StorageTek ACSLS libraries to manage backup and restore operations, tape volume movement, and tape volume organization.

To configure StorageTek ACSLS library, ensure that it is properly installed and running before you start Enterprise Module Configuration.

For information about using StorageTek ACSLS libraries with Arcserve Backup, see the *Enterprise Module Guide*.

- **IBM 3494**--The IBM 3494 configuration option lets you configure or modify the configuration of an IBM 3494 library. With this option, you can use the full capabilities of Arcserve Backup with the large tape volume capacities of the IBM® TotalStorage® Enterprise Automated Tape Library 3494.

To configure IBM 3494 Libraries, ensure that the following configurations are complete before you start Enterprise Module Configuration:

- All libraries are properly attached to your network.
- IBM 3494 Automated Tape Library software is installed on the primary server.

For information about using IBM 3494 libraries with Arcserve Backup, see the *Enterprise Module Guide*.

- **Image Option**--The Image Option configuration option lets you install a driver on target systems to enable Image Option capabilities. With this option, you can perform high-speed backups by bypassing the file system, creating a snapshot image of the drive, and reading data blocks from the disk.

For information about backing up and restoring data using the Image Option, see the *Enterprise Module Guide*.

To configure devices using Enterprise Module Configuration

1. From the Windows Start menu, click Start, point to Programs (or All Programs), CA, ARCserve Backup, and click Enterprise Module Configuration.

The Enterprise Module Configuration, Options dialog opens.







2. Click the button for the device that you want to configure, follow the prompts on the subsequent dialogs, and complete all required information.

Device Configuration

With Device Configuration, you can create file system devices, data mover file system devices, and deduplication devices using a wizard application, or you can modify the configuration of existing devices within the Windows environment. Devices are configured to a folder on a specific shared drive. When you specify the device as your backup destination, each session is stored as an individual file within that folder.

Arcserve Backup lets you configure disk-based devices without stopping the Tape Engine. When you configure devices, you can change the credentials of devices used for remote access by clicking Security from the Disk-Based Device Configuration dialog.

From Device Configuration, you can add one or many devices. When you click Next, Arcserve Backup verifies the validity of information specified for all devices and displays the results as tool tips. Point your mouse to the icon in the Device Name column. Device status is indicated by the series of icons described in the following table:

Icon	Description	Function
	Pending	Displayed while a device is being created or edited.
	Verifying	Displayed while a device is being verified.
	Passed	Displayed when a device passes verification.
	Failed	Displayed when a device fails verification.
	Warning	Displayed when a device passes verification, but requires corrections.
	Ready	Displayed when a device does not change and is ready for use.

When the status displayed is Failed:

- Ensure the paths specified for the Location are unique for each device.
- Ensure the security credentials are accurate.
- Check that the volume is shared.
- Check that the path specified for the location is valid.

Staging backup operations can quickly consume a large amount of free disk space on file system devices. Due to the maximum file size limitations of FAT 16 and FAT 32 file systems, you should not use these file systems on file system devices designated for staging operations.

You can specify the Location of the FSD using any of the following formats:

- To specify a path to a local folder, use the format that follows:

`c:\fs_drive`

- To specify a path to a folder that resides on a mapped drive, use the format that follows:

`k:\fs_drive`

Arcserve Backup converts the path and prompts you for credentials when you are finished editing.

- To use a shared folder as a File System Device that you want to access through the network, specify the Universal Naming Convention (UNC) path of the shared folder, for example, `\\SERVER1\fs_drive\`. For UNC path, we encourage users to always put the correct credentials in the Security Window. Incorrect credentials may lead to unexpected results on some platforms, such as Windows 2008, Vista, and others.

When you use a mapped drive as an FSD:

- Arcserve Backup can use the mapped drive, but you must have logged in to the mapped drive previously.
- If you use a mapped drive for an FSD, Arcserve Backup converts the mapped drive to a universal naming convention (UNC) path and prompts you to provide log in credentials when you click Finish.
 - The log in credentials provided must enable full access to the mapped drive.
 - By default, Arcserve Backup uses the Arcserve Backup System Account to gain access to each remote FSD. You can change the credentials that you use with the selected file system device by using Security.
 - You do not need to provide credentials when you create an FSD using a local disk.

Important! Arcserve Backup does not support sharing an FSD with multiple Arcserve Backup servers. When an FSD is shared, the Arcserve servers using the FSD can overwrite the other server's backup data.

Note: Arcserve Backup supports configuring an aggregate total of 255 FSDs and DDDs (only if the number of physical devices configured is 0).

Create File System Devices

File System Devices (FSDs) can be used as backup destinations in normal or staging jobs. If you want to create deduplication devices, see [Create Deduplication Devices](#) (see page 386).

Note: When you create FSDs on NTFS volumes that have NTFS compression enabled, Arcserve Backup turns compression off on the FSD folder. Arcserve Backup behaves in this manner to help prevent problems with NTFS compression and the large files that can be created when backing up data.

To create file system devices

1. Open the Arcserve Backup Manager Console.

From the Navigation Bar, expand Administration and click Device Configuration.

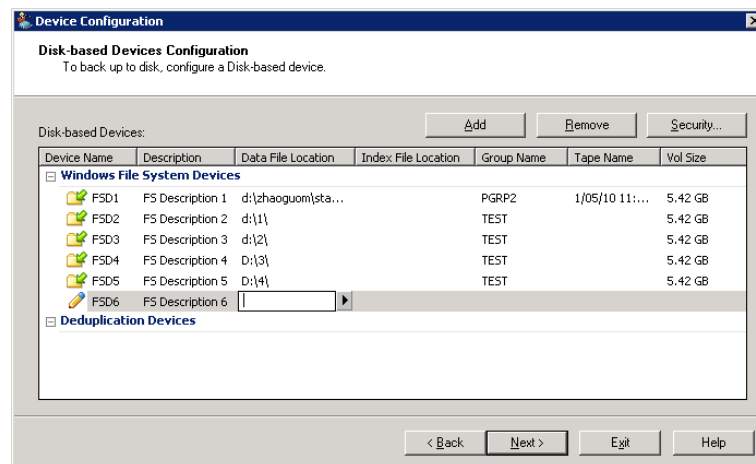
Device Configuration opens.

2. On the Welcome to Device Configuration dialog, select the Disk-based Devices option and click Next.
3. From the Login Server dialog, provide the required security credentials for the Primary Server and click Next.
4. From the second Login Server dialog, select the desired server and click Next.

The Disk-Based Device Configuration dialog opens, showing separate branches in the tree for Windows File System Devices and Deduplication Devices.

5. Click the branch for the type of device you wish to create, for example, Windows File System Devices, and then click Add.

A new, blank device is added to the appropriate branch of the tree.



6. Complete device configuration, as follows:
 - a. In the Device Name field, enter a name or accept the default.
 - b. In the Description field, enter a description or accept the default.
 - c. In the Data File Location field, type a location or click the Browse button to search for one.
 - d. In the Group Name field, enter a name.

Note: The Index File Location field applies only to deduplication devices and is not available when creating FSDs. The Tape Name and Vol Size fields auto-fill when verification completes successfully.

Add More than One File System Device to a Group

To add multiple file system devices to the same device group, the file device type you specify in Device Configuration should be the same for each device you want to include in the group. You can also use Configure Groups to place multiple file system devices in the same group after the file system devices have been created.

Note: You can place only one deduplication device in a deduplication device group.

Device Commands for File System Devices

The device commands that are available for file system devices are:

- **Format**--Deletes the sessions from that folder.
- **Erase**--Deletes the sessions and writes a blank header file on that folder.

The device commands that are not available for file system devices are:

- Retension
- Compression
- Eject
- Long erase

How to Create Deduplication Devices

Deduplication Devices (DDD) can be used as backup destinations in normal or staging jobs. For more information, see [Create Data Deduplication Devices](#) (see page 735).

Note: If you wish to create file system devices, see [Create File System Devices](#) (see page 385).

Remove Disk-Based Devices

When file system or deduplication devices break down, or when you no longer wish to use a device, you can remove the device from Arcserve Backup.

Note: The following procedure applies to file system and deduplication devices.

To remove disk-based devices

1. Launch Device Configuration and choose Disk-Based Devices.
2. Click Next.
The Login Server screen opens.
3. Specify the domain name, primary server name, and authentication type.
Enter your user name and password and click Next.
Specify the server on which the disk-based device is to be removed and click Next.
The Disk-Based Device Configuration screen opens.
4. From the desired branch of the tree (File System or Deduplication Device), click the device you wish to remove to select it.
Click Remove.
The selected device is flagged for removal. If you change your mind and wish to keep the device, click Cancel Remove.
5. Click Next and review results. The device you removed is shown as Delete - Successful in the Report column.
6. Click Next to remove more devices or click Exit to end device configuration.
The disk-based device is removed.

Change Disk-Based Devices

Arcserve Backup lets you change a disk-based device name, description, and data or index file locations (deduplication devices only). However, you may not modify the device's group or tape names here.

Note: The following procedure applies to File System and Deduplication Devices.

To change disk-based devices

1. From the Arcserve Backup Manager Console, launch Device Configuration.
The Device Configuration screen opens.
2. Select Disk-Based Devices and click Next.
The Login Server screen opens.

3. Specify the domain name, primary server name, and authentication type.

Enter your user name and password and click Next.

Specify the server on which the disk-based device you wish to change is connected and click Next.


The Disk-Based Device Configuration screen opens.

4. From the appropriate branch of the tree, click the device you wish to change to select it.

For File System Devices, you may change the Device Name, Description or Data File Location.

For Deduplication Devices, you may change the Device Name, Description, Data File or Index File Locations.

Note: You cannot change the Group and Tape Names from this screen. Use Device Group Configuration to change the Group Name.



5. Specify new security settings. If you changed remote locations for the Data File or Index File Locations, click Security and provide the required User Name, Domain and Password. Retype the password to confirm.
6. Click Next.
7. Click Exit to leave Device Configuration if you are done modifying devices.

Configure Device Groups

Group Configuration is a utility that allows you to create, rename, and delete groups, assign or remove devices from groups, and set the properties of groups, such as those needed for staging or deduplication jobs.

From the Device Group Configuration wizard, there are two options:

- [Configure groups](#) (see page 389) -- use this option to create, rename, or delete groups, or assign or remove devices from groups.
- [Set Device Group Properties](#) (see page 392) -- use this option to set properties for deduplication groups or staging groups.

Configure Groups

The Configure groups option of the Device Group Configuration wizard lets you create, rename and delete groups as well as assign or remove devices from groups.

When disk-based devices are created, they are automatically added to groups. However, for times when you need to reassign devices, such as for hardware maintenance or replacement, you can add a new (empty) group and assign devices to it later, or swap devices among existing groups.

Add a New (Empty) Disk-Based Device Group

You can create a new regular group and then assign a disk-based device to it, making the regular group a deduplication device group or a file system device group, accordingly. If an existing device is busy or damaged, but the group has been specified in various backup jobs, you can remove the device from a particular group and put a new device in its place.

Note: You can assign multiple file system devices to a group, but only one deduplication device to a group.

To add a new empty group

1. From the Arcserve Backup Manager Console, launch Device Group Configuration. Device Group Configuration appears.
2. Click Next.
3. On the Login Page, provide credentials and click Next.
4. On the Options dialog:
 - a. select the server you would like to configure,
 - b. choose Configure Groups,
 - c. click Next.

5. From the Device Group Configuration dialog, click New.

The New Group dialog opens.

6. On the New Group dialog,
 - a. Enter a Name for the new group,
 - b. click OK.

The new regular group appears in the Empty Groups list, but contains no device. You may now assign an available disk-based device to this group.

Assign Disk-Based Devices to Groups

You can assign devices to groups using Device Group Configuration. If you wish to assign a deduplication device to a group, the desired group must be empty. Only one deduplication device can be assigned to a group.

If Device Group Configuration is not already running, launch it from the Arcserve Backup Manager Console.

To assign disk-based devices to groups

1. From Device Group Configuration, click a Group from the list of available groups on the left.
2. From the list of Available Devices on the right, click a device to select it.

Note: If there are no Available Devices to choose from, remove a device from another group or delete an existing device groups. Devices in deleted groups are moved to the Available Devices list.

3. Click Assign.
The device is added to the group you selected.
4. Click Finish when you are done assigning devices.
5. Click Exit to leave Device Group Configuration.

You can convert an existing File System Device group to a Deduplication Device group by removing the FSD devices from it and adding a deduplication device to it. Similarly, you can convert a deduplication group to a normal FSD group in the same manner.

Remove Disk-Based Devices from Groups

You can remove file system or deduplication devices from groups to reassign the devices elsewhere.

To remove disk-based devices from groups

1. From Device Group Configuration, click a Group from the list of available groups on the left.
2. Select the device inside the group to select it.

3. Click Remove.

The device is removed from the Group and added to the list of Available Devices.

4. Click OK when you are done removing devices.

You may reassign devices you have removed to other device groups.

Delete Disk-Based Device Groups

You may delete file system or deduplication device groups. Devices assigned to deleted groups are moved to the list of available devices for reassignment.

To delete disk-based device groups

1. Launch Device Group Configuration.

Device Group Configuration appears.

2. Click Next.

3. On the Login Page, complete the required fields and click Next.

4. On the Options dialog, select the server you would like to configure, choose Configure Groups and then click Next.

5. From the Device Group Configuration dialog, click a Group from the Groups list to select it.

6. Click Delete.

A confirmation message appears.

7. Click OK to continue.

The selected group is deleted. The device assigned to the deleted group is moved to the Available Devices List.

8. Click Finish if you are done deleting groups.

9. Click Exit to leave Device Group Configuration.

10. Click Yes to clear the confirmation message.

Rename Disk-Based Device Groups

If you wish to rename an existing File System or Deduplication device group, you may do so from Device Group Configuration.

To rename disk-based device groups

1. Launch Device Group Configuration.

Device Group Configuration appears.

2. Click Next.

3. On the Login Page, complete the required fields and click Next.
4. On the Options dialog, select the server you would like to configure, choose Configure Groups and then click Next.
5. From the Device Group Configuration dialog, click a Group from the Groups list to select it.
6. Click Rename.
The Rename Group dialog opens.
7. From the Rename Group screen, specify a new name for the device group.
Click OK
The name is changed and the device previously assigned to the group remains unchanged.

Set Device Group Properties

Use the Set Device Group Properties option to set properties for groups. You can set staging properties for File System Devices or deduplication properties for Deduplication Devices.

To set disk-based group properties

1. Launch Device Group Configuration and click Next.
2. Specify the Primary Server and Authentication Type, enter the required security credentials and click Next.
3. Choose the Set Device Group Properties option and click Next.
4. For File System Device Groups, click the Enable Staging option to set staging properties. For Deduplication Device Groups, set the deduplication device properties.
5. Click Finish.

Disk-Based Device Group Properties

From the list of Groups displayed, choose the device group you wish to configure and then complete the following fields.

Depending on the group selected, there are several types of properties you can configure:

Deduplication Group Properties

The following options apply to devices configured as deduplication devices.

- **Max Threshold**--Specifies the maximum amount of space that can be used on a disk before a job fails. When the maximum threshold is reached, Arcserve Backup fails the jobs.

Default Value: 80%

The Max Threshold is represented as either a percentage of the total capacity used on the disk, or as the total number of GB or MB used.

- **Max # Streams**--Specifies the maximum number of simultaneous streams to the device.
- **Pause data migration**--Instructs Arcserve Backup to halt the data migration process. This option applies only to deduplication groups used in staging operations.

Default Setting: Disabled

- **Allow optimization in Data Deduplication backups**--Directs Arcserve Backup to examine file header parameters first. The process of identifying natural boundaries and performing hash calculations is performed only on files whose header details have changed since the last backup, greatly enhancing backup throughput.

Default Value: Enabled

Note: You cannot optimize stream-based data (for example, MS SQL or Oracle). If you use Optimization, ensure the Reset Archive bit for backup to deduplication device option on the Operation tab in Global Options is enabled. Failure to reset the archive bits after a backup job means that Optimization considers all files to be changed, even if no changes actually took place. We recommend disabling optimization in rare situations where applications running on the machine being backed up reset the file archive bit and file attributes like modified time.

- **Enable Global Deduplication**--Lets you to perform deduplication on the C:\ drives of different machines.

Note: Arcserve Backup lets you perform Global Deduplication operations on Oracle RMAN sessions.

- **Delayed Disk Reclamation**--Lets you reclaim disk space created by the deduplication process. Delayed disk reclamation reduces the risk of disk fragmentation.
- **Expedited Disk Reclamation**--Lets you immediately reclaim disk space created by the deduplication process. Although expedited disk reclamation improves the performance of disk reclamation, it can introduce disk fragmentation to the device. This option is enabled by default to improve the performance of disk reclamation.

Staging Device Group Options

The following options are available when the option, Enable Staging, is selected:

- **Max Threshold**--Specifies the maximum amount of space that can be used on a disk before a job fails. When the maximum threshold is reached, Arcserve Backup fails the jobs.

Default Value: 80%

The Max Threshold is represented as either a percentage of the total capacity used on the disk, or as the total number of GB or MB used.

- **Purge data when the used disk space exceeds the Max threshold**-- Instructs Arcserve Backup to delete old sessions when the disk space used for backup exceeds the Max Threshold value.
- **Purge to Threshold**--Available only when the Purge data when the used disk space exceeds the Max threshold option is enabled.
- **Max # Streams**--Specifies the maximum number of simultaneous streams to the device.
- **Enable SnapLock for this group**--Available only on devices that support SnapLock technology, this option prevents data purge, or over-write until the specified retention time elapses.
- **Pause Data Migration**--Instructs Arcserve Backup to halt the data migration process. This option applies only to deduplication groups used in staging operations.

Default Value: Disabled

Cloud Group Properties

The following options apply to cloud-based devices:

- **Max # Streams**--Specifies the maximum number of simultaneous streams to the cloud device.

Default Value: 4

More information:

[Back Up Data Using Global Deduplication](#) (see page 750)

Deduplication Device Management

Use Device Configuration to create data deduplication devices, to remove existing deduplication devices, or to change the properties on an existing device. You must have the proper security access to create a data deduplication device on a remote server.

Note: You may also create deduplication devices using the Create Disk-based Devices option from the Device Manager.

More information:

[Remove Disk-Based Devices](#) (see page 387)

[Change Disk-Based Devices](#) (see page 387)

[Protect Deduplication Devices with Arcserve Replication](#) (see page 395)

[Disk-Based Device Group Properties](#) (see page 393)

[Create Data Deduplication Devices](#) (see page 735)

Protect Deduplication Devices with Arcserve Replication

When deduplication devices are installed locally, the deduplication data files are excluded from Arcserve Backup jobs. If you wish to protect the deduplication device itself, you can do so using Arcserve Replication.

With Arcserve Replication, you can create a scenario that replicates the index and data file paths for a deduplication device. For more information, see the topic, [Create Arcserve Replication Scenarios for Deduplication Devices](#). (see page 746)

Device Manager

The Device Manager provides information about storage devices that are connected to your system, the media in these devices, and the status of these devices. When you highlight a storage device or the adapter card it is configured to, summary information is displayed about the adapter card or the storage device, such as the vendor, model name, and board configuration.

If you have more than one storage device connected to your machine, Arcserve Backup lets you separate them into groups. Establishing device groups is the key to the flexibility and efficiency of Arcserve Backup.

By default, Arcserve Backup is installed with each storage device assigned to its own group. If identical storage devices (same make and model) are detected, it automatically places them in the same group. You can use Device Group Configuration to:

- Create a new device group
- Assign a device to a device group (including a RAID group)
- Remove a storage device from a device group
- Rename or delete a device group
- Use a RAID tape set as one unit

Maintenance Tasks

Using the Device Manager, you can perform the following maintenance tasks on your media:

- [Format media](#) (see page 396).
- [Erase data](#) (see page 398).
- [Retension tapes](#) (see page 399).
- [Compress data](#) (see page 400).
- [Eject media](#) (see page 400).
- [Online and Offline drives](#) (see page 401).
- [Rebuild media - RAID devices only](#) (see page 401).
- [Scan device - USB storage devices only](#) (see page 402).

Important! Before you use these options, especially formatting and erasing, make sure you have the right media selected.

Format Media

Although Arcserve Backup automatically formats blank media during a backup job, you can use this option to manually format your media. Formatting writes a new label at the beginning of the media, effectively destroying all existing data on the media.

Note: Use this option with care. After the media is formatted, Arcserve Backup can no longer restore the data and any job sessions associated with the media.

Low level formatting, which is required on most hard drives and some mini cartridge device drives, is not required for drives that Arcserve Backup supports.

To format media

1. Click the Format toolbar button on the Device Manager window.

The Format dialog opens. It displays specific details about the media in your library slots. For example, unformatted media appears as <Blank Media>, and slots reserved for cleaning media do not appear.

Important! File System Devices (FSD) that are part of a staging group cannot be formatted using the Format toolbar button. To prevent accidental formatting of an FSD before the data is migrated to a final destination media, Arcserve Backup disables the Format toolbar button on the Device Manager window. If you want to format the FSD, you can either use the command line (`ca_devmgr`) or disable the staging option for the selected FSD.

2. Select the slot containing the media that you want to format. Assign a New Media Name and an Expiration date to the media that you want to format.

Note: When you assign a New Media Name to a slot, the light icon next to the selected slot turns green. Slots with write-protected media appear in red. These media cannot be formatted. You must specify a New Media Name before formatting the media.

Repeat this step if you want to specify more media.

3. If you want to use the media in a media pool, select a slot with the green light icon and check the Use Rotation option. Then, from the Media Pool drop-down list, select the media pool that you want to use the newly formatted media in. In the Serial No. field, you can accept the default serial number or specify a user-defined serial number. (If no media pool name is defined and the media has an assigned bar-coded serial number, then Arcserve Backup does not overwrite that serial number during the format procedure.)

Note: Click the Apply to all button if you want to use all of the formatted media in a media pool and assign all the media to the same media pool.

4. Assign a name and an expiration date to the media you want to format. You must specify a New Media Name before formatting the media. For more information, see [How to Choose Expiration Dates](#) (see page 404).

5. Click OK.

The Format dialog closes and the following message appears:

"Formatting will erase ALL of your data from the media. Do you want to format the media?"

6. Do one of the following:
 - To start the formatting process, click OK.
Arcserve Backup formats the media.
 - To cancel the formatting process, click Cancel.
Arcserve Backup does not format the media.

Erase Media

Use this option to erase all data from a single media or from multiple media. Arcserve Backup also erases all references to the contents of this media (if any) from the database. When you reformat this media, its physical history (read and write passes) is retained.

You should verify that you have selected the correct media before using the Erase option. Erased data cannot be retrieved. When erasing media, you can choose from the following options:

- **Quick Erase**--Quick Erase effectively erases media. It avoids the time a Long Erase would take (minutes to hours) by overwriting the media label. The media history remains available to Arcserve Backup for tracking purposes.
- **Quick Erase Plus**--This option performs the same operation as Quick Erase, and also erases bar codes and serial numbers. For more information about bar code and serial number cataloging, see Mount and Dismount Option.

Note: If the media you are erasing does not have a serial number or bar code, this option functions in the same manner as the Quick Erase option.

Media erased using the Quick Erase Plus option can no longer be tracked by Arcserve Backup, and information such as the expiration date is no longer carried forward.

- **Long Erase**--Long Erase completely removes all data from media. It takes much longer than a Quick Erase, but the media is literally blank. For security reasons, use the Long Erase option to ensure that all data on your media is erased completely.

Note: The long erase process consumes more time than the quick erase process. This will be apparent when you erase large capacity libraries. Exercise caution when using this option on large capacity libraries.

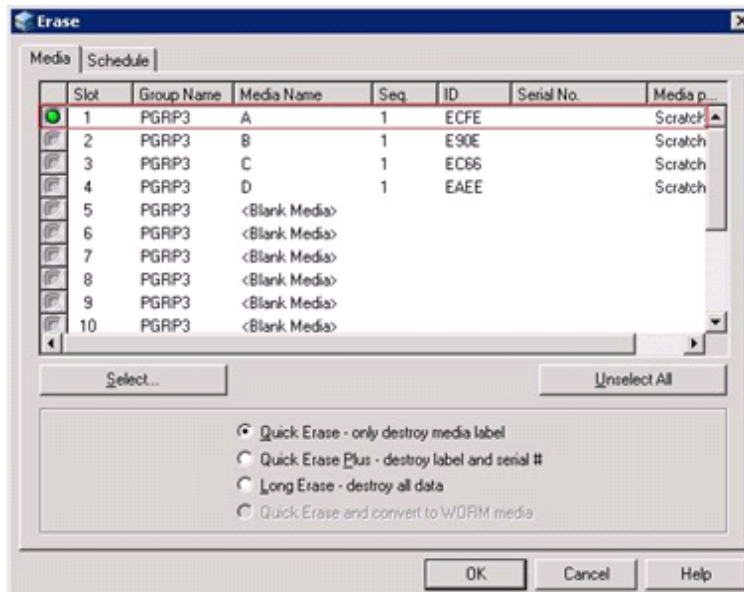
- **Quick Erase and convert to WORM**--This option quickly erases all data from the media. In addition, Arcserve Backup converts the media to Write Once - Read Many (WORM) media.

To use this option, Arcserve Backup must detect DLTWORM capable media in the library or in a stand-alone drive.

To erase media

1. Click the Erase toolbar button.

The Erase dialog opens.



Note: Slots reserved for cleaning media do not appear in the Erase dialog.

2. Select the slot you want to erase. When you select media, the light icon next to the media turns green.

You can press the Shift key to select multiple contiguous media. Press the Ctrl key to select multiple noncontiguous media. You can also click and drag the light icon to select multiple contiguous media.

3. Select an erase method, click OK, and then click OK to confirm.

Arcserve Backup erases the media.

Retension Tapes

The Retension option helps to ensure that tapes are tensioned properly to avoid errors, jamming, or breaking. You should retension media if you are having trouble writing to it or reading from it.

Note: This feature applies only to quarter inch cartridge tapes.

To retention tapes

1. Insert the tape into a storage device.
2. Select that tape.

In the left pane of the Device Manager, expand the tree under the storage device containing the tape.

Then highlight the tape.

Click the Retention on the toolbar.

Click OK.

Arcserve Backup retentions the tape.

Compress Media

Arcserve Backup lets you compress the backup data that is stored on the media. Use the Compression option only if your storage device supports tape compression. If it does not, the Compression toolbar button will be disabled.

Note: Set the Compression option off only if you plan to use a media in another drive that does not support compression. In this case, the drive that does not support compression will not be able to read the compressed data on the media.

To turn compression on or off

1. Open the Device Manager and browse to the library that you want to configure.
2. Select the device drive in the Device Management tree.

If the device drive supports compression, then Arcserve Backup enables the Compression toolbar button. To verify if the device supports compression, select the Detail tab when the device is highlighted.

3. Click Compression on the toolbar.
4. Click OK to set the Compression Mode to Off (if it is On) or On (if compression is Off).

Eject Media

Use this option to eject media from library storage drives and return the media to their home slots (the slot with which the media was associated during the inventory process).

To eject the media from all drives in a library or a single drive

1. Open the Device Manager window.
2. From the Device Manager's devices directory tree, do one of the following:
 - To eject the media from all drives in a library, select the library.
 - To eject the media from a single drive, select the individual drive.

3. To eject the media, do one of the following:
 - Right-click the library or drive and select eject from the pop-up menu.
 - Click the Eject toolbar button.
4. Click OK to confirm.

Arcserve Backup ejects the media.

Online and Offline Drives

You can status library drives as offline or online from the Device Manager by right-clicking on the drive and selecting offline or online, depending on the current state of the drive.

This capability can be useful for marking defective drives in a library as offline, which prevents Arcserve Backup from using the drive until it is repaired and in an online status.

Note: If there is media inside the drive you want to mark as online or offline, eject the media prior to marking the drive offline. Arcserve Backup cannot access the media inside an offline drive.

To online and offline drives

1. Open the Device Manager and browse to the server connected to the library containing the drive that you want to status as online or offline.
2. Expand the library, right-click the drive, and select Online or Offline from the pop-up menu.

The drive status changes to offline or online.

Note: Drives display in a disabled mode when they are in an offline state.

Rebuild Media

Arcserve Backup lets you rebuild one missing or unusable tape containing backup data in a RAID level 5 environment. Due to the architecture of RAID Level 5 (striping with parity), you cannot rebuild more than one missing or defective tape.

To rebuild RAID tape drives

1. Eject the incomplete RAID set, using the eject option from Arcserve Backup Device Manager (choosing 'Eject' while highlighting the RAID ejects all tapes in RAID).
2. Insert a tape that the user wants to use as the replacement for the missing tape in one of the tape drives.
3. Chose Erase from the Device Manager.
Arcserve Backup erases the tape.
4. Insert the incomplete RAID set in the other tape drives and click the Rebuild toolbar button.
Arcserve Backup rebuilds the media.

To rebuild RAID tape libraries

1. If there are no blank tapes in the library, import one tape, or erase an unused tape in the library.
2. Choose the RAID set that you want to rebuild and click the Rebuild toolbar button.
Arcserve Backup rebuilds the media.

Scan Devices

Use the Scan Device option to enumerate USB storage devices that are connected directly to the Arcserve Backup server.

Note: This option applies to USB storage devices only.

To scan a USB storage device using the Scan Device option

1. Open the Device Manager.
2. Connect the USB storage device to the Arcserve Backup server.
3. Select the USB controller icon in the device directory tree and click the Scan Device toolbar button.
Arcserve Backup detects and enumerates the device in the Device Manager, device directory tree.

Important! If the drivers for the USB storage device are not Plug and Play (PnP) compatible, Arcserve Backup may not be able to detect and enumerate the storage device. To resolve this, you must configure the USB storage device by stopping and restarting the Tape Engine.

To scan a USB storage device by stopping and starting the Tape Engine


1. Stop the Tape Engine by doing the following:

- a. From the Quick Start menu, select Server Admin.

The Server Admin Manager opens.

- b. From the server tree, locate and select the primary server.

The Arcserve Backup services appear in the right side of the window.



Name	Status	Up Time (days:hours:minutes)	Description
CA ARCserve Communication Foundation	Started	2 : 23 : 35	Provides data used by CA ARCserve Backup Dashboard.
CA ARCserve Database Engine (ODBC)	Started	2 : 23 : 35	Provides database services for ARCserve Backup products. If thi...
CA ARCserve Discovery Service	Started	2 : 23 : 35	Enables the discovery of all ARCserve Backup products on the ne...
CA ARCserve Domain Server	Started	2 : 23 : 34	Provides the management of domains and authentication service...
CA ARCserve Job Engine	Started	2 : 23 : 35	Maintains and executes jobs from the ARCserve Job Queue. If t...
CA ARCserve Management Service	Started	2 : 23 : 35	Provides remote services for command line utilities.
CA ARCserve Message Engine	Started	2 : 23 : 35	Allows remote management of other ARCserve Services.
CA ARCserve Service Controller	Started	2 : 23 : 35	Enables remote start/stop of ARCserve Backup services.
CA ARCserve Tape Engine	Started	2 : 23 : 34	Manages the configuration and operation of backup devices for ...

- c. Right-click Arcserve Tape Engine and select Stop on the pop-up menu.

The Tape Engine stops.

Note: Do not close the Server Admin Manager.

2. Attach the USB storage device directly to the Arcserve Backup server.

3. Restart the Tape Engine by doing the following:

- a. From the server tree, locate and select the primary server.

The Arcserve Backup services appear in the right side of the window.

- b. Right-click Arcserve Tape Engine and select Start on the pop-up menu.

The Tape Engine starts.

After the Tape Engine starts, Arcserve Backup detects and enumerates the device in the Device Manager device directory tree.

More information:

[Configure USB Storage Devices](#) (see page 425)

Schedule Device Management Jobs

Under Device Management, you can submit a scheduled Format or Erase job. In the Format or Erase dialog, choose Run Now to run and submit the device command now or choose Schedule to submit a device command job to the Arcserve Backup queue and run later. Specify the date and time you want to run the device command.

For more information about the Run Now option, refer to the section Job Queue Tab.

More information:

[How to Manage Jobs Using the Job Queue Tab](#) (see page 323)

Device Management Functions for Libraries

Using the Device Manager, you can perform the following management tasks for your libraries.


- [Inventory the slot range](#) (see page 407).
- [Mount and Dismount a magazine](#) (see page 409).
- [Import and export media](#) (see page 410).
- [Clean media](#) (see page 412).
- [Offline and online removable drives](#) (see page 415).
- [Configure library device groups](#) (see page 419).

How Arcserve Backup Labels Media with Bar Codes or Serial Numbers

Labeling media allows the library to quickly recognize and differentiate one media from the others. Bar code recognition is a library-specific feature. Each media comes from its manufacturer with a bar code label affixed to the outer edge of the media cartridge. This label has a predefined serial number in letters and numerals, which is used as the media serial number when the media is formatted.

If you select a media pool name and the media has an assigned bar coded serial number, that serial number is preserved and the media pool range is ignored.

Note: When the serial number or bar code does not exist on the media, Arcserve Backup displays N/A (not available) on the media description in the Device Manager.



How to Choose Expiration Dates

The expiration date tracks how long media should be in service. The life of media is generally based on passes. A pass is defined as the storage drive head passing over a given point on the media. For example, a backup without verification constitutes one pass, whereas a backup with verification constitutes two passes.

Tape manufacturers rate their tapes' useful lives from about 500 to 1500 passes. This does not mean that the tape is unusable after it reaches the maximum number of passes, only that it is more susceptible to errors at this point.

You should choose expiration dates based on how you plan to use the tape. If you plan to use the tape often, (for example, a few times a week), you should set the expiration date to one year, or less, from the date of formatting. By contrast, if you plan to use the tape only once or twice a month, you can set the expiration date to two or three years from the current date.

When media reaches its expiration date, Arcserve Backup notifies you that you cannot overwrite to expired media. To remedy this condition, you can specify to append the backup data to the expired media by doing the following:

1. Open the Backup Manager and click Options on the toolbar.

The Global Options dialog opens.


2. Select the Backup Media tab.

In the First Backup Media section, click Append, click OK, and resubmit the job.

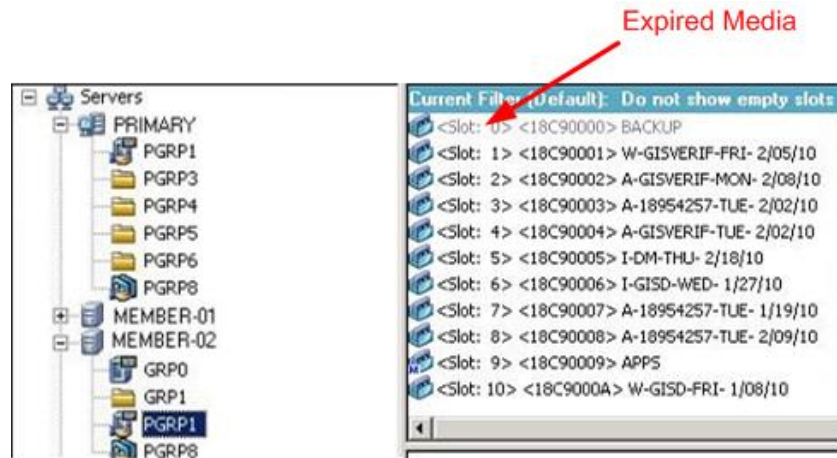
How Expired Media Appears in the Backup Manager and Device Manager

Arcserve Backup can detect expired media.

The following screen illustrates how expired media displays in the Backup Manager, Destination tab.



The following screen illustrates how expired media displays in the Device Manager.



Expiration Dates for New Media

Arcserve Backup formats media using the following guidelines:

- If you are formatting new, blank media, the default expiration date is three years from the current date.
- If you are reformatting media, the expiration date that appears is the date you specified the first time the media was formatted.

How Arcserve Backup Logs Expired Media

Arcserve Backup logs messages in the Activity Log that relate to media that is expired or will expire in certain number of days.

- When the backup job appends the backup data to an expired media, a warning message is displayed as shown below:

```
This job is appending to an expired media. (MEDIA=media_name[S/N:serial_number], ID=media_id,
SEQ=sequence_number)
```

- When a backup job chooses a media to overwrite or append the backup data, it checks the alert period of the media expiration and displays the following message:

```
This job is using media that will expire after <# of days> (MEDIA=media_name[S/N:serial_number],
ID=media_id, SEQ=sequence_number).
```

Where <# of days> represents a specific number days (for example, 3, 5), Media_name represents the name of the media (for example, tape1, Media_id represents the media ID (for example, 3d3c), and Sequence_number represents the sequence number.

Note: This operation applies to both first tape and spanning tape.

- The alert period of the media expiration is 30 days by default. You can change this by adding the DWORD AlertPeriodForTapeExpiration to the following registry key to set the alert period (number of days):

```
\\HKEY_LOCAL_MACHINE\ComputerAssociates\Arcserve
Backup\Base\Task\Backup\AlertPeriodForTapeExpiration
```

Note: This approach only applies to tape media, and you cannot overwrite to an expired media.

Inventory Slots

The Inventory Slots option checks the library slots and reads the media header. It then associates the media header with the slot in which it was found (called the home slot). In this way, the Tape Engine can track changes made to media in the library. For example, media added or removed from a magazine or moved to a different slot.


Be aware of the following considerations when using the Inventory Slots option to inventory media:

- Each media that you load into the storage drives in the libraries must have a unique serial bar code number.
- You should only add and remove media when the Tape Engine is running so that you can immediately inventory your slots.
- For media that was created using a previous Arcserve release, this release automatically creates a new media pool with the same name given to media that was inventoried and used in a media pool in the previous Arcserve release.

To inventory slots

1. Right-click a slot and select Inventory/Offline Slots from the pop-up menu.

The Inventory/Offline Slots dialog opens.



2. Select the slot you want to inventory. Press the Shift key to select multiple contiguous media. Press the Ctrl key to select multiple non-contiguous media. The light icon next to selected media turns green.

Note: You can also click and drag the light icon to select multiple contiguous media.

3. Choose one inventory method:

- **Quick Inventory**--The Tape Engine matches the bar code number to the media serial number, if the library supports bar codes, and the bar code option is enabled. You can only use this method if you are using the bar code option.

- **Regular Inventory**--The Tape Engine reads all the media information from the media.

Note: This method is also known as a Manual Inventory.

- **Offline**--Dismounts the selected slots.

4. Click OK.

Arcserve Backup inventories the slots.

Mount and Dismount Magazines

Use this option to mount (load) or dismount (remove) a magazine from the library. Mounting a magazine initiates an inventory of the slots in the magazine. Dismounting a magazine returns all media to their home slots and prepares the magazine for removal. The time this process requires varies based upon the number of media in the magazine you are mounting or dismounting. Additionally, the time required to mount and dismount magazines can vary from vendor to vendor.

This option checks the library slots and reads the media header. It then associates the media header with the slot in which it was found (its home slot). This enables the Tape Engine to keep track of any changes made to media in the library (media added to or removed from a magazine or moved to a different slot).

If you are using bar codes, each media that you load into a storage drive in the libraries must have a unique serial bar code number. If you purchased two media having identical serial numbers, you must use one of the media in a different backup session.

You should add and remove media only when the Tape Engine server is running, so that you can immediately inventory your slots.

To mount and dismount a magazine

1. Click the Mount toolbar button.

The Mount/Dismount Magazine dialog opens.



2. From the Magazines drop-down list, select the magazine you want to mount or dismount.

Depending on the operation that you want to perform, click one of the following:

- Mount
- Dismount

Arcserve Backup mounts or dismounts the magazine.

Import and Export Media

Arcserve Backup lets you import media and retrieve media information from the media or the Arcserve Backup database. You can also import or export multiple media to or from your library slots.

If the library has mail slots, Arcserve Backup lets you move tapes into and out of the library. You can:

- Import one or more media from mail slots to library slots.
- Export one or more media from library slots to mail slots.

When importing media, you can choose one of the following methods:

- **Quick Import**--Arcserve Backup imports the media and attempts to use the media's bar code information to retrieve the corresponding information from the Arcserve Backup database.

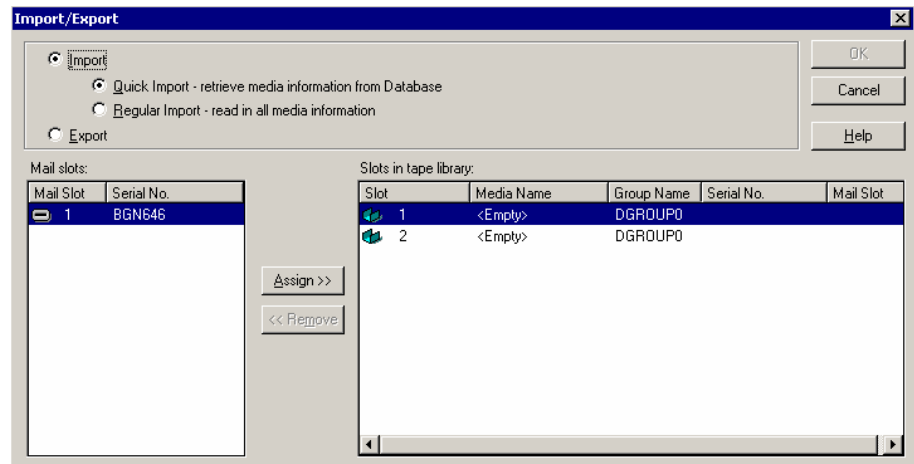
Note: You can only use this method if you are using the bar code option.

- **Regular Import**--Reads all media information from the media itself.

To import media to libraries

1. Open the Device Manager and browse to the library that you want to configure.
Select the library in the Device Management tree.
2. Click Import/Export from the tool bar.

The Import/Export dialog opens.



- Choose Import to view all the available empty slots in a media library.

Note: If your library has a bar code reader, the Serial No. field displays the bar code number of your tape. You can use the Serial No. field to identify tapes located in a specific mail slot.

Select the mail slot containing the media you want to assign to your library.

Select the empty slot to which you want to import the media and click Assign.

Arcserve Backup imports the media into the library.

Note: The best practice is to import cleaning tapes to slots that you have designated as cleaning slots, or set the barcode as a cleaning tape prefix. You can do this from the Library Properties window of the Device Manager. If you import it to a different slot, you may receive unrecognized media errors.

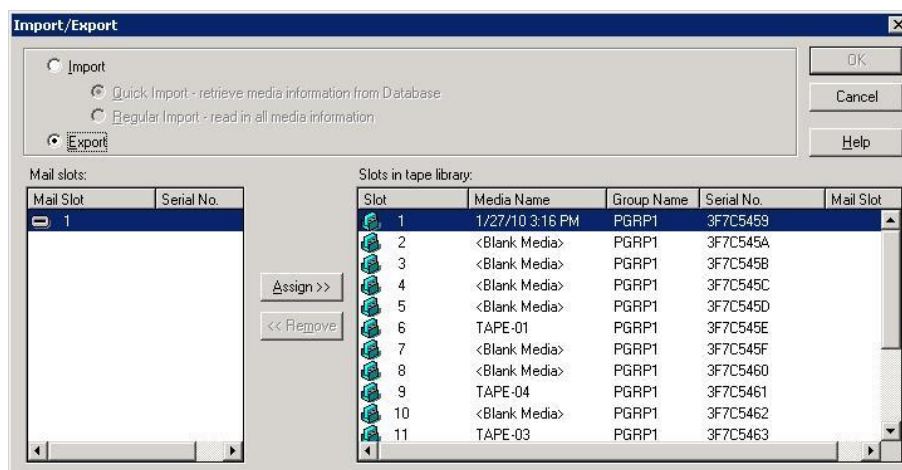
- Repeat the previous step for each media you want to import.
- Choose an Import method and click OK.

The media is imported to the library.

To export media from libraries

- Open the Device Manager and browse to the library that you want to configure.
- Select the library in the Device Management tree.
- Click Import/Export on the toolbar.

The Import/Export dialog opens.



- Select the Export option to view all the occupied slots in a library.

Highlight the media you want to export.

Select the mail slot to export to, and click Assign.

Arcserve Backup exports the media from the library.

5. Repeat the previous step for each media you want to export.
6. Click OK.
The media is exported from the library.

Clean Media


Use this option to clean the tape heads of any media drive in your library.

Note: To use this option, you must have at least one cleaning tape configured in your library.

To clean media (tape heads)

1. Click the Clean toolbar button.

The Clean Tape Head dialog opens.



Note: Offline drives do not appear in the Device Name and Description lists.

2. From the Choose the cleaning tape slot drop-down list, select the cleaning slot that you want to use.

From the Clean tape heads within tape device list, select the drive whose head you want to clean.

Click OK.

Arcserve Backup cleans the tape heads.

More information:

[Configure Libraries](#) (see page 374)

How to Configure Cleaning Slots

This section describes how you can configure more than one cleaning slot.

If supported by your library, you can use Arcserve Backup to specify more than one cleaning slot. You can designate a slot based on the following:

- **Slot number**--This option lets you designate specific slots as cleaning slots. You can specify one or more cleaning slots and they do not need to be in a successive order.
- **Bar code prefix**--This option lets you designate slots based on a bar code prefix.

Example 1: If your cleaning tape bar code number is CLN123, specify "CLN*" as the bar code prefix.

Example 2: If you are using more than one cleaning tape, and their bar codes start with ABC, specify "ABC*" as the bar code prefix.

More information:

[Add Cleaning Slots Based on Slot Number](#) (see page 413)

[Remove Cleaning Slots Based on Slot Number](#) (see page 414)

[Configure Cleaning Slots Based on Bar Code Prefix](#) (see page 414)

Add Cleaning Slots Based on Slot Number

Arcserve Backup lets you add cleaning slots based on the slot number.

To add cleaning slots based on slot number

1. Open the Device Manager and browse to the library that you want to configure.
2. Right-click the library and select Library Properties from the pop-up menu.

The Library Properties dialog opens.

3. Select the Cleaning tab.

The cleaning options display.

4. Select the Clean by Slot option.

From the Available Slots list, select the slot that you want to designate as a cleaning slot and click the Add button.

The available slot is added to the Clean Slots list.

5. Repeat the previous step to add more cleaning slots.
6. Click OK.

You have successfully added cleaning slots based on their slot number.

More information:

[How to Configure Cleaning Slots](#) (see page 413)

Remove Cleaning Slots Based on Slot Number

Arcserve Backup lets you remove cleaning slots based on the slot number.

To remove cleaning slots based on slot number

1. Open the Device Manager and browse to the library that you want to configure.
2. Right-click the library and select Library Properties from the pop-up menu.
The Library Properties dialog opens.
3. Select the Cleaning tab.
The cleaning options display.
4. Select the Clean By Slot option.
From the Available slots list, select the slot that you want to remove.
Click the Remove button to exclude the slot from use as a cleaning slot.
The available slot is removed from the Cleaning Slots list.
5. Repeat the previous step to configure more cleaning slots.
6. Click OK.
The cleaning slots are removed based on their slot numbers.

More information:

[How to Configure Cleaning Slots](#) (see page 413)

Configure Cleaning Slots Based on Bar Code Prefix

The Clean By Bar Code function lets you specify cleaning slots for your library based on a specific bar code or range of bar codes using a prefix and a wildcard character.

To configure cleaning slots based on bar code prefix

1. Open the Device Manager and browse to the library that you want to configure.
2. Right-click the library and select Library Properties from the pop-up menu.
The Library Properties dialog opens.
3. Select the Cleaning tab.
The cleaning options display.

4. Select the Clean by Barcode option.

Specify the bar code prefixes into the Clean Barcode Prefix(s) field.

Note: The asterisk is a wildcard character.

Click OK.

The cleaning slots are set based on their bar code prefixes.

Examples: Clean Bar Code Prefixes

The bar code on your cleaning tape is CLN123. In the Clean Barcode Prefix(s) field, specify CLN123.

There are several cleaning tapes in your library. The bar code prefix for the cleaning tapes is ABC. In the Clean Barcode Prefix(s) field, specify ABC*.

There are several cleaning tapes in your library. The cleaning tapes' bar code prefixes are ABC, CLN1, and MX. In the Clean Barcode Prefix(s) field, specify ABC*; CLN1*; MX*.

More information:

[How to Configure Cleaning Slots](#) (see page 413)

Offline and Online Removable Drives

Arcserve Backup automatically detects removable drives that are connected via Universal Serial Bus (USB) or Serial Advanced Technology Attachment (SATA) to a Arcserve Backup primary or member server. Before you can back up data to a removable drive, you must perform a one-time configuration and then bring the removable drive online.

After you perform a one-time configuration, you can specify removable drives as online or offline from Device Manager by right-clicking on the drive and selecting online or offline (depending on the current state of the drive) from the pop-up menu.

Note: If there is media inside the drive that you want to mark as online or offline, eject the media prior to marking the drive offline. Arcserve Backup cannot access the media inside a drive that is in an offline state.

Arcserve Backup automatically detects and configures removable drives that are connected via USB or SATA to a Arcserve Backup primary or member server. Before you can back up data to a removable drive, you must configure the removable drive and then bring the removable drive online.

To specify a removable drive as online

1. Ensure that the removable drive is attached to a Arcserve Backup domain primary or member server.
2. Open the Device Manager window and expand the Servers object.
Browse to and select the server to which the removable drive is connected.
Arcserve Backup presents you with a list of devices attached to the selected server.
3. From the list of devices attached to the server, select and right-click the removable drive that you want to bring online.
From the pop-menu, select **Online**.
Arcserve Backup prompts you to confirm that you want to configure the device.
Note: This message displays only the first time that you bring the removable drive online.
4. Click OK.
Arcserve Backup prompts you to confirm that you want to bring the device online.
5. Click OK.
The removable drive is now configured and is in an online state.

Use the **Offline** removable drive option when you want to perform maintenance, repairs, or detach a drive from your Arcserve Backup environment. For example:


- You do not want to use the removable drive for a period of time and you do not want to overwrite the media in the drive.
- You want to detach the removable drive from the Arcserve Backup server so that you can replace it with an identical removable drive, or remove it completely from your Arcserve Backup environment.

Important! When the removable drive is offline, jobs associated with the removable drive may fail.

To specify a removable drive as offline

1. Open the Device Manager window, expand the Servers object.
Browse to and select the server to which the removable drive is attached.
Arcserve Backup presents you with a list of devices attached to the selected server.
2. From the list of devices attached to the server, select and right-click the removable drive that you want to take offline.
From the pop-up menu, select **Offline**.
Arcserve Backup prompts you to confirm that you want to take the removable drive offline.
3. Click OK.
The removable drive is now in an offline state in the Device Manager window.

Note: After you mark the removable drive as offline, <Offline> displays in the Device Manager window.



How Device Replacement Works

Situations may arise that require you to repair or replace a device that is connected directly to your Arcserve Backup server (for example, a single-drive library, tape drive, CDROM, and so on).

When you replace a device, Arcserve Backup demonstrates the following behavior:

- After you replace the device with a device that is **different** from the original device and start the Tape Engine, Arcserve Backup assumes that the device is a new device and creates a new device group for the device. Since the replacement device is not associated with the original device group, jobs associated with the original device group will fail.

To remedy the failed jobs, you must reconfigure the jobs associated with the original device group and then resubmit the jobs.

- After you replace the device with a device that is the **same** as the original device and start the Tape Engine, Arcserve Backup assigns the device to the device group where the original device was assigned.

This behavior ensures that jobs associated with the original device group do not fail.

Limitations:

- The replacement device must be a product from the same manufacturer as the original device.
- The replacement device must be the same type of device as the original device (for example, a single-drive library, a tape drive, and so on).
- The replacement device must be connected to the same adapter and channel as the original device.
- The original device must not be assigned to a RAID device group.
- The Arcserve Backup server, where the original device is connected, must not be a member of a SAN domain.

Configure VTLs to Function as Arcserve Backup VTLs

Virtual tape libraries (VTLs) are disk-based devices that are designed to behave like physical tape libraries. To allow Arcserve Backup to back up data to these devices, VTLs, you must configure the devices to function as Arcserve Backup VTLs.

Important! You should not configure standard physical libraries to function as a Arcserve Backup VTLs. The library's backup and data migration performance can be adversely affected when configured to function as a VTL.

Prerequisite Tasks

Before you can configure devices to function as VTLs, verify that the following prerequisite tasks are complete:

- The Tape Library Option is licensed.
- The VTLs are properly configured using Device Configuration.
- Arcserve Backup detects the VTLs.

To configure VTLs to function as Arcserve Backup VTLs

1. From the Administration menu in the Navigation Bar on the Home Page, select Device.

The Device Manager window opens.

2. From the Server directory tree, locate the VTL.

Right-click the VTL and select Library Properties from the pop-up menu.

The Library Properties dialog opens.

3. Select the General tab.

In the VTL (Virtual Tape Library) section, check the Library is a VTL check box and click OK. Arcserve Backup ignores specified media expiration dates when you select this option.

The library is identified as a VTL in the Backup Manager, Destination tab.

The screenshot shows the Device Manager interface. On the left, a tree view shows the server hierarchy: Servers > MABZH02-2K8-2 > PGRP1. The main area displays the configuration for the selected group, PGRP1. At the top, a filter is set to 'Do not show empty slots'. Below this, a grid lists 16 slots with their IDs and contents. A sidebar on the left contains configuration options: 'Configure Groups...', 'Create Disk-based Devices...', 'Set Disk-based Device Group Properties...', and 'Refresh'. The main area shows a table with group and device information.

Group Information	
Group Name	PGRP1
Group Type	VTL Library Group

Device Information	
Library	HP MSL G3 Series
Drive1	HP Ultrium 2-SCSI
Drive2	HP Ultrium 2-SCSI

Note: If you do not want to identify a library as a VTL, repeat the above steps and remove the check mark from the Library is a VTL check box.

Media Movement

When you insert a media into a magazine slot or remove a media from a slot, you must either inventory this slot or remount the magazine.

Important! If you are manually inserting media into a library, always insert media into slots, never into its library drives.

Device Group Configuration Using the Device Manager

Arcserve Backup lets you separate the slots in your library into groups. Grouping slots lets you run several types of jobs at the same time. Additionally, if you have several slots in a group, you can let the library span the media in the group for you.

By default, the first time you start the Tape Engine, all the slots in each library you have attached to your machine are automatically assigned to that group.

After you start Arcserve Backup, you can use the Device Manager to:

- [Create new groups](#) (see page 420).
- [Assign slots to groups](#) (see page 421).
- [Remove slots from groups](#) (see page 423).
- [Delete groups](#) (see page 423).
- [Rename groups](#) (see page 424).

Example: Library Configuration Using the Device Manager

For example, if you have two libraries attached to your machine, you will have two groups—all the slots in the first library are assigned to GROUP0, and all the slots in the second library are assigned to GROUP1. You can retain these group names, or you can regroup and rename them. Since each slot in a library is viewed as a virtual storage drive, each slot can be assigned its own group.

Create New Library Groups

To create a new library group, you must first open the Device Group Configuration dialog. The following are methods you can use to open the Device Group Configuration dialog.

- From the Device Manager, click the Device menu and then click Manage Device Groups.
- From the Device Manager window or the Staging Location tab in the Backup Manager window, click Manage Device Groups located in the device properties preview pane.



Note: To access the Device Group Configuration dialog using a wizard-like application, from any manager window, click the Configuration menu and select Device Group Configuration.

To create new library groups

1. From the Device Manager, click the Device menu and select Manage Device Groups.

The Device Group Configuration dialog opens. Existing groups, and the slots assigned to each group, are listed here. If you have reserved one of your slots for cleaning media, it cannot be assigned to a group and it does not appear in this dialog.

2. Click New.

The New Group dialog opens.

3. Enter a name for the library group and click OK.

The new library group displays in the Groups field. You can now begin assigning slots to this group.

Assign Slots to a Library Group

Arcserve Backup lets you assign specific slots to a library group.

To assign slots to a library group

1. From the Administration menu in the Navigation Bar on the Arcserve Backup Manager Console, select Device Group Configuration.

The Device Group Configuration Welcome dialog opens.

2. Click Next.

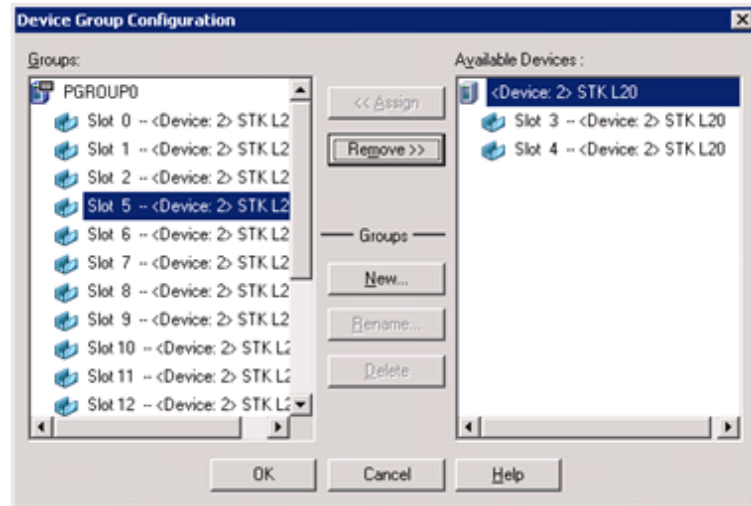
The Login Page dialog opens.

3. Complete the required fields on the Login Page dialog and click Next.

The Options dialog opens.

4. Select the server that you want to configure, click the Configure Groups option, and then click Next.

The Device Group Configuration dialog opens. Library devices and their corresponding slots (available for assignment) display in the Available Devices list.



5. From the Available Devices list, select the slots that you want to assign to a group. You can select one available slot at a time, or you can select the library to assign all of its available slots to a group.
6. From the Groups list, select the group to which you want to assign the slot.
7. Click Assign.

Arcserve Backup removes the slot from the Available Devices list and places it in the Groups list, below the group to which it was assigned.

8. Repeat Steps 5 through 7 to assign more slots to groups.

Note: If there are no slots available, you can remove them from their currently assigned group to make them available to other groups. To do this, from the Groups list, select the slot that you want to make available to other groups and click Remove. The slot is now available to other groups. You can now perform Steps 5 through 7 to assign the slot to a different group.

9. Click Finish and then click Exit to exit Device Group Configuration.

The slots are assigned to the library groups.

Remove Slots from a Library Group

Arcserve Backup lets you remove (unassign) specific slots from a library group.

To remove slots from a library group

1. From the Device Manager, click the Device menu and select Configure Groups.
The Device Group Configuration dialog opens.
2. Highlight the slot you want to remove. Slots are listed in the Groups list beneath the name of the group to which they were assigned.
3. Click Remove.

The slot is removed from the group to which it was assigned in the Groups list and placed in the Available Devices list.

4. Repeat Steps 2 and 3 to remove more slots from groups.
5. Click OK.

The slots are removed from the library groups.

Delete Library Groups

If you no longer require a specific library group, Arcserve Backup lets you delete library groups.

To delete library groups

1. From the Device Manager window, click Configure Groups (from the list of functions) or click the Groups toolbar button.
The Device Group Configuration dialog box opens.
2. Select the group you want to delete.
3. Click Delete, and then click OK to confirm.

The group is removed from the Groups list. Any slots that were assigned to the group are placed in the Available Devices list.

Rename Library Groups

Arcserve Backup lets you rename specific library groups.

To rename library groups

1. From the Device Manager window, click Configure Groups (from the list of functions) or click the Groups toolbar button.

The Device Group Configuration dialog opens.

2. Select the group you want to rename and click Rename.

The Rename Group dialog opens.

3. Specify a new name for the group and click OK.

The new group name is appears in the Groups list.

Universal Serial Bus (USB) Storage Devices

Arcserve Backup can detect the following types of Universal Serial Bus (USB) storage devices that are connected to the Arcserve Backup server:

- Tape drives
- Media changers
- USB removable drives

After you connect the USB storage devices to your Arcserve Backup server, you can use them for all of your backup and restore operations.

Note: If you disconnect USB devices from the Arcserve Backup server and do not restart the Tape Engine after disconnection, you can manually assign the disconnected devices to new groups. These assignments are activated after you reconnect your devices to the server and restart the Tape Engine. If you restart the Tape Engine after disconnecting USB devices from the Arcserve Backup server, you cannot manually assign the disconnected devices to new groups.

Configure USB Storage Devices

Use the Scan Device option to enable Arcserve Backup to detect and enumerate USB storage devices. You can start the Scan Device option by clicking the Scan Device button on the Device Manager toolbar.

Important! The drivers for the USB storage device must be installed on the Arcserve Backup server in order for Arcserve Backup to detect and communicate with the devices.

Note: For more information about configuring USB storage devices, see Scan Device Option.

More information:

[Scan Devices](#) (see page 402)

Prerequisites for Backing Up to Removable Drives

Before you can back up to removable drives, you must:

- Ensure that the media is formatted to the NTFS or FAT32 file system.
Note: If you need to format or reformat the media, consult the manufacturer's documentation for formatting guidelines, or use a Windows-based application to format the media.
- Attach the removable drive to a Arcserve Backup domain primary or member server.
- Status the removable drive as online.

Format Removable Media

After Arcserve Backup detects your drive, you must format the removable storage media as a Arcserve Backup storage media. In the Arcserve Backup graphical user interface, removable media is represented as if it is tape media. This is not an error. Arcserve Backup treats removable media in the same manner as tape media.

Note: Various manufacturers provide you with pre-formatted media that needs to be manually formatted before you can use it. For more information about how to format the media for the drive you are using, see the manufacturer's documentation.

To format removable media

1. Open the Device Manager window and expand the Servers object.
2. Browse to the server to which the removable drive is connected.

3. Select and right-click the removable drive.
4. From the pop-up menu, select Format media.
Arcserve Backup formats the media.

How You Can Configure Removable Device Groups

You configure removable device groups through the Device Management feature. Using this feature, you can perform the following tasks:

- Create or delete new removable device groups.
- Rename removable device groups.
- Assign or remove individual devices from a device group.

Note: You cannot assign a removable drive into a group of media drives. You must create a new group for the removable devices.

Filter Libraries

Arcserve Backup lets you use filters to configure the Device Manager to display only the information you need, thereby increasing data manageability and application performance.

To filter libraries

1. Open the Device Manager window and select Preferences from the View menu.
The Preferences dialog opens.

2. Select the Library Filter tab and specify the filter options that are appropriate to your needs:
 - **Show Write Protected Media in Format / Erase dialogs**--Lets you view information about write-protected media in all Format and Erase dialogs.
 - **Show device name as Vendor ID and Serial Number**--Lets you view device names as the Vendor ID and the serial number.
 - **Show Empty Slots**--Select this option to view the empty slots in the library.
 - **Show Slots Between**--Specify the range of slots to be displayed in the current manager. To define the range, enter the minimum and maximum number of slots allowed.
 - **Show Blank Media Only**--Select this option to view the blank media in the library.
 - **Show Tapes within Media Pool**--Select this option to view the tapes within a particular media pool. Wild cards ("*" and "?") are accepted in the media pool.
 - **Show Tapes Matching Serial #**--Select this option to view the tapes that match a certain serial number. Wild cards ("*" and "?") are accepted in the serial number.

If a filter was applied to the current manager, the status bar will indicate it by displaying FILTER in the second panel and it will be detailed in the right panel of the view.

Note: Click Clear to clear all the fields of their information and remove all library filter criteria.

3. Optionally, click the Save as Default button after you have entered the criteria for your library filter to apply the filtering criteria to all Device Manager views.
4. Click Apply.

The filtering criteria are applied to the current view.

Note: Click the **Cancel** button to discard the changes to your filtering options.

Removable Drive Support

Arcserve Backup supports SCSI and USB removable devices allowing you to back up data, restore data, scan sessions, merge removable sessions, and manage removable media on your removable devices. The Backup Manager identifies and treats the removable media as tape media.

Note: To access the most up-to-date list of certified devices, click the Technical Support link on the Arcserve Backup Home Page.

How Arcserve Backup Supports Write Once Read Many (WORM) Media

Arcserve Backup allows you to back up your data either to rewriteable media or WORM media. WORM media, with significantly longer shelf life than magnetic media, manifests secure, long-term storage for data you do not want to erase.

Arcserve Backup lets you mix WORM and non-WORM media in a library. From the Device Manager, you can identify WORM media by an icon with the letter "W" inside a red circle. In addition, Arcserve Backup lets you specify WORM media for custom backup jobs.

The Backup Manager contains three options for Daily, Weekly, and Monthly WORM media rotations for use with GFS rotations. You can locate these options on the Backup Manager, Schedule tab, when you specify the Use Rotation Scheme option.

■ **WORM Media Supported**

Arcserve Backup supports backing up data to the following WORM media:

- DLT WORM (DLTice)
- STK Volsafe
- IBM 3592 WORM
- LTO3 WORM
- SAIT WORM

■ **WORM Media Considerations**

The following list describes situations that can occur when using a DLT WORM device with DLT WORM media and how Arcserve Backup responds to such situations.

- When a backup job spans tapes and the media is WORM media, Arcserve Backup needs WORM media to complete the job.
 - If a blank WORM media is not available, and a blank DLT WORM-capable media is available, Arcserve Backup automatically converts blank DLT media to DLT WORM media and then completes the backup job.
 - If WORM media is not available for a WORM job to continue, Arcserve Backup does not convert non-blank media to WORM media.
- When you are running a backup job that specifies Use WORM Media and there is no WORM media available, Arcserve Backup may convert blank WORM-capable media to WORM media for the job.

Note: For these scenarios, the available WORM media must be DLT SDLT-II or higher.

- **WORM Media Limitations**

If you use WORM media, certain Arcserve Backup features, specifically those involving media pools, reformatting, and overwriting or reusing media, are disabled because of the nature of the media. These limitations include the following:

- You cannot erase WORM media.
- You cannot submit an Overwrite job to WORM media.
- You cannot format WORM media unless the media is blank.
- You cannot use WORM media for multiplexing jobs.
- Arcserve Backup does not automatically assign WORM media to the Scratch Set in a media pool. WORM media cannot be recycled and as such, is always assigned to the Save Set in a media pool.
- Arcserve Backup cannot use WORM media with file system devices and Arcserve Backup tape RAID devices.
- In cross-platform SAN environments, UNIX does not support WORM media.

DLTSage Error Handling

DLTSage is an error-monitoring, reporting, and alerting technology developed by Quantum for use on SuperDLT tape drives. To receive tape drive alerts, you must use SuperDLT tape drives with DLTSage firmware.

Arcserve Backup interfaces with the firmware on SuperDLT tape drives to analyze critical tape drive and media performance parameters collected for each track, segment, Magneto Resistive (MR) channel, and optical band. Arcserve Backup uses the information collected to:

- Diagnose information such as threshold conditions and tape drive history.
- Identify high-risk tape drives and media that are approaching or have reached their end of life.
- Predict tape drive cleaning needs.
- Analyze tape drive environmental conditions.
- Generate media and hardware error messages.

Note: For more detailed information, see [How DLTSage Error Handling Works](#) (see page 430).

Resolving DLTSage Errors

For information about resolving DLTSage errors, see the following:

- [DLTSage Errors with the Tape Library Option](#) (see page 431)
- [DLTSage Errors without the Tape Library Option](#) (see page 432)

How DLTSage Error Handling Works

Arcserve Backup queries DLTSage using a SCSI Log Sense. If a hardware or media error occurs as a backup job starts, during a backup job, or after a backup job ends, Arcserve Backup uses the information captured from the SCSI Log Sense to generate tape drive error messages that display in the Tape Log and the Activity Log.

An error message displays if any of the following conditions exist:

- The tape drive is experiencing difficulties reading from or writing to a tape.
- The tape drive cannot read from or write to a tape, or the media performance is severely degraded.
- The media exceeded its life or maximum number of passes expectancy.
- The tape drive may have a clogged head or needs cleaning.
- The tape drive has a cooling problem.
- There is a potential tape drive hardware failure.

If an error condition is detected, Arcserve Backup may attempt to automatically correct the problem and complete the job. However, you must install the Arcserve Backup Tape Library Option to use the capabilities of uninterrupted inline cleaning, drive usage balancing, and error-preventive drive selection features. For more information about automated error resolution, see the *Tape Library Option Guide*.

If the Arcserve Backup Tape Library Option is not installed, you must cure the error condition or problem area manually. See the manufacturer's documentation, as necessary.

DLTSage Errors with the Tape Library Option

If Arcserve Backup detects a problem with a tape drive, there are three mechanisms Arcserve Backup uses to cure tape drive errors:

Drive Error Prevention

Before performing a backup job, Arcserve Backup evaluates the tape drive's health. If DLTSage detects a problem with a tape drive, Arcserve Backup performs the following analyses and actions:

- If another tape drive is available (empty) Arcserve Backup moves the tape to the next available tape drive and runs the backup job.
- If there are no available drives, Arcserve Backup moves the tape to a tape drive that is not locked, and runs the backup job.
- If all other drives are locked, Arcserve Backup runs the backup job in the current tape drive.

Drive Usage Balancing

Arcserve Backup balances tape drive usage by rotating the usage of tapes evenly between all tape drives in your library. When Arcserve Backup starts a backup job, it detects the last drive used and uses the next available drive in the library.

To minimize the movement of tapes from tape drive to tape drive, Arcserve Backup performs the following analyses and actions:

- If the media targeted for the job is loaded, allow it to remain in its current drive and run the backup job.
- If the media targeted for the job is not loaded, identify the last used drive and load the tape into the next available, empty drive, and then run the backup job.

Uninterrupted Drive Cleaning

A contaminated tape drive condition is usually discovered when you are running a backup job. A significant number of tape drive and media errors can be remedied by cleaning the tape drive.

For Arcserve Backup to perform uninterrupted drive cleaning, you must have a cleaning tape installed in the tape cleaning slot specified during setup, and specified a cleaning schedule. If you did not specify a cleaning schedule, Arcserve Backup defaults to a 100 hour period between scheduled tape cleaning operations.

If Arcserve Backup detects a contaminated tape drive condition during a backup job, and a cleaning slot is configured, Arcserve Backup automatically performs the following analyses and actions:

- If Arcserve Backup detects a write error during a backup, and the symptoms relate to a contaminated tape drive or media, Arcserve Backup makes a second attempt to write to the tape drive.
- If the second write attempt fails, Arcserve Backup cleans the tape drive if one or more of the following conditions exist:
 - The tape drive was never previously cleaned.
 - DLTSage detected the need to clean the tape drive and drive usage exceeds one fourth of the scheduled cleaning.
 - Tape drive usage exceeds one third of the cleaning.
 - The user specified ForceClean the tape drive.

If Arcserve Backup determines that a tape drive must be cleaned to continue a job, the following actions take place:

- Arcserve Backup pauses the job.
- The library returns the tape to its home slot and locks the tape drive.
- Arcserve Backup directs the cleaning operation.
- The library reloads the tape into the cleaned drive and aligns the tape with the buffer.
- Arcserve Backup resumes the job.

Note: These mechanisms run automatically upon error detection. No user intervention is required. After Arcserve Backup cures the error, the backup job continues. For more information about automated error resolution, see the *Tape Library Option Guide*.

DLTSage Errors without the Tape Library Option

If you have one or more stand-alone Quantum SuperDLT tape drives with DLTSage firmware, Arcserve Backup will diagnose read and write errors reported by DLTSage on those devices if the Tape Library Option is not installed. However, if the option is not installed, you must cure the error condition or problem area manually. To receive DLTSage error reporting on multi-drive tape libraries, the option must be installed.

When an error occurs, Arcserve Backup performs a second attempt to complete the job. If the error persists, Arcserve Backup stops the backup job and the tape drive relays the information about the error condition to Arcserve Backup. You can view the details about the error condition in the Activity Log.

After determining the cause of the error condition and curing the problem, you must resubmit the job.

How Arcserve Backup Cures Tape Drive Errors

If an error condition occurs, Arcserve Backup makes a second attempt to complete the job. If the error persists, Arcserve Backup stops the backup job. The tape drive then relays the information about the error condition to Arcserve Backup. You can view the details about the error condition in the Activity Log.

After determining the cause of the error condition and curing the problem, you must restart the job.

How Arcserve Backup Integrates with Secure Key Manager

Secure Key Manager (SKM) is encryption technology that lets hardware vendors, such as HP and Quantum, secure data that is stored on storage devices. To support the encryption key management capabilities of these vendors, Arcserve Backup integrates with SKM technology.

If you are backing up data to devices that support SKM, as best practice, you should use the encryption features provided by the devices instead of the encryption features provided by Arcserve Backup. We recommend this approach because hardware-based encryption provides a higher level of security than software-based encryption.

Integration with SKM technology lets Arcserve Backup behave in a manner that is transparent to the user.


In some instances, devices may not be able to read from the media because the media is not recognized or the encryption key is not available. These conditions cause the device to appear as if it is offline or not functioning. If a device appears to be offline or not functioning, Arcserve Backup behaves as follows:

Note: The following behaviors apply to single-drive libraries and multiple-drive libraries that support SKM technology.


Device Manager

The Device Manager demonstrates the following behavior when Arcserve Backup detects that SKM is installed on the device and the SKM application is offline or not functioning:

- Encryption type**--For SKM controlled devices, Unknown Encryption appears in the Device Manager.




- Format and Erase operations**--For SKM controlled devices, Unknown Encryption appears in the Device Manager.



Messages

The messages that follow appear when Arcserve Backup detects that SKM is installed on the device and the SKM application is offline or not functioning:

- Failed Format and Erase Operations**--The message that follows opens when Format and Erase operations fail on SKM controlled devices:



- Activity Log messages**--Arcserve Backup generates the Activity Log messages highlighted in the following screen when Format and Erase operations fail on SKM controlled devices:

Generic Logs

E8021	100-LL-260-2	07/14/2009...	Failed to erase media in slot:2 (EC=HW ENCRYPTION ERROR) [Device:2][HP MSL G
E6001	100-LL-260-2	07/14/2009...	Error Erasing:HW ENCRYPTION ERROR
E6114	100-LL-260-2	07/14/2009...	The tape is encrypted and can not be used by ARCserve. (Serial Number=[TT0247L4]
E6112	100-LL-260-2	07/14/2009...	Failed to read the tape's header. (Serial Number=[TT0247L4])
E6114	100-LL-260-2	07/14/2009...	The tape is encrypted and can not be used by ARCserve. (Serial Number=[TT0247L4]
Information	100-LL-260-2	07/14/2009...	[JOBQUEUE]: Pruned [Job No: 3] [Description: Backup [Custom]]

New Error strings defined for SKM errors, if Format and Erase fails.

The following Activity Log message appears when Arcserve Backup cannot decrypt the encryption that is detected on SKM controlled media:

E3834	100-LL-260-2	07/13/2009...	26	Unable to find any suitable media that can be used in this job.
E3703	100-LL-260-2	07/13/2009...	26	Unable to open media. (MEDIA=HELLO, EC=MEDIA MISCOMPARE ERROR)
W6500	100-LL-260-2	07/13/2009...	26	Client connect to tape failed (Tape Name = HELLO, Random ID = AFF0, Seq = 1)
E6112	100-LL-260-2	07/13/2009...	26	Failed to read the tape's header. (Serial Number=[TT0247L4])
E6114	100-LL-260-2	07/13/2009...	26	The tape is encrypted and can not be used by ARCserve. (Serial Number=[TT0247L4])
Inform...	100-LL-260-2	07/13/2009...	26	Description: Backup [Custom].
Inform...	100-LL-260-2	07/13/2009...	26	Start Backup Operation. (QUEUE=1, JOB=3)
Inform...	100-LL-260-2	07/13/2009...	26	Run Backup Job Now.

In addition to the existing error messages, E6114 will be logged to indicate that the job failed, due to encryption not decryptable.

How to Ensure that Arcserve Backup Spans Media in a Single Drive Autoloader

When Arcserve Backup backs up data to a single drive autoloader, and detects there is no blank media at spanning, Arcserve Backup pauses the job at spanning to let you insert blank tapes into the drive. This behavior is designed to help ensure that backup data spans media properly.

When Arcserve Backup spans media and detects no blank media in a single drive autoloader, the events that follow occur:

1. Arcserve Backup pauses the job and prompts you to insert blank media into the autoloader.

Note: If there are no empty slots, you can replace the older media with blank media. However, you must not replace or remove the media from the recently spanned slots. While spanning tapes, Arcserve Backup locks the affected slots during the inventory process. As a result, removing or replacing the spanning tapes prevents Arcserve Backup from updating the slot information properly. If you insert media that is not blank, you may inadvertently erase the data from the media using Device Manager.

2. After you close the door to the autoloader, Arcserve Backup inventories all media in the slots.

If the inventory process does not start automatically, you can inventory the media manually using Device Manager.

Note: You must wait for the inventory process to finish, which can take several minutes to complete.

After the inventory process is complete, the backup job resumes using the blank media, after you click OK on the message box that prompted you to replace the media.

To ensure that Arcserve Backup spans media in a single drive autoloader

1. Open the Device Manager and browse to the autoloader.

Right-click the device and select Library Properties from the pop-up menu.

The Library Properties dialog opens.

2. Click the General tab.

Clear the Set unknown bar code media to not inventoried during initialization check box.

Click OK.

Arcserve Backup inventories the blank media automatically when spanning is required.

Important! If you do not perform these steps, you must inventory the media manually using Device Manager.

Media Assure

From the Media Assure & Scan Utility, you can select the Media Assure button to display the Media Assure Option dialog. This dialog allows you to enable a media assure operation, which helps you ensure that the sessions on the media are restorable. A Media Assure job scans sessions randomly based on specified criteria.

After a Media Assure & Scan job completes, check the Activity Log in the Job Status Manager for errors. Based on the nature of the errors, you can take corrective actions to remedy the error.

Note: For the Media Assure feature, it depends on the session records in the Arcserve Backup database. So, if there are no records for the media in the database or the session records for this media have been destroyed, media assure will not scan any sessions.

- **Enable Media Assure**--Check this option to enable a media assure scan job that will select some sessions to scan randomly. Otherwise, it is a regular scan job.
- **Scan all data in a session**--Select this option to scan all session details.
- **Scan only session headers for each session**--Select this option to scan session headers only and not the session details. This is quicker than scanning all data in a session, however it may be harder to find the problem.
- **Scan sessions that match the following criteria**
 - **Sessions that were backed up in the last (number) Days**--Specifies the number of days that sessions were backed up to include in the Media Assure operation. The default is 7 days. So, all sessions that were backed up in the last 7 days will be scanned.
 - **Choose session no more than**--Limits the number of sessions scanned, because there are too many sessions that can fit a scan condition. The default is 20%. A percent or numerical value can be selected.
 - **Specify the Nodes for which the sessions should be scanned (Using ", " to separate)**--Indicates the sessions to scan in the specified nodes. This can be a wildcard match. For example, if you specify the node name ARC*, sessions will be selected from the node name ARC001 and the node name ARC002. If you don't specify any node name, then any session in all nodes may be selected. By default, all sessions in all nodes may be selected.

How Uninterrupted Drive Cleaning Works

A contaminated tape drive condition is usually discovered when you are running a backup job. A significant number of tape drive and media errors can be remedied by cleaning the tape drive.

For Arcserve Backup to perform uninterrupted drive cleaning, you must have a cleaning tape installed in the tape cleaning slot specified during setup, and a specified cleaning schedule. If you did not specify a cleaning schedule, Arcserve Backup defaults to a 100 hour period between scheduled tape cleaning operations.

If Arcserve Backup detects a contaminated tape drive condition during a backup job, and a cleaning slot is configured, Arcserve Backup automatically performs the following analyses and actions:

- If Arcserve Backup detects a write error during a backup, and the symptoms relate to a contaminated tape drive or media, Arcserve Backup makes a second attempt to write to the tape drive.
- If the second write attempt fails, Arcserve Backup cleans the tape drive if one or more of the following conditions exist:
 - The tape drive was never previously cleaned.
 - DLTSage detected the need to clean the tape drive and drive usage exceeds one fourth of the scheduled cleaning.
 - Tape drive usage exceeds one third of the cleaning.
 - The user specified ForceClean the tape drive.

If Arcserve Backup determines that a tape drive must be cleaned to continue a job, the following actions take place:

1. Arcserve Backup pauses the job.
2. The library returns the tape to its home slot and locks the tape drive.
3. Arcserve Backup directs the cleaning operation.
4. The library reloads the tape into the cleaned drive and aligns the tape with the buffer.
5. Arcserve Backup resumes the job.

Implement Logical Block Protection

Logical Block Protection (LBP) is functionality offered by manufacturers of stand-alone tape drives, such as IBM and HP, that lets Arcserve Backup detect read and write errors that occur on drives under the following conditions:

- An error occurs during the process of writing the backup data to the media. For example, the data is corrupt.
- The backup data changed as it transmitted through the network or storage area network.

When an error occurs, the drive reports a SCSI error in the Tape Engine log file and the backup will fail.

The Arcserve Backup implementation of LBP does not support the following backups:

- LBP on Read operations
- Arcserve Tape RAID backups
- NAS (Network attached storage) backups
- Arcserve Backup UNIX/Linux Data Mover backups
- File system device backups

Example: Data deduplication and staging backups

To implement LBP, you can create the registry key described in the following procedure. The registry entry can be created under the DEVICE<#> key in the registry for the stand-alone tape drive.

Follow these steps:

1. Log in to the Arcserve Backup server
2. Open the following registry key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA ARCServe  
Backup\Base\TapeEngine\DEVICE<#>
```

Note: DEVICE# represents the number of the device where you want to implement LBP.

3. Create the following DWORD:
LBP
4. Set the DWORD value of LBP to 1 and save the changes.

Note: To disable LBP, set the DWORD value of LBP to 0.

How to Optimize Tape Usage

Suppose you have a scenario where you have multiple disk staging backup jobs or multiple GFS Rotation backup jobs, and each job formats its own tape for incremental or differential backups. If the incremental or differential size of the data is lesser than the capacity of the tapes, the tape usage will not be optimized and space on the tapes will be wasted. In addition, using more tapes will increase the requirement for the number of slots in a tape library and can also result in the need to ship more tapes off-site.

There are two approaches to resolving this problem: [Media Maximization](#) (see page 439) and [Consolidation During Migration](#) (see page 441).

Media Maximization


Media maximization is a process that helps optimize disk and tape usage in GFS and rotation jobs. In a GFS or rotation job, when data is backed up on a scheduled basis to the same media pool, Arcserve Backup automatically appends the newly backed up data to a partially filled tape, instead of formatting a new tape each time. Using media maximization, you can optimize disk and tape space and also reduce the number of tapes needed to store your GFS rotation job data.

Media maximization can be used with the following types of jobs:

- GFS jobs
- Disk staging GFS jobs
- Disk staging rotation jobs
- Custom disk staging jobs using media pools to append backup data

Note: Arcserve Backup applies media maximization to a GFS Rotation job only if the specified media pool prefix is the same for those sets of jobs that are intended to use media maximization. For example, you can consolidate data from Job 1 and Job 2 on to the same tape in media pool A, and you can also consolidate data from Job 3 and Job 4 to another tape in media pool B.

However, in the media maximization process, Arcserve Backup does not back up data to a media that is already being used by an active backup job. So, you must ensure that the backup job schedule or migration schedule (in a staging job) is configured so that the backup or migration of data is sequential. If Arcserve Backup detects that the media is currently in use, it will revert to formatting a new tape for the second job, rather than wait for the first job to complete.



Examples: How Media Maximization Works

- **GFS rotations**--Multiple backup servers are processing GFS backup jobs. Arcserve Backup will store the backup data on the same media, only if you specify the same media pool prefix for all of the jobs.
- **Staging backups**--Multiple backup servers (Server A and Server B) are processing backup or migration jobs. The job on Server B starts while the job on Server A is in progress. Arcserve Backup writes the data for the job on Server B to a different tape than the job on Server A. Arcserve Backup demonstrates this behavior because multiple backup servers cannot write data to the same media simultaneously. However, if the job on Server B starts after the job on Server A completes, Arcserve Backup writes the data to the same tape that was used by Server A.

Note: As a best practice, you can increase the Timeout for First Media value to control the length of time that the job waits before it selects a different tape to store the backup data. For more information, see [Backup Manager Backup Media Options](#) (see page 164).

Consolidation During Migration

Consolidation during migration is a process to help optimize tape usage in staging jobs. Consolidation during migration can be used in a custom job, rotation job, or GFS rotation job.

In a staging job, when data is migrated (or copied) from the staging area to the same media destination (same media pool prefix), the consolidation during migration option allows you to append migrated data onto a partially filled tape, instead of formatting a new tape each time. Through the use of the consolidation during migration option, you can optimize tape space and also reduce the number of tapes needed to store your migrated data.

The consolidation during migration option is similar to the media maximization feature and data will not be migrated to a media that already has an active migration job in process. However, with this option you no longer have the responsibility of scheduling each job so that the next migration job is not started before the previous migration job has been completed. If you select this option, Arcserve Backup will automatically detect if the media is currently in use, and if it is, will wait for the current migration job to complete before starting the next migration job. To consolidate data during migration you need to specify the same exact target media prefix and target media pool prefix so that data belonging to different jobs can be consolidated to the same exact tape.

The Consolidate data across jobs while copying option, which is a Miscellaneous option that appears on the Policies tab, lets you specify if you want to consolidate the data from different jobs onto a single tape during migration.

Example: How Consolidate Data During Migration Works

You can consolidate data from Job 1 and Job 2 onto the same tape, and you can also consolidate data from Job 3 and Job 4 to another tape. In this scenario you would need to do the following:

- When submitting a backup of Job 1, choose consolidation. Specify the media prefix as AAA and the media pool as MP1.
- When submitting a backup of Job 2, choose consolidation. Specify the media prefix as AAA and the media pool as MP1.
- When submitting a backup of Job 3, choose consolidation. Specify the media prefix as BBB and the media pool as MP2.
- When submitting a backup of Job 4, choose consolidation. Specify the media prefix as BBB and the media pool as MP2.

If you want data to be consolidated from Jobs 1 and 2 and from Jobs 3 and 4, you must specify the exact same media prefix and the exact same media pool in each submitted job. Since the consolidation has to be done to the same tape set you must also choose the same tape library group in the target destination. In addition, the jobs which are supposed to consolidate must run on the same backup server.

Important! If any of these four parameters are different (Media Prefix, Media Pool, target destination, and backup server), the data will NOT be consolidated to the same tape set.

You can also specify the copy method as either to overwrite the data on a tape or to append to the data on an existing tape.

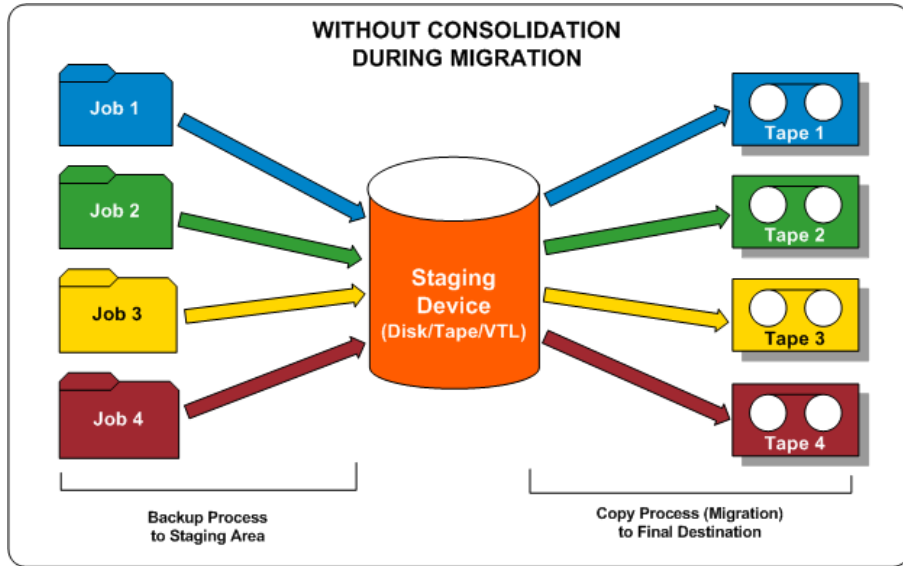
- **Overwrite**--If you have a requirement to consolidate data across multiple jobs and ship the tapes on a daily basis, you should choose the "Overwrite" option. This will ensure that a tape is formatted on a daily basis and all the data backed up on that day would be migrated to a single tape.

For example, if you have two jobs (Job 1 and Job 2) and you want to ship the tapes offsite on a daily basis. In this scenario you would choose Overwrite. When the backup job finishes on Monday, Arcserve Backup would format a final tape for Monday and copy the data from the staging tapes of Jobs 1 and 2 to the final tape. Then, after the backup finishes on Tuesday, Arcserve Backup would format a final tape for Tuesday and copy the data from staging tapes of Jobs 1 and 2 to the final tape. This way a tape is formatted every day and helps you to ship the tapes offsite on a daily basis.

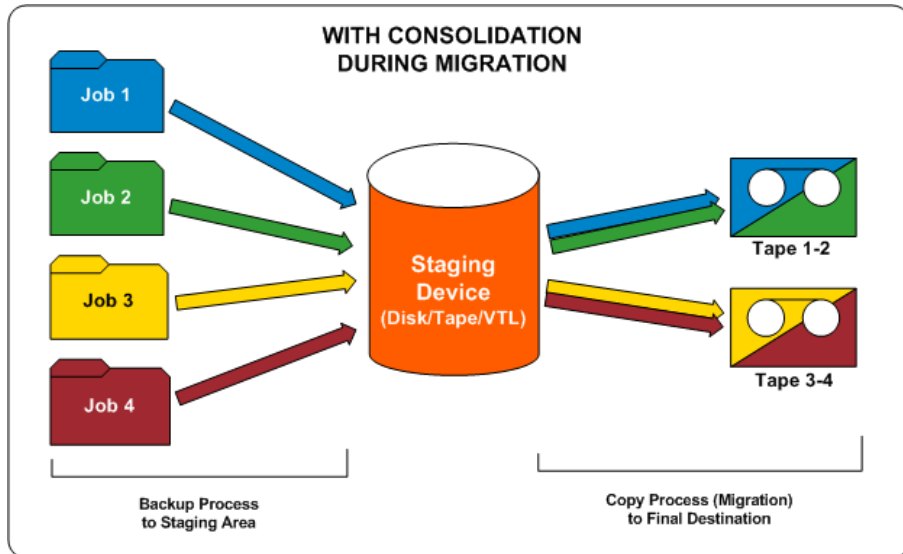
- **Append**--If you have a requirement to consolidate data across multiple jobs (for daily backups) for the whole week to a single tape and ship the tapes on a weekly basis, you should choose the "Append" option. This will ensure that for example, in 5-day GFS rotation jobs, all the incremental or differential data (belonging to different jobs) that are backed up on Monday, Tuesday, Wednesday, and Thursday is consolidated to one tape set. The full backups that happen (for different jobs) on Friday would be consolidated to another tape set.

For example, if you have two jobs (Job 1 and Job 2) and you don't want to ship the tapes offsite on a daily basis. In this scenario you would choose Append. When the backup job finishes on Monday, Arcserve Backup would format a final tape for Monday and copy the data from the staging tapes of Jobs 1 and 2 to the final tape. Then, after the backup finishes on Tuesday, Arcserve Backup would copy and append the data from staging tapes of Jobs 1 and 2 to the final tape from Monday. A new tape would not be formatted and only one tape would be formatted for the entire week of daily backups. This helps you utilize your tapes more efficiently.

The following diagram shows the tape usage requirements if you do not enable the consolidate during migration option:



The following diagram shows the tape usage requirements if you enable the consolidate during migration option:



How Media Pools Work

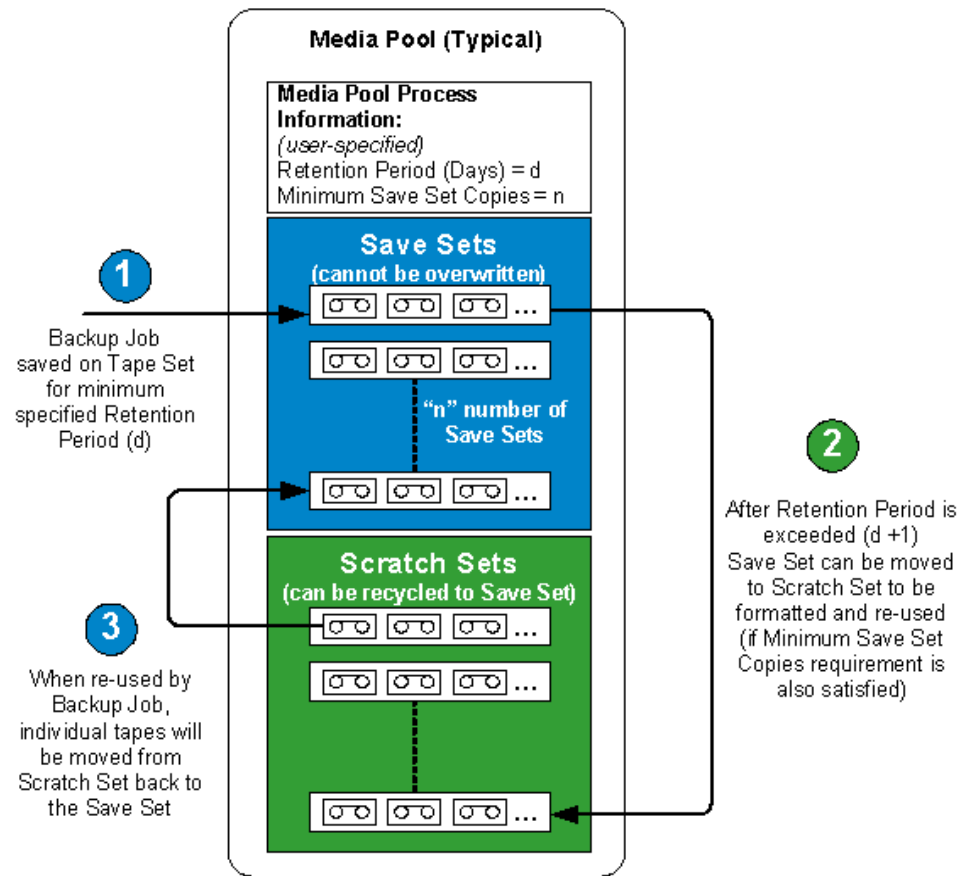
Each media pool is divided into Save Sets and Scratch Sets. These sets are used in conjunction with each other to control the preservation of backup data on tapes until your specified criteria has been met and then allows you to recycle these tapes for reuse. The two user-defined retention criteria are:

- the minimum number of media that must be contained in the Save Set
- the retention period (in days)

Example: Media Pool Used in a Rotation

During a 5-day work week, daily backups are performed on Monday, Tuesday, Wednesday, and Thursday. Each of these daily backups has their own set of backup media (daily Save Sets) that are retained for four days (the user-specified retention period). On the fifth day (Friday), a weekly Save Set is created and the daily Save Set from the previous Monday becomes part of the Scratch Set, so that it can be reused (overwritten). In other words, on the next Monday, the daily media pool from the previous Monday becomes part of the Scratch Set and can be reused for this Monday's backups. After the new Monday backup is completed, the Scratch Set for that day becomes the Monday Save Set and is retained all week.

The following diagram shows how a typical media pool processes a backup job and the movement of Save Sets and Scratch Sets within a media pool:



Save Sets

The media pool Save Set is a set of media that cannot be overwritten until the media pool's retention requirements that you specify have been met. You can modify Save Set information for all Custom backup jobs, move media from the Save Set to the Scratch Set, or you can move media from one media pool Save Set to another media pool Save Set.

You define the minimum number of media that must be contained in the Save Set and the retention period (in days). These settings determine how long media will be held. After both of these criteria have been satisfied, Arcserve Backup releases the oldest media in the Save Set back into the Scratch Set, where it can be recycled and re-used (overwritten).

- The retention period is the number of days in which a media has not been used (written to) before it is moved into the Scratch Set. For example, if you specify a retention period of 14 days, a media remains in the Save Set if it has been used within that specified time. If the media has not been used for 14 days, it is moved to the Scratch Set.
- The minimum number of media contained within the Save Set is the number of media that must be retained in the Save Set before the older media are recycled to the Scratch Set. This is a safeguard for preventing data loss in case backups are not done for extended periods of time.

Note: You will receive a warning if you attempt to format or erase media that is contained in a Save Set.

Scratch Sets

The media pool Scratch Set is a set of media that has been recycled from the Save Set after its specified retention criteria has been satisfied. The media from the Save Set that can be re-used and overwritten are placed in the Scratch Set after they have met the specified criteria (the minimum number of media to save and retention period). The oldest media in the Scratch Set, those that have not been used for the longest period of time, are used first.

Each time a media in the Scratch Set is used, it moves from the Scratch Set to the Save Set. The media moves back to the Scratch Set once the specified retention criteria have been met. If the media meets these retention criteria, Arcserve Backup prompts for a blank tape or accepts media from the Scratch Set.

Arcserve Backup performs media pool maintenance at the beginning of a job, and will not allow media in the Save Set to be moved to the Scratch Set until the two retention criteria are met. When you select a media pool Scratch Set in the left pane of the Media Pool Manager, the right pane will display the media pool name, the set name, the owner name, and the date the Scratch Set was created.

Save Sets and Scratch Sets

The set of media containing important data that cannot be overwritten is called the Save Set. You can move media from the Save Set in one media pool to the Save Set in another media pool. Media that has not been formatted for the longest period will be used first.

Note: You will receive a warning if you try to format or erase media in a Save Set.

When the media meets certain criteria in a Save Set (minimum number of media in Save Set and retention period) they are recycled to the Scratch Set. Each time a media in the Scratch Set is written to, it moves from the Scratch Set to the Save Set. Additionally, if Arcserve Backup detects non-blank media in the Scratch Set, the Media Pool Manager controls the usage of the media such that WORM media containing data is not used.

The retention period is the number of days in which a media has not been used before it is moved into the Scratch Set. For example, if you specify a retention period of 14 days, a media remains in the Save Set if it has been used within that specified time. If the media has not been used for 14 days, it is moved to the Scratch Set.

You define the minimum number of media contained within the Save Set. This is the number of media to be retained in the Save Set before the older media are recycled to the Scratch Set. This is a safeguard for preventing data loss in case backups are not done for extended periods of time.

Media pools apply to every media, regardless of which backup type and method were selected. Arcserve Backup performs media pool maintenance at the beginning of a job, and will not allow media in the Save Set to be moved to the Scratch Set until two criteria are met:

- The oldest tape in the Save Set is compared and exceeds the retention time.
- The minimum required number of media is in the Save Set.

If the media meets these criteria, Arcserve Backup prompts for a blank tape or accepts media from the Scratch Set.

Serial Numbers

The serial number of a media is one way to categorize media pools. You cannot change the serial number of media, but you can create a serial number for media by:

- **Bar code**--A number is read from a bar code label and this number becomes the serial number. A changer with a bar code reader is required for this method. This will override any previously defined media pool settings.
- **Automatic**--Arcserve Backup automatically assigns a serial number for the media based on the base and range of serial numbers set when the pool was created.
 - **Base**--This is the base number, which Arcserve Backup will use when automatically assigning serial numbers. The first media formatted will have the same serial number as the base number. Each media's serial number thereafter will be increased by one.
 - **Range**--You can specify the range (up to 31 digits) from which the media pool serial numbers will be categorized.

GFS Media Pools

Grandfather-Father-Son (GFS) Rotation media pools are based on basic media pooling architecture.

GFS Rotation jobs use three media pools: Daily, Weekly, and Monthly, which are based on the information you enter in the Media Pool Name Prefix field when submitting the job.

When a GFS Rotation job runs, Arcserve Backup automatically formats and names your media according to the backup type, media pool, and date using the following syntax:

(backup type)-(user-defined media pool prefix)-(day-of-the-week)-(date)

Where..	Is...
backup type	F - full backup I - incremental backup D - differential backup W - weekly backup M - monthly backup A - all daily backups (full, incremental, and differential) when you use the Media Maximization option (enabled by default) and enable the Append Media option. For more information on the Media Maximization option, see the section Media Maximization in GFS Rotation Jobs.
user-defined media pool prefix	The name you assigned to the media pool for your GFS Rotation scheme.
day of the week	An abbreviation for the day of the week on which the job was performed.
Date	The date on which the backup was performed in mm/dd/yy format.

This media naming convention allows you to easily identify backup media. For example, the media used for the first full backup in your rotation scheme will have the following name: F TP MON 11/1/05.

Note: Arcserve Backup prevents you from using the underscore character (_) and the hyphen character (-) when specifying Media Pool names.

Five-day rotation schemes have the following retention times for each media pool:

- **Daily (_DLY)**--six days (daily media in seven-day Rotation Schemes have a retention time of eight days)
- **Weekly (_WLY)**--five weeks
- **Monthly (_MLY)**--343 days

The following are the formulas used for calculating the number of media in the Save Sets and the retention times for the GFS media pools:

- **Daily pool**--This pool holds the media for daily backup jobs. The default retention period is six days and the number of Save Set media is based on the number of daily media in the GFS Rotation minus one [# of daily media -1].
- **Weekly pool**--This pool holds the weekly media. The retention period equals the number of weekly media times seven, minus one [(# of weekly media * 7) - 1]. The number of save media is based on the number of weekly media in the GFS setup minus one [# of weekly media - 1].
- **Monthly pool**--This pool holds the monthly media. The retention period equals the number of monthly media times 29 minus five [(# of monthly media * 29) - 5]. The number of save media is based on the number of monthly media in the GFS setup minus one [# of monthly media - 1].

More information:

[Media Maximization in GFS Rotation Jobs](#) (see page 450)

Media Maximization in GFS Rotation Jobs

By default, Arcserve Backup enables the Media Maximization option, which lets you submit multiple GFS backup jobs using the same media pool. By sharing the same media pool, you can append multiple jobs to the same tape sets rather than creating new tape sets for each job. This reduces the amount of media you use when processing GFS Rotation jobs.

Important! To help ensure that Arcserve Backup writes GFS rotation backup data to the same tape, you must specify the same media pool prefix for the desired jobs in the Backup Manager window.

Note: To disable the Media Maximization option, set the NT registry DWORD value EnableMediaMaximization to 0. This registry key is as follows:

HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\Arcserve Backup\Base\Task\Backup

More information:

[GFS Media Pools](#) (see page 449)

How You Can Maximize the Use of Media

To take full advantage of the Media Maximization option, use the following guidelines when submitting GFS backup jobs using the same media pool:

- **Use the same Rotation Scheme**--GFS jobs that use different rotation schemes may need different tape names. To ensure that multiple GFS jobs will share the media, use the same rotation scheme.
- **Start GFS jobs on the same day**--The first day of a GFS job is a full backup. Jobs that start with different dates may not be able to share media during the first week. To ensure that multiple GFS jobs will share media during the first week, start GFS jobs on the same day. Otherwise, media sharing will begin after the weekend.
- **If you want to modify multiple GFS backup jobs to use a new media pool, modify them on the same day**--This ensures that all jobs will share the media right away. Otherwise, media sharing will begin after the weekend.
- **Modify existing GFS jobs to use the same media pool as other GFS jobs**--If the existing GFS jobs you modified use the same rotation scheme, media sharing should begin right away. However, if any of the jobs have been running for less than one week, media sharing may begin after the weekend.

Media Maximization Methods

There are two different methods you can use to maximize your media usage. The method depends on whether you enable the Append Media feature when submitting your GFS backup job. Both methods significantly reduce the amount of media required. The following is a description of each method.

GFS Rotation Jobs without Append Media Enabled

If you submit GFS rotation jobs without the Append Media feature enabled, you can maximize media usage by submitting multiple jobs using the same media pool.

For example, if you submit three GFS rotation jobs all using the same media pool and 5-day rotation scheme, all three jobs share the same set of tapes. On each day of the rotation scheme, all three jobs append to the same tape:

- Monday = One tape that includes full backup data from job 1(day 1), job 2(day 1), and job 3(day 1)
- Tuesday = One tape that includes incremental backup data from job1(day 2), job 2(day 2), and job 3(day 2)

- Wednesday = One tape that includes incremental backup data from job 1(day 3), job 2(day 3), and job 3(day 3)
- Thursday = One tape that includes incremental backup data from job 1(day 4), job 2(day 4), and job 3(day 4)
- Friday = One tape that includes weekly backup data from job 1(day 5), job 2(day 5), and job 3(day 5)

This results in five tapes for the week.

Without the Media Maximization option, each job would require its own tape:

- Monday = Three full backup tapes. One tape for job1(day 1), one tape for job 2(day 1), and one tape for job 3(day 1)
- Tuesday = Three incremental backup tapes. One tape for job1(day 2), one tape for job 2(day 2), and one tape for job 3(day 2).
- Wednesday = Three incremental backup tapes. One tape for job1(day 3), one tape for job 2(day 3), and one tape for job 3(day 3).
- Thursday = Three incremental backup tapes. One tape for job1(day 4), one tape for job 2(day 4), and one tape for job 3(day 4).
- Friday = Three Weekly backup tapes. One tape for job1(day 5), one tape for job 2(day 5), and one tape for job 3(day 5).

Without the Media Maximization option, you need 15 tapes for the week.

Note: When submitting multiple GFS Rotation Jobs with the same media pool without Append Media enabled, tapes can be shared only if the same Backup Method is used. For example, a tape that has data from a full backup job can be shared only with data from another full backup job. It cannot be shared with data from incremental, differential, weekly, or monthly backup jobs.

GFS Rotation Jobs with Append Media Enabled

Similar to submitting GFS Rotation jobs without the Append Media feature enabled, you can maximize media usage when you enable Append Media by submitting multiple jobs using the same media pool. In addition, enabling Append Media also lets you maximize media usage by allowing you to share tapes among different jobs, regardless of the backup method that was used. (The only exceptions to this is weekly and monthly backup jobs. Weekly and monthly backup jobs can never share tapes with full, incremental, and differential backup jobs.)

For example, when submitting multiple GFS rotation jobs with the same media pool without Append Media enabled, a tape that has data from a full backup job can be shared only with data from another full backup job. If you enable Append Media, a tape that has full backup data can be shared with full, incremental, and differential data.

To share tapes among different jobs with different backup methods, Arcserve Backup uses the same GFS rotation naming syntax, but it uses a different naming convention for backup types when the Append Media feature is enabled:

(backup type)-(user-defined media pool prefix)-(day-of-the-week)-(date)

Without Append Media	With Append Media
F - full backup	A - full backup
I - incremental backup	A - incremental backup
D - differential backup	A - differential backup
W - weekly backup	W - weekly backup
M - monthly backup	M - monthly backup

If you submit GFS rotation jobs with the Append Media feature enabled, you can maximize media usage by submitting multiple jobs using the same media pool and use the previous day's tape within the current week.

For example, if you submit three GFS rotation jobs all using the same media pool and 5-day rotation scheme, all three jobs share the same set of tapes. In addition, multiple days can share the same tape, drastically reducing the amount of tapes you use:

- Monday, Tuesday, Wednesday, Thursday = One tape that includes full backup data from job 1(day 1), job 2(day 1), and job 3(day 1), and incremental backup data from job 1 (days 2, 3, and 4), job 2(days 2,3, and 4), and job 3(days 2,3, and 4).
- Friday = One tape that includes weekly backup data from job 1(day 5), job 2(day 5), and job 3(day 5)

This results in two tapes for the week.

Without the Media Maximization option, each job requires its own set of tapes. Among these tapes, only the ones that include data from the same Backup Method can be shared:

- Monday = Three full backup tapes. One tape for job1(day 1), one tape for job 2(day 1), and one tape for job 3(day 1)
- Tuesday, Wednesday, Thursday = Three incremental backup tapes. One tape for job1(days 2, 3, and 4), one tape for job 2(days 2, 3, and 4), and one tape for job 3(days 2,3, and4).
- Friday = Three weekly backup tapes. One tape for job1(day 5), one tape for job 2(day 5), and one tape for job 3(day 5).

This results in nine tapes for the week.

Note: If you submit a GFS Rotation job with Append Media enabled and Arcserve Backup cannot use the previous day's media for some reason, it will format a media in the Scratch Set or a blank media using the "With Append Media" naming convention. To minimize the likelihood of this situation occurring, see Media Maximization Rules in this chapter.

Overlapping Media Rules

Because the Media Maximization option allows multiple GFS jobs using the same media pool to share tapes, you may encounter a situation where a media is busy because it is being used by another GFS job. If this occurs when submitting a GFS backup job without Append Media enabled, the job waits for the tape to become available before appending. The default wait time is 10 minutes. If the media is still busy after 10 minutes, the job uses another tape.

If this occurs when submitting a GFS backup job with Append Media enabled, Arcserve Backup tries to append to a previous day's media. If that media is busy, it waits 10 minutes. If the media is still busy after 10 minutes, the job uses the current date to generate another media and attempts to use it. If the new media is busy, the job waits 10 minutes. If this media is still busy after 10 minutes, the job formats another media with a new name.

Note: You can change the wait time by entering a new value for the Windows registry key GFSwaittime. This value is stored in the following registry key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\Arcserve Backup\Base\Task\Backup
```

Media Pool Manager

The Media Pool Manager allows you to create and maintain media pools. Arcserve Backup creates a catalog file on your media to improve performance for merge jobs and database backup jobs. The Media Pool Manager helps you to perform the following tasks:

- **Create a new media pool**--To assign media to a media pool, you first need to create the media pool. A media pool name can consist of up to 16 uppercase characters.
- **Delete an existing media pool**--To delete a media pool, you must first re-assign the media to another media pool.
- **Move media in a pool**--You can move media from one set to another. You can also move media from a Scratch Set to a Save Set and vice versa by using the Assign Media and Remove Media options.

- **Perform location maintenance**--You can enter information about a new location, modify information about an existing location, or assign media to a location.
- **Assign media to a media pool**--You can assign media to a media pool during the process of formatting. When you format media using Device Management, you define certain media pooling information that will be associated with the media.
- **Remove media from a media pool**--You can remove media from a media pool.

Note: Media Pool operations, backup jobs using the Overwrite option, Tape Erase operations, and backup jobs involving Media Pools (such as GFS rotation jobs) are not supported on WORM (Write Once Read Many) media. These operations are either blocked or disabled in WORM support updates.

Create Media Pools

You can use automatic rotation schemes to control the media you use during backups. However, if you choose not to use the automatic features, the Media Pool Manager is an indispensable tool to efficiently schedule the maintenance and recycling of media. The Media Pool Manager helps you to organize your media into media pools, similar to those used in rotation schemes. Just as in rotation schemes, the media pools you create are collections of rewriteable storage media managed as a single unit.

Note: If you are using WORM media, media pool options are disabled. By definition, WORM media cannot be overwritten, so you cannot recycle it in a rotation scheme or a media pool.

For more information about media pools, see "Managing Devices and Media."

To create media pools

1. From Administration menu in the Navigation Bar on the Home Page, click Media Pool.

The Media Pool Manager opens.

2. From the Media Pool Manager, click New.

The Media Pool Configuration dialog appears.

Note: Arcserve Backup detects and assigns media serial numbers when media is formatted and placed in a specific media pool.

3. Enter a name for the media pool in the Pool Name field. Fill in the remaining fields appropriately.

The Minimum # of media in Save Set, Retention Time, Prune Retention Time, Base Serial Number, and Serial Number Range fields contain default values that you can change if desired.

Note: The Next Serial Number field cannot be set.

4. Click OK when you are finished.

The new media pool you created appears in the Media Pool Manager. You can now assign media to the Save Sets and Scratch Sets of this media pool.

How You Can Create a Rotation

To create a rotation, select the Schedule object in the left pane of the MM Admin window, double-click it, and select a schedule from the list. Double-click the schedule to access the Rotation object. Right-click the Rotation object and select Create. The Create Rotation dialog appears where you can set the following:

- **Sequence Number**--MM Admin automatically generates a sequence number for your rotation. Vault cycles start with the lowest sequence number. The default for a new rotation is 10 and the next new rotation that follows is 20. If you would rather assign a particular sequence number, select the Sequence Number option and select a number.
- **Vault Name**--A vault name must be specified for each rotation. You can select the name of a vault from the drop-down vault list.
- In the Retention fields, set any of the following conditions:
 - **Hold Days**--The number of days you want tape volumes to be retained.
 - **Keep for Cycles**--The number of vault cycles you want tape volumes to be retained in this rotation.
 - **Days Elapsed from First Format Date**--Starting from the day the tape volumes were first formatted, enter the number of days you want tape volumes to remain in this rotation.
 - **By Date**--Tape volumes are retained in this rotation until the date you enter here is reached.
 - **By Tape Expiration Date**--Tape volumes are retained in this rotation until their expiration date passes.
 - **Permanent**--All tape volumes are retained in this rotation permanently.

If a tape volume meets one of these conditions, it remains in the same rotation. None of these conditions have priority over the other so if any condition is true, the media will stay in the vault--even if conditions appear to conflict. For example, if you select 60 in the Hold Days field but enter a date that is only 30 days away in the By Date field, the tape volume will stay for 60 days.

When you click Add, the new rotation is saved and added to the Rotation branch in the MM management window.

When the Retention period for a tape volume expires, the tape volume is unvaulted and returned to Tape Service to be re-used.

Media Management Administrator (MM Admin)

Note: To use the MM Admin, you must install the Enterprise Module.

MM Admin lets you protect, control, and manage your media resources. Using MM Admin, you can organize tape movement to off-site storage locations, define retention policies to ensure that your tapes are protected from being prematurely overwritten, secure access to tape-resident files, and maintain a comprehensive inventory of tape library resources.

MM Admin activities are recorded in the Activity Log. This includes information, warnings, and errors. This important function allows you to centrally keep track of all media management operations.

To manage media using MM Admin, you must create a vault, create a schedule, select a vault criteria descriptor, and define a rotation. The following sections include information on each of these steps and cover all topics associated with managing your media using MM Admin.

Media Management and Tape Service

In data centers with off-site storage locations, tape volumes are typically cycled out of the central tape library to more secure storage areas (vaults), and then cycled back into the central library. MM Admin works with Tape Service to provide additional media control, rotation, slot number assignment, and reporting on vaulted tape volumes so that you can physically route these tape volumes to off-site storage locations and back to the data center, as necessary.

You can define vaulting criteria using MM Admin. The criteria for holding tape volumes in vaults can be different for each schedule and for each vault. As tape volumes meet these criteria, they are checked out of Tape Service with the proper vault code and reports are generated indicating the current location and destination to where the tape volumes must be moved.

Media Management Administrator Terms

The following are important terms associated with the MMO:

- **Vault**--Any identifiable storage area or location you define.
- **Slot**--Virtual slots in a vault are assigned when a tape volume is vaulted. Each slot is used to store one tape volume. By default, there are 32000 slots in a vault, but you can designate a different maximum number of slots as you create a vault.
- **Schedule**--Determines when a tape volume is to be placed in or removed from a vault.
- **Rotation**--Determines when to move tape volumes, and is associated with a schedule. Each rotation you define points to a vault.
- **Vault Criteria Descriptor (VCD)**--Defines the controlling data set you want to use for the selected tape volume. You can choose the controlling data set by media name or file name, or you can select an individual media as the controlling data set.
- **Vault cycle**--The actual movement of tape volumes. You must describe the vault, the tape volumes, and the rules for tape volume movement under the MMO by creating a Vault Criteria Descriptor (VCD) record. The MMO uses this descriptive information to execute a vault cycle when movement is scheduled.
- **Reports**--Each time you execute a vault cycle or an estimated vault cycle, Arcserve Backup generates several reports before another vault cycle can be initiated. The Vault Selection Report contains a list of tape volumes to be selected for moving into the vaults through the VCD. The Shipping Report and the Receiving Report provide a reliable record of the result of the vault cycle and the current location of your tape volumes.

The Shipping Content Report and the Receiving Content Report provide you with basic session details—in addition to the information contained within the Shipping Report and the Receiving Report—such as the session number, source path, start date, size, and number of files.

An Inventory Report is also available, which you can generate at any time.

MM Admin Interface

The MM Admin interface is designed to make vault creation, scheduling, VCD creation, rotation, and report creation easy. The tools provided by MM Admin allow you to establish the vaulting policy needed for complete Media Management.

The MM Admin workspace includes a menu bar, the main MM Admin toolbar, and the MM management window. The left pane of the MM Management window displays the MM primary management server in a tree structure for easy navigation. The right pane displays information related to the object selected in the left pane. It also displays any output messages and reports generated during your MM Admin session.

MM Admin Toolbar

The following table describes tasks that you can perform using the Media Management Administrator (MM Admin). Click the corresponding toolbar button to start the task.

Button	Task
Initialize MM Database	Lets you Initialize the MM database.
Retrieve Data	Lets you retrieve data and display the latest information if the database fails.
Refresh	Lets you refresh and update the information displayed in the MM Admin window.
Start Vault Cycle	Lets you start the vault cycle process.
Simulate Vault Cycle	Lets you produces a Vault Selection Report that predicts how many tape volumes will be moved without actually updating location information.
Find Media in Vault	Lets you search for media by Tape Name or Serial Number.
Property	Lets you view the server's properties.
Print	Lets you print the information displayed in the right pane of the MM Admin window.
Print Preview	Lets you preview information before printing.

MM Admin Window

The objects in the left pane of the MM Admin window are arranged in an expandable tree. To view related information, double-click the branch you need. After you access a branch, you can add, modify, or delete objects from the tree structure using the available pop-up menus. Right-click any object to access pop-up menus.

When you open MM Admin, the MM primary management server is displayed at the top of the tree. Double-click the branch to expand it and access the following objects:

- **Current Server**--Displays information about the server you are currently using.
- **Vault**--Provides information about previously created vaults.
- **Schedule**--Lists the names of the previously created schedules, and allows you to access the Vault Criteria Descriptor and Rotation objects.
- **Reports**--Provides access to the seven available reports.
- **Status**--Allows you to view the status of the most recent operation.
- **Find Media in Vault**--Lets you access the Find Media dialog to locate a particular media.

Schedule Object

The Schedule object provides information about previously defined schedules and allows you to create new schedules. You must create a schedule before you define the Vault Criteria Descriptor and Rotation that determine selection and retention policies for your vault.

When you select the Schedule object, the right pane of the MM Admin window displays the names of previously defined schedules. These schedules are also listed under the Schedule object in the left pane. Right-click the Schedule object to create a new schedule. Right-click a specific schedule to delete it. For more information about creating or deleting a schedule, see the section How You Can Schedule Tape Volume Movement.

After you have named and created a schedule, the Vault Criteria Descriptor (VCD) and Rotation objects appear in the left pane of the MM Admin window.

Vault Criteria Descriptor Object

The Vault Criteria Descriptor (VCD) allows you to set source information that governs the tape volumes assigned to a vault. You can select a media pool name or a file name as the controlling data set. If you want to assign only one tape, select the Assigned by User option as the controlling data set. If you use this option, you must enter command line information. When this data set is vaulted, the tape volumes are assigned to slots in the vault.

When you select the Vault Criteria Descriptor object, the right pane of the MM Admin window displays columns listing the following information for existing VCDs:

- **VCD Name**--The name of the Vault Criteria Descriptor.
- **VCD Type**--Indicates whether the controlling data set is defined by media pool, file name, or by user.
- **Media Pool**--If the controlling data set is a media pool, the name of the media pool appears.
- **Host Name**--If the controlling data set is a file name, the host where the file resides appears in this column.
- **Path/File Name**--If the controlling data set is a file name, the full path and file name appear in this column.
- **Create Date**--The date the VCD was created.

In the right pane of the MM Admin window, right-click an existing VCD to update or delete it. Right-click the Vault Criteria Descriptor object in the left pane to create a new VCD. For more information about creating, updating, or deleting a VCD, see the section [How You Can Manage Tape Volumes and VCDs](#).

More information:

[How You Can Manage Tape Volumes and VCDs](#) (see page 469)

Rotation Object

Media management relies upon user-defined rotation policies to determine when and where tape volumes should be moved. Use the Rotation object to set or update the retention policies that determine when tapes will be moved or released from the vault and returned to Tape Service.

When you select the Rotation object, the right pane of the MM Admin interface lists the following information about previously defined rotations:

- **Rotation Name**--The name of the rotation.
- **Vault Name**--The name of the vault the rotation is associated with.
- **Retention Hold Days**--Starting from the Last Write date (the date the media was last written to), this indicates the number of days that the tape volumes will be held in this rotation.

Note: To view the Last Write date, expand the Vault object and highlight a media name in the top right-hand pane. The Last Write date appears in the lower right-hand pane.

- **Retention Keep for Cycles**--Indicates the specific number of vault cycles and tape volumes that are held in this rotation.

- **Retention Days Elapsed from First Format Date**--Indicates that tape volumes are held in this rotation until a specified number of days have elapsed since they were first formatted.
- **Retention Permanent**--Indicates that tape volumes will remain in this rotation permanently.
- **Retention By Tape Expiration Date**--Indicates that tape volumes remain in this rotation until the tape expiration dates have passed.
- **Retention By Date**--Indicates that tape volumes remain in this rotation until the specified date has passed.
- **Create Date**--The date the rotation was created.
- **Description**--A user-defined description of the rotation.

Existing rotations are also listed in the right pane of the MM Admin window under the Rotation object.

- To update an existing rotation, right-click the rotation name and select Update from the pop-up menu.
- To create a new rotation, right-click the Rotation object and select Create.

Reports Object

Although tape volume location information in the database is updated when you initiate a vault cycle, the physical movement of tape volumes is done manually. MM Admin generates reports indicating the current location and destination where the tape volumes must be moved so that you can route them to other storage locations and back to the data center, as necessary.

The Reports object provides access to the reports generated by the vault cycle process and the Inventory reports, which can be generated at any time. Expand the Reports object in the left pane of the MM Admin window to view the following report types:

- **Vault Selection Report**--Contains a list of tape volumes to be selected for moving into the vaults through the Vault Criteria Descriptor (VCD).
- **Shipping Report**--Contains a list of tape volumes to be pulled from each of the vaults.
- **Shipping Content Report**--Contains a list of tape volumes and sessions in each tape volume to be pulled from each of the vaults.
- **Receiving Report**--Contains a list of tape volumes to be distributed to the vaults.
- **Receiving Content Report**--Contains a list of tape volumes and sessions in each tape volume to be distributed to the vaults.
- **The Inventory Report, By Vault**--Lists tape volumes grouped by the vault where they reside.
- **The Inventory Report, By Media**--Lists tape volumes grouped by vault and shows Media name in front.

When you select a report type in the left pane of the MM Admin window, the right pane displays the contents, listing the available reports identified by date. Click a report to view it in the right lower pane. You can print any of these reports using the Print button on the MM Admin toolbar. At the time it is generated, you can also select to send a report by email if you configure the alert notification system to use Microsoft Exchange. For more information on sending reports using email, see [How the Media Management Process Works](#) (see page 465). For more information on using alerts, see "Using the Alert Manager."

The Inventory Reports are based on information in the Slot table, and can be generated at any time. The Shipping and Receiving Reports are based on movement records generated during a vaulting cycle, and are updated after each vault cycle process completes.

The Vault Selection Listing is produced each time the Start Vault Cycle command is executed. For each VCD processed, this listing identifies the first tape volume in the tape volume set and the controlling data set. This information is provided for all tape volume sets selected for the vaulting cycle.

Find Media in Vault Object

The Find Media in Vault object provides the quickest way to search vaults for a specific media, if, for example, you require that media to execute a restore job. You can choose to search for the media using its Tape Name or its Serial Number (case sensitive).

To open the Find Media in Vault dialog, right-click the Find Media in Vault object, and choose Find from the pop-up menu. Using this dialog you can set the criteria for your media search.

Status Object

MM Admin can run only one vault cycle at a time. To monitor the progress of the vault cycle, or to obtain current online status, double-click the Status object in the left pane of the MM Admin interface to view the following information:

- **Current Status**--The status of the current operation is displayed as either Active or Finished.
- **Last Operator**--The owner of the last operation executed.
- **Last Operation Type**--Operation types can be Ready, Vault Cycle, Commit, Browsing, Update, and Reset.
- **Last Operation Started At**--The date and time the last operation began.
- **Last Operation Finished At**--The date and time the last operation ended.

Reset the Status of Vault Processing

Use MM Admin to manually reset the status of Vault Processing if something goes wrong during the vault cycle, such as corruption of the MM Admin database.

You can use the `ca_mmo` command line utility to reset the status. After the status is reset, you can restart another vault cycle.

Note: For more information about the `ca_mmo` command line utility, see the *Command Line Reference Guide*.

How the Media Management Process Works

The Media Management process includes setting a vaulting policy, scheduling tape volume movement, selecting tape volumes, defining retention policies, executing the vault cycle, and moving the media to the proper location.

After you set a vaulting policy and retention policies, the vaulting rotation process begins. You should run vault cycles as often as you run backup operations. For example, if you back up your data every day, you should also run a vault cycle every day. If you back up your data once a week, run a vault cycle once a week after your backup operation is complete.

The vault cycle process updates location information for tape volume sets, indicating movement into a vault or from a vault back to the Tape Service. You must initiate the process by clicking Start from the Vault Cycle menu on the MM Admin toolbar. You can also initiate the vault cycle using the `ca_mmo -start` or `-startAll` command at the DOS prompt.

Notes:

- When using Media Management Administrator (MM Admin), the vault cycle processes tapes for the primary server and all of the member servers.
- You must click the Start Vault Cycle button every time you want current information on the MM location of the media.

Execute the Start Vault Cycle process to generate reports detailing the movement of the tape volumes and location information. The slots that already contain tape volumes and the new slots that will be vaulted are grouped together by their common schedule. Beginning with the first rotation in the schedule, tape volume sets are assigned to a vault and its slots based on the expiration criteria. Slots are automatically created and tape volumes automatically vaulted during this process.

When the first rotation is satisfied, the next rotation in the schedule is processed, and so on through the entire schedule until all rotations have been exhausted. Media management then generates reports indicating the current location and destination where the tape volumes must be moved. If you do not want to remove these tapes manually, you can use the `ca_mmo -export` command at the DOS prompt so MM Admin automatically exports them. See [Device Manager](#) (see page 395) for more information about command line utilities for media management.

You can use the Simulate Vault Cycle command to produce a Vault Selection Report. Use this command at any time to predict how many tape volumes will be moved without actually updating the location information. If you want to send the Vault Selection Report to someone by email, make sure your system is configured to send alerts using Microsoft Exchange and, from the Configuration menu, enable the option Send the report by E-mail. For more information on configuring alerts, see "Using the Alert Manager."

The vault cycle generates the Shipping and Receiving Reports, listing the old and new locations of the tape volume set, to provide you with the information you need to manage your media. These reports provide the following information:

- **Shipping Report**--tells you what media to pull manually, and where to send it.
- **Shipping Content Report**--lists all tape volumes and sessions in each tape volume to be pulled from each of the vaults.
- **Receiving Report**--tells you what media will be coming in to each particular vault.
- **Receiving Content Report**--lists all tape volumes and sessions in each tape volume to be distributed to the vaults.

If you want to send the shipping and receiving reports to someone by email, make sure your system is configured to send alerts using Microsoft Exchange and, from the Configuration menu, enable the option Send the report by E-mail. For more information on configuring alerts, see "Using the Alert Manager."

When a tape volume comes under Media Management control, Tape Service updates the tape volume's location status to OFF_SITE. To prevent a tape volume from being used while under Media Management control, the tape volume is automatically checked out, and the location is updated to reflect this. Because all vaulted tape volumes are placed in checked out status, if you need to retrieve tape volumes, they must be checked into Tape Service before they can be used.

Vault Management

The first step in establishing a vaulting policy is to create a vault. You can create vaults using MM Admin.

This section contains the following topics:

- [Create Vaults](#) (see page 466).
- [Modify Vaults](#) (see page 467).
- [Delete Vaults](#) (see page 468).

Create Vaults

When you create a vault, location information is automatically updated and integrated with the Location Maintenance feature in Arcserve Backup. If you select a vaulted tape through Arcserve Backup, vault location information appears. Location information is also updated in the Media Pool Manager. If you select a vaulted tape for restore, a message appears indicating that the tape is OFF_SITE.

To create vaults

1. From the Arcserve Backup Home page, open the Media Management Administrator window.
2. Right-click the Vault object and select Create from the pop-up menu.
The Create Vault dialog opens.
3. Enter a name and description for the new vault.
4. Select the Use in Local option if this vault will not be moved to another location. If the tape volumes in this vault are to be maintained off site, do not select this option.
5. Click Add to save and add the vault to the Vault branch in the Media Management Administrator window.
The vault is created.

Modify Vaults

Use the following steps when you want to modify the Vault Name, Vault Description, or Use in Local option.

To modify vaults

1. From the Administration menu in the Navigation Bar on the Home Page, click MM Admin.
The Media Management Administrator window opens.
2. Browse to and double-click the Vault object in the left pane of the Media Management Administrator window.
A list of existing vaults displays.
3. Right-click the vault you want to modify from the list, and select Modify from the pop-up menu.
The Edit Vault dialog opens
4. Make your changes and click OK.
The modified settings for the vault are saved.

Delete Vaults

Use the following steps to delete vaults from the MM Admin.

Note: Before you delete a vault, you must remove all media from the vault and ensure that there are no rotations are associated with the vault.

To delete vaults

1. Open the Media Management Administrator window and browse to the vault that you want to delete.
2. Right-click the vault name and select Delete from the pop-up menu.
A delete confirmation message box opens.
3. If you are sure that you want to delete the vault, click Yes.
The vault is deleted.

Create Schedules

Media Management relies upon a user defined schedule to determine the tape volumes to move, and when and where to move them. When you select the Schedule object, you can view existing schedules in the right pane of the Media Management Administrator window or you can define new rotation policies and vaulting criteria.

To create schedules

1. Open the MM Admin, right-click the Schedule object in the left pane of the Media Management Administrator window, and select Create from the pop-up menu.
The Create Schedule dialog opens
2. On the Create Schedule dialog, specify a name for the schedule and click Add.
The new schedule is saved and added to the Schedule branch in the Media Management Administrator window.
After you create a schedule, the Vault Criteria Descriptor (VCD) and Rotation objects appear in the left pane of the Media Management Administrator window. These objects allow you to select media and retention policies.

Modify Schedules

After a schedule is created and the corresponding VCD and Rotations are configured, you can then modify each of these subordinate settings. For more information, see [Modify Vault Criteria Descriptors \(VCD\)](#) (see page 470) and [Modify Rotations](#) (see page 474).

Delete Tape Volume Movement Schedules

Before you can delete a schedule, you should first make sure that any VCD and Rotation for the schedule have been deleted.

To delete schedules

1. Expand the list of schedules below the Schedule object.
2. Right-click the Schedule object that you want to delete and select Delete from the pop-up menu.

A confirmation window appears.

3. Verify that this is the Schedule that you want to delete and then click Yes.

The schedule is deleted.

How You Can Manage Tape Volumes and VCDs

To assign media to vaults you must specify a VCD and rotation. You can select a media pool, file name, or an individual media for the controlling data set. When this data set is vaulted, its tape volume set is placed in slots in the vault. The assignment of slot numbers is based on the rotation records you defined.

More information:

[Vault Criteria Descriptor Object](#) (see page 460)

Create Vault Criteria Descriptors (VCD)

After you have created a schedule, you must describe the rules for media selection by creating vault criteria descriptors (VCDs).

To create vault criteria descriptors

1. From the Administration menu in the Navigation Bar on the Home Page, click MM Admin.

The Media Management Administrator window opens.

2. Expand the Schedule object, expand a schedule, right-click Vault Criteria Descriptor object, and select Create from the pop-up menu.

The Create Vault Criteria Descriptor dialog appears.

3. Choose one of the following options:
 - **Media Pool Name**--To use a media pool name as the controlling data set, enter the name of the media pool or use the drop-down list to select a media pool name from the pool list. Only the media within the Save Set of the media pool can be vaulted. The media in the Scratch Set cannot be vaulted.
 - **File name**--To use a file name as the controlling data set, select the File Name option and enter the host name and the full path and file name from your backup, such as C:\DOC\Readme.txt, in the appropriate fields. Browse through the Database or Restore Manager to obtain path or file information. The MM Admin finds all tapes used for the backup of this directory or file.
 - **Assigned by User**--If you want to use an individual media as the controlling data set, select the Assign by User option. This is useful in emergencies when you need to use a specific tape. Because MM Admin lets you start a vault cycle only with local media, the media icon appears in yellow if the vaulted media is not a local media with a remote host name. To start a vault cycle with a remote media and members servers, you must execute the ca_mmo command line utility using the -startall argument.

Note: For more information about command line utilities for media management, see the *Command Line Reference Guide*.

4. Click Add.

The VCD is added to the Vault Criteria Descriptor branch in the Media Management Administrator window.

Modify Vault Criteria Descriptors (VCD)

Use the following steps to modify the media pool name, the file name, the assigned by user option associated with the vault criteria descriptor (VCD).

To modify vault criteria descriptors

1. Open the MM Admin, expand the list of schedules under the Schedule object and select a schedule from the list.

Expand the schedule to display the Vault Criteria Descriptor and Rotation objects.

Right-click the Vault Criteria Descriptor object and select Modify from the pop-up menu.

The Edit Vault Criteria Descriptor dialog opens.

2. Modify the Media Pool Name, the File Name, the Assigned by User option associated with the VCD and click OK.

The modified value is applied.

Delete Vault Criteria Descriptors (VCD)

To delete a schedule, you must first delete the associated rotation and vault criteria descriptor (VCD).

To delete vault criteria descriptors

1. From the Schedule object, select the specific VCD from the list under the Vault Criteria Descriptor.
2. Right-click and select Delete from the pop-up menu.
3. Click OK.

Tape Volume Retention Policies

After you create a schedule, you must set the policies governing tape volume retention for your vault. To do this, use the Rotation object.

Note: The Rotation object appears in the left pane of the MM Admin window only after you create schedule.

Special Tape Volume Movement

Special circumstances may arise in which you need to move a particular tape volume. If this situation occurs, you have three options—Temporary Check In, Manual Check In, and Manual Check In and Retire. You also have the option to permanently vault a volume so that it does not return to Tape Service. The following sections describe each of these options.

Temporary Check In

The Temporary Check In option is useful for tracking media movement if you want to temporarily move a tape volume from a vault to use for a restore job, but want to return it back to the vault when the job is finished.

All tape volumes that are vaulted are in *checked out* status. Use the Temporary Check In option to change this status to *checked in* so that you can keep track of your tape volume while it is temporarily being used for a restore job. When you finish using the tape volume, the next vault cycle returns it to the vault and changes the status back to *checked out*.

Note: The Temporary Check In option is only for tracking tapes that are temporarily returned from the vault, and is not a requirement for the actual tape movement; if you do not use this option, you can still manually move a tape volume from a vault and return it when a job is finished. However, you should use this option because, if you do not use it and move a tape volume, there will be a discrepancy between the status of the tape volume that appears in the MM Admin and the actual location of the tape.

To use the Temporary Check In option, in the left pane of the MM Admin window, double-click the Vault object to see a list of existing vaults. Select a vault to display its information in the right pane. Select the name of the media you want to move, right-click, and select Temporary Check In.

Example: Temporary Check In

For example, to perform an emergency restore operation using a tape volume from one of the vaults, use the Temporary Check In feature to temporarily check the tape volume in to Tape Service, execute the restore operation, and then run a vault cycle to return the tape volume to the vault.

Check in Tape Volumes Manually

Use the Manual Check In option to check a tape volume back into Tape Service before the time it is scheduled to be checked in. When you manually check a tape volume back into Tape Service, it does not return to the vault.

To check in tape volumes manually

1. From the left pane of the MM Admin window, double-click the Vault object.
A list of existing vaults displays.
2. Select a vault to display its information in the right pane.
Select the name of the media you want to move, right-click, and select Manual Check In from the pop-up menu.
The tape volume is checked in.

Manual Check In and Retire

Use the Manual Check In and Retire option to check a tape volume back into Tape Service before the time it is scheduled to be checked in, and retire it so it is no longer used.

To use the Manual Check In and Retire option, in the left pane of the MM Admin window, double-click the Vault object to see a list of existing vaults. Select a vault to display its information in the right pane. Select the name of the media you want to move, right-click, and select Manual Check In and Retire.

Permanent Retention

Use the Permanent Retention option to permanently vault slots and the tape volumes they contain. If you use this option, when a tape volume is vaulted, it does not return to Tape Service. The only way to return it is to change the vault status back to the default.

To permanently check tape volumes out of Tape Service, select the Permanent Retention option on the Create Rotation dialog.

Create Rotations

Media Management relies upon rotation policies to determine when and where tape volumes should be moved. You can define rotation policies using the Media Management Administrator interface.

Before you create a rotation, you must create a schedule. If no schedule has been created, you cannot access the Rotation object.

To create a rotation, perform the following procedure:

1. In the Media Management Administrator windows, double-click the Schedule object.
2. From the list, select the schedule and double-click.
3. Right-click the Rotation object.
4. Select Create from the pop-up menu.
The Create Rotation dialog appears.
5. (Optional) Enter a Sequence Number.
6. Enter a Vault Name for your rotation. You can select a vault name from the list.
7. Enter Retention conditions.

If a tape volume meets one of these conditions, it remains in the same rotation. The Retention conditions are:

- **Hold Days**--The number of days you want tape volumes to be retained.
 - **Keep for Cycles**--The number of vault cycles you want tape volumes to be retained in this rotation.
 - **Days Elapsed from First Format Date**--Starting from the day the tape volumes were first formatted, enter the number of days you want tape volumes to remain in this rotation.
 - **By Date**--Tape volumes are retained in this rotation until the date you enter here is reached.
 - **By Tape Expiration Date**--Tape volumes are retained in this rotation until their expiration date passes.
 - **Permanent**--All tape volumes are retained in this rotation permanently.
8. Click Add.

The new Rotation is saved and added to the Rotation branch in the Media Management Administrator windows.

Modify Rotations

Use the steps that follow to modify the movement of tape volumes associated with a schedule.

To modify rotations

1. Expand the Schedule object in the left pane of the MM Admin window and select a schedule from the tree.
2. Double-click the schedule to access the Rotation object.
3. Double-click the Rotation object and select a rotation in the right pane.
4. Right-click the rotation and select Modify.

The Edit Rotation dialog appears

5. Apply your changes and click OK.

The new settings are saved.

Delete Rotations

If you want to delete a schedule, you must first delete the associated rotation and VCD.

To delete rotations

1. From the left pane of the MM Admin window, expand the schedule object, and the expand schedule for the rotation that you want to delete.

The rotation displays in the right pane of the MM Admin window.

2. From the right pane of the MM Admin window, select the rotation from the list.
3. Right-click the rotation that you want to delete and select Delete from the pop-up menu.

A confirmation message appears

4. Click Yes.

The rotation is deleted.

Slot Detail and Status Information

When tape volumes have been assigned to slots in a vault, MM Admin displays slot information for the vault. Select the Vault object in the left pane of the MM Admin window and expand it. When you select a particular vault from the list, the right pane of the MM Admin window displays a view of the vault and its slots.

This view provides the following information:

- **Media Name**--Lists the media name, ID, sequence number, and serial number.
- **Slot Status**--Either Active, Unvaulted, Temporary Check In, Manual Check In, or Manual Check In and Retire:
 - **Active**--The media has been sent to this vault.
 - **Unvaulted**--The media has not yet been sent to this vault.
 - **Temporary Check In**--The media will be checked in temporarily during the next vault cycle.
 - **Manual Check In**--The media will be checked in during the next vault cycle.
 - **Manual Check In and Retire**--The media will be checked in and retired during the next vault cycle.
- **Slot Name**--Lists the vault name and slot number.
- **Media Export Status**--Either Ready, Success, or Fail:
 - **Ready**--The default status. The media has been assigned to the vault, but has not been exported from the tape library to the mail slot yet.
 - **Success**--Appears once the media is successfully exported to the mail slot.
 - **Fail**--Appears if MM Admin failed to export the media to the mail slot.
- **Local**--Either Yes or No. Yes appears if the media belongs to a local machine. No appears if it belongs to a remote machine.
- **Create Date**--The date the slot was created.

When you highlight a media name, additional information appears in the Properties pane in the lower right-hand corner of the page. This information includes the media name, serial number, random ID, host name, slot status, slot name, media export status, media type, media class, last write, last read, and slot creation date. Because MM Admin lets you start a vault cycle only with local media, the media icon appears in yellow if the vaulted media is not a local media with a remote host name. If you want to start a vault cycle with a remote media and member servers, use the `ca_mmo -startall` command line utility.

Note: For more information about command line utilities for media management, see the *Command Line Reference Guide*.

Because slots are automatically created when a tape volume is vaulted, you typically have no reason to update slot information.

Find Specific Media in a Vault

To help you locate media in your vaults, MM Admin provides the Find Media in Vault feature. This feature is the fastest way to locate media in your vaults if you know the tape name or serial number of the tape volume you need. If you do not know this information, you can use the Database Manager to find the media.

To find a specific media in a vault

1. From the Administration menu in the Navigation Bar on the Home Page, click MM Admin.

The Media Management Administrator window opens.

2. From the Media Management Administrator window, right-click the Find Media in Vault object and select Find from the pop-up menu.

The Find Media in Vault dialog opens.

3. Select one of the following methods to find your media:

- **Find by Tape Name**--Lets you enter the tape name, the random ID, and the sequence number to identify the tape you want Arcserve Backup to find.
- **Find by Serial Number**--Lets you enter the serial number of the desired media.

Arcserve Backup completes the Find by Serial Number task using case-sensitive values. For example, the serial number ABC123 is different from the serial number abc123.

4. Click Find.

When the search is finished, the vault and slot information appears in the right pane of the Media Management Administrator window.

How to Manage Cloud Storage

Arcserve Backup provides you with various tools and options that you can use to manage cloud storage. This section includes information about how you can create cloud connections, devices, and cloud-based device management.

Note: For reports on backing up data to a cloud storage, see the following dashboard reports in the *Dashboard User Guide*:

- Backup Data Location Report
- Data Distribution on Media Report
- Recovery Point Objective Report

This section contains the following topics:

[Cloud Storage Limitations](#) (see page 477)

[Configure Cloud Connections](#) (see page 478)

[Create Cloud-Based Devices](#) (see page 480)

[Browse Cloud Folders](#) (see page 483)

[Format Cloud-Based Device](#) (see page 485)

[Erase Cloud-Based Device](#) (see page 486)

[Online and Offline Cloud-Based Device](#) (see page 486)

[Specify Compression on Cloud-Based Device](#) (see page 487)

[Migrate Data to Cloud Storage in a Staging Backup Job](#) (see page 487)

[Configure Throughput to Eucalyptus-Based Cloud Media](#) (see page 489)

Cloud Storage Limitations

Consider the following limitations when backing up data to cloud storage:

- Arcserve Backup does not support normal backups to cloud-based devices. Currently, Arcserve Backup allows the migration to cloud-based devices.
- Arcserve Backup does not support Tapecopy using the cloud-based device feature.
- Arcserve Backup does not allow multiple cloud-based devices to point to a single cloud folder or group at the same time. This may cause data corruption.

Configure Cloud Connections

A cloud connection is a user-defined configuration that contains the information that Arcserve Backup uses to communicate with the specified cloud vendor. Cloud connections let you access cloud storage services through applications that are on the current server. The cloud connection helps ensure that your backup data is stored under the account that you created with the cloud vendor.

For example, in Arcserve Backup you would only specify the cloud connection name when creating one cloud-based device. By providing this information, all cloud related items are transparent for Arcserve Backup users.

Note: To eliminate potential clock skew errors when connecting to the cloud, verify that your computer has the correct time zone specified, and the clock is synchronized with the global time. As a best practice, you should always check the time of your computer against the GMT time. If the time of your computer is not synchronized with the global clock time (within 5 to 10 minutes), cloud connection will not work. If necessary, reset the correct time for your computer and then resubmit the archive job.

Follow these steps:

1. Open the Arcserve Backup Manager Console.
From the Navigation Bar, expand Administration and click Device Manager.
The Device Manager window opens.
2. From the Servers directory tree, locate the server you want to configure.
Select Manage Cloud Connection from the toolbar to configure your cloud connections.

Note: You can also select Manage Cloud Connection by right-clicking the specific server in the Servers directory tree and select Manage Cloud Connections or click the Manage Cloud Connection link in the Servers details section.

The Cloud Connection Configuration dialog opens.

3. Click the desired cloud vendor and click Add.

The Add Connection dialog opens.

4. Complete the required fields on the Add Connection dialog and click OK.

Note: The fields that appear on the Add Connection dialog vary based on the cloud vendor that you are using.

The Add Connection dialog contains the following fields:

- **Connection Name**--Specify the name of the cloud connection.
- **Access Key ID or Query ID**--Specify the account name for the cloud connection.
- **Secret Access Key or Secret Key**--Specify the password for the cloud connection.
- **Vendor URL**--Specify the vendor's URL. Typically, this field defaults to the URL of the public vendor you selected. If you selected a private vendor, it is then required for you to enter the URL.

- **Description**

- **Do not use proxy**--Select this option if you want a cloud vendor connection to bypass the settings from the Proxy Settings dialog. Arcserve Backup directly connects to the cloud vendor instead of going through the proxy.

- **Advanced**

- **Bucket Name**--Specify a unique bucket name to store your data.

Note: Bucket Names can be from 3 through 63 characters in length and cannot contain the following:

-- Underscores, for example: _

-- Uppercase characters, for example: AB

-- Adjacent periods, for example: ..

-- End with a hyphen, for example: name-

-- An adjacent hyphen and period, for example: my-.bucket.com

-- Non-English characters

-- End with a period, for example: name.

- **Bucket Region**--Specify a region where your bucket resides.

- **Enable Reduced Redundancy Storage**--Select this option to reduce your cost by storing noncritical, reproducible data at lower levels of redundancy than the standard storage.

5. Click OK to close the Add Connection dialog to save the configuration and return to the Cloud Connection Configuration dialog.

6. (Optional) On the Cloud Connection Configuration dialog, click Proxy Settings to configure the proxy settings for all cloud vendor connections. You can automatically detect, automatically configure the script, or set the proxy server manually.
7. From the Cloud Connection Configuration dialog, click Add to add more cloud connections or click close to apply the configurations.

Create Cloud-Based Devices

A cloud-based device is a virtual, Arcserve Backup device that you create in association with a Arcserve Backup cloud connection. The cloud-based device lets Arcserve Backup store data with the cloud vendor specified by the Arcserve Backup cloud connection.

Arcserve Backup lets you create cloud-based devices that link to cloud folders. You can link cloud-based devices to new or existing cloud folders. If the cloud folder is configured with password protection, you must specify a password otherwise you cannot link the cloud-based device to the cloud folder.

Note: The cloud folder must reside in the subdirectory under the root path of the cloud account or connection. For example, you can enter "directoryA" as the cloud folder. You cannot enter a cloud folder that contains '\' character.

Arcserve Backup does not allow the following characters when creating a cloud folder:

- asterisk (*)
- question mark (?)
- forward slash (/)
- back slash (\)
- less-than sign (<)
- greater-than sign (>)
- vertical line (|)
- semicolon (;)
- blank space
- ampersand (&)
- all non-English characters

Important! You cannot link multiple cloud-based devices from different servers under a single cloud folder at the same time. This may cause data corruption.

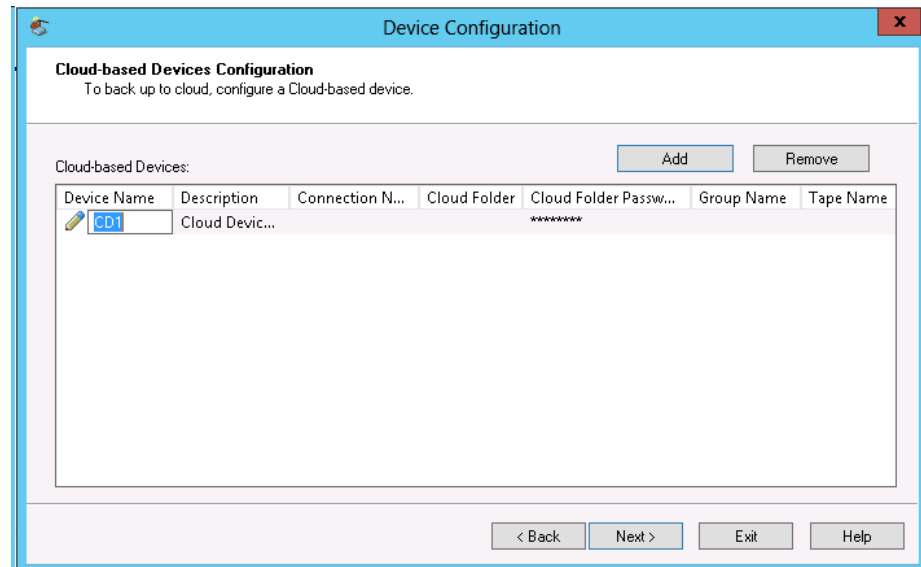
Follow these steps:

1. Open the Arcserve Backup Manager Console.
From the Navigation Bar, expand Administration and click Device Configuration to open Device Configuration.
2. On the Welcome to Device Configuration dialog, select the Cloud-based Device Configuration option and click Next.
3. From the Login Server dialog, provide the required login credentials for the Primary Server and click Next.
4. From the second Login Server dialog, select the server that you want to log in to and click Next to open the Cloud-Based Device Configuration dialog opens. The Cloud-Based Device Configuration dialog displays a list of Devices and any corresponding cloud information.

Note: You create the cloud connections before you create cloud-based devices. For more information, see [Create Cloud Connection Configuration](#) (see page 478).

5. Click Add.

A new, blank device is added.



6. Complete the following fields on the Cloud-Based Devices Configuration dialog:
 - **Device Name**--In the Device Name field, enter a name, or accept the default.
 - **Description**--In the Description field, enter a description, or accept the default.
 - **Connection Name**--In the Connection Name field, select a cloud connection name from the drop-down.
 - **Cloud Folder**--In the Cloud Folder field, enter the folder name where the cloud-based device resides, or click the arrow button next to the field to launch a new window to browse for a cloud folder for the specific cloud-based device.
Note: You can also browse for an existing cloud folder directly from the Device Manager screen. For more information, see [Browse Cloud Folder](#) (see page 483).
 - **Cloud Folder Password**--(Optional) In the Cloud Folder Password field, enter a password to protect the cloud folder from unauthorized access and to create a cloud-based device.
 - **Group Name**--(Optional) In the Group Name field, enter a name for the group.
7. Click Next to add the device.

When the configuration completes, a list of your cloud-based devices appears with a status of your cloud connection.

Note: You can also launch the Cloud-based Device Configuration through Device Manager and Device Configuration:

■ **Device Manager:**

Click the Manage Cloud-based Device button from the toolbar;



■ **Device Configuration:**

1. From the Manager Console, click the Navigation Bar and expand the Administration menu.
2. Click Device Configuration.
The Welcome screen for Device Configuration opens.
3. Select the Cloud-based Device Configuration option.

Browse Cloud Folders

Arcserve Backup lets you browse existing cloud folders directly from the Device Manager home screen.

The Browse Cloud Folder launches a window to browse for a cloud folder for a specific Cloud Connection.

To browse for a cloud folder

1. Open the Arcserve Backup Manager Console.

From the Navigation Bar, expand Administration and click Device Manager.


The Device Manager window opens.

2. From the Device Manager, there are three ways you can launch the Browse Cloud Folder feature:
 - Right-click the specific server in the Servers directory tree and select Browse Cloud Folder;
 - Click the Browse Cloud Folder button from the toolbar;

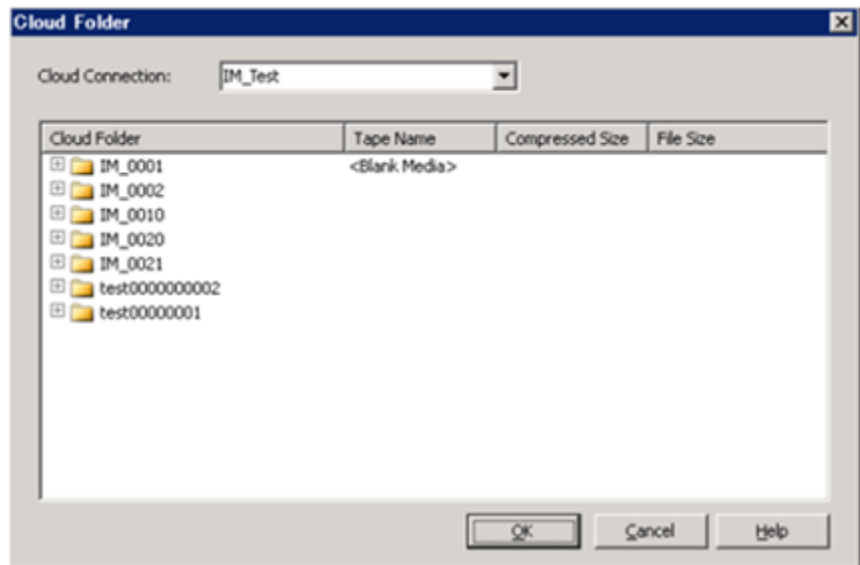


OR

- Click the Browse Cloud Folder link in the Servers details section.



The Browse Cloud Folder window appears.



3. Select a Cloud Connection Name to display the list of Cloud Folders where your Cloud-based Device resides.

The Cloud Folder displays the tape name of the cloud-based device, the compressed size and file size of the cloud folder.

Be aware of the following:

- You can also browse for a cloud folder when creating a cloud-based device. For more details, see [Create Cloud-Based Devices](#) (see page 480).
 - From the Cloud-Based Device Configuration screen, you can select a Cloud Folder to mount as a Cloud-based Device, then click OK.
- You can right-click a folder or file to delete. This helps you manage the Arcserve Backup objects created in the cloud.

Format Cloud-Based Device

Although Arcserve Backup automatically formats blank devices during a migration job, you can use this option to manually format your cloud-based device. Formatting writes a new label at the beginning of the device, effectively destroying all existing data on the cloud-based device.

Note: Use this option with care. After the cloud-based device is formatted, Arcserve Backup can no longer restore the data and any job sessions associated with the media.

To format cloud-based devices

1. Open the Device Manager and select the specific cloud-based device from the Servers directory tree.
2. Right-click the cloud-based device and click Format on the pop-up menu or click Format on the toolbar.

The Format dialog opens.

3. Assign a New Media Name to the cloud-based device that you want to format.
4. Click OK.

The Format dialog closes and the following message appears:

"Formatting will erase ALL of your data from the media. Do you want to format the media?"

5. Do one of the following:
 - To start the formatting process, click OK.
Arcserve Backup formats the cloud-based device.
 - To cancel the formatting process, click Cancel.
Arcserve Backup does not format the cloud-based device.

Erase Cloud-Based Device

Use this option to erase one cloud-based device at a time. Arcserve Backup also erases all references to the contents of this cloud-based device (if any) from the database. When you reformat this cloud-based device, its physical history (read and write passes) is retained.

Verify that you have selected the correct cloud-based device before using the Erase option. Erased data cannot be retrieved. When erasing cloud-based devices, you can choose from the following options:

- **Quick Erase**--Quick Erase effectively erases the cloud-based device. It avoids the time a Long Erase would take (minutes to hours) by overwriting the device label. The history remains available to Arcserve Backup for tracking purposes.
- **Quick Erase Plus**--This option performs the same operation as Quick Erase, and also erases bar codes and serial numbers. For more information about bar code and serial number cataloging, see the Mount and Dismount Option.

Note: If the cloud-based device you are erasing does not have a serial number or bar code, this option functions in the same manner as the Quick Erase option.

If the Quick Erase Plus option is used, it can no longer be tracked by Arcserve Backup, and information such as the expiration date is no longer carried forward.

To erase cloud-based devices

1. Open the Device Manager and select the specific cloud-based device from the Servers directory tree.
2. Right-click the cloud-based device and click Erase on the pop-up menu or click Erase on the toolbar.

The Erase dialog opens.

3. Select an erase method, click OK, and then click OK to confirm.

Arcserve Backup erases the cloud-based device.

Online and Offline Cloud-Based Device

You can set cloud-based devices as offline or online from the Device Manager by right-clicking on the drive and selecting offline or online, depending on the current state of the drive.

This capability can be useful for marking defective devices as offline, which prevents Arcserve Backup from using the cloud-based device until it is repaired and in an online status.

To online and offline cloud-based devices

1. Open the Device Manager and browse to the server connected to the library containing the cloud-based device that you want to status as online or offline.
2. Expand the library, right-click the cloud-based device, and select Online or Offline from the pop-up menu.

The status changes to offline or online.

Note: Offline state appears disabled, when selected.

Specify Compression on Cloud-Based Device

Arcserve Backup lets you compress the backup data that is stored on the cloud-based device.

Note: The following steps describe how to turn compression on and off.

To specify compression on cloud-based devices

1. Open the Device Manager and select the specific cloud-based device from the Servers directory tree.

If the device supports compression, then Arcserve Backup enables the Compression toolbar button. To verify if the device supports compression, select the Detail tab when the device is highlighted.

2. Right-click the cloud-based device and click Compression on the pop-up menu or click Compression on the toolbar.
3. Click OK to set the Compression Mode to Off (if it is On) or On (if compression is Off).

Migrate Data to Cloud Storage in a Staging Backup Job

The process of setting up a staging backup job lets you specify how long you want to retain the data on the staging device. Arcserve Backup lets you purge the data from the staging device or migrate the data to a final destination device after the retention time of a session expires. The final destination device can be cloud storage.

To migrate data to cloud storage in a staging backup job

1. Open the Backup Manager and click the Start tab.

From the Start tab, click Normal Backup and Enable Staging.

The Staging Location and Policies tabs appear in the Backup Manager.

2. Click the Source tab.
From the Servers directory tree, specify the data that you want to back up.
3. Click the Staging Location tab and expand the Staging Servers object.
 - a. Browse to and select the group you want to choose as the staging group for this backup job.
Note: Cloud group cannot be selected as a staging group.
 - b. Click the Policies tab to specify copy policies.
 - c. Specify the staging policies for full, differential, and incremental backups required for your job.
4. Click the Destination tab and expand the Servers object.
 - a. Browse to and select the cloud group you wish to use as the final destination for this backup job.
Note: For information about creating cloud-based devices, see [Create Cloud-Based Devices](#) (see page 480).
 - b. Click Cloud Policy to open the Cloud Purge Policy dialog.
 - c. Specify the cloud purge policy for backups required for the job.
 - d. Click OK.
5. Click the Schedule tab and specify the schedule that you want to use for the backup job.
6. Click Options on the toolbar.
The Global Options dialog opens.
7. Specify the global options that you want to apply to the job and then click OK.
Note: For more information about Global Options, see [Global Backup Options](#) (see page 154).
8. Click Submit on the toolbar to submit the job.
The Submit job dialog opens.
9. Complete the required fields on the Submit job dialog and then click OK.
The job is submitted.

More information:

[Submit a Backup Job](#) (see page 136)

[Global Backup Options](#) (see page 154)

[Specify Copy and Purge Policies for Disk Staging Backups](#) (see page 213)

Configure Throughput to Eucalyptus-Based Cloud Media

You can define the block size of data that Arcserve Backup migrates to Eucalyptus-based cloud storage media. This capability is useful in that you can define a block size that suits the capabilities of your network. For example, when the network bandwidth is high, you define a higher block size. Conversely, when the network bandwidth is low, you define a lower block size. Keep in mind, when communication between the source server and the Eucalyptus-based cloud storage media fails, Arcserve Backup retransmits the chunks of data that were transmitted previously after communication resumes.

The default chunk size is 10 MB. You can define the value of the chunk size as from 1 MB to 50 MB.

Follow these steps:

1. Open the following .ini file on the Arcserve Backup server:

```
$ARCSERVE_HOME\CCIconfig\CCIconfig.ini
```

2. Locate the following section in the CCIconfig.ini file:

```
[Eucalyptus_Config]  
ChunkSize=1048576
```

3. Specify the value of the chunk size that you require (in bytes).

Examples:

```
5 MB = 5242880 bytes
```

```
25 MB = 26214400 bytes
```

4. Close the CCIconfig.ini file.

Chapter 8: Administering the Backup Server

This section provides you with information that you can use to administer, manage, and maintain the Arcserve Backup server.

This section contains the following topics:

[How Arcserve Backup Engines Work](#) (see page 491)

[Configure Arcserve Backup Engines](#) (see page 499)

[Additional Server Admin Functions](#) (see page 520)

[Authentication Levels for Arcserve Backup Services, Components, and Applications](#) (see page 527)

[Arcserve Backup Domains](#) (see page 538)

[Manage User Profiles Using the User Profile Utility](#) (see page 554)

[Restore the Arcserve Backup Job Queue](#) (see page 557)

[BSBAB--Manage Arcserve Servers Using the Server Configuration Wizard-AdminW](#) (see page 559)

[How Arcserve Backup Protects Active Directory Data on Domain Controller Servers](#) (see page 579)

[Install and Uninstall Arcserve Backup Server Based Options](#) (see page 590)

[Discovery Configuration](#) (see page 591)

[Arcserve Backup Maintenance Notifications](#) (see page 598)

[Apply Arcserve Backup Component Licenses](#) (see page 601)

[Managing Firewalls](#) (see page 602)

How Arcserve Backup Engines Work

The Arcserve Backup Server consists of three functional engines:

- **The Job Engine**--This engine processes your jobs at their designated date and time. It scans the job queue for a job that is ready to run, then sends it to the appropriate handler.
- **The Tape Engine**--This engine communicates with, and controls, your storage devices. The Tape Engine selects the device needed for a job.
- **The Database Engine**--This engine maintains a history of:
 - Information about jobs processed by Arcserve Backup, such as the job type, the final result, the start and end time, submitter, and description.
 - Media used by Arcserve Backup, such as its type, its name, the date it was first formatted, its expiration date, and the sessions on it.
 - Files, directories, drives, and machines that Arcserve Backup has backed up or copied.

You can control these Arcserve Backup engines in the Server Admin. To view information about an individual engine, open the Server Admin from the Quick Start menu in the Navigation Bar on the Home Page. From the Arcserve domain directory tree, select the primary server, member server, or stand-alone where you want to obtain engine status information.

Important! To manage and configure Arcserve Backup engines, you must be logged in to Arcserve Backup with the caroot password or a Arcserve Backup Administrator account.

- **Job Engine**--Displays information about the jobs submitted, such as the total number of jobs and the number of ACTIVE, READY, HOLD, and DONE jobs. It also shows the queues, which ones are being scanned, and the scanning interval.
- **Tape Engine**--Displays information about jobs using the Tape Engine, such as the type of job, and who submitted it. It also displays information on media groups.
- **Database Engine**--Displays pruning information related to the Arcserve database.

How Engine Status Affects Arcserve Backup Operations

A stopped engine is an engine that is completely offline. This may be caused by errors, manual shutdown, or a new installation. Whatever the reason, it means that the services of that engine are not available.

The Arcserve Backup engines are designed to run independently of each other. For example, if you stop the Tape Engine, the Database Engine and the Job Engine are not affected. They continue to run, performing their services as configured. The Database Engine continues to log pertinent Arcserve Backup information in the database, and the Job Engine continues to scan the job queue and start jobs as required. If a job requires a storage device, the Job Engine launches the job, but the job fails because the Tape Engine is not able to communicate with the storage device. The Database Engine then logs this information.

Note: Although Arcserve Backup still functions if one or two engines are not running, Arcserve Backup needs all three engines running simultaneously to achieve complete functionality.

Service State Icons

The toolbar at the top of each Arcserve Backup manager displays an icon for each of the back-end services--Job Engine, Tape Engine, and Database Engine, as shown by the following illustration:



Depending upon the color, the icons indicate one of the following states:

- **Green**--Indicates that the service is running.
- **Red**--Indicates that the service is not running.
- **Gray**--Indicates that the service cannot be connected to or is in an unknown state.
- **Blue**--Indicates that the service is paused.

Stopping and Starting Arcserve Backup Services

The following sections describe the methods that you can use to stop and start the Arcserve Backup services on primary, stand-alone, and member servers.

This section contains the following topics:

[Stop and Start All Arcserve Backup Services Using Batch Files](#) (see page 493)

[Stop and Start Individual Services Using the Command Line](#) (see page 496)

[Stop and Start Arcserve Backup Services Using the Server Admin](#) (see page 497)

Stop and Start All Arcserve Backup Services Using Batch Files

There are two methods that you can use to manually stop and start Arcserve Backup services, such as the Job Engine, the Tape Engine, and the Database Engine.

One method is to open the Server Admin, select the server name from the domain tree, select the individual service that you want to stop or start, and then click the Stop or Start toolbar buttons. However, circumstances may arise that require you to stop all Arcserve Backup services. For example, you need to apply a patch or fix that was released by Arcserve Support.

The `cstop` and the `cstart` commands let you shut down and restart all Arcserve Backup services sequentially, based upon their dependencies to the other Arcserve Backup services. This process ensures that there is no loss of data while the services are shutting down, and that all Arcserve Backup services are running properly when your system restarts.

To stop or start all Arcserve Backup services using a single command, use the file `cstop.bat` or `cstart.bat` located in the Arcserve Backup home directory.

cstop.bat

When you run `cstop.bat`, Arcserve Backup stops the services in the following order:

1. Arcserve Communication Foundation (Global)
2. Arcserve Dashboard Sync Service
3. Arcserve Central Remoting Server
4. Arcserve Communication Foundation
5. Arcserve Management Service
6. Arcserve Tape Engine
7. Arcserve Job Engine
8. Arcserve Database Engine
9. Arcserve Message Engine
10. Arcserve Discovery Service
11. Arcserve Domain Server
12. Arcserve Service Controller
13. Arcserve PortMapper
14. Alert Notification Server
15. Arcserve Universal Agent

cstart.bat

When you run cstart.bat, Arcserve Backup starts the services in the following order:

1. Alert Notification Server
2. Arcserve Discovery Service
3. Arcserve PortMapper
4. Arcserve Service Controller
5. Arcserve Domain Server
6. Arcserve Database Engine
7. Arcserve Message Engine
8. Arcserve Tape Engine
9. Arcserve Job Engine
10. Arcserve Management Service
11. Arcserve Universal Agent
12. Arcserve Communication Foundation
13. Arcserve Central Remoting Server
14. Arcserve Dashboard Sync Service
15. Arcserve Communication Foundation (Global)

Be aware of the following behavior, as it relates to stopping and starting Arcserve Backup Global Dashboard services:

- Arcserve Backup Global Dashboard requires the following services for Central Primary server configurations:
 - Arcserve Communication Foundation (Global)
 - Arcserve Dashboard Sync Service
 - Arcserve Central Remoting Server
 - Arcserve Communication Foundation
- Arcserve Backup Global Dashboard requires the following services for Branch Primary server configurations:
 - Arcserve Dashboard Sync Service
 - Arcserve Communication Foundation
- When you execute cstop.bat and cstart, Arcserve Backup stops and starts the services that correspond to the type of primary server that you installed (Central Primary server or Branch Primary server).

Stop and Start Individual Services Using the Command Line

Circumstances may arise that require you to stop and start only one or two Arcserve Backup services. Arcserve Backup lets you use the command line to stop individual services.

To stop and start the Arcserve Backup services using the command line

1. Start the Windows command line.
2. After the command line opens, enter one of the following commands:
 - NET START [enginename]
 - NET STOP [enginename]

Substitute one of the following for [enginename]:

- Arcserve Communication Foundation (Global)
Arcserve Communication Foundation (Global)
- Arcserve Dashboard Sync Service
CADashboardSync
- Arcserve Central Remoting Server
Arcserve_RemotingServer
- Arcserve Communication Foundation
Arcserve Communication Foundation
- Arcserve Management Service
CASMgmtSvc
- Arcserve Tape Engine
CASTapeEngine
- Arcserve Job Engine
CASJobEngine
- Arcserve Database Engine
CASDbEngine
- Arcserve Message Engine
CASMessageEngine

- Arcserve Discovery Service
CASDiscovery
- Arcserve Domain Server
CasUnivDomainSvr
- Arcserve Service Controller
CasSvcControlSvr
- Arcserve PortMapper
CASportmap

Note: If stop and restart the CA Remote Procedure Call service (CASportmap) using the Command Line (or the Computer Management console), the service cannot communicate with its port assignments properly. This can prevent a user account with caroot equivalence from logging in to the Arcserve Backup domain. To remedy the inability to log in to the Arcserve Backup domain, run the cstop command and then run the cstart command. This enables the service to communicate properly and lets the user account with caroot equivalence log in to the Arcserve Backup domain.

- Alert Notification Server
"Alert Notification Server"
- Arcserve Universal Agent
CASUniversalAgent

Note: Repeat this step to start and stop each Arcserve Backup service.

Stop and Start Arcserve Backup Services Using the Server Admin

Using the Server Admin, you can stop and start individual Arcserve Backup services that are running on a primary, stand-alone, and member server.

Use this method when you need to stop one or two Arcserve Backup services for a short time. For example, you need to stop and start the Tape Engine on the primary server so that it can detect a newly installed library.

When you need to stop and start all Arcserve Backup services, you should use the cstop and cstart batch files. These batch files let you stop and start all Arcserve Backup services sequentially, based on their dependencies to other Arcserve Backup services. For more information, see [Stop and Start All Arcserve Backup Services Using Batch Files](#) (see page 493).

Be aware of the following behavior regarding stopping all Arcserve Backup services:

- If you use the Server Admin to stop all services, the service status displays as unknown.
- The Stop all Services option lets you stop all Arcserve Backup services except the Arcserve Service Controller service. Arcserve Backup behaves in this manner because the Arcserve Service Controller service controls the starting of Arcserve Backup services.

To stop and start Arcserve Backup services using the Server Admin

1. From the Quick Start menu in the Navigation Bar on the Home Page, click Server Admin.

The Server Admin opens.

2. Expand the domain directory tree and select the server where you want to stop or start Arcserve Backup services.

The Name, Status, Up Time, and Description of the Arcserve Backup services display in the Server Admin window.

3. Select the service that you want to stop or start.

- If the status is Started, click Stop on the toolbar.
- If the status is Stopped, click Start on the toolbar.

The Arcserve Backup service stops or starts.

4. (Optional) To stop all Arcserve Backup services running on a Arcserve Backup server, right-click the server and click Stop all services on the pop-up menu. To restart all Arcserve Backup services on the server, right-click the server and click Start all services on the pop-up menu.
5. (Optional) To stop all Arcserve Backup services running on all Arcserve Backup servers in a domain, right-click the domain and click Stop all services in domain on the pop-up menu. To restart all services on all servers in a domain, right-click the domain and click Start all services in domain on the pop-up menu.

Configure Arcserve Backup Engines

The Arcserve Backup Server Admin allows you to configure each engine to suit your needs.

Important! To manage and configure Arcserve Backup engines, you must be logged in to Arcserve Backup with the caroot password or a Arcserve Backup Administrator account.

To configure Arcserve Backup engines

1. Open the Arcserve Backup Server Admin by clicking Server Admin in the Quick Start menu.

The Server Admin window opens.

2. Click Configuration on the toolbar.

The Server Admin Configuration dialog opens.

3. Select the desired engine tab and specify the settings that suit your needs.

More information:

[Job Engine Configuration](#) (see page 500)

[Tape Engine Configuration](#) (see page 503)

[Database Engine Configuration](#) (see page 514)

[Alert Configuration](#) (see page 518)

Job Engine Configuration

The Arcserve Backup Job Engine controls the execution time of jobs in the job queue. It scans the job queue regularly, starting jobs as their execution dates and times are reached. Arcserve Backup provides the following job engine options:

- **Job Queue Scanning Interval (seconds)**--The Job Engine constantly scans the job queue for jobs that should execute. By default, the job queue is scanned every 10 seconds. To change the time interval, specify a number from 1 through 9999.
- **Retention Time for DONE Job (hours)**--Jobs with a final status of DONE remain in the job queue for the time specified in this field. By default, Arcserve Backup keeps DONE jobs for 24 hours before they are deleted from the queue. To change the time, specify a number from 0 through 999.

Note: Single occurrence staging jobs (disk to disk to tape and disk to tape to tape) will be removed from the job queue after the migration phase of the job is complete and the length of time specified for this option has elapsed.

- **Database Polling Interval (minutes)**--The Job Engine periodically polls the Arcserve Backup database to discover copied and purged sessions on staging enabled devices. The value specified in this field determines the time interval between polls. The default value for this field is five (5) minutes, and the minimum value is one (1) minute.
- **Message Type in Activity Log**--The Activity Log contains information about all Arcserve Backup activities. By default, notes, warnings, and errors that occur when running Arcserve Backup appear in its Activity Log. To change the types of messages, specify one of the following values:

None

No messages appear.

Errors

Only errors that occur while running Arcserve Backup appear.

Warnings & Errors

Warnings and errors that occur while running Arcserve Backup appear.

Notes, Warnings & Errors (default)

Includes all notes, warnings, and errors that occur while running Arcserve Backup.

Debug

Includes debugging information and all notes, warnings, and errors that occur while running Arcserve Backup.

- **Network Shares**--By default, Arcserve Backup opens Use All Shares in the Browser. This means that both Default Shares and User Shares are available for selection as either your source or destination for a job. To change the type of shares that are displayed in the Browser, specify one of the following:

Use Default Shares Only

Only administrative shares are available.

Use Users Shares Only

Only shares that have been specifically set by users are displayed.

- **Buffer Size (K Bytes)**--Defines the buffer size used by Arcserve Backup.

Default value: 256 KB

All computers behave differently. Factors affecting their behavior can be related to backup server hardware, the total size of the backup job, and the number child jobs that a backup job spawns. You can increase or decrease the size of the buffer to optimize your system's performance while performing a backup.

Increasing or decreasing the buffer size does not necessary improve backup and restore performance. For example, if the backup server has abundant system resources, such as a large amount memory, access to a fast network, and fast disk I/O, increasing the buffer size can increase the system's backup and restore performance. Conversely, if the backup server has limited system resources, reducing the buffer size can increase the system's backup and restore performance.

Note: For a typical server, the best practice is to apply the default value of the buffer size.

- **Backup**--Allows you to customize additional options in your backup jobs:

Record Hard Links for NTFS Volumes

If you back up hard links files, this information is included and preserved by default.

Default value: ON

Confirm when Overwriting Media

Any time a media is to be overwritten, Arcserve Backup can prompt you to confirm that you really want to overwrite the media. By default, this option is disabled (OFF). If you set this option, a confirmation dialog is displayed. If you do not respond within five minutes, the job is canceled.

Default value: OFF

Backup Registry key details when an entire machine is selected

You can turn on the option to back up the Registry key details for target machines by clicking the Backup registry key details when an entire machine is selected check box.

Default value: OFF

Enable Media Maximization

Lets you optimize disk and tape usage in GFS and rotation jobs. For more information, see [Media Maximization](#) (see page 439).

You cannot change this value from a member server. Member servers inherit the value specified for this option from the Arcserve Backup primary server.

Default value: ON

- **Retry Crashed Jobs after Job Engine Restart**--This option is a checkpoint mechanism. Arcserve Backup attempts to restart a crashed job if this box is checked. It should only be turned on if a cluster environment is configured to allow for failover.
- **Submit Makeup Jobs on HOLD**--Use this option to place a hold status on a job rather than a ready status.
- **Block pop-ups when data migration jobs finish**--When a staging migration job is finished, pop-up messages display to inform you if a job was successful, failed, and so on. Specify this option when you do not want pop-up messages to appear after the migration job is finished.
- **Block pop-ups when any job finishes**--When a job is finished, pop-up messages display to inform you if a job was successful, failed, and so on. Specify this option if you do not want pop-up messages to appear after a job is finished.

- **Enable archive CSV file**--Specify this option to create a CSV file for the archived files. By default, this option is disabled.
- **Archive CSV file folder**--This field lets you specify the folder where you want to store the Archive CSV file. By default, Arcserve Backup stores the archive CSV files in BAB_HOME\Archived files.

Be aware of the following:

- The Archive CSV file folder field is enabled only if the Enable archive CSV file option is specified.
- The folder that you specify must reside on the Arcserve Backup server.
- Arcserve Backup includes the CSV file folder with the Arcserve Backup database protection job.
- Arcserve Backup backs up the CSV file folder in any backup job that includes the Arcserve Backup Catalog Database when specified as a global backup option. For more information, see [Backup Manager Operation Options](#) (see page 167).
- Arcserve Backup compresses CSV files on NTFS platforms.

More information:

[Job Status Types](#) (see page 325)

Tape Engine Configuration

The Arcserve Backup Tape Engine identifies all the backup devices that are connected to your system. The default configuration log options can be changed; for example, when you want to troubleshoot a hardware or Tape Engine specific problem.

To modify any of the options, settings, and parameters described in the following sections, start the Arcserve Backup Server Admin and select the Tape Engine tab.

Tape Engine Message Log Options

The following lists describe tape engine message log options:

- **Level**--If you keep the default (Summary), you do not need to specify any other options. The available values are:
 - None--No information is logged. Tape Engine logging is halted and the Tape Engine Log does not appear in the Job Status Manager.
 - Summary--(default) Logs critical messages and reduces the size of the tape log by excluding unnecessary information. For this option, the Tape.log is present in the Job Status Manager. The Tape.log file, by default, is generated and stored in the Arcserve Backup\Log folder. If the log path needs to be changed, you can do so by creating an alternate log path entry in the registry file. For more information on creating an alternate log path entry, see [Alternate Path to the Tape Engine Log](#) (see page 509).
 - Detail--This option logs all commands sent to the attached backup devices by Arcserve Backup. Reads/Writes and Test Unit Ready commands are excluded. Tape Engine specific information, which may be used by Arcserve Support to help troubleshoot backup and restore issues, is also logged. The Tape.log file, by default, is generated and stored in the Arcserve Backup\Log folder. If the log path needs to be changed, you can do so by creating an alternate log path entry in the registry file.

The Tape.log file for this option can be viewed in the Job Status Manager by accessing the Tape Log tab.

- Detail with Read/Writes--Logs all commands sent to attached backup devices by Arcserve Backup. Unlike the "Detail" option, this option includes Reads/Writes and Test Unit Ready commands. Tape Engine specific information, which may be used by Arcserve Support to help troubleshoot backup and restore issues, is also logged. The Tape.log file, by default, is generated and stored in the Arcserve Backup\Log folder. If the log path needs to be changed, you can do so by creating an alternate log path entry in the registry file.

For more information on creating an alternate log path entry, see [Alternate Path to the Tape Engine Log](#) (see page 509). The Tape.log file for this option can be viewed in the Job Status Manager by accessing the Tape Log tab.

Note: You may incur a potentially large log file size due to the Read/Write capability. The Read/Write logging may impede performance on the machine.

- **Output**--If you specified either "Summary," "Detail," or "Detail with Reads/Writes," you can define where you want the messages sent. Specify one of the following:
 - Both Screen and File--The messages are recorded in the Tape Engine Log as well as to a DOS box (the Tape Engine Message window).
 - Screen Only--The messages are sent to the Tape Engine Message window only.
 - File Only--(default) The messages are recorded in the Tape Engine log only. You can view the Tape Engine log in the Job Status Manager.

Important! If you select either option, Both Screen and File or Screen Only, you must configure the Arcserve Tape Engine service such that it can interact with your desktop and display the contents of the tape log in a DOS window. For more information, refer to the section [Enable Interaction with the Desktop](#) (see page 512).

Specify Tape Engine Log Options

The Limit Log Size section of the Tape Engine tab on the Server Admin Configuration dialog lets you direct how Arcserve Backup controls the behavior of the Tape Engine's log files.

To specify Tape Engine log options

1. From the Arcserve Backup Manager interface, select Server Admin from the Quick Start Menu in the Navigation Bar on the Home Page.

The Server Admin window opens.

2. From the Domain/Server directory tree, select the server that you want to configure.

Click the Configuration toolbar button.

The Configuration - *Server Name* dialog opens.

3. Click the Tape Engine tab.

From the **Limit log size** section, specify the following options as applicable to your requirements:

- **Limit log size by**--Check the Limit log size by check box to enable Circular Logging. In the Limit Log Size By field enter the value that you want to specify as the maximum total size of all chunked TAPE.LOG files.

The Limit log size by value divided by the Maximum log file count value represents the maximum size of all chunked log files. For example, if you specify a Limit log size by value of 100 MB and a Maximum log file count of 10, Arcserve Backup chunks TAPE.LOG when it reaches 10 MB (100/10 = 10).

The default value for the limit log size by option is 100 MB, and the range is between 1 and 2000 MB.

To disable Circular Logging, clear the limit log size by check box.

- **Prune logs older than**--Use this option to specify the number of days that must elapse before Arcserve Backup prunes the log files.

The default value for the Prune logs older than option is 100 days, and the range is between 1 and 365 days.

- **Log file split criterion**--The options in this section define the behavior of how Arcserve Backup splits the log files.

- Maximum log file count--Specifies the number of chunked log files Arcserve Backup retains.

The default value for the Maximum log file count option is 10, and the range is between 3 and 32.

Note: You can modify this setting only if the Limit log file size by option is specified.

- Maximum single log file size--This option works in conjunction with the Prune logs older than option. When you specify the Maximum single log file size and Prune logs older than options, Arcserve Backup switches to Circular Logging mechanisms when the TAPE.LOG reaches its maximum size, and deletes chunked log files when their age is greater than the value specified under the Prune logs older than option.

The default value for the Maximum single log file size option is 10000 KB, and the range is between 1 and 100000 KB.

Note: You can modify this setting only if the Prune logs older than option is specified.

4. Click OK to apply the Tape Engine log options.

The Tape Engine Log options are applied.

Note: Click Cancel to discard your changes.

Tape Engine General Options

Arcserve Backup lets you specify the following general options:

- **Use Global Scratch Set**--Lets Arcserve Backup use a Global Scratch Set. This option is enabled by default.

The Global Scratch Set treats all of the scratch tapes in all media pools as one large Scratch Set. This helps ensure that backup jobs do not fail when a scratch tape is not available in its own media pool.

When this option is enabled, the Media Pool Manager shows only the Save Set for each pool (not the Scratch Set), but adds an object named GlobalScratchSet. This object contains all of the media available in the scratch sets of all your media pools. If you right-click GlobalScratchSet and select Assign Media, you can move media from an unassigned set to the Scratch Set.

When you select a media in the Global Scratch Set, two extra properties appear as column headings on the top right-hand pane and on the Properties tab in the lower right hand pane: Medium Type and Media Pool. If you click a column heading on the top right-hand pane, you can sort the list by that column. If the media you select in the Global Scratch Set is vaulted, it appears in a different color to indicate that it is inactive.

Note: If you enable the Global Scratch Set and submit a backup job using a specific media pool, Arcserve Backup first attempts to find media in that media pool's Scratch Set. If no media is available, the Global Scratch Set media will be used. Also, if you specify a media pool and submit a backup job that spans tapes, media in the Global Scratch Set can be used.

- **Show Tape Log on Job Status Manager**--Lets you view the Tape Log in the Job Status Manager. If the Activity Log is open when you enable this option, you must click Refresh to update the manager.

Note: This option is available only on Windows computers.

- **Use TapeAlert**--Lets Arcserve Backup detect and report TapeAlert flags asserted by your tape drives and libraries. If you do not want to receive TapeAlert-related messages, disable this option.
 - With this option enabled, Arcserve Backup queries all devices connected to Arcserve Backup for TapeAlert flags in one minute intervals. If Arcserve Backup detects a TapeAlert flag, it reports real-time details about the flag in the Activity Log and the Tape.log file.
 - With this option disabled, Arcserve Backup does not maintain a separate thread-querying mechanism for detecting and reporting TapeAlert flags. As a result, Arcserve Backup will not query for TapeAlert flags until the job is running and a SCSI error occurs. If Arcserve Backup detects a TapeAlert flag while the job is running, it reports the details about the flag in the Activity Log and the Tape.log file.

Event Log Configuration (Windows Servers)

The Log tab allows you to enable or disable confirmation messages and to specify which messages may be written to the Windows Event Log.

- **Enable Message Logging into Event Log**--By default, all messages are written only into the Arcserve Backup Activity Log. If you check this box, the following groups of check boxes become enabled:
- **Exclude Message Type From Logging Check Boxes**--Use these check boxes to select which type of message should be excluded from the Event Log.
- **Exclude Message Logging From Check Boxes**--Use these check boxes to exclude all messages from a particular Arcserve Backup module.

How Arcserve Backup Records Events in the Windows Event Viewer

Event Viewer is a Windows administrative tool that lets you monitor events that relate to application, security, and system logs. The information stored in Event Viewer can vary, based on the role of the computer to your environment and the applications that are running on the computer.

Note: To open Event Viewer, click Start on the Windows toolbar, select Programs, Administrative Tools, and click Event Viewer.

The Server Admin lets you specify the type of Arcserve Backup event information that you want to record in Event Viewer. For more information, see [Event Log Configuration](#) (see page 508).

The list that follows describes the event codes for Arcserve Backup information, warning, and error events that appear in Windows Event Viewer.

- **500**--Most information events and agent information events
- **600**--Agent warning events
- **700**--Agent error events
- **900**--Audit events
- **Unique event codes**--resource ID of the message

The diagram that follows displays Arcserve Backup events in Windows Event Viewer.

Type	Date	Time	Source	Category	Event	User	Computer
Information	9/22/2008	11:45:19 ...	CA Arcserve Backup	(55)	500	N/A	SHADA07-...
Information	9/22/2008	11:45:19 ...	CA Arcserve Backup	(55)	500	N/A	SHADA07-...
Information	9/22/2008	11:45:07 ...	CA Arcserve Backup	(55)	500	N/A	SHADA07-...
Information	9/22/2008	11:45:07 ...	CA Arcserve Backup	(55)	500	N/A	SHADA07-...
Information	9/22/2008	11:45:07 ...	CA Arcserve Backup	(55)	500	N/A	SHADA07-...
Information	9/22/2008	11:45:01 ...	CA Arcserve Backup	(55)	500	N/A	SHADA07-...
Information	9/22/2008	11:44:37 ...	CA Arcserve Backup	(55)	1406	N/A	SHADA07-...
Warning	9/22/2008	11:33:17 ...	CA Arcserve Backup	(50)	600	N/A	SHADA07-...
Warning	9/22/2008	11:33:17 ...	CA Arcserve Backup	(30)	600	N/A	SHADA07-...
Information	9/22/2008	11:33:17 ...	CA Arcserve Backup	(55)	500	N/A	SHADA07-...
Information	9/22/2008	11:33:17 ...	CA Arcserve Backup	(25)	500	N/A	SHADA07-...
Error	9/22/2008	11:33:17 ...	CA Arcserve Backup	(20)	700	N/A	SHADA07-...
Error	9/22/2008	11:33:17 ...	CA Arcserve Backup	(55)	700	N/A	SHADA07-...
Information	9/22/2008	11:28:54 ...	CA Arcserve Backup	(30)	7101	N/A	SHADA07-...
Error	9/22/2008	11:28:24 ...	CA Arcserve Backup	(15)	1303	N/A	SHADA07-...
Error	9/22/2008	11:27:58 ...	CA Arcserve Backup	(15)	1301	N/A	SHADA07-...
Information	9/22/2008	11:24:09 ...	CA Arcserve Backup	(50)	900	N/A	SHADA07-...
Information	9/22/2008	11:23:08 ...	CA Arcserve Backup	(50)	900	N/A	SHADA07-...
Information	9/22/2008	11:00:17 ...	CA Arcserve Backup	(50)	900	N/A	SHADA07-...

More information:

[Event Log Configuration \(Windows Servers\)](#) (see page 508)

Alternate Path to the Tape Engine Log

You can change the default tape log path if, for example, you want to move the log to a volume with more space. You can create an alternate location for the file by configuring a registry setting in the Windows NT registry. Create a String Value called “LogPath” under the following key:

HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\Arcserve Backup\Base\TapeEngine\Debug

Set the value to the local drive path you want to use as the new log file location (for example, D:\temp\log). After the log path is established, you can enable it by either restarting the Tape Engine or changing one of the logging options described earlier. To change the tape log path back to the default, you can remove the “LogPath” value and restart the Tape Engine.

Note: Alternate paths can only be local drives as mapped drives are not supported for redirecting the log.

Circular Logging

Circular Logging is a process that lets you control the size and behavior of the Tape Engine log file. Using this feature, you can set a size limit that directs Arcserve Backup to chunk the log file into smaller log files when a user-specified size limit is exceeded. Additionally, you can specify a retention period, total count, or both for log files. After the retention period elapses, Arcserve Backup deletes the chunked log files.

The Tape Engine log file is labeled TAPE.LOG. It can be found in the CA\ARCserve Backup\LOG directory.

To configure and use Circular Logging, start the Server Admin from the Quick Start menu on the Arcserve Backup Home Page. For more information see [Specify Circular Logging Settings](#) (see page 511).

Log File Names

If you do not specify Circular Logging settings, Arcserve Backup uses the default file name, TAPE.LOG. If you do specify settings, TAPE.LOG is still generated, but it is chunked into smaller files and the smaller files are named using the following format:

TAPE.LOG.####

where #### represents the sequential log number created on a given day.

Example: Log File Names

For example, on a given day, the Tape Engine generates three log files based upon a file size limit of 100 MB. The log file names are as follows:

TAPE.LOG
TAPE.LOG.0001
TAPE.LOG.0002

How Arcserve Backup Labels Log Files

Arcserve Backup labels the log files using the following guidelines:

1. If TAPE.LOG reaches a specified value, Arcserve Backup renames TAPE.LOG to TAPE.LOG.0001, and creates a new TAPE.LOG file.
2. If TAPE.LOG reaches a specified value for the second time, Arcserve Backup renames TAPE.LOG.0001 to TAPE.LOG.0002, renames TAPE.LOG to TAPE.LOG.0001, and creates a new TAPE.LOG file.

3. If TAPE.LOG reaches a specified value for the third time, Arcserve Backup renames TAPE.LOG.0002 to TAPE.LOG.0003, renames TAPE.LOG.0001 to TAPE.LOG.0002, renames TAPE.LOG to TAPE.LOG.0001, and creates a new TAPE.LOG file.

This process continues in a cyclical manner. Arcserve Backup always retains the latest three log files.

Important! Arcserve Backup calculates the value in which a new log file is created based upon the amounts that you specify in the Limit Log Size By and Log File Count options. For example, if you specify a Log Limit Size By amount of 500 MB and a Log File Count of 10, Arcserve Backup creates a new log file when the current log size exceeds (500 divided by 10) 50 MB.

Specify Circular Logging Settings

Circular Logging lets you customize the characteristics of log files generated by the Tape Engine.

To specify Circular Logging settings

1. From the Arcserve Backup Home Page, click the Quick Start menu and select Server Admin.
The Arcserve Backup Server Admin dialog opens.
2. From the Admin menu, select Configuration.
The Configuration dialog opens.
3. Click the Tape Engine tab.
4. To enable Circular logging, click the Limit Log Size By option in the Limit Log Size section of this dialog, and then specify then maximum size in MB. This amount represents the maximum size of all log files.
5. In the Log File Count field, select the number of log files that you want Arcserve Backup to retain. This amount represents the maximum number of TAPE.LOG files that Arcserve Backup will retain.
6. Click OK to apply the settings.

Note: After the log file count exceeds the number specified using the Log File Count option, Arcserve Backup deletes the oldest log files.

Prune Log Files

To specify log file pruning only:

1. Disable the Limit Log Size By option.
2. Click the Prune Logs Older Than option and specify the number of days that you want to elapse before Arcserve Backup prunes log files.

3. (Optional) In the Single Log File Size field, enter a size in KB to specify a size limit for a single log file. If you do not specify a value in the Single Log File Size field, Arcserve Backup uses the default value, 10000 KB, as the size limitation for each single log file.
4. Click OK to apply the settings.

Important! If you enable both Limit Log Size options (Limit Log Size By and Prune Logs Older Than), Arcserve Backup prunes log files if either the total number of log files exceeds the Log File Count, or the date of the log files exceeds the number of days specified under the Prune Log Files Older Than option. You cannot specify a Single Log File Size—Arcserve Backup uses the formula Total Log Size divided by Log File Count to calculate the Log File Size setting.

Enable Interaction with the Desktop

This section describes how to enable the Arcserve Backup Tape Engine to interact with the desktop. However, these steps can be used when you want to allow any Arcserve Backup service or engine to interact with the desktop.

To enable interaction with the desktop

1. From the Windows Start menu, select Programs (or All Programs), Administrative Tools, and select Component Services.
The Component Services dialog opens.
2. From the object tree, select the Services (Local) object.
In the Services list, locate, right-click Arcserve Tape Engine (for example), and select Properties from the pop-up menu.
The Arcserve Tape Engine Properties (Local Computer) dialog opens.
3. Select the Log On tab.
Under Local System account, select the Allow service to interact with desktop option and click Apply.
Click OK to close the Arcserve Tape Engine Properties (Local Computer) dialog.
4. Stop and then restart the Arcserve Tape Engine service.
The Tape Engine can interact with the desktop.
5. Close the Windows Component Services dialog.

Increase Virtual Memory Allocation to Improve Tape Engine Performance

The Tape Engine is a Arcserve Backup component that controls all jobs that involve backing up, restoring, and recovering data. In large backup environments, the following conditions can have adverse affects on Tape Engine performance:

- High quantities of backup jobs running simultaneously.
- High quantities of backup jobs that use deduplication, that may or may not run simultaneously.
- Multistreaming jobs, multiplexing jobs, or both with a high number of streams specified running simultaneously.

Diminished Tape Engine performance can occur due to the amount of virtual memory that is allocated to applications by the operating system. By default, Windows 32-bit operating systems allocate 2 GB of virtual memory for applications. The default allocation on 32-bit operating systems can prevent computers that contain 3 or more GB RAM from utilizing all of the RAM that is installed on the computer.

Note: Windows 64-bit operating systems can allocate up to 4 GB of virtual memory.

Windows operating systems let you increase the amount of virtual memory allocated for applications. If your backup server demonstrates slow performance in a large backup environment, you can increase the amount of virtual memory available to applications to improve the performance of the Tape Engine.

Note: You should increase virtual memory only if your backup server contains a Windows 32-bit operating system with at least 2.5 GB RAM.

The following sections describe how to increase virtual memory allocation on computers running 32-bit versions of Windows Server 2003 and Windows Server 2008.

To increase virtual memory allocation on 32-bit Windows Server 2003 systems

1. From the desktop, right-click My Computer and select Properties on the pop-up menu.

The System Properties dialog opens.

2. On the System Properties dialog, click the Advanced tab.

The Advanced properties appear.

3. In the Startup and Recovery area, click Settings.

The Startup and Recovery dialog opens.

4. In the System startup area on the Startup and Recovery dialog, click Edit

The Windows boot.ini file opens in Notepad.

5. Add /3GB to the end of the startup string below [Operating Systems], as illustrated by the following:

[operating systems]

```
multi(0)disk(0)rdisk(0)partition(1)WINDOWS="Windows Server 2003, Enterprise" /noexecute=optout /fastdetect switch: /3GB
```

6. Close Notepad.

When Notepad prompts you to save your changes, click Yes.

7. Close all open dialogs and then restart the computer.

To increase virtual memory allocation on 32-bit Windows Server 2008 systems

1. Open a Command Line window.
2. From the Command Line, execute the following command:

```
BCDEDIT /Set IncreaseUserVa 3072
```

3. After the command executes, restart the computer.

Database Engine Configuration

The Arcserve Backup Database Engine stores the following types of statistical information for all jobs processed.

- Files and directories that have been backed up, copied, and restored.
- Jobs that Arcserve Backup has processed.
- Storage devices and media used for Arcserve Backup operations.

Arcserve Backup provides the following database engine options:

- **Enable auto pruning**--When database pruning is enabled, information about the files and directories that were backed up or copied in a session is deleted. By default, this option is selected to free up space in the database file. It can be useful to disable this option to maintain the detailed information for restoring purposes, but be aware that your database can become very large if you do not prune it.
 - **Run Pruning at**--This field is active only if the Enable Database Pruning option is on. Specify when you want the pruning operation to run.
Default value: If enabled, will occur at 12:00 p.m.
 - **Prune database job records older than**--This field is active only when the Enable auto pruning option is specified. This option lets you specify how long job records should be retained in the database before Arcserve Backup prunes them.
Default value: If pruning is enabled, 180 days.
Range: 1 to 999 days.
 - **Prune other database records older than**--This field is active only when the Enable auto pruning option is specified. This option lets you specify how long other records (for example, session detail records) should be kept in the database before Arcserve Backup prunes them.
Default value: If pruning is enabled, 30 days.
Range: 1 to 999 days.
- **Delete Re-Formatted or Erased Media-Related database Records when Pruning**--When you reformat or erase a media, Arcserve Backup also deletes the records in the database that pertain to the media. Performing this extra step, however, can be a time-consuming process. Select this option to postpone deleting these records until pruning is performed.
- **Prune activity logs older than**--Specify how long activity logs should be kept in the database before Arcserve Backup prunes them.
Default value: 14 days
Range: 1 to 999 days.
- **Prune catalog files older than**--Specify how long catalog files should be kept in the database before Arcserve Backup prunes them.
Default value: 60 days.
Range: 1 to 9999 days.

- **Database Maintenance Operations**--The following options apply to maintenance operations that can be performed on the Arcserve Backup database.

When you enable the following options, Arcserve Backup performs the specified task the next time the Database Pruning Job runs. If the Database Pruning Job is scheduled to run on a daily basis, the specified operations are performed when the pruning job runs. To schedule the database maintenance operations to run without the dependence of the Database Pruning Job, you can use the Job Scheduler Wizard to create individual jobs that use the `ca_dbmgr` command line utility to facilitate the database maintenance operations.

Note: For more information, see [How You Can Use the Job Scheduler Wizard to Schedule Jobs](#) (see page 341) or the *Command Line Reference Guide*.

- **Update statistics**--This option lets Arcserve Backup update table and index statistics. With correct and up-to-date statistical information, SQL Server and SQL Server 2008 Express can determine the best execution plan for queries, which improves query performance.

You should update the statistics on a daily basis.

- **Re-build indexes**--This option lets Arcserve Backup remove fragmentation (by compacting the pages based on the specified or existing fill factor setting) and reorder the index rows in contiguous pages. As a result, Arcserve Backup improves query performance and reclaims disk space.

You should rebuild the indexes on a weekly basis.

- **Check DB integrity**--This option lets Arcserve Backup check the allocation, structural, and logical integrity of all the objects in the Arcserve database.

You should check the integrity of the database on a weekly basis and allocate a sufficient amount of time for this task to run.

- **Reduce DB size**--This option lets Arcserve Backup reclaim disk space on your system by reducing the size of the data files in the Arcserve database.

You should reduce the size of the database on an as-needed basis.

- **Submit Prune Job**--Select this option to submit the pruning job now.
- **Submit Arcserve DB protection job**--This option lets you recreate the Arcserve Backup Database Protection Job because the original job was deleted. For more information, see [Recreate the Arcserve Backup Database Protection Job](#) (see page 625).

- **Catalog Database**

- **Catalog database folder**--This option let you specify the location of the Arcserve Backup catalog database folder. Click the ellipsis button to browse and select a different location for the catalog database folder.

By default, the catalog database folder will be located on the Primary Server at:

C:\Program Files\CA\ARCserve Backup\CATALOG.DB\

- **Compress catalog transfer on the following member servers**--This option lets Arcserve Backup compress catalog information when the data is transferred from a member server to the primary server.

If the Primary Server has any associated Member Servers, the "Compress catalog transfer on the following member servers" field will be enabled, displaying the names of the Member Servers.

By default, this option is disabled. With this option disabled, Arcserve Backup will not compress the catalog information when it is transferred from the Member Server to the Primary Server.

- **Minimum disk free space threshold**--This option lets you specify the minimum percentage of free disk space when Arcserve Backup deletes catalog files.

Default value: 10 %

Range: 1% to 99%

Note: Arcserve Backup periodically checks the free disk space percentage on the volume where the catalog database folder is located. If the detected free space is lower than the specified percentage, a warning message will be sent to the activity log and it would automatically begin to delete catalog database files (minimum of 7 days old and starting with the oldest first) from the disk until the detected free space percentage is greater than the threshold setting.

Example: If the detected free space is lower than the 10%, a warning message is sent to the activity log and it would automatically begin to delete catalog database files (minimum of 7 days old and starting with the oldest first) from the disk until the detected free space percentage is greater than 10%.

- **Enable Media Pool Maintenance**--When selected, all media scheduled to be moved from a media pool's Save Set to its Scratch Set are automatically moved any time a prune job is run.

- **Maximum Database Server Memory**--This applies to Microsoft SQL Express only. Used to ensure that the size of the Microsoft SQL Express memory usage does not exceed this limit.

Default value: 1024 MB

Range: 256 MB to 1024 MB

More information:

[How to Protect the Arcserve Backup Database](#) (see page 607)

[How the Catalog Database Works](#) (see page 650)

Alert Configuration

Alert is a notification system that sends messages to people in your organization using various methods of communication. Alert does not generate its own messages. You must tell Alert what information you want to communicate and where you want to send it.

If you configure Alert from Server Admin, you can generate notifications for non-job related events, such as starting or stopping the Tape Engine. To do this, enter the words or phrases you want to communicate in exactly the same format as they appear in the Activity Log and click Add.

Or, if you want to send all activity log messages, enter an asterisk and click Add. Alert generates notification messages and sends them to the appropriate recipients. For information on selecting recipients and configuring methods to transmit Alert notifications, see "Using the Alert Manager."

Add and Remove Alert Notifications

Arcserve Backup lets you configure notifications Alert notifications for non-job related events (for example, starting and stopping the Tape Engine and the successful completion of operations) using the Server Admin Manager.

The procedure that follows describes how to add Alert notifications for non-job related events.

To add alert notifications

1. Open the Arcserve Backup Server Admin Manager.

Click the Arcserve Backup primary server or stand-alone server in the directory tree, and then click Configuration on the toolbar.

The Configuration dialog opens.

2. Click the Alert tab.

In the Alert list field, enter the text for the event for which you want to receive an Alert notification.

You can enter a complete phrase, or only a part of it (even a single keyword). The Alert engine tries to match the keyword or phrase to the text of each generated event. However, you should be as specific as possible, in order to avoid receiving unwanted Alert notifications.

Examples:

- To receive an Alert notification when the Tape Engine starts, enter engine in the Alert list to have Alert detect the event. However, the Alert engine then sends notifications for any events that related to engine, such as Database engine started.
- To receive Alert notifications for all activity log messages, enter '*'.
- To receive Alert notifications for Audit log events, enter the following in the Alert list field:
 - [Auditlog]--Sends an Alert notification for all Audit Log events.
 - [Auditlog][Success]--Sends an Alert notification for all successful Audit Log events.
 - [Auditlog][Failure]--Sends an Alert notification for all failed Audit Log events.

Note: The keywords for Audit Log events are case-sensitive and square brackets are required.

3. Click Add to add the search text.

4. Click OK.

When the text of an event matches one of the keywords you have entered, the Alert engine generates Alert notifications for all recipients previously configured using the Alert Manager.

The procedure that follows describes how to remove Alert notifications.

To remove Alert notifications

1. Open the Arcserve Backup Server Admin Manager.

Click the Arcserve Backup primary server or stand-alone server in the directory tree, and then click Configuration on the toolbar.

The Configuration dialog opens.

2. Click the Alert tab.

From the Alert list, click the event for which you no longer want to receive Alert notifications, and then click Delete.

The Alert notification is deleted.

Note: To delete all Alert notifications from the Alert list, click Delete All.

3. Click OK to close the Configuration dialog.

Additional Server Admin Functions

You can use the Server Admin to perform the following functions:

- Change the system account
- Configure multiple NIC cards
- Manage licenses centrally

Modify the Arcserve Backup System Account

The Arcserve Backup server requires a valid system account on the host Windows computer (initially entered during installation). You can modify the login credentials for the system account at any time, using the Server Admin.

If you are using a Windows domain user account to serve as the credentials for the Arcserve Backup server system account, you must update Arcserve Backup with the new password when you modify your Windows domain password.

To modify the Arcserve Backup system account

1. From the Arcserve Backup Home Page, open the Server Admin by selecting Server Admin from the Quick Start menu.

The Arcserve Backup Server Admin window opens.

2. Select Arcserve Backup System Account from the Server Admin menu.

The Arcserve Backup System Account dialog opens.

3. Complete the following fields, as required:

- Server
- User Name
- Password
- Domain

4. Click OK.

Note: If you are using a remote deployment of Microsoft SQL Server to host the Arcserve Backup database, and the login credentials for the Arcserve Backup system account (User Name and Password) are identical to the login credentials for the remote SQL Server account, a message appears to notify you that modifying the login credentials for the system account also modifies the login credentials for the remote SQL Server account. If you are sure that you want to modify the login credentials for the remote SQL Server account, click OK.

The login credentials for the Arcserve Backup system account are modified.

Reconfigure Node Tier Assignments

You can use the Arcserve Backup Server Admin or the Central Agent Admin to change the assigned priority classifications of your Arcserve Backup nodes. These tiers are used to filter the information displayed on the Arcserve Backup Dashboard by the priority level of the monitored nodes.

The Node Tier Configuration dialog contains three priority categories (High Priority, Medium Priority, and Low Priority), and is automatically populated when a node is added to your system and browsed. By default, a High Priority tier is configured to include all Arcserve Backup servers (Primary and Member) and any nodes with Arcserve Backup application agents installed (such as Oracle, Microsoft Exchange Server, Microsoft SQL Server, Microsoft Sharepoint Server, and so on), and a Low Priority tier is configured to include all other nodes (having file system agents installed). The Medium Priority tier is not configured to include any nodes, and is available for customized use.

The node assignments for each tier can be reconfigured and customized to meet your individual needs by using the Node Tier Configuration dialog, which is accessed from the Arcserve Backup Server Admin or from the Backup Manager (right-click Windows Systems in Source tab) or from the Central Agent Admin (right-click Windows Systems).

To reconfigure node tier assignments

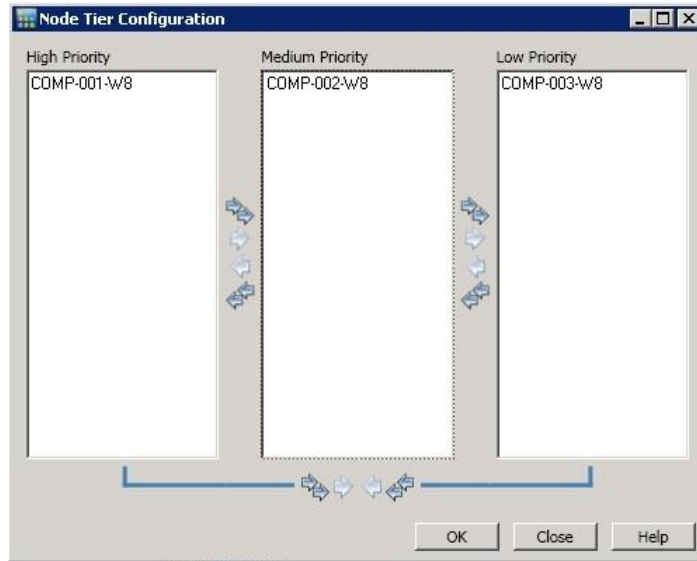
1. Open the Arcserve Backup Server Admin by clicking Server Admin on the Quick Start menu on the Home Page.

The Server Admin window opens.

2. Expand the domain directory tree and select a server that you want to view or reconfigure the node tier assignments.

3. Select the Node Tier Configuration option from the Admin menu.

The Node Tier Configuration dialog opens, displaying the nodes assigned to each Tier category (High Priority, Medium Priority, Low Priority).



4. Select one or more nodes that you want to reassign to a different tier category and click on the corresponding arrow icon to move the selected nodes from one tier to another.

Note: Multiple nodes can be selected for tier assignment by using the "CTRL" or "SHIFT" key combinations.

- Click the single arrow icon to move only the selected nodes.
 - Click the double arrow icon to move all the node in the tier.
5. Click OK when done.

The node tier assignments have been changed to meet your individual needs.

Manage Arcserve Backup Component Licenses

The Arcserve Backup Server Admin lets you perform the following license management tasks:

- View the Arcserve Backup products installed on a primary server, stand-alone server, member servers, and agent servers in a Arcserve Backup domain.
- Identify the total number of licenses applied and the number of active licenses for each component in a Arcserve Backup domain.
- View the names of the servers using the component licenses in a Arcserve Backup domain.
- Release licenses from servers to make the licenses available to other servers in your domain.


Note: For information about releasing licenses from servers, see [Release Licenses from Servers](#) (see page 525).

To manage Arcserve Backup component licenses

1. From the Arcserve Backup Manager Console, open the Server Admin by clicking Server Admin in the Quick Start menu.

The Server Admin opens.

The Arcserve Backup primary server and its member servers display in a directory tree structure as illustrated by the following:



2. To view the Arcserve Backup products installed on a primary server and a member server, select the server in the directory tree.

The components and licenses for the selected server display in the properties view, as illustrated by the following:

The screenshot shows the Arcserve Backup Server Admin console for server CRW002. The left sidebar contains navigation options: Arcserve Backup System Account, Configuration..., Stop all services, Add/View Licenses..., Manage Licenses..., Install/Uninstall Options..., and Node Tier Configuration... The main area displays 'Server Information' with a dropdown menu. Below this, a table titled 'Products Installed: 6' lists the installed products, their versions, and build numbers.

Product Name	Version	Build
Arcserve Backup	17.0	7490
Tape Library Option	17.0	7490
Agent for Microsoft SQL Server	17.0	7490
Enterprise Module	17.0	7490
Disaster Recovery Option	17.0	7490
Agent for Virtual Machines	17.0	7490

3. To view the component and licensing relationships in a Arcserve Backup domain, right-click the primary server and select Manage Licenses from the pop-up menu.

The License Management dialog opens.

The License Management dialog provides you with the following information:

- **Version**--Specifies the release number of the license for the selected component.
- **Active Licenses**--Specifies the number licenses that are currently active for the selected component. The total includes purchased licenses and trial licenses.
- **Available Licenses**--Specifies the number of licenses available for use for the selected component. The total includes only purchased licenses.
- **Total Licenses**--Specifies the total number of licenses purchased for the selected component.
- **Licenses Needed**--Specifies the number of additional licenses that you need to use the selected component.


Examples:

- You are using one purchased license and one trial license for a component. Arcserve Backup recommends purchasing one license to replace the trial license so that you have uninterrupted use of selected component.
- You are protecting six Windows computers using the Client Agent for Windows. You purchased four Client Agent for Windows licenses. In the past, backups may have failed due to an insufficient number licenses. Arcserve Backup recommends purchasing two additional licenses to help ensure that you have uninterrupted use of the Client Agent for Windows.

- **Licensed machines**--Specifies the names of the computers using the active licenses for the selected component.

Example:

- The following dialog illustrates that there are 10 active licenses and zero available licenses for the Tape Library Option. The host names of the computers using the Tape Library Option licenses display in the Licensed machines field.



More information:

[Central License Management](#) (see page 53)

Release Licenses from Servers

Arcserve Backup licensing functions on a count-based mechanism. Count-based licensing lets you grant a single overall license to the application with a predetermined number of active license rights included in the overall license pool. Each server that uses the license is granted an active license from the pool, on a first-come basis, until the total number of available license rights has been reached. If all the active license rights have already been applied and you need to add a license to a different member server, you must remove the license rights from one of servers to reduce the count before the different member server can use the license.

To release licenses from servers

1. From the Arcserve Backup Manager Console, open the Server Admin by clicking Server Admin in the Quick Start menu.

The Server Admin opens.

2. From the server directory tree, right-click the primary server and select Manage Licenses from the pop-up menu.

The License Management dialog opens.

3. From the License status section, select the component containing the license that you want to release.

The machines using the license display in the Licensed machines field.

4. Clear the check box next to the machine name with the license that you want to release and click Apply.

The active license is released from the selected server. The license is now available to other servers running the Arcserve Backup product in your Arcserve domain.

Note: After you click the Apply button, the selected machine no longer appears in the Licensed machines field.

Configure Multiple Network Interface Cards

If the Arcserve Backup server has more than one network connection, Arcserve Backup can be configured to use a specific network interface card for its backup and restore operations. If you configure Arcserve Backup in this manner, it will not interfere with the other cards that are attached to the system.

You can configure Arcserve Backup to use a set of network interface cards, which it effectively uses when performing multistreaming backups. You can also configure Arcserve Backup to use an appropriate card from the configured set of network interface cards when connecting to a client agent.

To configure multiple network interface cards

1. From the Arcserve Backup Home Page, open the Server Admin by selecting Server Admin from the Quick Start menu.

The Arcserve Backup Server Admin window opens.

2. Select Multiple Network Cards from the Server Admin menu.

The Multiple Network Cards settings dialog opens.

3. Select one of the following options:
 - **OS decide which network card to be used (default)**--Lets the operating system decide which network interface card is to be used.
 - **Use selected network card**--Lets you specify the network interface cards that you want to use from the list. When configured in this manner, any job that Arcserve Backup executes defaults to the first network interface card. When you are using multistreaming, and more than one process is created, each subsequent process will use the next configured network interface card.
4. Click OK.

The network card settings are applied.

Authentication Levels for Arcserve Backup Services, Components, and Applications

You must log in to Windows Vista, Windows Server 2008, and Windows Server 2012 systems using an administrative account or an account with the highest available permissions to run various Arcserve Backup services, components, and applications. The binaries corresponding to these services, components, and applications contain Arcserve Backup specific functionality that is not available to a basic user account. As a result, Windows will prompt you to confirm an operation by specifying your password or by using an account with administrative privileges to complete the operation.

This section contains the following topics:

[Arcserve Backup Services, Components, and Applications that Require Administrative Privileges](#) (see page 527)

[Arcserve Backup Services, Components, and Applications that Require the Highest Available Privileges](#) (see page 533)

Arcserve Backup Services, Components, and Applications that Require Administrative Privileges

The administrative profile or an account with administrative privileges has read, write, and execute permissions to all Windows and system resources.

The following table describes Arcserve Backup services, components, and applications require administrative privileges:

Component	Description
_HTMSETUP.EXE	Shows the html page during Setup.

Component	Description
<CD_ROOT>\IntelINT\Exchange.DBA\SETUP.EXE	Lets Arcserve Backup launch the following executables: <ul style="list-style-type: none"> ■ IntelINT\Exchange.DBA\Exchange.DBA\SETUP.EXE--Installs the Agent for Microsoft Exchange Server for database level backups. ■ IntelINT\Exchange.DBA\ExchangeD.DBA\SETUP.EXE--Installs the Agent for Microsoft Exchange Server for document level backups.
<CD_ROOT>\IntelINT\EO\SETUP.EXE	Lets Arcserve Backup launch the Windows installer labeled MSIEXEC.EXE to install the MSI package.
<CD_ROOT>\SETUP.EXE	Lets Arcserve Backup launch the CD browser so that you can install or upgrade Arcserve Backup.
AGENTDEPLOY.EXE	Agent Deployment application.
AGIFPROB.EXE	Dashboard SRM (storage resource management) back-end service for collecting the agent storage resource information for Dashboard.
AGPKIMON.EXE	Dashboard SRM client component for collecting SRM information from agents running on Windows nodes.
ALADMIN.EXE	Alert administration application.
ARCSERVECFG.EXE	Server Configuration Wizard.
ASDBInst.exe	Lets the installation process install Microsoft SQL Server Express Edition when you specify Microsoft SQL Server Express Edition during the installation process.
ASRECOVERDB.EXE	Utility that lets you recover the Arcserve Backup database.
ASREMSVC.EXE	Lets you install Arcserve Backup on a remote system.
AUTHSETUP.EXE	Authentication Setup command line utility.
BABHA.EXE	Lets you configure Arcserve Backup for high availability. This component is commonly used with Microsoft Cluster Service installations.
BACKINT.EXE	Backup integration module for the Agent for SAP R/3 for Oracle.
BACKINTCONFIG.EXE	Configuration utility for the Agent for SAP R/3 for Oracle.

Component	Description
BCONFIG.EXE	Lets Arcserve Backup configure server information (for example, a primary server, a member server, or a stand-alone server) as it relates to the type of Arcserve database and the caroot password. This component runs when you are installing or upgrading Arcserve Backup.
BDELOBJ.EXE	Lets the uninstallation process delete temporary and dynamic files from a system where you are uninstalling Arcserve Backup. The uninstallation process copies this application to the target system.
BDELOBJ_BAB.EXE	Lets the installation process delete temporary and dynamic files from a system where you are upgrading Arcserve Backup from a previous release. The uninstallation process copies and replaces this application to the target system.
BRANCHCFG.EXE	Opens the Branch Manager window.
BRANCHSERVICE.EXE	Arcserve Backup Global Dashboard. Allows communication between a branch site and the central site.
C:\Program Files (x86)\Microsoft SQL Server\100\Setup Bootstrap\Release\setup.exe	Microsoft SQL Server installation file. This executable is not an Arcserve binary file.
C:\Program Files (x86)\Microsoft SQL Server\100\Setup Bootstrap\Update Cache\KB968369\ServicePack\setup.exe	Microsoft SQL Server patch installation file. This executable is not an Arcserve binary file.
C:\Program Files (x86)\Microsoft SQL Server\100\Setup Bootstrap\Update Cache\KB968369\ServicePack\x86\setup\1033\pfiles\sqlservr\100\setup\release\setup.exe	Microsoft SQL Server patch installation file. This executable is not an Arcserve binary file.
CA.ARCserve.CommunicationFoundation.WindowsService.exe	Provides data used by Arcserve Backup Global and Local Dashboard.
CABATCH.EXE	cabatch command line utility.
CadRestore.exe	CA Active Directory Object Level Restore utility. Lets you browse .NTDS AD DB files and restore Active Directory objects to the current Active Directory.
CALICENSE.EXE	Arcserve Backup License Application for monitoring each option's License account number.
CALicnse.exe	CA license check application.
CARUNJOB.EXE	Local backup and restore utility.
CENTRALSERVICE.EXE	Arcserve Backup Global Dashboard (Arcserve Central Remoting Server). Allows a branch site to synchronize data to the central site database.

Component	Description
CHGTEST.EXE	Test Changer utility.
DBACONFIG.EXE	<p>Lets Arcserve Backup configure database instances during the installation process for the following agents:</p> <ul style="list-style-type: none"> ■ Agent for Microsoft SQL Server ■ Agent for Oracle ■ Agent for SAP R/3 for Oracle ■ Agent for Informix ■ Agent for Sybase ■ Agent for Lotus Notes <p>Lets you configure database instances after the installation process is complete for the following agents:</p> <ul style="list-style-type: none"> ■ Agent for Microsoft SQL Server ■ Agent for Oracle ■ Agent for SAP R/3 for Oracle ■ Agent for Informix ■ Agent for Lotus Notes
DELETEME.EXE	Lets the installation process delete temporary files from remote systems when installing Arcserve Backup components on a remote system. This component runs on the local and remote system.
DEPLOYDUMMY.EXE	Agent Deployment application preload module.
DSCONFIG.EXE	Discovery Configuration utility.
dosboot.exe	Disaster Recovery utility.
EMCONFIG.EXE	Enterprise Module Configuration utility.
ETPKI_SETUP.EXE	ETPKI encryption/decryption library installation utility.
HDVSSCOM.exe	Lets VSS hardware backup jobs import metadata.
jucheck.exe	JAVA application that checks for upgrades.
LandingPage.exe	Microsoft SQL Server 2008 Express Edition installation file. This executable is not a Arcserve binary file.
LICCHECK.EXE	Lets the Arcserve Backup Agent for Lotus Domino and the Arcserve Backup Agent for Informix verify the status of their licenses.
MASTERSETUP.EXE	Lets the installation process launch Windows Installer 3.1 and VC8 SP1 redistribute package.

Component	Description
MASTERSETUP_MAIN.EXE	Lets the installation process display the installation wizard dialogs, configure, and call individual products when you are installing Arcserve Backup.
MEDIASVR.EXE	Proxy for Tape Engine Communication.
MergeIngres2Sql.exe	Arcserve Backup Merge Ingres Database to SQL Database tool. This tool lets Arcserve Backup migrate database data from an Ingres database to a Microsoft SQL Server database when you upgrade a UNIX or Linux server to a data mover server.
ORAUPGRADE.EXE	Upgrades Oracle Agent used during the upgrade of an old version of Arcserve Backup to the current version.
qphmbavs.exe	Microsoft SQL Server 2008 Express Edition installation file. This executable is not an Arcserve binary file.
RAIDTEST.EXE	Lets you configure and test RAID devices using a Windows command line utility. You can configure and tape RAID and tape changer RAID devices, not actual RAID.
rdbgsetup.exe	Microsoft SQL Server 2008 Express Edition component. This executable is not an Arcserve binary file.
RMLIC.EXE	Module that uninstalls License Modules.
SDOInst.exe	Installs prerequisite components when you deploy Arcserve Backup using SDO.
SETUP.EXE	Installation Wizard.
setup100.exe	Microsoft SQL Server 2008 Express Edition installation file. This executable is not an Arcserve binary file.
SETUPFW.EXE	Windows Firewall Configuration utility.
SETUPSQL.EXE	Lets Arcserve Backup create the Arcserve database with Microsoft SQL Server when you are installing or upgrading Arcserve Backup.
SETUPSQL_EXP.EXE	Lets Arcserve Backup create the Arcserve database with Microsoft SQL Server 2008 Express Edition when you are installing or upgrading Arcserve Backup.
SILENT.EXE	Arcserve Backup License Application.

Component	Description
SIMULATE.EXE	Lets Arcserve Backup configure fictitious SCSI devices (for example, tape drives and tape libraries) based on their file system. This Arcserve Backup component is a command line utility.
SMPLEMON.EXE	Dashboard component for collecting storage resource utilization at the Agent node.
SPS012UPGRADE.EXE	Lets Arcserve Backup upgrade a previous version of the Agent for Microsoft SharePoint Server to the current version.
SPADMIN.EXE	Agent for Microsoft SharePoint Server 2003 installation wizard.
SQLAGENTRMTINST.EXE	Agent for Microsoft SQL Server installation wizard.
SQLCONFIG.EXE	Agent for Microsoft SQL Server configuration utility.
SQLdiag.exe	Microsoft SQL Server 2008 Express Edition installation file. This executable is not an Arcserve binary file.
SqlWtsn.exe	Microsoft SQL Server 2008 Express Edition installation file. This executable is not a Arcserve binary file.
TAPEENG.EXE	Arcserve Backup Tape Engine.
TAPETEST.EXE	Test Tape Drive utility.
Uninstall.exe	Utility that lets you uninstall Arcserve Backup.
UNINSTALLER.EXE	Application that uninstalls the ETPKI component.
UPDATECFG.EXE	Saves the configuration file during upgrade.
UPDATECFG.EXE	Lets Arcserve Backup back up the registry and file entries from previous Arcserve Backup installations when you are upgrading from a previous release.

Arcserve Backup Services, Components, and Applications that Require the Highest Available Privileges

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 Arcserve Backup services, components, and applications require an account with the highest available privileges:

Component	Description
ACSCFG.EXE	Volume Configuration utility for StorageTek ACSLS Library.
ADMIN.EXE	Arcserve Backup Agent Admin utility.
adrasr.exe	Disaster Recovery utility.
AdrLogViewer.exe	Disaster Recovery utility.
adrmain.exe	Disaster Recovery tool.
adrstart.exe	Disaster Recovery utility.
AgPkiMon.exe	SRM component that lets Arcserve Backup collect and monitor agent PKI (performance key indicators). For example, CPU, Memory, NIC, and disk information.
ALERT.EXE	Arcserve Backup Alert Service.
ARCSERVEMGR.EXE	Arcserve Backup Manager Console.
ASWANSYNC.EXE	Lets Arcserve Replication and High Availability interface with the Arcserve Backup Client Agent for Windows.
ATLCFG.EXE	Volume Configuration utility for IBM 3949 Library.
BAOFCONFIGMIGRATION.EXE	Upgrade tool for migrating the configuration for the previous version of the Agent for Open Files to the current version (Unicode format).
BDAEMON2.EXE	Raima DB Daemon Application.
CA_AUTH.EXE	ca_auth command line utility.
CA_BACKUP.EXE	ca_backup command line utility.
CA_DBMGR.EXE	ca_dbmgr command line utility.
CA_DEVMGR.EXE	ca_devmgr command line utility.
CA_JOBSECMGR.EXE	ca_jobsecmgr command line utility.
CA_LOG.EXE	ca_log command line utility.

Component	Description
CA_MERGE.EXE	ca_merge command line utility.
ca_msvmpopulatedb.exe	Arcserve Hyper-V Configuration Tool.
CA_QMGR.EXE	ca_qmgr command line utility.
ca_recoverdb.exe	ca_recoverdb command line utility. Lets you recover the Arcserve Backup database.
CA_RESTORE.EXE	ca_restore command line utility.
CA_SCAN.EXE	ca_scan command line utility.
CAADVREPORTS.EXE	caadvreports command line utility.
CAAGSTART.EXE	Lets the Universal Agent start processes that launch database agents. This is in internal utility that is not exposed to the end user.
CAAUTHD.EXE	Authentication service.
CACLURST.EXE	caclurst utility.
CADIAGINFO.EXE	Lets the Diagnostic wizard collect diagnostic information about Arcserve Backup from remote systems. This utility stores the collected diagnostic data in a file with a .caz file extension.
CADIAGSUPPORT.EXE	Lets Arcserve Support personnel and end users open and view diagnostic information saved in diagnostic data (.caz) files.
CADIAGWIZ.EXE	Lets Arcserve Backup collect Windows system and network information from local and remote systems as it relates to Arcserve Backup. Arcserve support can use the collected information to troubleshoot an Arcserve server.
CADISCOVD.EXE	Domain Server application.
CADVWIZ.EXE	Device Configuration Wizard.
CADVWIZE.EXE	Device Wizard for configuring devices.
CALICNSE.EXE	License check application.
CAMINFO.EXE	License information display application.
CAREPORTS.EXE	careports Report Writer Command Line utility.
CASDSCSVC.EXE	Discovery Service.
CASERVED.EXE	Service Controller.
CASISCHK.EXE	Single instance support application.

Component	Description
CATIRPC.EXE	Arcserve Portmapper.
CAVER.EXE	Lets Arcserve Backup display the version and build number details of the Arcserve Backup base product in a graphical user interface.
CDBMERGELOG.EXE	Lets Arcserve Backup merge activity log details from the local cache to the Arcserve database.
CONFIGBAF.EXE	BAF (Bright Agent frame) configuration utility. The installation wizard launches this utility to register agents into a configuration file that is used by the Universal agent.
CONFIGENCR.EXE	Encryption configuration utility.
CSTMSGBOX.EXE	Lets Arcserve Backup show message boxes. This is in internal utility that is not exposed to the end user.
DBACFG.EXE	Lets Arcserve Backup configure account details for database agents.
DBENG.EXE	Arcserve Backup Database Engine.
DBTOSQL.EXE	Lets Arcserve Backup migrate Raima VLDB database information and data to Microsoft SQL Server databases.
DBTOSQL_EXP.EXE	Lets Arcserve Backup migrate Raima VLDB database information and data to Microsoft SQL Server 2008 Express Edition databases.
drcreate.exe	Lets Arcserve Backup create a disaster recovery boot kit. For example, MSD (machine specific disk and bootable media).
DUMPDB.EXE	Lets Arcserve Backup dump session or export session passwords stored in the database to a specified target file or import session password stored in file to database.
DVCONFIG.EXE	Device Configuration utility.
ELOConfig.exe	Lets Arcserve Backup configure the Storage Area Network (SAN) Option.
ERRBOX.EXE	Custom pop-up error display application.
EXPTOSQL.EXE	Microsoft SQL Server 2008 Express Edition to Microsoft SQL Server conversion utility.

Component	Description
GROUPCONFIG.EXE	Device Group Configuration utility.
IMPORTNODEINFO.EXE	Imports node information to Arcserve Backup database during upgrade.
INSTALLALERT.EXE	Used for installing Alert Modules.
JOBENG.EXE	Arcserve Backup Job Engine.
JOBWINDOW.EXE	Job window configuration utility.
JOBWINUTIL.EXE	Job window.
JOBWIZARD.EXE	Generic Job Scheduler Wizard.
LDBSERVER.EXE	ONCRPC service to handle Database Engine queries.
LIC98LOG.EXE	A license service.
LIC98SERVICE.EXE	A license service.
LIC98VERSION.EXE	A license service.
LICDEBUG.EXE	Enables debugging for the license application.
LICRCMD.EXE	Enables remote command execution for the license application.
LOGWATNT.EXE	License application for providing the license event log management function.
LQSERVER.EXE	ONCRPC service to handle Job Queue queries.
MERGEALIC.EXE	License management component.
MERGEALICAT.EXE	Merge Catalog utility.
MERGEALOLF.EXE	License management component.
MERGEALROLF.EXE	License management component.
MMOADMIN.EXE	Media Management Option user interface.
MSGENG.EXE	Arcserve Backup Message Engine.
Ofant.exe	Agent for Open Files service.
ofawin.exe	Agent for Open Files console utility.
PFC.EXE	Preflight Check utility.
RMANCFG.EXE	Lets you configure Oracle databases that you are protecting using the Agent for Oracle.
SERVERMIGRATION.EXE	Lets Arcserve Backup migrate BrightStor ARCserve Backup r11.x database information to the current database configuration.

Component	Description
ServerMigrationDR.exe	Lets Arcserve Backup migrate disaster recovery information to the primary server.
SETMANPC.EXE	Disaster Recovery utility.
SETUPRD.EXE	RAID configuration command line utility.
SETUPSQL.EXE	Builds Arcserve Backup database (creates ODBC connection, sets user and password in database for the database and creates tables for the database).
SQLCLEAN.EXE	Lets Arcserve Backup clean all destroyed media information in a Microsoft SQL Server database. This component runs after you format or erase media that contains SQL Server data.
SQLCLEAN_EXP.EXE	Lets Arcserve Backup clean all destroyed media information in a Microsoft SQL Server 2008 Express Edition database, when the media is formatted or erased.
SQLTOSQL.EXE	Lets Arcserve Backup migrate database information from BrightStor ARCserve Backup r11.x and older versions of the Microsoft SQL Server database to the current release.
TAPECOMP.EXE	Tape Compare command line utility.
TAPECOPY.EXE	Tapecopy command line utility.
UNIVAGENT.EXE	Arcserve Backup Universal Agent.
UPGRADEUTIL.EXE	Lets Arcserve Backup back up and restore configuration files and registry entries during the installation process. This component runs when you perform a build to build upgrade.
vmdbupd.exe	Arcserve VMware auto-populate update utility.
VSERVICE.EXE	Lets Arcserve Backup validate customer access privileges when you install Arcserve Backup Agent for Microsoft Exchange Server on Exchange Server 2003 systems. Note: VSERVICE.EXE is an internal application that is not exposed to the end user.
W95AGENT.EXE	Client Agent for Windows on Windows 95 and Windows 98.

Arcserve Backup Domains

Arcserve Backup domains are a logical grouping of Arcserve Backup domain primary and member servers that allow easier administration of Arcserve Backup servers and users. In addition to providing a single sign-on to multiple Arcserve Backup servers, it also provides the same access level (privileges) on all the servers for the same user.

A Arcserve Backup domain has a name and a collection of one primary and one or more member servers. This allows you to manage any server from the Arcserve Backup domain to perform database management, tape and device management, and backup policy and schedule management without requiring logging in to each Arcserve Backup server separately.

Primary servers dispatch instructions about jobs and tasks to member servers in a Arcserve Backup domain. If a primary server becomes disabled or unavailable for a period time, tasks such as executing scheduled jobs and authenticating licenses on member servers will not function properly.

Each domain has a name, a mandatory designated primary server, and optional member servers. From the primary server, you can start and stop Arcserve Backup services on any member server in the domain.

When configuring the primary and member servers in a domain, the Arcserve Backup domain name must be the same on all computers in the domain. You must define the Arcserve Backup domain name when you install the primary server. You can reconfigure the domain name using the Server Configuration Wizard to change the domain membership. This wizard configures the Arcserve Backup domain name for all of the domain.

Note: Arcserve Backup server names and Arcserve Backup domain names cannot exceed 15 bytes. A name totaling 15 bytes equates to approximately 7 to 15 characters.

Manage Domain Users and Groups Using the `ca_auth` Command Line Utility

To manage the domain user and groups, Arcserve Backup provides a command line utility called `ca_auth`.

For more information about domain user management, type `ca_auth` under the command prompt, or see *Command Line Reference Guide*.

Create caroot Equivalence

By default, Arcserve Backup creates a caroot equivalency for the administrator user on the primary and all member servers during setup. However, it does not create this equivalency for any other users on the member servers and all other member users. Hence, prior to using the command line utilities in a Arcserve Backup domain, you must create this equivalency.

The caroot password can consist of any combination of alphanumeric and special characters, but may not exceed 15 bytes. A password totaling 15 bytes equates to approximately 7 to 15 characters.

By creating an equivalence list, all clients can use Arcserve Backup without the user logging into the Domain. Arcserve Backup can validate if the current user has equivalent access to the domain. The access rights to the operating system ensure a particular access level to the Arcserve Backup domain.

For more information about creating equivalence, see the section about `ca_auth` in the *Command Line Reference Guide*.

More information:

[How Arcserve Backup Equivalence Works](#) (see page 36)

[Equivalency and the System Account](#) (see page 37)

How to Manage Multiple Domains Using the Job Status Manager

Arcserve Backup lets you manage one or more Arcserve Backup using the Job Status Manager, which lets you monitor and manage the job queues relating to all Arcserve Backup domains in your enterprise.

A Arcserve Backup domain consists of the following configurations:


- A primary backup server that is connected to one or more member servers.
- A single, stand-alone backup server.

The first time that you open the Job Status Manager, the domain directory tree displays the Arcserve Backup domain that you are currently logged in to, as illustrated by the following screen.

The screenshot shows the Job Status Manager interface. On the left, a domain tree displays 'Arcserve Backup Domain' with a sub-entry 'DOMAIN-001 (DEFAULT)' containing 'PRIMARY SERVER 1'. A red line points to 'PRIMARY SERVER 1' with the label 'Default domain'. On the right, the 'Job Queue' table is displayed with the following data:

Job Name	Backup Server	Job No.	Job ID	Status	Execution T.
Database protection job	PRIMARY SERVER1	2	40	▶ READY	8/24/2009
Database protection job (#1 Makeup J...	PRIMARY SERVER1	4		⏸ HOLD	8/21/2009
Database pruning job	PRIMARY SERVER1	1	41	▶ READY	8/22/2009

At any other time, you can Add and Delete domains from the Job Status Manager.



After you add domains to the Job Status Manager, you can perform the following tasks:

- **Manage job status**--Modify job status, such as Ready, Hold, Run now, Stop Job, and refresh the data for the jobs.
- **Manage job protection**--Modify the user name for the job and the encryption password for the job.
- **Manage job maintenance**--Modify job schedules, add jobs, delete jobs, preflight check jobs, Quick search jobs, generate logs, and print information about jobs.

The capability to perform these tasks is limited by the role assigned to account for the specified domain.

Example: Role Assignment

The current Arcserve Backup domain is DomainA. The user "Arcserve User" adds two Arcserve Backup domains: DomainB and DomainC. The roles for each domain are as follows:

- DomainA: caroot user
- DomainB: Backup Operator
- DomainC: Monitor Operator

Arcserve User can perform the following tasks in each Arcserve Backup domain:

- DomainA: Modify, delete, stop all jobs in the domain.
- DomainB: Control jobs submitted using the Arcserve User account.
- DomainC: View Job information about jobs in this domain.

Add Domains to the Job Status Manager

Arcserve Backup lets you add Arcserve Backup domains to the Job Status Manager.

To add domains to the Job Status Manager

1. Open the Arcserve Backup Manager Console.
From the Navigation Bar, expand Quick Start and click Job Status.
The Job Status Manager opens.
2. Right-click Arcserve Backup Domains and select Add Domain from the pop-up menu.
The Default Server Information dialog opens.
3. From the Arcserve Backup Primary Server drop-down, select the Arcserve Backup domain that you want to add.
Complete the following:
 - **Authentication Type**--From the drop-down list, select Arcserve Backup Authentication or Windows Authentications.
 - **User**--Specify the account required to log in to the Arcserve Backup server.
 - **Password**--Specify the password corresponding with the user password.
 - **(Optional) Remember the security information**--Lets Arcserve Backup remember your user name and password.
4. Click OK.
The Arcserve Backup domain specified appears in the Job Status Manager below Arcserve Backup domains.

Delete Domains from the Job Status Manager

Arcserve Backup lets you delete domains from the Job Status Manager.

To delete domains from the Job Status Manager

1. Open the Arcserve Backup Manager Console.
From the Navigation Bar, expand Quick Start and click Job Status.
The Job Status Manager opens.
2. From the domain directory tree expand Arcserve Backup Domains and locate the domain that you want to delete.
Right-click the Domain and click Delete Domain on the pop-up menu.
The Arcserve Backup domain is deleted from the Job Status Manager.

How to Process Computer Name Changes in an Arcserve Domain

The computer name is a name that your computer uses to identify itself in a network or a domain. In a centralized management environment, an Arcserve domain can consist of a primary server and one or more member servers, or a stand-alone server. You establish the names of the Arcserve domain, the computer name of the primary server, and the computer names of the member servers when you install Arcserve Backup.

Arcserve Backup uses the computer names of the primary server and the member servers to establish communication between the servers. Arcserve Backup specifies the computer name of the primary server in the Discovery.cfg configuration file. The Discovery.cfg configuration file resides on the primary server and the member servers.

Note: The Arcserve domain name and the computer name of the primary server can be different. However, both names must not exceed 15 bytes. A name totaling 15 bytes equates to approximately 7 to 15 characters.

When you change the computer name of the primary server or the member servers, the servers cannot communicate with each other in the Arcserve domain.

In an Arcserve domain, the following scenarios exist when you change the computer name of an Arcserve server:

- The computer name of the primary server in an Arcserve domain was changed.
To ensure that the primary server and the member servers can communicate, see [Change the Computer Name of the Primary Server on the Primary Server](#) (see page 543), and [Change the Computer Name of the Primary Server on a Member Server](#) (see page 548).
- The computer name of a member server in an Arcserve domain was changed.
To ensure that the member server can communicate in the Arcserve domain, see [Change the Computer Name on a Member Server](#) (see page 550).
- The computer name of a stand-alone server was changed.
To ensure that a stand-alone server can communicate in an Arcserve domain, see [Change the Computer Name on a Stand-alone Server](#) (see page 553).
- The computer name of a server that is running the Manager Console was changed.
To ensure that a server that is running the Manager Console can communicate in an Arcserve domain, see [Change the Computer Name on a Server that is Running the Manager Console](#) (see page 553).

More information:

[Discovery.cfg Configuration File](#) (see page 549)

Change the Computer Name of the Primary Server on the Primary Server

The following procedure helps ensure that the primary server and member servers in an Arcserve domain can communicate after you change the computer name of the primary server.

You must change the computer name of the primary server before you complete these steps.

Note: You can use this procedure when you change the computer name of a stand-alone server.

If you are using only Microsoft SQL Server 2008 Express for Arcserve Backup without any other SQL instance installed, you may also need to:

- Install Microsoft SQL Server Management Studio Express (SSMSE) onto this machine, if it is not already installed. SSMSE is a graphical tool for managing SQL Server 2008 Express Edition, and for managing instances of the SQL Server Database Engine created by any edition of SQL Server 2005. For more information, see Microsoft SQL Server Management Studio Express on the Microsoft Download Center website.
- Be familiar with the sqlcmd utility, which is used to enter Transact-SQL statements, system procedures, and script files at the command prompt. For more information, see sqlcmd Utility on the Microsoft Developer Network website.

For more information about renaming systems hosting Microsoft SQL Server databases, see the following topics on the Microsoft Developer Network website:

- How to Rename a Computer that Hosts a Stand-Alone Instance of SQL Server 2005
- How to Rename a SQL Server 2005 Virtual Server
- Installing SQL Server (SQL Server 2000)
- Renaming a Server

Follow these steps:

1. Restart the target system to complete the Windows computer name change process.
2. Log in to the primary server.

Note: Do not open the Manager Console or log in to Arcserve Backup.

3. Open the Windows Command Line and change the directory to the following directory:

%ARCSERVE_HOME%

Execute the following command, to stop all Arcserve services:

```
cstop
```

All Arcserve services stop.

Note: Do not close the Windows Command Line.

4. Using a text editing application, such as Notepad, open the discovery.cfg configuration file located in the following directory on the primary server:

%ARCSERVE_HOME%\config\discovery.cfg

In the PRIMARY field, change the name of the primary server as required for your environment.

Important! Do not modify the Arcserve Domain Name in the discovery.cfg configuration file. When you change the Arcserve Domain Name in the discovery.cfg configuration file, the password for the caroot account is deleted. Use the discovery.cfg configuration file only for the purposes of changing the host name of the primary server, a member server, and a stand-alone server.

Close the file and save your changes.

For more information, see [Discovery.cfg Configuration File](#) (see page 549).

5. From the Windows Command Line that you opened earlier, execute the following command to start all Arcserve services:

```
cstart
```

All Arcserve services start.

Note: Do not close the Windows Command Line.

6. From the Windows Start menu, select All Programs, CA, ARCserve Backup, and click Server Configuration Wizard.

The Server Configuration Wizard opens and the Select Options dialog appears.

7. From the Select Options dialog, click the Select Database option, and click Next.

The Check caroot dialog opens.

8. Click Next.

Important! You must specify the caroot password to complete this task.

The System Account dialog opens.

9. Complete the following fields on the System Account dialog and click Next.

- **User Name**--Specify the Windows user name required to log in to the primary server.
- **Domain**--Specify the Windows domain name or host name of the new primary server.
- **Password**--Specify the password for the Windows user name required to log in to the primary server.

10. From the Select Database Options dialog, complete the fields and follow the prompts, as required, for your current database installation and click Next.

The subsequent dialogs that open will vary, based on whether you are running Microsoft SQL Server or Microsoft SQL Server 2008 Express in your current environment.

Note: For the Select Database option, if the server is a Central Primary Server in a Global Dashboard domain, and the new selected database is Microsoft SQL Server Express or Microsoft SQL Server 2000 (which are not supported by a Global Dashboard Central Primary Server), consider exporting and retaining the Global Dashboard information prior changing the database. After the Select Database operation is completed, the Global Dashboard information will be lost because the server will no longer function as a Central Primary Server. If you want to retain the grouping configuration and the registered branch information, you need to export this Global Dashboard information to a temporary location before performing the Select Database operation. For more information about exporting and importing Global Dashboard information, see the *Dashboard User Guide*.

Important! The Server Configuration Wizard prompts you to overwrite the existing Arcserve_DB instance, and by default, the option is enabled. To retain your previous data, such as job history, activity logs, and so on, you must clear the checkmark from the Overwrite the existing "Arcserve_DB" instance option.

11. After the Server Configuration Wizard completes the updates, click Finish.

12. From the Windows Command Line that you opened earlier, execute the following commands to stop and restart all Arcserve services:

```
cstop  
cstart
```

All Arcserve services stop and restart. The primary server functions using the new computer name.

Note: Do not close the Windows Command Line.

13. You must now create equivalence for the caroot user account.

From the Windows Command Line, execute the `ca_auth` command using the following syntax:

Note: Do not include angle brackets `<>` with your arguments.

```
ca_auth -cahost <new primary server host name> -equiv add <user name> <new primary server host name>
caroot caroot <password>
```

Equivalence is applied to the caroot user account.

14. If your Arcserve domain consists of member servers, complete the steps in [Change the Computer Name of the Primary Server on a Member Server](#) (see page 548).
15. If you are running Microsoft SQL Server 2008 Express as the Arcserve Backup database, note that SQL Express is installed as a named instance. As described in Microsoft document MS143799, execute the following commands using SSMSE to link the named instance to the new computer name:

```
sp_dropserver <old_name\instancename>
GO
sp_addserver <new_name\instancename>,local
GO
```

Restart the SQL Server instance.

16. Run the Microsoft SQL Agent Account Configuration utility to update the ODBC communication settings if any of the following conditions are met:

- The server is the primary server with a locally installed Arcserve Backup database,
- The server is a standalone server with a locally installed Arcserve Backup database,
- The server is the primary server, a standalone server, or a member server AND Arcserve Backup database or the Agent for Microsoft SQL Server is installed on the same machine.

To start the Microsoft SQL Agent Account Configuration utility, click Start on the Windows Taskbar, choose All Programs, CA, ARCserve Backup, and Microsoft SQL Agent Account Configuration.

After you start the utility, follow the prompts and accept all settings.

17. Verify the renaming operation.

To verify the renaming operation has successfully completed, select information from either `@@servername` or `sys.servers`. The `@@servername` function returns the new name, and `sys.servers` table shows the new name.


Note: After a computer has been renamed, any connections that used the old computer name must connect using the new name.

18. Release all Arcserve Backup licenses registered to the former primary server.

Note: For more information, see [Release Licenses from Servers](#) (see page 525).

19. Update the Arcserve Backup System account. To do this, open the Arcserve Backup Manager Console and then open the Server Admin Manager.

Select the Arcserve Backup server and click Arcserve Backup System Account as illustrated by the following:



The Arcserve Backup System Account dialog opens.

20. Complete the fields that follow:
 - Microsoft Windows User Account
 - Password
 - Microsoft Windows Domain

Click OK.

21. Open the Job Status Manager and complete the following tasks:
 - Delete and re-create the Database Pruning Job.
Note: For more information, see [Re-create the Arcserve Backup Database Pruning Job](#) (see page 328).
 - Modify the Database Protection Job and any other backup jobs that are set to run on the renamed server, to update the Staging and Destination locations.
Note: For more information, see [Modify or Create a Custom Database Protection Job](#) (see page 614).
22. Perform a full backup of the Arcserve Backup database.

More information:

[BSBAB--Manage Arcserve Servers Using the Server Configuration Wizard-AdminW](#) (see page 559)

Change the Computer Name of the Primary Server on a Member Server

The following procedure ensures that the primary server and member servers in an Arcserve domain can communicate after you change the computer name of the primary server.

Be aware of the following considerations:

- You must change the computer name of the primary server before you complete this task.

Note: For more information, see [Change the Computer Name of the Primary Server on the Primary Server](#) (see page 543).

- You must complete this task on all member servers in the Arcserve domain.
- The Server Configuration Wizard may display the following messages while you are completing this task:
 - Arcserve Backup is unable to connect to the original primary server. You can safely click Continue to clear this message.
 - Arcserve Backup is unable to unregister the member server. You can safely click Yes to clear this message.

To change the computer name of the primary server on a member server

1. Log in to the member server.

Note: You do not need to start Arcserve Backup to complete this task.

2. From the Windows Start menu on the member server, select All Programs, CA, ARCserve Backup, and click Server Configuration Wizard.

The Server Configuration Wizard opens.

3. From the Select Options dialog, click Move this server to another Arcserve Backup domain and click Next.

The Check caroot dialog opens.

4. Specify the password for the caroot account and click Next.

5. In the Add to another Arcserve Backup domain dialog, specify the new hostname for the primary server, specify the password for the caroot account and click Next.


The System Account dialog opens.

6. In the System Account dialog, complete the following fields:
 - **User Name**--Lets you specify the Windows user name that is required to log in to the primary server.
 - **Domain**--Lets you specify the Windows domain name or host name of the new primary server.
 - **Password**--Lets you specify the password for the Windows user name that is required to log in to the primary server.

Click Next and follow the on-screen instructions to complete the configuration.
7. Repeat Steps 1 through 6 for all member servers in the Arcserve domain.

Discovery.cfg Configuration File

The discovery.cfg configuration file specifies the name of the Arcserve domain and the computer name of the primary server, as illustrated by the following example:



```

# Sample Discovery Service Configuration File
#
# $ARCSERVE_HOME/config/discovery.cfg
#
#
# Please use BAB tool ARCServeCfg.exe to change the configuration.
# Don't change it manually.
#
#
# PRIMARY required: primary discovery server host name
#PRIMARY CASPrimaryServer
#
# DOMAIN the name of ARCServe domain. DOMAIN is required when localhost is
# either the PRIMARY Discovery Server or the Member Discovery Server
#
# Note : This is a logical name, not associated with any actual machine name.
#
#
DOMAIN ARCSERVE_01
PRIMARY AS_PRIMARY

```

The discovery.cfg configuration file is located in the following directory on the primary and member servers:

```
%ARCSERVE_HOME%\config\discovery.cfg
```

Important! Do not modify the Arcserve Domain Name in the discovery.cfg configuration file. When you change the Arcserve Domain Name in the discovery.cfg configuration file, the password for the caroot account is deleted. Use the discovery.cfg configuration file only for the purposes of changing the host name of the primary server, a member server, and a stand-alone server.

Change the Computer Name of a Member Server

The following procedure helps ensure that the member servers in a Arcserve Backup domain can communicate with the primary server after you change the computer name of the member server.

You must change the computer name of the member server before you complete this procedure.

To change the computer name of a member server

1. Log in to the member server.

Note: Do not open the Manager Console or log in to Arcserve Backup.

2. Open the Windows Command Line and change the directory to the following directory:

```
%ARCSERVE_HOME%
```

Execute the following commands, to stop and start all Arcserve services:

```
cstop  
cstart
```

All Arcserve services stop and restart.

Note: Various Arcserve Backup services will not start after the cstart command completes. This is expected behavior that will not adversely affect this procedure.

Do not close the Windows Command Line.

3. From the Windows Start menu, select All Programs, CA, ARCserve Backup, and click Server Configuration Wizard.

The Server Configuration Wizard opens and the Select Options dialog appears.

4. Click the Move this server to another Arcserve Backup domain option and click Next.

The Add to Another Arcserve Backup Domain dialog opens.

5. On the Add to Another Arcserve Backup Domain dialog, complete the following fields and click Next.
 - **Primary Server Name**--Specify the name of the primary server.
 - **Password**--Specify the password for the caroot user account.

The System Account dialog opens.

6. Complete the following fields on the System Account dialog:

- **User Name**--Specify the Windows user name required to log in to the member server.
- **Domain**--Specify the Windows domain name or host name of the new member server.
- **Password**--Specify the password for the Windows user name required to log in to the member server.

Click Next.

Note: At this time, a pop-up message may appear to inform you that various Arcserve Backup services will not start. This is expected behavior that will not adversely affect this procedure.

7. Click OK.

The Arcserve Backup Data Migration dialog opens.

8. On the Arcserve Backup Data Migration dialog, click Next.

The Migrate Server Data Dialog opens.

9. On the Migrate Server Data Dialog, click Start.

The Complete dialog opens after the data migration process starts and completes.

10. Click Finish on the Complete dialog.

Note: At this time, jobs do not transfer to the new member server. Continue to the next step to transfer the jobs to the new member server.

11. Open the Job Status Manager.

Locate a job associated with the old member server.

Right-click the job and click Modify Job on the pop-up menu.

From the Backup Manager, click the Destination tab.

Specify the new member server as the destination for the job.

Submit the job with a Hold status.

Close the Backup Manager.

Note: You cannot modify and transfer jobs if the source data for the jobs reside on the old member server. As such, you must delete jobs with this configuration and then re-create them on the new member server.

12. After the member server name change is complete, the old (invalid) member server name remains in the Arcserve Backup Manager. To remove the invalid member server name from the Manager, do the following:

- a. Open the command line window and browse the Arcserve Backup installation directory.
- b. Execute the following command:

```
bab -cahost <primary server> -removehost <invalid member server>
```

Example: The following syntax describes a primary server named A, and an invalid member server named B.

```
bab -cahost A -removehost B
```

Note: If your Arcserve Backup implementation contains more than one invalid member server, repeat this step (b) for each member server.

- c. Log in to the primary server or the server hosting the Arcserve Backup database to verify the status of the Arcserve Backup database.

(Optional) Open Microsoft SQL Server Management Studio and open to the Arcserve Backup database instance using Windows authentication.

For example, the path to a Microsoft SQL Server Express Edition database is as follows:

```
<server name>\ARCServe_DB
```

- **Windows authentication is required**--If you must log in to the Arcserve Backup database using Windows authentication, execute the following command (applies to Microsoft SQL Server and Microsoft SQL Server Express Edition databases):

```
osql -S <server_name[instance_name]> -E -d asdb -Q "delete from ashost where rhostname = '<member server name>'"
```

Example: The following syntax describes a Arcserve Backup database named asdb and the database requires Windows authentication:

```
osql -S A -E -d asdb -Q "delete from ashost where rhostname = 'B' "
```

- **Windows authentication is not required**--If Windows authentication is not required to log in to the Arcserve Backup database, execute the following command (applies only to Microsoft SQL Server databases):

```
osql -S <server_name[instance_name]> -U <login_user> -d asdb -Q "delete from ashost where rhostname = '<member server name>'"
```

Example: The following syntax describes a Arcserve Backup database named asdb, an invalid member server named B, a Microsoft SQL Server user name sa, and a password 123.

```
osql -S A -U sa -d asdb -Q "delete from ashost where rhostname = 'B' " password: 123
```

Note: If your Arcserve Backup implementation contains more than one invalid member server, repeat this step (c) for each member server.

If you can view the details for the instance, the database is functioning properly. Close Microsoft SQL Server Management Studio and continue to the next step.

If Microsoft SQL Server Management Studio displays pop-up messages, the database instance is not functioning properly. You must try to resolve the problems indicated on the pop-up messages and then verify the status of the Arcserve Backup database.

If the above commands complete successfully, the following message appears:

```
n rows affected
```

If the above commands did not complete successfully, verify that the server name, the user name that is logged in to Arcserve Backup, and database name are correct, and then repeat this step.

13. To verify the changes, open the Manager Console, open the Backup Manager, and select the Source tab.

Expand the Windows Systems object in the Source directory tree.

The member server, with its new host name, appears under the Windows Systems object.

More information:

[BSBAB--Manage Arcserve Servers Using the Server Configuration Wizard-AdminW](#) (see page 559)

Change the Computer Name of a Stand-alone Server

A stand-alone server is an Arcserve server that resides in an Arcserve domain that does not manage member servers.

The procedure to change the computer name of a stand-alone server is identical to that of changing the computer name of a primary server.

Note: For more information, see [Change the Computer Name of the Primary Server on the Primary Server](#) (see page 543).

Change the Computer Name of a Server that is Running the Manager Console

When you change the computer name of a server that is running the Manager Console, you do not need to process modifications to the primary server, a stand-alone server, a member server, or the server that is running the Manage Console.

Manage User Profiles Using the User Profile Utility

The Arcserve Backup User Profile utility lets the Arcserve Backup administrator control user access to Arcserve Backup.

This section contains the following topics:

[Add a User Using the User Profile Utility](#) (see page 554)

[Delete a User Using the User Profile Utility](#) (see page 555)

[Change a User Password Using the User Profile Utility](#) (see page 555)

[Assign Roles to Users Using the User Profile Utility](#) (see page 556)

[Suspend Users Using the User Profile Utility](#) (see page 556)

Add a User Using the User Profile Utility

A default user is created when Arcserve Backup is installed. The default user name is caroot.

To add a user profile using the User Profile Utility

1. Open the Arcserve Backup Manager Console.

From the Quick Start Menu select Administration and click User Profile.

The User Profile utility opens.

2. Click Add on the toolbar.

The Add User dialog opens.

3. Click the General tab.

4. In the User Name field, specify a name for the user.

Select one of the following options:

- **Windows Authentication**--Lets you specify a Windows user name that the user will use to log in to Arcserve Backup.
- **Arcserve Backup Authentication**--Lets you specify a non-Windows user name that the user will use to log in to Arcserve Backup.

Note: If you specified Arcserve Backup Authentication, you must complete the following fields:

- **Password**
- **Confirm Password**

Click OK.

The user is added.

Delete a User Using the User Profile Utility

The User Profile Utility lets you delete Arcserve Backup users.

Note: You cannot delete the default Arcserve Backup user name (caroot).

To delete a user using the User Profile Utility

1. Open the Arcserve Backup Manager Console.
From the Quick Start Menu select Administration and click User Profile.
The User Profile utility opens.
2. Select the user that you want to delete and click Delete on the toolbar.
Click OK to confirm that you want to delete the user profile.
The user is deleted.

Change a User Password Using the User Profile Utility

The User Profile Utility lets you change Arcserve Backup user passwords.

To change a user password using the User Profile Utility

1. Open the Arcserve Backup Manager Console.
From the Quick Start Menu select Administration and click User Profile.
The User Profile utility opens.
2. Select the user that you want to modify and click Properties on the toolbar.
The User Properties dialog opens.
3. Click the General tab.
Complete the following fields:
 - **Password**
 - **Confirm Password**Click OK.
The user password is changed.

Assign Roles to Users Using the User Profile Utility

The User Profile Utility lets you assign Arcserve Backup roles to Arcserve Backup users.

To assign roles to users using the User Profile Utility

1. Open the Arcserve Backup Manager Console.
From the Quick Start Menu select Administration and click User Profile.
The User Profile utility opens.
2. Select the user that you want to modify and click Properties on the toolbar.
The User Properties dialog opens.
3. Click the Roles tab.
Check the check box next to the Arcserve Backup roles that you want to assign to the user.
Note: For more information, see [Roles and Permissions](#) (see page 83).
Click OK.
The roles are applied to the user.

Suspend Users Using the User Profile Utility

The User Profile Utility lets you temporarily suspend users from logging in to and using Arcserve Backup.

To suspend users using the User Profile Utility

1. Open the Arcserve Backup Manager Console.
From the Quick Start Menu select Administration and click User Profile.
The User Profile utility opens.
2. Select the user that you want to modify and click Properties on the toolbar.
The User Properties dialog opens.
3. Click the General tab.
In the Status field, click Suspend.
Click OK.
The user is suspended.

Note: To reactivate users, perform the above steps and click Active in the Status field.

Restore the Arcserve Backup Job Queue

You can protect the Job Queue by backing it up using the following methods:

- Back up the Job Queue using the Database Protection Job.
- Back up the Arcserve Backup primary server or stand-alone server and include the directory that contains the Arcserve Backup database with the backup.
- Back up data and specify the Job Scripts option on the Operations section of the Global Options dialog.

In all of these methods, the Job Queue is one of the last few objects backed up during the job. If you used the Database Protection Job, or included the Arcserve Backup database in the same backup job, the Arcserve Backup database and Job Queue will reside on the same backup media, and the Job Queue sessions will be between one and six sessions before the Arcserve Backup database.

In the event the Arcserve Backup Job Queue is damaged or is deleted in error, use the following steps to restore the Job Queue to the last backup.

To restore the Arcserve Backup Job Queue

1. Ensure that there are no jobs running on any Arcserve Backup server in the affected Arcserve Backup domain.
2. Open the Restore Manager and select Restore by Tree from the methods drop-down list.

Expand the Windows Systems object and locate the Arcserve Backup server, or the remote server where the Arcserve Backup database is located.

Expand the Arcserve Backup server, locate and click the Arcserve Backup database.

From Properties on the Restore Manager window, locate and notate the Media Name, the ID, and the Session Number for the Arcserve Backup database.

Arcserve Database (Microsoft SQL Express instance)

Agent Option...

Properties

Additional Information	
Media Name	MEDIA_011
ID	5467
Sequence No.	1
Session No.	8
Backed Up On	10/22/08 3:39 PM
Source Path	\\COMP-011
Session Status	Finished
Session Type	Arcserve Database (Microsoft SQL Express instance)

Select Restore by Session from the restore methods drop-down list.

The Session directory tree opens.

3. From the Session directory tree, locate and expand the Backup Media containing the Arcserve Backup database backup data.

Locate the session that contains the Arcserve Backup database, and start searching upward from there. The session path of the Job Queue backup session will end in 00000001.QSD, and the Session Type will be Arcserve Job Queue. Locate the session that contains the Job Queue backup session, as illustrated below.



Click the check box next to the session containing the Job Queue backup session.

4. Click the Destination tab.
5. Clear the check mark next to Restore files to their original location, and specify an alternate location to restore the Job Queue backup session.

Note: The Client Agent for Windows must be installed on the system containing the alternate location and the alternate location must be an empty directory (For example, C:\Temp). The best practice is to specify directory on the Arcserve Backup primary or stand-alone server.

Submit the job.

6. Closed the Restore Manager window.
7. After the restore job is complete, open the Server Admin and stop all Arcserve Backup services by doing the following:
 - a. Locate and select the primary server or stand-alone server.
 - b. Right-click the Arcserve Backup server and select Stop All Services from the pop-up menu.

All services stop on the primary or stand-alone server.

8. Browse to directory where you restored the Job Queue backup session.

Copy all Job Queue files under the folder that you restored to the following directory:

ARCSERVE_HOME\00000001.qsd

Note: <ARCSERVE_HOME> represents the directory where you installed Arcserve Backup. By default, Arcserve Backup is installed in the directory that follows:

C:\Program Files\CA\ARCserve Backup

9. Open the Server Admin and restart all Arcserve Backup services by doing the following:
 - a. Locate and select the primary server or stand-alone server.
 - b. Right-click the Arcserve Backup up server and select Start All Services from the pop-up menu.

All Arcserve Backup services restart on the primary or stand-alone server.

Open the Job Queue Manager and you will see that the job queue has been restored to its original form. The Arcserve Backup Job Queue is restored, and you can resume normal operation.

BSBAB--Manage Arcserve Servers Using the Server Configuration Wizard-AdminW

The Server Configuration Wizard lets you manage how Arcserve Backup servers function. Using the Server Configuration Wizard, you can perform the following tasks:

- Manage the roles of the servers in your Arcserve Backup domain. For example, you can:
 - Promote a Arcserve Backup member server to a Arcserve Backup primary server.
 - Demote a Arcserve Backup primary server to a Arcserve Backup member server.
 - Allow a member server to separate from one Arcserve Backup domain and join a different Arcserve Backup domain.

- Select the application that you want to use to manage the Arcserve Backup database.

For Microsoft SQL Server 2008 Express installations, the database must be installed on the primary server. If you require remote database communication, you must use Microsoft SQL Server to host the Arcserve database.

- Move the Arcserve Backup database to other systems or use a different SQL Server database instance in your environment.
- Repair the Arcserve database connection to a primary server and member servers.

- Register a member server with a Arcserve Backup domain primary server.

To register a member server with a domain primary server, you must provide valid credentials (for example, the user name and password). After Arcserve Backup authenticates your credentials, the member server is registered into the Arcserve Backup database.

Arcserve Backup lets you register the member server with the Arcserve Backup primary server when you install Arcserve Backup. If the registration process fails when you are installing Arcserve Backup, Setup displays messages to notify you that an error occurred.

- Specify the Arcserve Backup Domain Administrator (caroot) password on a primary server.
- Correct installation failures.

When you install Arcserve Backup, the installation process can fail under the following scenarios:

- Arcserve Backup cannot communicate or authenticate properly with the Arcserve Backup database.
- Arcserve Backup cannot authenticate the caroot account or a system account.

If a database communication error or user authentication error occurs, the installation wizard displays an error message. To remedy the problem, run the Server Configuration Wizard.

More information:

[Tasks You Can Perform Using the Server Configuration Wizard](#) (see page 561)

[Start the Server Configuration Wizard](#) (see page 564)

Tasks You Can Perform Using the Server Configuration Wizard

Using the Server Configuration Wizard you can perform the following tasks:

■ Primary Server and Stand-alone Server Tasks:

You can perform the following tasks on primary and stand-alone servers:

- Modify the Arcserve Backup Domain Administrator (caroot) account password.
- The caroot account password lets you log in to the Arcserve Backup Manager Console to perform administrative tasks.
- Specify the application that you want to use to host the Arcserve Backup database.
- You can specify Microsoft SQL Server 2008 Express or Microsoft SQL Server as the Arcserve database application. SQL Server 2008 Express must be installed locally to the Arcserve Backup primary server. SQL Server can be installed locally or remotely to the Arcserve Backup primary server.
- Move the Arcserve Backup database to a different system, instance, or both.
- Repair database connections with member servers.
- Re-initialize the Arcserve Backup database.
- Specify SQL Server collation to ensure that you can search and sort backup data that contains Unicode-based characters.
- Correct installation failures.
- Demote a primary server to a member server.

Note: Arcserve Backup does not support migrating Arcserve Backup database information from multiple Arcserve Backup domains into a single Arcserve Backup domain. Although you can demote a primary server and allow it to join a different Arcserve Backup domain, joining a different domain will result in the loss of the backup job history from the demoted primary server, and you will not be able to view media and session details in the Restore Manager on the demoted server. For more information, see [Data Migration Limitations in an Arcserve Domain](#) (see page 562).

■ Member Server Tasks:

You can perform the following tasks on member servers:

- Assign the member server to a different Arcserve Backup domain.
- Promote a member server to a primary server or stand-alone server.

Note: To enable central management capabilities on the promoted server, you must install the Central Management Option on the new primary server after the promotion process is complete.

- Repair the database connection.

- Correct installation failures.
- Register virtual clusters as Arcserve Backup domain member servers.

Note: Use the Server Admin Manager to modify the Arcserve Backup System Account on a primary server and a member server (for example, user name, password, and so on). For more information, see [Change or Modify the Arcserve Backup System Account](#) (see page 520).

Data Migration Limitations in a Arcserve Backup Domain

The Server Configuration Wizard lets you define the roles of the servers in an Arcserve domain and specify the application that you want to use to host the Arcserve database instance.

Arcserve Backup lets you migrate Arcserve database instance data as described by the following scenarios.

Scenario 1:

You exchange the roles of the primary server and a member server in an Arcserve domain. You can successfully migrate data under the following conditions:

- The original primary server hosted the Arcserve database instance using Microsoft SQL Server 2008 Express Edition and new primary server is hosting the Arcserve database instance using Microsoft SQL Server 2008 Express Edition.
- The original primary server hosted the Arcserve database instance using Microsoft SQL Server and new primary server is hosting the Arcserve database instance using Microsoft SQL Server.

Important! Arcserve Backup does not support data migration when the original primary server hosted the Arcserve database with Microsoft SQL Server and the new primary is hosting the Arcserve database with Microsoft SQL Server 2008 Express Edition.

To accomplish a successful data migration, complete the following steps:

1. From the primary server that you want to demote, back up the Arcserve database using the Database Protection Job.

Note: Allow the Database Protection Job to finish before continuing.

2. Promote the member server to a primary server.
3. Demote the original primary server and allow it to join new primary server's domain.

4. From the [Agent Restore Options dialog](#) (see page 635) on the new primary server, specify the following options:
 - Use current Arcserve database as original location.
 - Preserve current Arcserve domain memberships.
5. Restore the original Arcserve database to the new primary server.

Scenario 2:

You modify the application hosting the Arcserve database from Microsoft SQL Server 2008 Express Edition to Microsoft SQL Server.

Note: This scenario applies to Arcserve primary server and Arcserve stand-alone server installations.

To accomplish a successful data migration, complete the following steps.

1. Run the [Server Configuration Wizard](#) (see page 564) on the primary or stand-alone server and specify the Select database option.

After the database modification and configuration process is complete, the Server Configuration Wizard prompts you to migrate the data from the old database instance to the new database instance.

2. Migrate the data from the Microsoft SQL Server 2008 Express Edition instance to the Microsoft SQL Server instance.
3. Uninstall Agent for Arcserve Database.
4. Reinstall the Agent for Arcserve Database to create the new database instance using the following executable:
`<ARCserve_HOME>\Packages\ASDBSQLAgent\SQLAgentRmtInst.exe`
5. Uninstall the ARCSERVE_DB instance of Microsoft SQL Server 2008 Express Edition.

Scenario 3:

You exchange the roles of the primary server and a member server in an Arcserve domain. The original primary server hosted the Arcserve database instance using Microsoft SQL Server 2008 Express Edition and new primary server is hosting the Arcserve database instance using Microsoft SQL Server.

Note: In this scenario you must convert the Arcserve database from a Microsoft SQL Server 2008 Express Edition instance to a Microsoft SQL Server instance on the primary server that you want to demote before you back up the Arcserve database instance.

To accomplish a successful data migration, complete the following steps:

1. Run the Server Configuration Wizard on the primary server and specify the Select database option.

After the database modification and configuration process is complete, the Server Configuration Wizard prompts you to migrate the data from the old database instance to the new database instance.

2. Migrate the data from the Microsoft SQL Server 2008 Express Edition instance to the Microsoft SQL Server instance.
3. From the primary server that you want to demote, back up the Arcserve database using the Database Protection Job.

Note: Allow the Database Protection Job to finish before continuing.

4. Promote the member server to a primary server.
5. Demote the original primary server and allow it to join new primary server's domain.
6. From the [Agent Restore Options dialog](#) (see page 635) on the new primary server, specify the following options:
 - Use current ASDB as original location.
 - Preserve current Arcserve domain memberships.
7. Restore the original Arcserve database to the new primary server.

Start the Server Configuration Wizard

The Server Configuration Wizard lets you manage how Arcserve Backup servers function.

To start the Server Configuration Wizard

1. From the Windows Start menu, select Programs (or All Programs), CA, ARCserve Backup, and click Server Configuration Wizard.

The Server Configuration Wizard opens.

2. Select the task that you want to perform, click Next, and follow the on-screen procedures to complete your configurations.

Promote a Member Server to a Primary Server

Using the Server Configuration Wizard, you can promote a Arcserve Backup member server to a Arcserve Backup primary server.

Before you promote a member server to a primary server, the following considerations apply:

- All jobs must be stopped on the member server before the upgrade process starts. Arcserve Backup detects all jobs with a Ready Status and places them in a Hold status for you. If there are jobs in progress, Arcserve Backup displays a message and the upgrade process pauses until all jobs in progress are complete.
- If the promoted primary server will be configured as the central server in a Global Dashboard domain and you want to continue to use the grouping configuration and registered branch information collected from the old central server, you must import this dashboard information into the server after it has been promoted. For more information about how to import this dashboard information, see the *Dashboard User Guide*.
- During the upgrade process, you will be prompted to specify a Arcserve Backup database application. You can specify Microsoft SQL Server 2008 Express Edition or Microsoft SQL Server.

Microsoft SQL Server 2008 Express Installations

- You must install the database local to the primary server.

Microsoft SQL Server Installations

- You can install the Arcserve Backup database local or remote to the primary server.
- Microsoft SQL Server does not support local installations when Arcserve Backup is installed in NEC CLUSTERPRO environments.
- For remote Microsoft SQL Server database installations, the primary server must have a system account that properly authenticates with SQL Server and communicates via ODBC before you start the upgrade process.

To specify ODBC communication, do the following:

1. Open the Windows Control Panel, select Administrative Tools, Data Sources (ODBC), and System DSN.
2. Add a System Data Source labeled as follows:
Name: ASNT
Server: MachineName\InstanceName
3. Follow the on-screen instructions to test and complete the configuration.

- To enable central management capabilities, you must install the Central Management Option on the new primary server after the promotion process is complete.
 - **Note:** Use the Server Admin to install Arcserve Backup options, such as the Central Management Option, on the new primary server after the promotion process is complete. For more information, see [Install and Uninstall Arcserve Backup Server Based Options](#) (see page 590).

Promote a member server to a primary server


1. From the Windows Start menu, click Start, select All Programs, CA, ARCserve Backup, and click Server Configuration Wizard.

The Server Configuration Wizard opens.

2. Click the **Promote this server to primary server** option and then click **Next**.
3. Follow the on-screen instructions to complete the configuration.

Note: If you protect data that contains Unicode-based characters from East Asian languages (for example, JIS2004) you must enable SQL collation to ensure that you can search and sort the data. To do this, click Language Support Options on the SQL Server Express Instance dialog and follow the on-screen instructions to complete the configuration.

If you are hosting the Arcserve Backup database with Microsoft SQL Server, you click Language Support Options on the Select Database Installation Path dialog.)



After the configuration is complete, you must install the Arcserve Backup database protection agent on the system hosting the Arcserve Backup database.

4. To install the Arcserve database protection agent, do one of the following:
 - If the SQL Server database is installed on the Arcserve Backup primary server, open Windows Explorer and browse to the following directory:
C:\Program Files\CA\ARCserve Backup\Packages\ASDBSQLAgent
 - If the SQL server database is not installed on the Arcserve Backup primary server, open Windows Explorer and browse to the following directory:
C:\Program Files\CA\ARCserve Backup\Packages\ASDBSQLAgent
Copy the contents of the ASDBSQLAgent directory to any location on the system hosting the SQL Server database installation.
5. In the ASDBSQLAgent directory, double-click the following file:
SQLAgentRmtInst.exe
The **Arcserve Backup Agent for SQL Setup** dialog appears.
6. Complete the following fields, as required, for your installation:
 - SQL Instance Name
Specify the name of the SQL instance that you want to protect.
 - Auth Mode
Specify the authentication mode that the agent will use to communicate with and protect the Arcserve database.
If you specify SQL Authentication as the authentication mode, complete the following fields:
 - SQL SA Name
Specify the SQL system account name.
 - SQL SA Password
Specify the SQL system account password.
7. Click **Install** and follow the on-screen instructions to complete the installation.

Demote a Primary Server or Stand-alone Server to a Member Server

Using the Server Configuration Wizard, you can demote a Arcserve Backup primary server and a Arcserve Backup stand-alone server to a Arcserve Backup member server.

The demotion process lets you transfer all Arcserve Backup database information that relates to jobs, media, devices, and so on for the primary server and the related member servers and data mover servers to a different Arcserve Backup domain.

Review the following considerations and best practices before you demote a primary server to a member server:

- All jobs must be stopped on the primary server before the demotion process starts. Arcserve Backup detects all jobs with a Ready Status and places them in a Hold status for you. If there are jobs in progress, Arcserve Backup displays a message and the demotion process pauses until all jobs in progress are complete.
- You must specify Arcserve Backup authentication credentials to allow the demoted primary server to join the domain of another primary server (for example, *caroot* and your Arcserve Backup password). The process of allowing a member server to join a Arcserve Backup domain does not support using Windows authentication.
- If the primary server to be demoted contains member server relationships, data mover server relationships, or both, the Server Configuration Wizard presents you with list of servers that the primary server is managing and the following options:
 - Demote the primary server.
 - Demote the primary server and allow the member servers and the data mover servers that it is managing to join the new domain.

As a best practice, you should move the member servers to different Arcserve Backup domains and register the data mover servers with different primary servers before you demote the primary server. Optionally, you can promote the member servers to primary servers or stand-alone Arcserve Backup servers. Likewise, you should promote the member servers before you demote the primary server.

- If the primary server to be demoted contains data mover server relationships, you must reconfigure the data mover server's file system device settings after you demote the primary server and register the data mover server with different primary server.
- If the primary server to be demoted is also configured as the central dashboard server in a Global Dashboard domain and you want to retain the grouping configuration and the registered branch information, you must first export this dashboard information to a temporary location until a new central dashboard server is configured. For more information about how to export this dashboard information, see the *Dashboard User Guide*.

- If the primary server to be demoted is joining a Arcserve Backup domain that is running a remote Microsoft SQL Server database installation, and the primary server communicates with the SQL Server database using Windows authentication, the new member server must have a system account that uses Windows authentication and communicates via ODBC before you start the demotion process.
- All registered licenses will be removed from the demoted primary server.
- If you set up jobs using a different caroot user account before the demotion, you must manage the migrated jobs on the primary server in the domain that the new member server joins using the original caroot account and password as the job owner for all migrated jobs.
- The table that follows describes database migration scenarios and the type of data that Arcserve Backup migrates from the demoted primary server to the domain that the new member server (demoted primary server) joins:

Database on Demoted Primary Server	Database on New Primary Server	Database Data Migrates?	Job and Job History Data Migrates?	Authentication Data Migrates?
Microsoft SQL Server	Microsoft SQL Server Express Edition	No	No	No
Microsoft SQL Server Express Edition	Microsoft SQL Server Express Edition	No	No	No

Database on Demoted Primary Server	Database on New Primary Server	Database Data Migrates?	Job and Job History Data Migrates?	Authentication Data Migrates?
Microsoft SQL Server	Microsoft SQL Server	Yes	Yes Note: If you set up jobs using a different caroot user account before the demotion, you must manage the migrated jobs on the primary server in the domain that the new member server joins using the original caroot account and password as the job owner for all migrated jobs.	No
Microsoft SQL Server Express Edition	Microsoft SQL Server	Yes	Yes Note: If you set up jobs using a different caroot user account before the demotion, you must manage the migrated jobs on the primary server in the domain that the new member server joins using the original caroot account and password as the job owner for all migrated jobs.	No

To demote a primary server or stand-alone server to a member server

1. From the Windows Start menu, click Start, select All Programs, CA, ARCserve Backup, and click Server Configuration Wizard.
The Server Configuration Wizard opens.
2. Click Demote this server to member server and then click Next.

3. Follow the on-screen instructions to complete the configuration.
4. (Optional) After the configuration is complete, you can uninstall the Arcserve database protection agent from the server that you demoted by doing the following:

- From the Windows Control Panel, open Add and Remove Programs.
- Browse to and select Arcserve Backup Agent for Microsoft SQL.
- Click the Remove button to uninstall the agent.

The Uninstall Agent message box appears.

5. Select the Agent for Arcserve Database option and click OK.

Follow the on-screen instructions to complete the uninstallation.

6. (Optional) To move the CA licenses from the demoted primary server to a different Arcserve Backup primary server, do the following:

- a. On the demoted primary server, locate the file labeled ca.olf in the following directory:

```
c:\program files\ca\SharedComponents\ca_lic
```

- b. Save ca.olf as ca.old.

- c. Copy ca.old from the demoted primary server to the following directory on the other primary server:

```
c:\program files\ca\SharedComponents\ca_lic
```

- d. On the other Arcserve Backup Primary server, open a Command Line window and open the following utility.

```
c:\program files\ca\SharedComponents\ca_lic\mergeolf.exe
```

For more information about using the MergeOLF command, see the *Command Line Reference Guide*.

7. Uninstall the Arcserve Backup server-based options from the demoted primary server.

You can use the Server Admin Manager to uninstall the following server-based options from the demoted primary server:

- Central Management Option
- Tape Library Option
- Storage Area Network (SAN) Option

Note: For more information, see [Install and Uninstall Arcserve Backup Server Based Options](#) (see page 590).

You must remove all other options from the demoted primary server (for example, Global Dashboard) using Windows Add and Remove Programs.

Note: For information about Arcserve Backup server-based options that you can install on Arcserve Backup servers, see "Types of Arcserve Backup Server Installations" in the *Implementation Guide*.

8. Restart the primary server in the domain that the member server (demoted primary server or stand-alone server) joined. This step helps to ensure that the information about backup data associated with the member server is accurate on the primary server.

Move a Member Server to a Different Arcserve Backup Domain

Using the Server Configuration Wizard, you can move a member server to a different Arcserve Backup domain.

Before you move a member server to a different Arcserve Backup domain, the following considerations apply:

- All jobs must be stopped on the member server before the move process starts. Arcserve Backup detects all jobs with a Ready Status and places them in a Hold status for you. If there are jobs in progress, Arcserve Backup displays a message and the move process pauses until all jobs in progress are complete.
- After the member server joins a different Arcserve Backup domain, the jobs associated with the previous domain will migrate to the new domain. However, all database information relating to the member server will remain with the previous domain.

To move a member server to a different Arcserve Backup domain

1. From the Windows Start menu, click Start, select All Programs, CA, ARCserve Backup, and click Server Configuration Wizard.

The Server Configuration Wizard opens.

2. Click the **Move this server to another Arcserve Backup domain** option and then click **Next**.
3. Follow the on-screen instructions to complete the configuration.

Change the Password for the Arcserve Backup Domain Administrator (caroot) Account

The caroot password can consist of any combination of alphanumeric and special character but may not exceed 15 bytes. A password totaling 15 bytes equates to approximately 7 to 15 characters. Using the Server Configuration Wizard, you can change the password for the Arcserve Backup Domain Administrator (caroot) account. The Domain Administrator account lets you log in to the Arcserve Backup Manager Console to perform administrative tasks.

The caroot password can consist of any combination of alphanumeric and special characters, but may not exceed 15 bytes. A password totaling 15 bytes equates to approximately 7 to 15 characters.

Before you change the password for the Arcserve Backup system account, you must be logged in to a Arcserve Backup primary server.

Note: Use the Server Admin to change the password to the system account on a member server. For more information, see [Change or Modify the Arcserve Backup System Account](#) (see page 520).

To change the password for the Arcserve Backup Domain Administrator (caroot) account

1. From the Windows Start menu, click Start, select All Programs, CA, ARCserve Backup, and click Server Configuration Wizard.

The Server Configuration Wizard opens.

2. Click the **Password for Backup Server Logon and Administration** option and then click **Next**.
3. Follow the on-screen instructions to complete the configuration.

Repair the Arcserve Backup Configuration

Installation errors can occur when you install and upgrade Arcserve Backup from a previous release on a primary server or a member server. For example, an incomplete installation occurred.

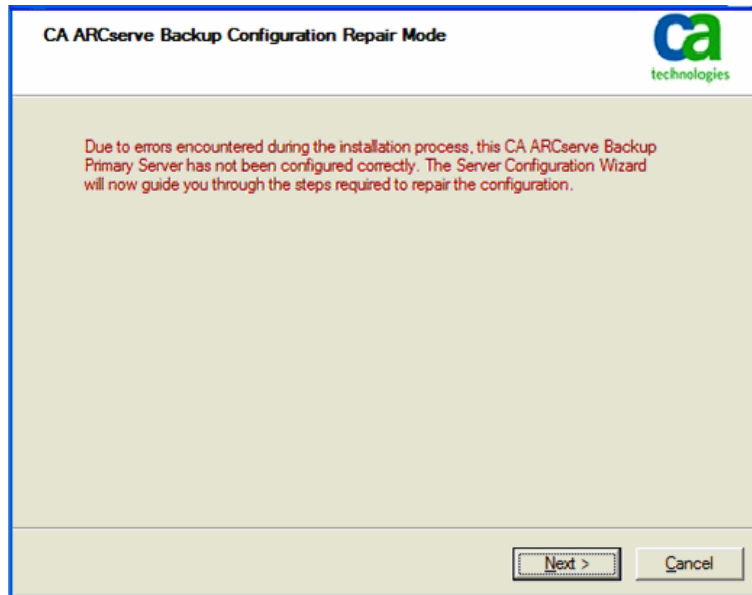
If the installation wizard detects errors, the Server Configuration Wizard prompts you to correct the installation errors.

The following procedure describes how to correct the Arcserve Backup configuration.

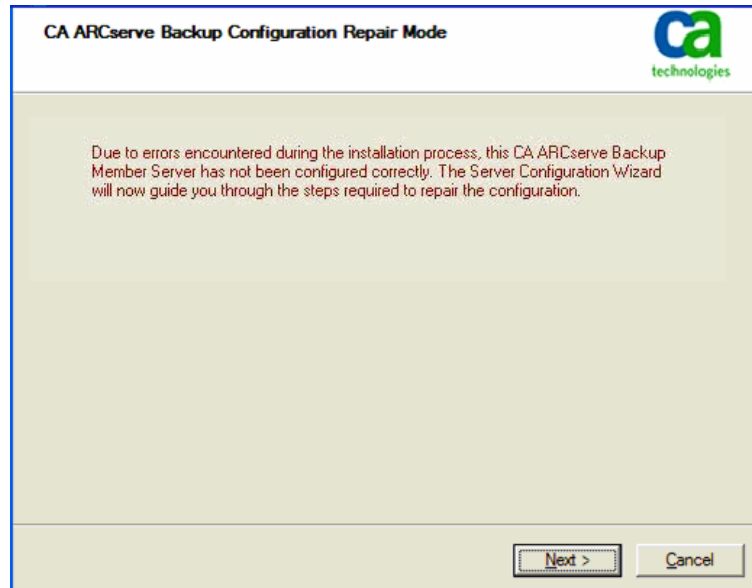
To repair the Arcserve Backup configuration

1. Install Arcserve Backup or upgrade Arcserve Backup from a previous release.

If the installation wizard detects installation errors with a primary server, the Arcserve Backup Configuration Repair Mode dialog appears as illustrated by the following screen:



If the installation wizard detects installation errors with a member server, the Arcserve Backup Configuration Repair Mode dialog appears as illustrated by the following screen:



2. Click Next.

The Server Configuration Wizard starts in repair mode.

Follow the prompts and complete the required fields on the subsequent dialogs to repair the Arcserve Backup configuration.

Repair the Arcserve Database Connection on a Primary Server

This task lets you repair Open Database Connectivity (ODBC) communication between a primary server and an Arcserve database instance that is hosted with Microsoft SQL Server, and register member servers with the primary server.

The Repair database connection option is disabled on stand-alone server installations or when you are hosting the Arcserve database using Microsoft SQL Server 2008 Express Edition.

ODBC is the most efficient method for the Database Engine to communicate with a Microsoft SQL Server instance that communicates through a network. Occasionally, network communication problems, Microsoft SQL Server communication settings problems, or both, can cause the Database Engine to communicate with the Arcserve database instance using Remote Procedure Call (RPC) communication. As a result, RPC communication will adversely affect the performance of the Arcserve database.

To remedy this problem, troubleshoot and repair the communication using the SQL Server Configuration Manager and then use the Server Configuration Wizard to repair ODBC communication between the Database engine and the Arcserve database instance.

To repair the Arcserve database connection on a primary server

1. Log in to the primary or stand-alone server where Arcserve Backup is installed.

Note: Do not open the Manager Console.

2. From the Windows Start menu, click Start, select All Programs, CA, ARCserve Backup, and click Server Configuration Wizard.

The Server Configuration Wizard opens.

3. Select the Repair database connection for member server(s) option and click Next.
4. Follow the prompts and complete all required fields on the subsequent dialogs to repair the database connection.

Note: When you repair the database connection on a primary server that is managing member servers, the Server Configuration Wizard attempts to repair the database connection on all member servers in the Arcserve domain.

Repair the Arcserve Database Connection on a Member Server

This task lets you repair Open Database Connectivity (ODBC) communication between a member server and an Arcserve database instance that is hosted with Microsoft SQL Server.

ODBC is the most efficient method for the Database Engine to communicate with a Microsoft SQL Server instance that communicates through a network. Occasionally, network communication problems, Microsoft SQL Server communication settings problems, or both, can cause the Database Engine to communicate with the Arcserve database instance using Remote Procedure Call (RPC) communication. As a result, RPC communication will adversely affect the performance of the Arcserve database.

To remedy this problem, troubleshoot and repair the communication using the SQL Server Configuration Manager and then use the Server Configuration Wizard to repair ODBC communication between the Database engine and the Arcserve database instance.

To repair the Arcserve database connection on a member server

1. Log in to the member server where Arcserve Backup is installed.

Note: Do not open the Manager Console.

2. From the Windows Start menu, click Start, select All Programs, CA, ARCserve Backup, and click Server Configuration Wizard.

The Server Configuration Wizard opens.

3. Select the Repair database connection option and click Next.
4. Follow the prompts and complete all required fields on the subsequent dialogs to repair the database connection.

Register Virtual Clusters as Arcserve Backup Domain Member Servers

Arcserve Backup lets you register virtual clusters with the Arcserve Backup primary server as domain member servers. You can register the virtual clusters as domain member servers when your backup and restore operations contain the following requirements:

- You want to protect applications that support virtual, cluster-aware configurations in Microsoft Cluster Server (MSCS) environments. For example, Microsoft SQL Server.
- You want to protect the application data (submit backup and restore jobs) using the virtual name of the service or application that is registered with MSCS.
- You want to add the virtual name of the service or application to the Backup Manager Source tab and Destination tab.
- You want to submit backup and restore jobs in a storage area network (SAN) instead of a local area network (LAN).

Note: To submit backups of this type, the physical nodes in the MSCS cluster must be connected to the SAN and a shared tape library.

To complete this task, the virtual clusters and the member server must reside on the same physical machine. In addition, the virtual clusters and the member server must be in an online state while you complete this task.

You can perform this task in cluster environments where MSCS is installed on any of the following Windows operating systems:

- Windows Server 2003
- Windows Server 2008
- Windows Server 2008 R2
- Windows Server 2012

Before you complete this task, verify that the following prerequisite tasks are complete:

- MSCS is installed on the nodes that you want to register with the primary server.
- The Arcserve Backup member server component and the application agent are installed on all of the nodes in the cluster.
- The physical nodes contained in the virtual cluster that you want to register with the primary server are connected to the SAN and a shared library.

To register virtual clusters as Arcserve Backup domain member servers

1. Log in to the Arcserve Backup member server.
2. From the Windows Start menu, select All Programs, CA, ARCserve Backup, and click Server Configuration Wizard.

The Server Configuration Wizard opens to the Select Options dialog.

3. Click Register virtual clusters as Arcserve Backup domain member servers and then click Next.

The Check caroot dialog opens.

4. Specify the carrot password and click Next.

The Check Cluster Nodes dialog opens.

5. Click the check boxes next to the Virtual Names that you want to register with the primary server and click Next.

6. When the registration completes successfully, the Complete dialog opens.

Click Finish to close the Server Configuration Wizard.

Note: If the registration process was unable to register one or more virtual clusters, the wizard provides you with a list of nodes that were not registered with the primary server.

7. Open the Backup Manager and click the Destination tab.

Note: If the Backup Manager was open while you completed this task, close and reopen the Arcserve Backup Manager Console.

The Servers directory tree appears.

8. Right-click the node that you registered with the primary server and click Refresh on the pop-up menu.

The virtual node displays the correct status.

Note: As a best practice, you should verify that you have a sufficient quantity of Arcserve Backup server component licenses to support backing up data to member servers. To submit jobs to member servers (virtual clusters), you must apply Arcserve Backup server component licenses to the member servers. The license count will be consumed based on the quantity of virtual clusters that you registered as member servers in the Arcserve Backup domain.

How Arcserve Backup Protects Active Directory Data on Domain Controller Servers

Active Directory is a hierarchical database that is stored on Domain Controller servers. The Active Directory includes static information about computer users, groups, printers, computer network configuration data, and so on.

Arcserve Backup lets you back up and restore the entire Active Directory on Windows Server 2003, Windows Server 2008, and Windows Server 2012 systems. On Windows Server 2003, Windows Server 2008, and Windows Server 2012 systems, you can restore the Active Directory at object level granularity.

You can restore the Active Directory files (*.dit and log files) to any domain controller server that meets the following conditions:

- The Client Agent for Windows is installed on the domain controller server.
- The domain controller server resides in the same domain as the server from which the backup was taken.
- The operating system running on the domain controller server is the same version, release, and service pack as the server from which the backup was taken.

Arcserve Backup lets you protect the Active Directory using the following approaches:

- **Restore the System State to its original location**--Arcserve Backup lets you restore the System State, which includes all objects in the Active Directory, to the server from which it was backed up. With this approach, you overwrite all of the objects contained in the Active Directory.

Use this approach when you need to restore the entire Active Directory to a previous point in time.

- **Restore the Active Directory to an alternate location**--Arcserve Backup lets you restore the Active Directory to an alternate location. This approach is a two-phase process that lets you restore Active Directory data at object level granularity on Windows Server 2003, Windows Server 2008, and Windows Server 2012 systems. With this approach, you restore the Active Directory to an alternate location using the Restore Manager, and then restore Active Directory objects using the CA Active Directory Object Level Restore utility.

The alternate location can reside on a server that does not function as a Domain Controller server. However, the best practice is to restore the Active Directory to an alternate location on the server from which it was backed up.

Note: Arcserve Backup cannot restore Active Directory objects at object level granularity on Windows Server 2008 and Windows Server 2012 systems that function as read-only domain controllers.

Example: When to Restore the Active Directory at Object Level Granularity

- A system administrator deleted a group of users, groups, or an object from the Active Directory in error.

Note: To protect Active Directory data, Arcserve Backup Client Agent for Windows must be licensed on the Domain Controller server.

Arcserve Backup lets you restore Active Directory data that was backed up using the following Arcserve Backup releases:

- Arcserve Backup r12. Includes the general availability release and all of the latest service packs.
- Arcserve Backup r12.5. Includes the general availability release and all of the latest service packs.
- This release of Arcserve Backup.

The CA Active Directory Object Level Restore utility lets you restore the following Active Directory objects:

- Organizational Unit
- User
- Group
- Computer
- Contact
- Connection
- Shared folder
- Printer
- Site
- Site container
- Site link
- Site link bridge
- Site settings
- Subnet container
- Trusted domain

- Configuration class
- Lostandfound class
- Builtindomain class
- Dnszone class
- Domain class
- Domaindns class
- Dmd class
- Organizationalunit class
- Containerrecifiers class

The CA Active Directory Object Level Restore utility cannot restore the following Active Directory objects

- System Schema
- Global Policy Object (GPO)

This section contains the following topics:

[Back up the Active Directory](#) (see page 582)

[Restore Active Directory Objects](#) (see page 583)


[Microsoft Exchange Server 2010 Mailboxes Appear as Legacy Mailboxes After Recovering the Active Directory](#) (see page 589)

[Reset Microsoft Exchange Server User Passwords After Recovering the Active Directory](#) (see page 589)

Back up the Active Directory

There are several approaches that you can use to back up a computer's System State.

- Create a backup job as you would create any other backup job and include the System State object for the computer with the source selections for the job.
- Create a backup job that includes only the computer's System State.



Note: The following steps describe how to submit a normal backup job. For information about submitting staging and deduplication backup jobs, see "Backing up Data."

To back up the Active Directory

1. Open the Backup Manager Window and click the Start tab.

The Backup job types display.

2. Click Normal backup to specify a normal backup job.

Note: For more information about the types of backup jobs, see the "Backing up Data."

Click the Source tab.

The backup source directory tree displays.

3. Browse to the computer that you want to back up.
Expand the volumes contained by the server and display the System State object.
Click the check box next to System State.
Note: The Backup Manager prevents you from selecting only the Active Directory.
4. Click the Schedule tab to define when and how frequently you want to back up the System State.
Note: For information about scheduling jobs, see "Customizing Jobs."
5. Click the Destination tab.
The available Device Groups display in the directory tree.
6. Select the Device Group where you want to store the backup data.
7. Click Options on the toolbar to define backup options for the job.
Note: For information about backup options, see "Backing up Data."
8. Click Submit on the toolbar to submit the job.
The job is submitted.

Restore Active Directory Objects

Arcserve Backup lets you restore the Active Directory at object level granularity. However, before you can restore Active Directory objects, you must back up the Active Directory as part of the computer's System State.

The process of recovering Active Directory objects consists of two phases:

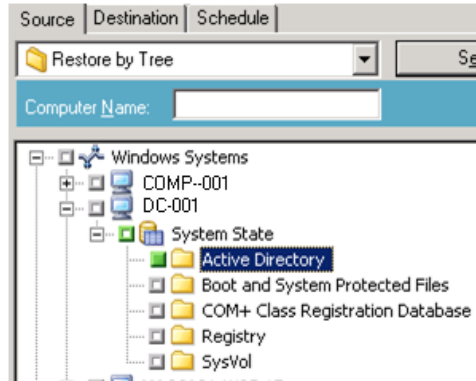
- Restore the Active Directory backup data to an alternate location using the Restore Manager. The alternate location should reside on the server where the System State was backed up.
- Recover the Active Directory object to the current Active Directory using the CA Active Directory Object Level Restore utility.

Restoring the Active Directory to its original location restores all of the objects contained in the Active Directory. The process for restoring the Active Directory to its original location is the same that of restoring files, directories, and so on. For more information, see "Restoring Data."

To restore Active Directory objects

1. Open the Restore Manager window, click Source, and expand the server and System State containing the Active Directory that you want to restore.

Click the check box next to Active Directory as illustrated by the following screen:



2. Click the Destination tab.

Clear the checkmark from Restore files to their original location(s).

In the location field, specify a path to an alternative location.

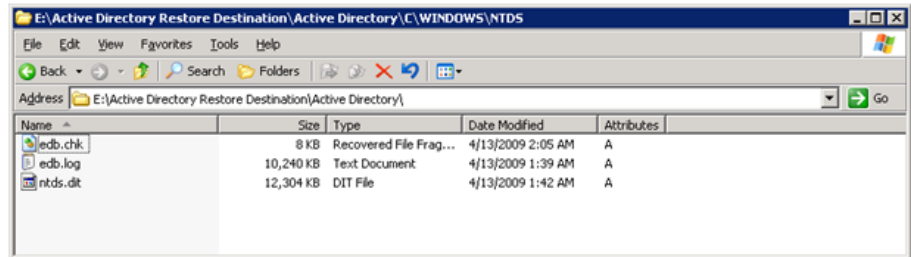
Example:

\\172.31.255.254\E:\Active Directory Restore Destination

Click Submit on the toolbar to submit the job.

Complete the required fields on the Submit dialog and click OK.

After the restore is complete, the recovered data appears in the alternate location.



3. (Optional) To restore the Active Directory to an alternative server, copy the files from the restore destination that you specified in the previous step to the alternative server.

Note: The best practice is to specify an alternative directory on the server where you are restoring the Active Directory. However, if you must restore the Active Directory to an alternative server, you can copy the restored Active Directory files from the alternative server to the source server or any other domain controller that is in the same domain as the source server. The restrictions on this capability are as follows:

- The Client Agent for Windows must be installed on the alternative server.
 - The operating system running on the alternative server must be the same version, release, and service pack as that of the server from which the backup was taken.
4. Log in to the Domain Controller server containing the restored, or copied, Active Directory data.

Open Arcserve Backup Agent Admin by doing the following:

From the Windows Start menu, select All Programs, CA, ARCserve Backup, and click Arcserve Backup Agent Admin.

Arcserve Backup Agent Admin opens.

Note: To open the Arcserve Backup Agent Admin, you must be logged in to the server using an account that has Domain Administrative privileges.

5. From the Options menu on the Arcserve Backup Agent Admin dialog, click AD Object Level Restore Utility.

The CA Active Directory Object Level Restore dialog opens.

6. Click Open on the CA Active Directory Object Level Restore utility.

The Open Active Directory Files dialog opens. The Active Directory backup restore points appear in the restore points list.


(Optional) Click the ellipsis to browse for more restore points.

Select a restore point and click OK.

Arcserve Backup populates the CA Active Directory Object Level Restore dialog with the Active Directory objects as illustrated by the following window.

- **Left pane**--Lets you view all objects included in the selected Active Directory ntds.dit database file.
- **Right pane**--Lets you view the attributes and child objects related to the item selected in the left pane.

Note: Arcserve Backup lets you restore only the Active Directory Objects that appear in the CA Active Directory Object Level Restore dialog. Arcserve Backup cannot restore system related objects.



7. (Optional) Click Option on the toolbar to open the Restore Options dialog.

Arcserve Backup lets you filter the active directory objects that you want to restore and specify a log level for the restore operation.

- a. Specify the restore options that you require:

- **Restore Renamed Objects**--Lets you restore objects that were renamed in the current Active Directory.
- **Restore Moved Objects**--Lets you restore objects that were moved to a different location in the current Active Directory.
- **Restore Permanently Deleted Objects**--Lets you restore objects that were deleted permanently from the current Active Directory.

- b. Specify the log level options that you require:

- **Log Level**--Lets you specify the level of details that you require in the debug log.

Default value: 0

Range: 0 to 3

0--Prints error messages to the log files.

1--Prints error and warning messages to the log files.

2--Prints error, warning, and information messages to the log files.

3--Prints error, warning, information, and debug messages to the log files.

Note: Level 3 is the highest level of logging details. If you encounter problems restoring active directory files, you should specify level 3 and then send the log files to Arcserve Support.

The Active Directory restore process generates the following log files:

- adrestorew.log
- CadRestore.exe.trc

Arcserve Backup stores the log files in the following directory:

C:\Program Files\CA\ARCserve Backup Client Agent for Windows

Click OK to close the Restore Options dialog.

8. (Optional) Click Filter on the toolbar to open the Filter Settings dialog.

Note: The best practice is to use filters when you are searching for a specific object.

Specify one of the following filter settings:

- **Show all types of objects**--Lets you display all objects in the Active Directory Object Restore dialog.

(Optional) To limit the number of child nodes, click Maximum number nodes under each parent nodes and specify a limit in the text box.

- **Show only the following types of objects**--Lets you display only objects of a particular type in the Active Directory Object Restore dialog.
- **Show only the following named objects**--Lets you display only objects with a particular name in the Active Directory Object Restore dialog.

Note: The Active Directory for a computer can contain a large number of objects. The best practice is to filter the objects using the Show only the following named objects filter and specify the name of the object that you want to restore.

Click OK to close the Filter Settings dialog.

9. From the Active Directory Object Restore dialog, expand the Active Directory tree and click the check box next to the objects that you want to restore.

Click Restore on the toolbar to restore the specified objects.

Arcserve Backup restores the Active Directory objects to the current Active Directory.

After the restore is complete the Restore Status message box opens.

Note: The Restore Status messages box describes the outcome of the job.

10. Click OK to close the message box.

The Active Directory objects are restored.

11. (Optional) Click Report on the toolbar to check the status of the restore.

Note: You may want to view the Job Report if Arcserve Backup reports that it could not restore the objects.

Microsoft Exchange Server 2010 Mailboxes Appear as Legacy Mailboxes After Recovering the Active Directory

Valid on Windows platforms running Microsoft Exchange Server 2010.

Symptom

The Recipient Type attribute for Microsoft Exchange Server 2010 user accounts appears as Legacy Mailbox instead of User Mailbox after you recover the Active Directory using CA Active Directory Object Level Restore Utility. In addition, the recovered user accounts are disabled after the recovery is complete.

Solution

This behavior occurs because Microsoft Exchange Server 2010 user accounts contain attributes relating to the Recipient Type that the CA Active Directory Object Level Restore utility cannot recover. As a result the Recipient Type appears as Legacy Mailbox instead of User Mailbox.

To remedy this problem, complete the following steps:

1. Recover the Active Directory using the [CA Active Directory Object Level Restore utility](#) (see page 583).
2. Log in to the Microsoft Exchange Server 2010 system.
3. Open Windows PowerShell.
4. Execute the following command:

```
Set-Mailbox -id [username or mailbox alias] -ApplyMandatoryProperties
```

Reset Microsoft Exchange Server User Passwords After Recovering the Active Directory

Arcserve Backup lets you recover the Active Directory on domain controller servers. Although the Active Directory contains data that relates to Windows user accounts, Windows does not store the passwords for user accounts in the Active Directory. As such, the Active Directory recovery process does not let you restore user passwords. Use the following guidelines for resetting user passwords after you recover the Active Directory:

- If the user account was present in the Active Directory before recovered the Active Directory, you do not need to reset the password for the user account.
- If the user account was not present in the Active Directory (for example, deleted), you must reset the password for the user account.

Install and Uninstall Arcserve Backup Server Based Options

From a primary and stand-alone Arcserve Backup server, you can use the Server Admin to install and uninstall the following Arcserve Backup options:

- Arcserve Backup Central Management Option
- Arcserve Backup Tape Library Option
- Arcserve Backup Storage Area Network (SAN) Option

Before you install and uninstall Arcserve Backup server based options, the following considerations apply:

- You can install and uninstall options only on a primary or stand-alone Arcserve Backup server.
- The Arcserve Backup options that display in the Install/Uninstall Options dialog will vary depending on the type of Arcserve Backup server you are configuring.
- If you are installing server based options, ensure that all external devices (for example, libraries) are connected to the primary servers, member servers, and the SAN in your environment. Arcserve Backup automatically detects supported devices and configures them for use automatically when the tape engine starts.

You must manually configure devices that Arcserve Backup does not automatically detect.


To install and uninstall Arcserve Backup server based options

1. From the Quick Start menu in the Navigation Bar on the Home Page, click Server Admin.

The Server Admin opens.

2. Expand the domain directory tree and click the primary or stand-alone server where you want to install or uninstall options.

The domain directory tree is illustrated by the following:



3. Right-click the server where you want to install and uninstall options and select Install/Uninstall Options from the pop-up menu.

The Install/Uninstall Options dialog opens.

4. From the Product Name list on the Install/Uninstall Options dialog, place a check mark next to the options that you want to install and clear the check mark next to the options that you want to uninstall.
5. Click OK and follow the on-screen instructions to complete the installation, uninstallation, or both.

Discovery Configuration

Discovery Configuration is a service that you can use to periodically discover computers in your network for newly added or upgraded Arcserve Backup software. A Discovery server runs as a background process that collects information from all other Discovery servers installed with Arcserve products across the corporate network.

Discovery Configuration allows you to distribute discovered network target information to remote servers. This capability allows administrators to decrease network traffic load created by Discovery servers to discover Windows domains or IP subnet addresses.

The Discovery Configuration allows you to perform the following tasks:

- Start or stop the discovery service
- Distribute tables with discovered network targets
- Add, remove or modify information in any of the three tables created by Discovery Configuration (IP subnets, IP subnet masks, and Windows domains)
- Set or modify Discovery Configuration parameters

You can open the Discovery Configuration at the command line or from Windows Explorer:

- **Command line**--Start dsconfig.exe from the following directory:

C:\Program Files\CA\SharedComponents\ARCserve Backup\CADS

- **Windows Explorer**--Double-click dsconfig.exe located in the Arcserve Backup Shared Components directory. For example:

C:\Program Files\CA\SharedComponents\ARCserve Backup\CADS

How the Discovery Service Detects Other Computers

A Discovery server is implemented as a Windows service. As soon as a Discovery server starts, it enumerates the list of products to create a behavior mask. Based on this mask, the Discovery server initializes the following required components:

- The Listen/Serialize component is initialized first and starts listening on a particular port (or Mailslot) for incoming packets (messages) from other Discovery servers. When a message is received, the Discovery server writes data (if any) into the repository (registry, for example) and then notifies the Query/Advertise component.
- The Query/Advertise component sends the message received from the Listen/Serialize component, (prepared with the product's list) directly to the Discovery server, which requested the data. The Query/Advertise component can also send messages as a broadcast message to the selected network targets (the list of IP Subnets or Windows Domains) if it is scheduled or initiated manually. It does this in order to query other Discovery servers across the network, and advertise its own list of Arcserve Backup products.

Discovery Service Configuration Options

You can specify the transport protocol used to broadcast queries by choosing configuration options that meet your needs. Select the Configuration button in the Discovery Configuration.

In the Network tab, you can choose the protocols to discover and define the TCP/IP subnet sweep. Choose to enable discovery of Arcserve Backup products if you want the discovery service to broadcast queries repeatedly with a specified interval.

Choose to enable network discovery if you want the Discovery server to run a process of discovering new Windows domains and IP subnets. By default, the Discovery server runs this process only when the discovery service is restarted. You can also modify the interval, depending on how dynamic your network environment is.

Note: It is not recommended that this option be run at all times because it continually broadcasts queries which can increase network traffic.

You can direct the Discovery Configuration application to clean its tables as the Discovery service starts. Discovery tables store information about computers with a Arcserve Backup product installed. When you enable this option, the Discovery service purges the data from its tables, discovers computers with a Arcserve Backup product installed, and then updates the tables with current, accurate data.

Important! The Clean Up Discovery Table on Startup option is enabled by default. If you disable this option and uninstall Arcserve Backup applications from systems from your environment, the details about these computers will remain in the tables when the Discovery service restarts.

Discovery Configuration Dialog

You can use the Discovery Configuration dialog to perform the following tasks:

- Specify the transport protocol used to broadcast queries.
- Enable Discovery of Arcserve Backup Products allows the Discovery Service (DS) to broadcast queries repeatedly, with a specified interval.

It is not recommended that this be run at all times because it continually broadcasts queries which could increase network traffic.

- Enable Network Discovery allows the DS to run a process of discovering new Windows domains and IP Subnets.

By default, the DS runs this process only when the Discovery Service is restarted. You can also modify the interval, depending on how dynamic your network environment is.

Network Tab

Use the Network tab to configure the following options:

- Specify the Protocols used to Discover.
- Specify the Protocols used to Reply.
- Enable Discovery of Arcserve Backup Products.
- Enable Discovery using TCP/IP Subnet Sweep.
- Enable Network Discovery and specify the time interval.
- Direct the Discovery Service to clean the discovery tables on startup .

Adapters Tab

- Select IP address to Run Discovery.

IP Subnets/Windows Domains Discovery

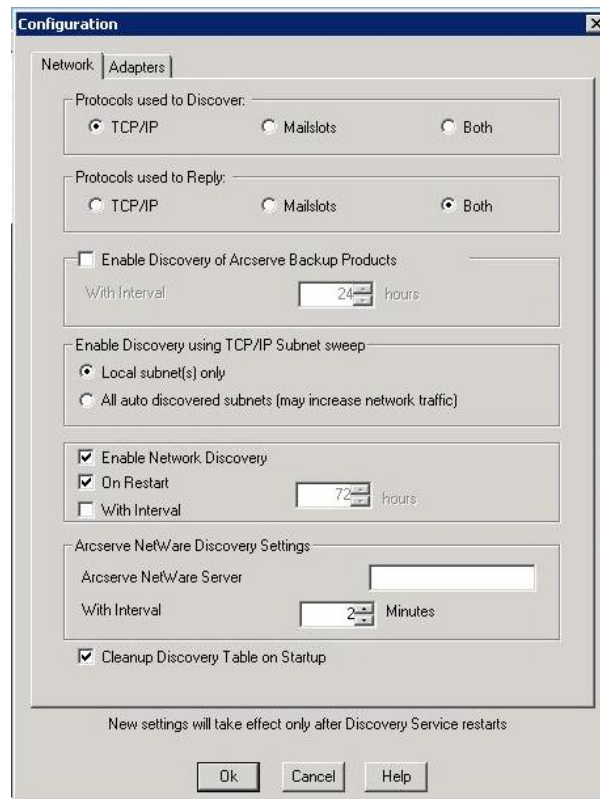
In order to query and advertise, the Discovery Configuration needs the list of network targets to broadcast. For example, you might want to discover Arcserve Backup products in only one IP subnet on a TCP/IP enabled network or in a list of subnets located physically in the same country.

In another scenario, you might want to enumerate the entire corporate network to find all possible subnets and then filter some of them out. The discovery process runs in the background, and enumerates Windows network's resources. A list of Windows domains for subsequent Mailslot broadcasting or a listing of IP subnets for UDP broadcasting is created.

Note: This process may take some time, depending on the size of your network. It is recommended that the IP subnets and domains discovery be run during a time of minimum network traffic.

Enable Discovery Using TCP/IP Subnet Sweep

You can configure to use a local or remote subnet sweep. The default setting is for a local subnet.



The discovery service broadcasts and retrieves all the information of the local subnet machines as well as manually-defined subnets and manually-defined machines.

To enable add a machine name (IP address) manually

1. Start Discovery and click the Add button on the Windows Domain tab.
2. Enter the IP address of the machine and click OK.

The discovery service can then ping, publish, and return a product list from the specified machine.

Note: You can limit the discovery range by disabling any auto-discovered machine or Windows domain, which can reduce network traffic. Check the Disable box for an existing machine or Windows domain.

To add a remote subnet

1. Start Discovery and click the Add button on the Subnet tab.


The Add a Subnet dialog opens.

2. Enter the Subnet and Subnet Mask and click OK.

This enables the discovery service publish its product list to each machine to the specific subnet and also return product information for every machine in this subnet.

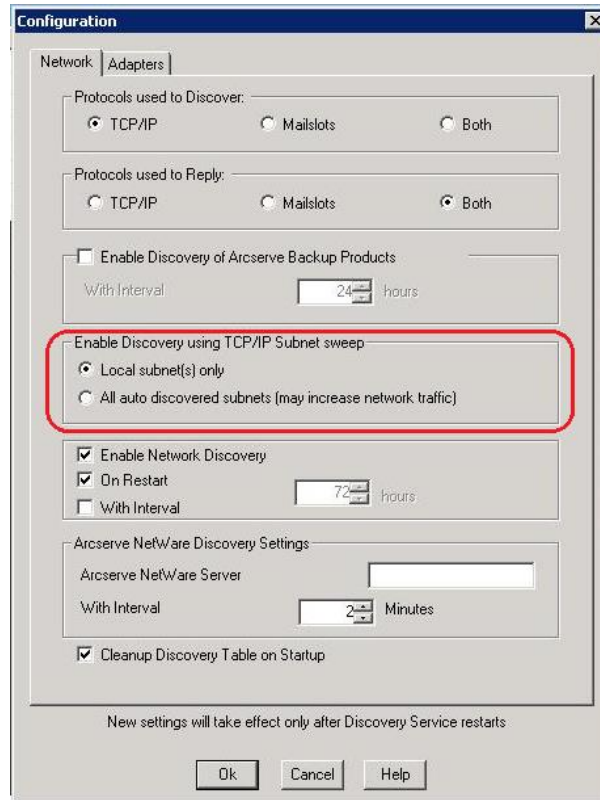
Note: You can limit the discovery range by disabling any auto-discovered machine or Windows domain, which can reduce network traffic. Check the Disable box for an existing machine or Windows domain.

If specific subnets use a different subnet mask you can add a subnet mask manually by starting Discovery and clicking the Add button on the Subnet Mask tab. The Add a Subnet Mask dialog opens as shown in the following example:



If you choose to perform an auto discovery, the discovery service will ping and publish a product list to each machine in each subnet listed in the Discovery Configuration Subnet tab and retrieve the product information from the remote machine.

3. To do this, start Discovery and click the Configure button on the Summary tab. The Configuration dialog opens as shown in the following example:



The Discovery server initiates an IP address sweep for remote subnets by using the auto-discovered subnets, subnet masks, machine, or Windows domains along with the manually-configured subnets, subnet masks, machines or Windows domains. Selecting this option may increase network traffic and can take a considerable amount of time to complete, depending on the size of your network. We recommend that you run this option during a time of minimum network traffic.

Note: If you choose to discover Arcserve Backup products in remote subnets, the discovery service does not rely on a UDP broadcast to locate remote instances. You need to know the size of the subnet and range of IP addresses using a subnet mask.

Discovery Configuration for the SAN Option

The Discovery service configuration for servers in a SAN environment require additional modifications in order for all servers to be discovered. To ensure that all SAN servers are enabled for discovery, one of the following options are available:

- Select the "All auto discovered subnets (may increase network traffic)" option in the Enable Discovery using TCP/IP Subnet sweep field in the Configuration dialog.
- Add other remote SAN machine names/IP addresses by accessing the Windows Domain tab.

Note: The configuration for each remote SAN server must be consistent to ensure accurate discovery. If you only configure one SAN server, other servers may still fail in discovery.

Discover Client Agent Systems with Non-default IP Addresses

If you select a non-default IP address for a client agent node, you must perform the following procedure to update the IP address to the new address, to allow you to add machines using Add/Import/Export Nodes.

To discover client agent systems with non-default IP addresses

1. On the agent machine, restart the discovery service using dsconfig.exe.
2. From the Backup Manager of the base server, delete the old machine object.
3. Click Add/Import/Export Nodes and add or import the nodes that you require.



Note: For more information, see [Add, Import, and Export Nodes Using the User Interface](#) (see page 334).

4. If you have a scheduled job in the queue, delete that job and recreate it to ensure that it runs properly.

Arcserve Backup Maintenance Notifications

Arcserve Backup displays a maintenance notification message near the Windows system tray to help ensure that you are running the most current version of Arcserve Backup. The maintenance notification message appears nine months after you install or update Arcserve Backup.

The maintenance notification message appears near the Windows system tray area of your computer's desktop, as illustrated by the following screens.

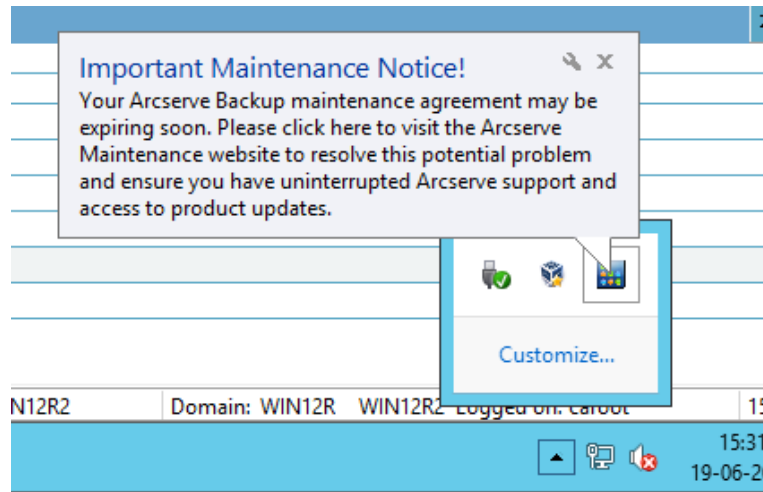
- Windows Server 2003 systems:



- Windows Server 2008 systems:



- Windows Server 2012 systems:



Maintenance Options

- **Update your Arcserve Backup maintenance contract**--To update Arcserve Backup, click the maintenance notification message.

After you click the maintenance notification message, the Arcserve Backup, Arcserve Maintenance Program web site opens. The Arcserve Maintenance Program web site lets you obtain product enhancements and updates, and install new releases of Arcserve Backup.

Nine months after you install Arcserve Backup, the maintenance notification message and the Arcserve tray icon open and close in 30-day intervals until you update your Arcserve Backup maintenance contract.

Note: Twelve months after you install Arcserve Backup, the maintenance messages open and close in six month intervals.

- **Do not update your Arcserve Backup maintenance contract**--If you do not wish to update Arcserve Backup at this time, click the X in the upper left-hand corner of the maintenance notification message to close the message. If you do not want to receive the messages, you must [disable](#) (see page 600) maintenance notification messages.

Alert Options

Arcserve Backup lets you [disable](#) (see page 600) and [enable](#) (see page 600) the maintenance notification message using the Arcserve Backup icon located in the Windows system tray.

Disable Maintenance Notification Messages

Arcserve Backup lets you disable the maintenance notification message that appears near the Window systems tray.

To disable the maintenance notification message

1. From the Windows system tray, right-click the Arcserve Backup icon. In the pop-up menu, click **Disable Alert**.



The maintenance notification message is disabled.

Enable Maintenance Notification Messages

Arcserve Backup lets you enable the maintenance notification message that appears near the Window systems tray, if it is currently disabled.

To enable maintenance notification messages

1. Open Windows Registry Editor and browse to the following registry key:

- **x86 platforms:**

HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA ARCserve Backup\Base\Admin\MessageEngine\MonthsPassedInLastMaintenance

- **x64 platforms:**

HKEY_LOCAL_MACHINE\SOFTWAREWow6432Node\ComputerAssociates\CA ARCserve Backup\Base\Admin\MessageEngine\MonthsPassedInLastMaintenance

Right-click MonthsPassedInLastMaintenance and click **Modify** on the pop-up menu.

The Edit DWORD Value dialog opens.

2. In the Value Data field, specify 0 (zero).

Click **OK**.

Close Windows Registry Editor.

Maintenance notification messages are enabled. The Arcserve Backup maintenance notification icon appears in the Windows system tray when maintenance is recommended.



Apply Arcserve Backup Component Licenses

Arcserve Backup lets you apply component licenses after you install Arcserve Backup. You would apply component licenses under the following scenarios:

- You installed Arcserve Backup components using trial licenses and you want to apply license keys to the components.
- You obtained additional licenses to support the growth of your backup environment.

To apply Arcserve Backup component licenses

1. Open the Arcserve Backup Manager Console.

From the Help menu, click About Arcserve Backup.

The About Arcserve Backup screen opens.

2. From the About Arcserve Backup screen, click Add/View Licenses.

The License Verification dialog opens.

Important! You must always apply license keys using the License Verification dialog.

The list box on the License Verification dialog displays the Component Name and the Install Type. The Install Type is Unlicensed Trial or Licensed.

3. From the list of components, do the following:

- a. Click a component.
- b. Click Use License Key.
- c. Enter the license key in the fields provided.

Click Add.

Note: If you attempt to apply an upgrade license key to a new Arcserve Backup component installation, the Upgrade Verification dialog opens. You must provide the license key for your previous installation to apply your upgrade license key.

(Optional) Repeat this step to apply license keys to other Arcserve Backup components.

4. When you are finished, click Continue.

The License Verification dialog closes and the license keys are applied to the Arcserve Backup components.

Managing Firewalls

This section contains the following topics:

[Allow Arcserve Backup Services and Applications to Communicate Through the Windows Firewall](#) (see page 602)

[How to Configure Your Firewall to Optimize Communication](#) (see page 602)

Allow Arcserve Backup Services and Applications to Communicate Through the Windows Firewall

During the installation or upgrade process, the installation wizard configures your Windows firewall such that Arcserve Backup services and applications can communicate properly. The installation wizard performs the configuration task only if the Windows firewall was in the On state when you installed Arcserve Backup.

If the Windows firewall was in the Off state when you installed Arcserve Backup, and then turned on the Windows firewall at any time after you installed Arcserve Backup, the Arcserve services and applications will not be able to communicate through the Windows firewall.

The following procedure lets you allow Arcserve Backup services and applications to communicate if the Windows firewall was in the Off state when you installed Arcserve Backup.

To allow Arcserve Backup services and applications to communicate through the Windows firewall

1. Open the Windows Command Line, and change to the following directory:

```
c:\Program Files\CA\SharedComponents\ARCserve Backup\
```

2. Execute the following command:

```
setupfw.exe /INSTALL
```

Arcserve Backup services and applications are added to the Windows firewall exception list. Arcserve Backup services and applications can now communicate through the Windows firewall.

How to Configure Your Firewall to Optimize Communication

For information about configuring firewalls to optimize Arcserve Backup communication, see the *Implementation Guide*.

Chapter 9: Managing the Database and Reporting

This section contains the following topics:

[How to Manage the Database and Reports](#) (see page 603)

[Database Manager](#) (see page 604)

[How to Protect the Arcserve Backup Database](#) (see page 607)

[How the Catalog Database Works](#) (see page 650)

[Using Microsoft SQL Server as the Arcserve Backup Database](#) (see page 658)

[Specify a Arcserve Backup Database Application](#) (see page 662)

[Arcserve Backup Logs and Reports](#) (see page 667)

[Arcserve Backup Diagnostic Utility](#) (see page 682)

[Arcserve Backup Infrastructure Visualization](#) (see page 687)

How to Manage the Database and Reports

The Arcserve Backup database maintains job, media, and device information on your system. Arcserve Backup stores the following types of information in the database:

- Detailed information about all jobs.
- Session details for all backup jobs.
- Information about the media used for all backup jobs.
- Detailed information about each file and directory that was backed up to media when you perform a restore.

When you want to restore a specific file, the database determines which media a file is stored on.

- Detailed information about media pools and media location.

The database information is also used to generate many types of reports.

Database Manager

The Database Manager lets you perform the following tasks:

- Track of the location of your media.
- Determine the session number of your backup.
- Determine if media should be retired.
- View log information about jobs you have run.
- Delete old records from the database.
- Visually compare the size of your database to the total available disk space.

Note: For Microsoft SQL Server databases, the total database size reported by the Arcserve Backup Database Manager is the size of the data device. You can obtain more information by browsing through the Microsoft SQL Server Enterprise Manager.

Database Views

When you open the Database Manager, the left panel displays the following options:

- **Summary**--Space the database used on your hard disk, database type, and other settings.
- **Job Records**--Jobs processed by Arcserve Backup.
- **Media Records**--Media used by Arcserve Backup.
- **Device Records**--Devices used by Arcserve Backup.

Sort Order

To change the sort order of records displayed in the Job, Media, and Device Records view, click on the field name you want to sort.

Database Pruning

You can set Arcserve Backup to remove old records from the database. For more information, see "Administering the Backup Server."

When You Should Rebuild the SQL Indexes

Note: This section only applies when using Microsoft SQL as the Arcserve Backup database.

You should rebuild the SQL Server index periodically to keep the index size manageable and at optimal performance. The best practice is to rebuild the index once or twice per month, or when the Arcserve database performs slowly.

The process of updating the SQL Server indexes can take a lot of time. If you do not have enough time to update all the indexes, you should update the key indexes: IX_astpdat_1, IX_astpdat_2, X_astpdat_3, K_pathname, and PK_filename. These indexes play an important role and affect the browsing speed in the Restore Manager and Database Manager.

For information about how to rebuild the SQL Server indexes, see the Microsoft SQL Server documentation.

Types of Errors Reported

The following statistical information is recorded in the database:

- **Media Errors**--Indicates data corruption occurred on the media preventing the read or write operation from successfully completing.
- **Soft Read Errors**--An error occurred while reading the media. Arcserve Backup attempted to correct the problem in real-time. A higher number of soft read errors indicate possible defective media. Media should be replaced for any future backups.
- **Soft Write Errors**--A write error occurred during the backup. Arcserve Backup is correcting the media problem in real time. A high number of soft write errors indicates the media should be replaced for future backups. Make sure the drive heads are cleaned after the current backup session is completed.

Device Error Records

If a drive has a critical error, the error log may contain some of the following information:

- **Time**--Time the error occurred.
- **Sense Info**--SCSI error code.
- **Media**--Number of media errors that occurred during the job.
- **Soft Write**--Number of soft write errors

- **Soft Read**--Number of soft read errors that occurred during the job.
- **Media Usage**--Amount of time the media was used during the job.
- **KB written**--Amount of data written to the media during the job.
- **Times Formatted**--Number of times media has been formatted.

Last Arcserve Backup Database Backup Information

Every time the Arcserve Backup database is backed up successfully, the backup media information is retained in a series of log files labeled ASDBBackups.txt and ASDBBackups.n.txt. The log files are stored in the Arcserve Backup home directory.

Each of these log files contains the detail information about the complete backup history of the Arcserve Backup database. For example, the log files contain detail information about the following items:

Note: This is not an exhaustive list.

- Tape name, serial number, and sequence number
- Session GUID, session ID, session type
- Backup method (for example, full, incremental, differential)
- Backup date
- Type of database
- Path of the instance

Note: The log files do not contain descriptive text or comments.

Arcserve Backup references the log files when you need to restore the Arcserve Backup database because it is not in a usable state or corrupt. The Arcserve Backup Disaster Recovery Option and Arcserve Database Recovery Wizard reference ASDBBackups.txt to determine the media that contains the latest Arcserve Backup database backup data.

Arcserve Backup manages the log files using the following logic:

1. The information for the most recent full, incremental, and differential backup of the Arcserve Backup database is always stored in the file labeled ASDBBackups.txt.
2. After the second full backup of the Arcserve Backup database is complete, ASDBBackups.txt is renamed ASDBBackups.1.txt and then a new ASDBBackups.txt is created.
3. After the third full backup is complete, ASDBBackups.1.txt is renamed ASDBBackups.2.txt, ASDBBackups.txt is renamed ASDBBackups.1.txt and then a new ASDBBackups.txt is created.
4. The log file renaming and creating process continues until Arcserve Backup creates a log file named ASDBBackups.10.txt.
5. If a log file labeled ASDBBackups.10.txt exists after the full backup is complete, Arcserve Backup deletes ASDBBackups.10.txt, renames the older log files, and then creates a new ASDBBackups.txt log file.

If ASDBBackups.txt is corrupt, you can rename any of the files labeled ASDBBackups.n.txt to ASDBBackups.txt and use the renamed log file to recover the Arcserve Backup database.

Enable Media Pool Maintenance

Using the Enable Media Pool Maintenance option you can allow daily maintenance of the media pool. Arcserve Backup performs media pool maintenance tasks according to the Prune job schedule.

To enable media pool maintenance

1. Open the Server Admin Manager and click the Configuration toolbar button.
The Configuration dialog opens.
2. Select the Database Engine tab.
The Database Engine dialog appears displaying the media pool maintenance option at the bottom of the dialog.
3. Click Enable Media Pool Maintenance and click OK.
Media pool maintenance is enabled.

How to Protect the Arcserve Backup Database

The following sections describe how to back up and restore the Arcserve Backup database.


Agent for Arcserve Database

The Arcserve Backup Agent for Arcserve Database is a form of the Arcserve Backup Agent for Microsoft SQL Server. The agent is either installed automatically when you install Arcserve Backup, or manually using a special utility, either after the location of the Arcserve Backup database is changed, or on multiple nodes of a cluster.

This utility, named SQLAgentRmtInst.exe, is placed in the Packages sub-folder of the Arcserve Backup home directory, in a folder named ASDBSQLAgent, when you install Arcserve Backup. If you need to install the agent on a computer that is not a Arcserve Backup server, you must copy the ASDBSQLAgent folder to the system where you are installing the agent, and run the SQLAgentRmtInst.exe utility on that machine.

By itself, the Agent for Arcserve Database allows you to back up and restore the Arcserve Backup database, and the system databases and Disaster Recovery Elements from the Microsoft SQL Server instance that contains the Arcserve Backup database. When installed with the Agent for Microsoft SQL Server, it allows the Agent for Microsoft SQL Server to recognize the presence of a Arcserve Backup database, and to work with Arcserve Backup to provide the special recovery mechanisms that are available for the Arcserve Backup database.

When upgrading from a previous release of Arcserve Backup, you must upgrade the Agent for Arcserve Database. This behavior is designed to help ensure that the current version of the Arcserve Backup database is protected by the current version of the agent. As a result, you cannot clear the check box next to Agent for Microsoft SQL Server in the product selection tree on the Components dialog.



Because the Agent for Arcserve Database is a form of the Agent for Microsoft SQL Server, it will appear as the Arcserve Backup Agent for Microsoft SQL Server in the system's installed programs list. If both are present, only a single entry will appear. If you need to uninstall one or the other, the installation sequence will prompt you to select which variant to remove.

You can use the stand-alone utility that installs the Agent for Arcserve Database in any of the following situations:

- When the Arcserve Backup database is moved
- To re-install the agent if it is accidentally uninstalled
- To install the agent to additional nodes of a cluster
- To install the agent on a remote computer, if the Arcserve Backup installer is unable to do so directly

Configure Backup and Restore Parameters for the Agent for Microsoft SQL Server Using the Central Agent Admin

Use the Central Agent Admin utility to configure the Agent for Microsoft SQL Server backup and restore parameters for supported versions of Microsoft SQL Server. The parameters include settings for Microsoft Virtual Device Interface (VDI) objects and remote communication.

To configure backup and restore parameters for the Agent for Microsoft SQL Server using the Central Agent Admin

1. From the Arcserve Backup Quick Start menu, choose Administration, Central Agent Admin.
Central Agent Admin opens.
2. In the Windows Systems tree, expand the server on which the agent is installed and then select the Agent for Microsoft SQL Server.
Click Configuration on the toolbar.
The Options Configuration dialog opens.
3. Click Agent for Microsoft SQL Server from the list on the left.
The Options Configuration shows the corresponding SQL Server settings.
4. Specify the synchronized recording under Agent Log Settings as follows:
Synchronized Recording--Forces the log messages to be written to the Activity Log as they are posted. You can disable this option to improve the performance on high-load systems by caching several messages and writing them as a group.

5. Select the Instance (ARCSERVE_DB) or the name of the instance for which you wish to change configuration for the Agent for Microsoft SQL Server.
6. Set the parameters under Virtual Device Configuration as follows:
 - **Number of Stripes**--Determines the number of CPUs used to perform backups. Set this value to match the number of CPUs in the database server for the fastest backup performance. The default setting is 1 and the maximum value is 32.
 - **Number of Buffers**--The total number of VDI buffers (of maximum transfer size) used to back up and restore. The default setting is 1. This number cannot be less than the number of stripes.
 - **Data Block Size (in bytes)**--All data transfer sizes are multiples of this value. Values must be a power of 2 between 512 bytes and 64 KB inclusive. The default is 65536 or 64 KB.
 - **Maximum transfer size**--The maximum input or output request issued by Microsoft SQL Server to the device. This is the data portion of the buffer. This parameter value must be a multiple of 64 KB. The range is from 64 KB to 4 MB. The default setting is 2097152 or 2 MB.
 - **Maximum VDI Wait Time - Backup (ms)**--The time, in milliseconds, a Virtual Device object waits for a response from Microsoft SQL Server during a backup operation. This setting is also used by the agent when waiting for parallel operations to synchronize or background operations to complete, including during some parts of restore operations. The default setting is 600000 ms (ten minutes).
 - **Maximum VDI Wait Time - Restore (ms)**--The time, in milliseconds, a Virtual Device object waits for a response from Microsoft SQL Server during a restore. Increase this time if the database to be restored contains very large data files. The default setting is 9000000 ms (2.5 hours).
7. Under Named Pipes Configuration, specify the Maximum Connection Wait Time (ms) time, in milliseconds, the Agent for Microsoft SQL Server should wait to close a named pipe if a remote connection fails. The default setting is 400 ms.
8. Set the parameters under Restore Post-Processing Wait as follows:
 - **Polling Period (seconds)**--The amount of time to wait between checks of the database status. The default setting is 60 seconds (one minute).
 - **Maximum Wait Timeout (minutes)**--The total amount of time to wait before abandoning the waiting process. If this timeout elapses and the job contains additional Transaction Log sessions to be restored, then those additional sessions may fail to restore because SQL Server is not yet ready. The default setting is 180 minutes (three hours).

Click Apply to Multiple to display a dialog from which you can select additional SQL Servers. Click OK to apply the settings and return to Configuration.
9. Click OK to end configuration.

Configure Backup and Restore Parameters for the Agent for Microsoft SQL Server Using the Backup Agent Admin

Use the Backup Agent Admin utility to configure the Agent for Microsoft SQL Server backup and restore parameters for supported versions of Microsoft SQL Server. The parameters include settings for Microsoft Virtual Device Interface (VDI) objects and remote communication.

To configure backup and restore parameters for the Agent for Microsoft SQL Server using the Backup Agent Admin

1. Log in to the computer where the Agent for Microsoft SQL Server is installed.
From the Windows Start menu, select All Programs, CA, ARCserve Backup, Backup Agent Admin.
The Arcserve Backup Agent Admin dialog opens.
2. Select Agent for Microsoft SQL Server from the drop-down list and click the Configuration icon.
The appropriate configuration dialog opens.
3. On the Common Settings tab, specify the Level of Detail and Synchronized Recording under Agent Log Settings, as follows:
 - **Level of Detail**--Controls the settings for level of detail of the agent's Activity Log and Debugging Log. For the Activity Log settings, a Level of Detail setting of Normal (0) includes basic information about agent activity. A setting of Detail (1) includes more detailed information about agent activity. A setting of Debug (2) enables the Debugging Log at a moderate level of detail. A setting of Trace (3) enables the Debugging Log at a very high level of detail. The Activity Log is localized for your reference. The Debugging Log is for Arcserve Support use, and is not available in multiple languages.
 - **Synchronized Recording**--Forces the log messages to be written to the Activity Log as they are posted. You can disable this option to improve the performance on high-load systems by caching several messages and writing them as a group.

4. Click Apply to apply the Common Settings.
 5. On the Instance Settings tab, select the Instance (ARCSERVE_DB) or the name of the instance for which you wish to change configuration for the Agent for Microsoft SQL Server.
 6. Set the parameters under Virtual Device Configuration as follows:
 - **Number of Stripes**--Determines the number of CPUs used to perform backups. Set this value to match the number of CPUs in the database server for the fastest backup performance. The default setting is 1 and the maximum value is 32.
 - **Number of Buffers**--The total number of VDI buffers (of maximum transfer size) used to back up and restore. The default setting is 1. This number cannot be less than the number of stripes.
 - **Data Block Size (in bytes)**--All data transfer sizes are multiples of this value. Values must be a power of 2 between 512 bytes and 64 KB inclusive. The default is 65536 or 64 KB.
 - **Maximum transfer size**--The maximum input or output request issued by Microsoft SQL Server to the device. This is the data portion of the buffer. This parameter value must be a multiple of 64 KB. The range is from 64 KB to 4 MB. The default setting is 2097152 or 2 MB.
 - **Maximum VDI Wait Time - Backup (ms)**--The time, in milliseconds, a Virtual Device object waits for a response from Microsoft SQL Server during a backup operation. This setting is also used by the agent when waiting for parallel operations to synchronize or background operations to complete, including during some parts of restore operations. The default setting is 60000 ms (ten minutes).
 - **Maximum VDI Wait Time - Restore (ms)**--The time, in milliseconds, a Virtual Device object waits for a response from Microsoft SQL Server during a restore. Increase this time if the database to be restored contains very large data files. The default setting is 9000000 ms (2.5 hours).
 7. Under Named Pipes Configuration, specify the Maximum Connection Wait Time (ms) time, in milliseconds, the Agent for Microsoft SQL Server should wait to close a named pipe if a remote connection fails. The default setting is 400 ms.
 8. Set the parameters under Restore Post-Processing Wait as follows:
 - **Polling Period (seconds)**--The amount of time to wait between checks of the database status. The default setting is 60 seconds (one minute).
 - **Maximum Wait Timeout (minutes)**--The total amount of time to wait before abandoning the waiting process. If this timeout elapses and the job contains additional Transaction Log sessions to be restored, then those additional sessions may fail to restore because SQL Server is not yet ready. The default setting is 180 minutes (three hours).
- Click Apply to apply the Instance Settings.
9. Click OK to end configuration.

How the Database Protection Job Works

Arcserve Backup lets you use Microsoft SQL Server 2008 Express Edition or Microsoft SQL Server for the Arcserve Backup database. Microsoft SQL Server 2008 Express Edition is a free, lightweight version of Microsoft SQL Server. Although these applications are quite different from each other, with respect to architecture and scalability, you can easily protect either version using the Arcserve Backup default Database Protection Job.

After you install Arcserve Backup, the Database Protection Job maintains a status of Hold. To protect the Arcserve Backup database, you must change the status of the Database Protection Job from Hold to Ready. For more information see, [Start the Arcserve Backup Database Protection Job](#) (see page 623).

If you accept the default Database Protection Job, the job schedule will contain the following values:

- **Schedule Name**--5-day weekly incremental backup, full backup on Friday
- **Execution Time**--11:00AM
- **Rotation Rules**--Append media
- **Media Pool Used**--ASDBPROTJOB

Note: The default retention time of six days allows you to have recovery points of at least a week. If you require more recovery points, you can increase the retention time of the media pool named ASDBPROTJOB manually.

Important! After you start the Database Protection Job, the Tape Engine will connect to a blank media in the first group that Tape Engine detects, and assign the media pool labeled ASDBPROTJOB. If the Tape Engine cannot connect to a blank media in the first group within five minutes, the Tape Engine will try to connect with blank media in the other groups sequentially. If the Tape Engine cannot connect to blank media, in any group, the job will fail.

More information:

[Modify, Create, and Submit a Custom Database Protection Job](#) (see page 614)

How to Back Up the Arcserve Backup Database

There are two approaches that you can use to back up the Arcserve Backup database.

- Create a backup job as you would create any other backup job and include the Arcserve Backup database objects with the source selections for the job.

This method requires that you know if you are running SQL Server 2008 Express or SQL Server in your environment. With this knowledge you must specify the proper source selections and Global Backup Operation options for the job to ensure that the required metadata and related items for each database type are backed up.

This approach lets you back up the affected databases, files, or both, when the backup up job is complete.

- Modify an existing Arcserve Backup Database Protection Job.

When you modify an existing Database Protection Job, Arcserve Backup detects the type of database that is running in your environment.

With this approach, Arcserve Backup selects the proper source selections and specifies the required Global Backup Operation options to ensure that the required metadata and related items for each database type are backed up.

Important! You should not protect the Arcserve database using multiple backup servers that do not reside in the domain that is using the Arcserve database.

Modify, Create, and Submit a Custom Database Protection Job

This section describes how to modify or create a custom Database Protection Job from an existing Database Protection Job. This task helps to ensure that your Arcserve Backup database is backed up and is protected.

Prerequisite Tasks

Before proceeding, ensure that the following prerequisite tasks are complete:

- Ensure that the Tape Engine can detect at least one device in your environment. For more information, see "Managing Devices and Media."
- Ensure that the default Database Protection Job exists in the Job Queue. If the Database Protection Job does not exist in the Job Queue, you must recreate it. For more information, see [Recreate the Arcserve Backup Database Protection Job](#) (see page 625).

To modify, create, and submit a custom database protection job

1. Open the Arcserve Backup Manager Console.

From the Quick Start menu on the Arcserve Backup Home Page, click Job Status.

The Job Status Manager window opens.

2. Click the Job Queue tab.

Locate and select the Database Protection job.

Note: If the Database Protection Job does not exist, you must recreate it. For more information, see [Recreate the Arcserve Backup Database Protection Job](#) (see page 625).

Right-click the Database Protection Job and select Modify from the pop-up menu.

The Backup Manager window opens displaying the Start, Destination, and Schedule tab.

Note: When you modify the Database Protection Job, Arcserve Backup detects the type of database that is running in your environment (SQL Express 2008 or SQL Server) and specifies the database objects that are required to protect the database.

3. Do one of the following:

- Click the Destination tab and specify a location, media, or both where you want to store the backup data. For more information, see [Options You Can Specify on the Backup Manager Destination Tab](#).
- From the Start tab, click the Enable Staging check box.

The Staging Location and Policies tabs appear.

Click the Staging Location tab and specify where you want to stage the backup data. For more information about using staging, see [Backup Staging Methods](#) (see page 200).

4. Click the Schedule tab and specify a schedule for the job. For more information, see [Rotation Schemes](#) (see page 307).
5. Click Options on the toolbar.

The Global Options dialog opens.

6. Click the Operation tab.

Specify Append Backup of Arcserve Backup data at the end of job options required for your Arcserve Backup database:

- **SQL Server 2008 Express Edition--Required Options**

For SQL Server 2008 Express databases, the following options are required and preselected for you:

- **Arcserve Backup database**--This option ensures that the Arcserve Backup database is backed up after jobs are complete.
- **Job scripts**--This option ensures that new and updated job scripts are backed up after jobs are complete.
- **SQL Server Disaster Recovery Elements for Arcserve Backup database**--This option ensures that the elements required to recover a SQL Server database from a disaster are backed up after jobs are complete.

- **SQL Server 2008 Express Edition--Optional Options**

For SQL Server 2008 Express databases, the following options are optional:

- **Catalog files**--This option ensures that the Catalog files are backed after backup jobs are complete.

- **SQL Server--Required Options**

For SQL Server databases, the following options are required and preselected for you:

- **Arcserve Backup database**--This option ensures that the Arcserve Backup database is backed up after jobs are complete.
- **Job scripts**--This option ensures that new and updated job scripts are backed up after jobs are complete.

- **SQL Server--Optional Options--**

For SQL Server databases, the following options are optional:

- **SQL Server Disaster Recovery Elements for Arcserve Backup database**--This option ensures that the elements required to recover a SQL Server database from a disaster are backed up after jobs are complete.
- **Catalog files**--This option ensures that the Catalog files are backed after backup jobs are complete.

Click OK.

The Global Options dialog closes and the Operation options are applied.

7. Click Submit on the toolbar to submit the job.

The Submit Job dialog opens.

8. From the Submit Job dialog, enter a description for your job and click OK.

The Database Protection Job is submitted.

Specify Microsoft SQL Server 2008 Express Backup Options for the Arcserve Backup Database

Arcserve Backup can use full and differential backup methods when backing up Microsoft SQL Server 2008 Express databases. This capability lets you use a rotation scheme or a schedule when backing up the Arcserve Backup database. Additionally, Arcserve Backup lets you check the consistency of the database before the backup job starts or after the backup job is complete.

To specify Microsoft SQL Server 2008 Express backup options

1. Open the Backup Manager windows, select the Source tab, and expand the Windows System object to locate the Arcserve Backup primary server.
2. Expand the Primary server, right-click the Arcserve Backup Database object, and select Agent Option from the pop-up menu.

The Agent Backup Option dialog opens.

3. From the Agent Backup Option dialog, specify the options that you require to protect the database.
 - Specify one of the following Backup Methods:
 - **Use Global or Rotation Options**--Select this option to perform a full or differential backup based on the global job method or rotation phase. The Incremental job method or rotation phase results in a Differential backup.
Note: This is the default backup option.
 - **Full**--Select this option to perform a full backup every time the job is run. When you perform a full backup, Arcserve Backup performs a full backup of the three system databases, the 24 Arcserve databases, and records a synchronization checkpoint. Arcserve Backup creates two backup sessions. One session will contain the disaster recovery elements. The other session will contain all of the data that is required to restore the Arcserve Backup database.
 - **Differential**--Select this option to perform a differential backup every time the job is run. When you perform a differential backup, Arcserve Backup performs a differential backup of the 24 Arcserve databases and records a synchronization checkpoint. Arcserve Backup creates one backup session containing all of the data that is required to restore the Arcserve Backup database.
 - Specify the Database Consistency Check (DBCC) options that you require. DBCC options let you check the allocation and structural integrity of all the objects in the specified databases.
 - **Before backup**--Select this option to check the consistency of the database before the backup starts.
 - **After backup**--Select this option to check the consistency of the database after the backup completes.
 - **Continue with Backup, if DBCC fails**--Select this option to continue the backup even if the check before backup operation fails.
 - **Do not check indexes**--Select this option to check only the system tables.
 - **Check the physical consistency of the database**--Select this option to detect torn pages and common hardware failures. In addition, it checks the integrity of the physical structure of the page and record headers, and the consistency between the page's object ID and index ID.

Example: How DBCC Options Work

The following example illustrates how DBCC options work with Override Global Options on the Agent Backup Options dialog.

- With Override Global Options specified, the DBCC options selected at the database level will be the only DBCC options specified.

- With Override Global Options not specified, all of the DBCC options specified for the database and all the DBCC options selected in the Global options will be applied together.

On the Global Options/Agent Options tab, the DBCC options that follow are specified:

- After backup
- Do not check indexes

On the Agent Backup Options dialog, Override Global Options is not selected and the DBCC options that follow are specified:

- Before backup
- Continue with backup, if DBCC fails

When you submit the backup job, Arcserve Backup applies the DBCC options specified in logical order: Perform the DBCC before the backup starts. If the DBCC fails, perform the backup. After the backup is complete, do not check the indexes.

- **Override Global Options**--Lets the agent ignore all checkbox-based options from the Global Agent Options tab for only the specified database.

Note: The Backup Method is not affected by this option because it can be overridden separately. This option is available only on a per-database basis.

4. Click OK.

Specify Microsoft SQL Server Backup Options for the Arcserve Backup Database

Arcserve Backup lets you protect the Arcserve Backup database using full, incremental, and differential backup methods. This capability lets you use rotation schemes and schedules to protect the database. Additionally, Arcserve Backup lets you back up only the transaction log and check the consistency of the database before the backup job starts or after the backup job completes.

Note: For information about how to protect Microsoft SQL Server databases that do not function as the Arcserve Backup database, see the *Agent for Microsoft SQL Server Guide*.

Follow these steps:

1. Open the Backup Manager windows, select the Source tab, and expand the Windows System object and locate the server that is hosting the Arcserve Backup database.

The server hosting the Arcserve Backup database can be a primary server, a member server, or a remote system. If the server hosting the Microsoft SQL Server database does not appear in the Backup Manager directory, you must add the server to the directory tree under the Windows Systems object before continuing. For more information, see [Back Up Remote Servers](#) (see page 197).

Note: To specify SQL Server backup options, you must authenticate using Windows or SQL Server credentials.

2. Expand the server, right-click the Arcserve Backup database object, and select Agent Option from the pop-up menu.

The Agent Backup Option dialog opens.

3. From the Agent Backup Option dialog, specify the options that you require to protect the database.

The following backup methods are provided:

- **Use Global or Rotation Options**--Backs up the database selected using the job's Global or Rotation Phase Backup Method. The Global or Rotation Options provides the following options:
 - The Full job method will result in a Full backup of the database.
 - The Differential job method will result in a Differential backup of the database, unless this database has not yet had a Full backup.
 - The Incremental job method will result in a Transaction Log backup With Truncation for databases using the Full and Bulk-Logged Recovery Models, and a Differential backup of databases using the Simple Recovery Model, unless this database has not yet had a Full backup.
 - The three main System databases are exempt from the Global or Rotation job method; selecting this option for databases [master], [model], or [msdb] will always result in a Full backup.
- **Full**--Full backup is performed. The files included in the Database Subset will be backed up in their entirety.
- **Differential**--Backs up data that has changed since the last Full backup. For example, if you ran a complete backup of your database on Sunday night, you can run a differential backup on Monday night to back up only the data that changed on Monday. This option is not available for the [master] database.
- **Transaction Log**--Backs up only the Transaction log. This option is only available for databases using the Full and Bulk-Logged Recovery Models.

The following Database Subset options are provided:

The database subset backs up selected files in a database. Use this option to back up a file or FileGroup when the database size and performance requirements do not allow you to perform a full database backup.

Note: The Database Subset options are disabled if the selected Backup Method is Transaction Log Only.

- **Entire Database**--Backs up the entire database.
- **Files and FileGroups**--Backs up selected files in a database. Use this option to back up a file or FileGroup when the database size and performance requirements make it impractical to perform a full database backup. This option is available only for databases using the Full and Bulk-Logged Recovery Models.
- **Partial Database**--Backs up the Primary FileGroup, and any other Read-Write FileGroups. For a Read-Only database, only the Primary FileGroup will be backed up. This option requires SQL Server 2005 or later.
- **Back up Transaction Log After Database**--Backs up the transaction log after the database is backed up. This allows you to perform a Full or Differential backup and a Transaction Log backup in the same job. This option is available only for databases using the Full and Bulk-Logged Recovery Models.

The following Log Truncation Options are provided:

- **Remove inactive entries from transaction log, after backup**--Truncates the log files. This is the default option.
- **Do not remove inactive entries from transaction log, after backup**--Retains inactive log entries after backup. These entries are included in the next Transaction log backup.
- **Backup only the log tail and leave the database in unrecovered mode**--Backs up the log and leaves the database in a restoring state. This option is available for Microsoft SQL Server 2000 or later. Use this option to capture activity since the last backup and take the database offline to restore it.

Important! Do not use the "Backup only the log tail and leave the database in unrecovered mode" log truncation option to back up the ARCserve Database. Performing a backup with this option causes the database to be placed in an offline status, and you can lose the ability to find the backups of the ARCserve Database in order to perform a restore and bring the database online. If you perform a backup of the ARCserve Database using this option, you can use ARCserve Database Recovery Wizard to recover the Arcserve Backup database and bring it back online.

The following database consistency options are provided:

A DBCC tests the physical and logical consistency of a database. DBCC provides the following options:

- **Before Backup**--Select this option to check the consistency of the database before the backup starts.
- **After Backup**--Select this option to check the consistency of the database after the backup completes.
- **Continue with backup, if DBCC fails**--Select this option to continue the backup even if the check before backup operation fails.
- **After restore**--Performs DBCC after the restore of the database.
- **Do not check indexes**--Select this option to check the database for consistency without checking indexes for user-defined tables.

Note: The system table indexes are checked regardless of whether you select this option.

- **Check the physical consistency of the database**--Select this option to detect torn pages and common hardware failures. In addition, it checks the integrity of the physical structure of the page and record headers, and the consistency between the page's object ID and index ID.

Example: How DBCC Options Work

The following example illustrates how DBCC options work in conjunction with Override Global Options on the Agent Backup Options dialog.

- With Override Global Options specified, the DBCC options selected at the database level will be the only DBCC options specified.
- With Override Global Options not specified, all the DBCC options specified for the database and all the DBCC options selected in the Global options will be applied together.

On the Global Options/Agent Options tab, the DBCC options that follow are specified:

- After backup
- Do not check indexes

On the Agent Backup Options dialog, Override Global Options is not selected and the DBCC options that follow are specified:

- Before backup
- Continue with backup, if DBCC fails

When you submit the backup job, Arcserve Backup applies the DBCC options specified in logical order: Perform the DBCC before the backup starts. If the DBCC fails, perform the backup. After the backup is complete, do not check the indexes.

The following miscellaneous options are provided:

- **Include Checksums generated by SQL Server**--Includes error checking information from Microsoft SQL Server, which can be used to validate the integrity of the backed-up data during restore. This option requires SQL Server 2005 or later.

All error messages that are generated during the DBCC are recorded in the Agent for Microsoft SQL Server log file named sqlpag.log. The log is located in the Backup Agent directory.

- **Override Global Options**--Enabling this option will override the Global option setting pertaining to the selected database.

Note: Backup Method and Transaction Log Truncation options are not affected by this option because they can be overridden separately. This option is available only on a per-database basis.

- **SQL Native Backup Compression**--This option applies to only SQL Server 2008 (Enterprise) and later versions. If enabled, this option lets Arcserve Backup use SQL Server database backup compression settings, which results in faster backup times and smaller sessions.

4. Click OK.

The Agent Backup Options are applied.

Start the Arcserve Backup Database Protection Job

The Arcserve Backup database maintains job, media, and device information on your system. After you install Arcserve Backup, the Database Protection Job maintains a status of Hold. To use the Database Protection Job to protect the Arcserve Backup, you must change the status of the Database Protection Job from Hold to Ready.

To start the Arcserve Backup Database Protection Job

1. Open the Arcserve Backup Manager Console.

From the Quick Start menu on the Arcserve Backup Home Page, select Job Status.

The Job Status Manager window opens.

2. Select the Job Queue tab and find the Database Protection Job.

Note: If the Database Protection Job was deleted, you can recreate the job using the steps in [Recreate the Arcserve Backup Database Protection Job](#) (see page 625).

Right-click the Database Protection Job and select Ready from the pop-up menu.

The status of the Database Protection Job changes from Hold to Ready. A full backup of the database will be performed at the next Execution Time.

3. (Optional) To start the Database Protection Job now, right-click the Database Protection Job and select Run Now from the pop-up menu.

The Database Protection Job starts now.

Important! After you start the Database Protection Job, the Tape Engine will connect to a blank media in the first group that Tape Engine detects, and assign the media pool labeled ASDBPROTJOB. If the Tape Engine cannot connect to a blank media in the first group within five minutes, the Tape Engine will try to connect with blank media in the other groups sequentially. If the Tape Engine cannot connect to blank media, in any group, the job will fail.

Access Requirements

When you submit a job that includes remote Windows database servers, Arcserve Backup prompts you for a default user name and password for the system on which the database resides. Arcserve Backup accesses the remote servers using this user name and password.

A Microsoft SQL Server native user name and password are also required to access some database servers. When prompted by the system, enter the Microsoft SQL Server user ID and the password of the system administrator (sa), or enter a user ID and password with equivalent privileges. This user may be a Windows user, depending on security settings.

Note that there are two different data transfer mechanisms available to the agent, and that they have different permission requirements. A backup using Named Pipes only requires the Backup Operator permission for the specific database being backed up, and the Database Creator role to back up the database. A backup using Virtual Devices, however, requires the System Administrator role.

Note: A user in the Backup Operator Group does not have rights to access the Arcserve Backup database. As a result, member servers are not visible to the user in the Backup Manager.

Delete the Arcserve Backup Database Protection Job

Use the following procedure to delete the default Arcserve Backup Database Protection Job.

Important! You should always back up the Arcserve Backup database. Failure to do so can result in unrecoverable backup data.

To delete the Arcserve Backup Database Protection Job

1. Open the Arcserve Backup Manager Console.
From the Quick Start menu on the Arcserve Backup Home Page, select Job Status.
The Job Status Manager opens.
2. Select the Job Queue tab and locate the Database Protection Job.
Right-click the Database Protection Job and select Delete from the pop-up menu.
A warning message appears.
3. If you are sure that you want to delete the Database Protection Job, click OK.
The Database Protection Job is deleted.

Note: For information about how to recreate the Arcserve Backup Database Protection Job, see [Recreate the Arcserve Backup Database Protection Job](#) (see page 625).

Re-create the Arcserve Backup Database Protection Job

Arcserve Backup lets you re-create the Arcserve Backup Database Protection Job in the event it was deleted intentionally or unintentionally.

To re-create the Arcserve Backup Database Protection Job

1. Open the Arcserve Backup Manager Console.
From the Quick Start menu on the Arcserve Backup Home Page, select Server Admin.
The Server Admin window opens.
2. Click the Configuration toolbar button.
The Configuration - <Server name> dialog opens.

3. Select the Database Engine tab and do the following:
 - a. Check the Submit ARCserve DB protection job check box.
 - b. In the Server field, specify the name of the Arcserve Backup server where you want the Database Protection Job to run. You can specify an Arcserve primary server or an Arcserve member server from the domain where you want to recreate the Database Protection job.
 - c. In the Group field, specify the name of the device group where you want to store the Database Protection Job data.

Click OK.

Arcserve Backup recreates the Database Protection Job.

4. Start the Arcserve Backup Database Protection Job.

Note: For more information, see [Start the Arcserve Backup Database Protection Job](#) (see page 623).

Re-initialize the Arcserve Backup Database

The following procedure describes how to re-initialize the Arcserve Backup database. The Arcserve Backup database may not initialize under the following scenarios:

- The Arcserve Backup database could not be recovered from a disaster using ARCserve Database Recovery Wizard.
- The Arcserve Backup database could not be started for various reasons.

When you execute this procedure, the Server Configuration Wizard overwrites the existing Arcserve Backup database instance, which allows you to re-initialize the Arcserve Backup database.

Important! The re-initialization process overwrites your Arcserve Backup database and you will lose backup data. Ensure that the Arcserve Backup database is corrupt before you complete this task.

You can perform this task with Microsoft SQL Server databases and Microsoft SQL Server 2008 Express Edition databases.

To re-initialize the Arcserve Backup database

1. From the Arcserve Backup primary or stand-alone server, start the Server Configuration Wizard.

Note: To start the Server Configuration Wizard, click Start, point to All Programs, CA, ARCserve Backup, and click Server Configuration Wizard.

The Server Configuration Wizard opens with the Select Options dialog.

2. Click Select Database and click Next.

The Check caroot dialog opens.

3. Specify the caroot password and click Next.

The System Account dialog opens.

4. Specify the system account information and click Next.

The Select Database dialog opens.

5. In the Select Database dialog, accept the default options and click Next.

A message appears warning you that certain information will not be migrated.

6. Click OK to clear the message.

7. Do one of the following:

- If the Arcserve Backup database is a Microsoft SQL Server 2008 Express Edition database, continue to the next step.
- If the Arcserve Backup database is a Microsoft SQL Server database, the SQL Database System Account dialog opens. Accept the default options on the SQL Database System Account dialog and click Next.

8. In the Select database installation path dialog, ensure that the Overwrite DB option is enabled.

Note: Enabling this option is essential for re-initializing the database.

Click Next.

The Arcserve Backup database is re-initialized.

9. Click Finish.

Note: After the re-initialization process is complete, all scheduled jobs in the Job Queue (for example, the Arcserve database pruning job and the Arcserve database protection job) will maintain a status of Hold. For the jobs to resume their schedule, you must change the status of each job from Hold to Ready. To change the job status, open the Job Status Manager, right-click the jobs and click Ready on the pop-up menu.

How to Restore the Arcserve Backup Database

There are several methods that you can use to restore the Arcserve Backup database. The following list describes these methods and includes a description of any special considerations or limitations associated with the method.

Important! Microsoft SQL Server database architecture and Arcserve Backup sessions are quite different from that of Microsoft SQL Server 2008 Express. If you attempt to restore Microsoft SQL Server data with Microsoft SQL Server 2008 Express backup sessions, or, conversely, restore Microsoft SQL Server 2008 Express with Microsoft SQL Server backup sessions, the recovery process can corrupt your data.

- **Standard Restore - Backed up in the Arcserve Backup domain that is using the database**--This restore method can be used in the following scenarios:
 - The Arcserve Backup database was backed up in the Arcserve Backup domain that is using the database.
 - The Arcserve Backup database is on line and functioning properly.
 - You want to restore the Arcserve Backup database to a particular point in time.
Note: You can restore the Arcserve Backup database to its original location or a different location.
- **Standard Restore - Backed up in a different Arcserve Backup domain**--This restore method can be used in the following scenarios:
 - The database was backed up in a Arcserve Backup domain that is different from the Arcserve Backup domain that is using the database.
 - The Arcserve Backup database is on line and functioning properly.
 - You want to restore the Arcserve Backup database to a particular point in time.
Note: For SQL Server 2008 Express installations, you must restore the Arcserve Backup database to its original location. For SQL Server installations, you can restore the Arcserve Backup database to its original or a different location.

Recover the Arcserve Backup Database Using Arcserve Database Recovery Wizard

ARCserve Database Recovery Wizard is a self-protection utility that lets you recover the Arcserve Backup database if it fails and was backed up by the Arcserve Backup domain that is using the database. Using the wizard you can recover the database from recent full or differential backups, or recover the database from full backup sessions that are stored on devices connected to the backup server.

Important! You cannot use ARCserve Database Recovery Wizard to recover a Arcserve Backup database that was backed up and used in a different Arcserve Backup domain.

To use ARCserve Database Recovery Wizard, ensure that your system meets the following prerequisite conditions:

- Agent for Arcserve Database is installed on the computer that is hosting the Arcserve Backup database.
- You have a Windows account with Administrative privileges (as a Local Administrator or a user in an Administrative group) on the computer that is hosting the Arcserve Backup database.
- You are running this wizard on the primary backup server or a stand-alone backup server.
- The Tape Engine is running on the Arcserve Backup server.
- The Arcserve Backup Manager Console is not running on the Arcserve Backup server.
- The Arcserve Backup Server Configuration Wizard is not running in the Arcserve Backup server.

When a Arcserve Backup server is configured as cluster-aware, all critical Arcserve Backup base-related services (not agent-related services) are monitored by the applicable cluster service (MSCS or NEC CLUSTERPRO). If a Arcserve Backup base-related service fails or needs to be shut down, the cluster service automatically tries to restart it or trigger a failover if the restart attempt fails. To run this task, you must stop Arcserve Backup services. However, in a cluster-aware environment, you must first manually stop the cluster service from continuing to monitor the service and attempting an automatic restart or failover. For information about how to stop HA service monitoring by the Cluster Service, see [Stop HA Service Monitoring by MSCS](#) (see page 917) or [Stop HA Service Monitoring by NEC CLUSTERPRO/ExpressCluster](#) (see page 930).

To recover the Arcserve Backup database using Arcserve Database Recovery Wizard

Note: If the Arcserve Backup server is not configured as cluster-aware, start at Step 4.

1. If the Arcserve Backup server is configured as cluster-aware, offline the following Arcserve Backup resources:
 - Arcserve ASDB
 - Arcserve HA
 - Arcserve Registry
2. If the Arcserve Backup server is configured as cluster-aware, start the SQL Server service (ARCSERVE_DB).
3. From the Windows Start menu, click Start, point to Programs, CA, ARCserve Backup, and click Database Recovery Wizard.

The Authentication dialog opens.

4. From the following scenarios, specify the credentials that are required to log in to the server:
 - Microsoft SQL Server or Microsoft SQL Server Express Edition is installed on the same computer as Arcserve Backup--Specify the Windows Domain\Account and Password for the Arcserve Backup server.
 - Microsoft SQL Server is installed on a remote server--Specify the Windows Domain\Account and Password for the server that is hosting the Microsoft SQL Server database.
 - You are using SQL Server authentication--Click SQL Server Authentication and specify the Login ID and Password that are required to log in to the SQL Server database.

Click Next.

The Restore Points dialog opens.

Note: The caroot Authentication dialog opens only if you logged in to Arcserve Backup using a Windows account and then started the ARCserve Database Recovery Wizard. On the caroot Authentication dialog, enter the password in the caroot Password field and click OK.

5. The Restore Points dialog retrieves information about available backup sessions from the Arcserve Backup database backup log files.

To retrieve more backup sessions, click More Recovery Points.

The Scan Media dialog opens.

To recover the Arcserve Backup database using the sessions that currently appear on the Restore Points dialog, select the session that you want to recover, click Next, and go to Step 6.

6. On the Scan Media dialog, do the following:
 - a. Specify a backup server and click Connect.
The devices connected to the specified server appear in the devices list.
 - b. Specify a backup device and click Scan.
The available Recovery Points stored on the specified backup device appear in the Detected Recovery Points list.
 - c. Specify the session that you want to recover and click Add to List.
The specified Recovery Points appear in the Selected Recovery Points list.

Note: To retrieve more Recovery Points, select a different device and repeat the Steps b and c.

Additional options:

- **Eject**--Lets you eject tapes and removable hard disk (RDX) media from the device.
Note: This option functions only on stand-alone tape drives and RDX media devices.
- **Refresh**--Lets you refresh the backup device list. You must click Refresh after you insert a new tape or RDX media into the device.

Note: If a backup session spans multiple media, the wizard prompts you to insert the related media.

Click OK.

The Restore Points dialog opens.

7. On the Restore Points dialog, select the session that you want to restore and click Next.
If the specified session is encrypted or contains password protection, the Session Password dialog opens.
8. From the Session Password dialog, enter the password in the Password field and click OK.

The Recovering the ARCserve Database dialog opens and the recovery process starts.

Important! You will be given three opportunities to provide the correct session password. If you cannot provide the correct password after three attempts, the recovery will fail. You must then click Back and repeat Step 5 and specify a different recovery point.

Note: The Messages field on the Recovering ARCserve Database dialog contains important information about the results of the recovery. To view detailed information about the recovery, see the following log file:

ARCserve_HOME\Log\ASrecoveryDB.log

9. When the recovery is complete, click Finish.
Arcserve Backup Database Recovery Wizard recovers the Arcserve Backup database and restarts all required Arcserve Backup services and engines.
10. If the Arcserve Backup server is configured as cluster-aware, bring the following Arcserve Backup resources back online:
 - Arcserve ASDB
 - Arcserve HA
 - Arcserve Registry

Recover the Arcserve Backup Database Using the `ca_recoverdb` Command

Each time you run a backup job, Arcserve Backup records information in its databases about the machines, directories, and files that have been backed up, and the media that was used. This allows you to locate files whenever you need to restore them. The database recovery command (`ca_recoverdb`) is a self-protection feature that allows you to recover a Arcserve Backup database if it is lost and was backed up by the Arcserve Backup domain that is using the database.

The `ca_recoverdb` utility invokes the `ca_restore` commands to implement the database recovery function. The `ca_recoverdb` utility automatically determines if the Arcserve Backup database is a SQL Server database or SQL Server 2008 Express Edition instance and provides the appropriate parameters for the `ca_restore` command.

Be aware of the following behavior:

- The first job that you run after you recover the Arcserve Backup database appears in the Job Status Manager with the same Job ID as the restore job for the Arcserve Backup database. Arcserve Backup demonstrates this behavior because the Job ID assigned to the Arcserve Backup database restore job is lost after you restore the Arcserve Backup database.
- When restoring the Arcserve Backup database in a disk staging environment, Arcserve Backup may attempt to purge data that was purged previously from the staging device. You will receive a warning message, however, the purge job will complete successfully.

Note: For more information about `ca_recoverdb` syntax and options, see [Syntax](#) (see page 633) and [Options](#) (see page 634).

Recovering the Database in Cluster-aware Implementations

When an Arcserve Backup server is configured as cluster-aware, all critical Arcserve Backup base-related services (not agent-related services) are monitored by the applicable cluster service (MSCS or NEC CLUSTERPRO). If an Arcserve Backup base-related service fails or needs to be shut down, the cluster service automatically tries to restart it or trigger a failover if the restart attempt fails. To run this task, you must stop Arcserve Backup services. However, in a cluster-aware environment, you must first manually stop the cluster service from continuing to monitor the service and attempting an automatic restart or failover. For information about how to stop HA service monitoring by the Cluster Service, see [Stop HA Service Monitoring by MSCS](#) (see page 917) or [Stop HA Service Monitoring by NEC CLUSTERPRO/ExpressCluster](#) (see page 930).

To recover the Arcserve Backup database using the `ca_recoverdb` command in a cluster-aware environment, do the following:

1. Offline the following Arcserve Backup resources:

- Arcserve ASDB
- Arcserve HA
- Arcserve Registry

2. Start the SQL Server service (ARCSERVE_DB).

3. Start all Arcserve Backup services using the `cstart.bat` batch file.

Note: The `cstart.bat` batch file is stored in the Arcserve Backup home installation directory.

4. Execute `ca_recoverdb.exe`.

Note: For more information see [Syntax](#) (see page 633) and [Options](#) (see page 634).

5. Bring the following Arcserve Backup resources back online:

- Arcserve ASDB
- Arcserve HA
- Arcserve Registry

Syntax

The `ca_recoverdb` command line syntax is formatted as follows:

```
ca_recoverdb [ -cahost <hostname> ]
           [-i [n]]
           -username <username> [-password <password>]
           [-dbusername <database username> [-dbpassword <database password> ] ] [-sessionpassword [session
           password] -session password [session password]...]
           [-waitForjobstatus <polling interval>]
```

Options

The `ca_recoverdb` provides various options for recovering a lost Arcserve Backup database.

The `ca_recoverdb` command includes the following options:

cahost <hostname>

Redirects default host from the backup log to the host specified by `cahost`.

For example:

HostA - The default host that existed in backup log, which will be used in `ca_restore`.

HostB - The host that you specify.

In these examples, if you do not specify the `cahost` switch, then the `ca_restore` command invoked by the `ca_recoverdb` utility will look as follows:

```
ca_restore -cahost HostA
```

If you do specify the `cahost` switch with the parameter HostB, then the `ca_restore` command invoked by the `ca_recoverdb` utility will look as follows:

```
ca_restore -cahost HostB
```

-i [n]

Specifies to use the interactive mode. If you include this switch, it allows you to specify a point in time from which to perform the Arcserve Backup database recovery by selecting which backup to use as a baseline. When the interactive mode is invoked, the `ca_recoverdb` displays the list of Arcserve Backup sequences for which it has log files. Each of the log files start with a Full database backup, and contains all of the other backups which are dependent on that Full backup to be restored (the Full backup is root of the “dependency chain” for those sessions).

The parameter *n* is used to specify the number of latest backup log sets (dependency chains) that you want to select from. The range of values for *n* is 1 to 99, and the default value is 10.

When you select a Full backup sequence, you will then be prompted to select which session to use as the restore point. After you select a session, the `ca_recoverdb` utility will determine the dependency chain for that sequence, and use `ca_restore` to submit a restore job for each session.

If you do not include the `-i` switch, the `ca_recoverdb` utility automatically uses the most recent backup as the specified selection, and builds the dependency chain for that session. This is helpful if you just want to recover to the latest point in time backup. However, if the most recent backup is lost or damaged, you can use the interactive mode to restore from an older session, and then merge tapes to re-integrate the latest information.

-username <username> [-password <password>]

Specifies the authenticating information for the database agent that will perform the actual recovery job. If you do not include the password option, it will default to no password required.

-dbusername <database username> [-dbpassword <database password>]

Specifies the authenticating information for the database. If you do not include the database username and corresponding database password, it will default to "dbusername" and "dbpassword" for authenticating purposes.

[-sessionpassword [session password] -sessionpassword [session password] ...]

Specifies the authenticating information for the sessions being set authenticating password.

[-waitForJobStatus <polling interval>]

Specifies the time interval (in seconds) that ca_recoverdb will wait until the job is completed and then exit with a return code that indicates the success or fail outcome of the job.

The <polling interval> value defines how often (in seconds) that the ca_recoverdb utility checks the job status with the Queue services. The default polling interval is 60 seconds.

Specify Microsoft SQL Server 2008 Express Edition and Microsoft SQL Server Database Restore Options

The Agents Restore Options dialog lets you specify how you want to restore Microsoft SQL Server 2008 Express Edition and Microsoft SQL Server database instances.

To specify Microsoft SQL Server 2008 Express Edition and Microsoft SQL Server database restore options

1. From the Quick Start menu in the Navigation Bar on the home page, click Restore.
The Restore Manager window opens.
2. From the restore method drop-down list, select Restore by Tree.
From the server tree, locate and expand the system hosting the Arcserve database instance.
Right-click the Arcserve Backup Database object and select Agent Option from the pop-up menu.
The Agent Restore Options dialog opens.
3. Complete the required fields for the Arcserve database instance.

More information:

[Agent Restore Options - Microsoft SQL Server Express Edition - Restore Options](#) (see page 636)

[Agent Restore Options - Microsoft SQL Server Express Edition - Restore Database File Options](#) (see page 638)

[Agent Restore Options - Microsoft SQL Server - Restore Options](#) (see page 638)

[Agent Restore Options - Microsoft SQL Server - Database File Options](#) (see page 645)

Agent Restore Options - Microsoft SQL Server Express Edition - Restore Options

Arcserve Backup lets you specify Microsoft SQL Server Express Edition restore options and the location to restore them.

The Restore Options tab lets you choose how your database is recovered. This tab contains the following selections:

Arcserve Backup Automatic Selection

Lets you automatically select all required sessions and options. This option is enabled by default for every restore job and applies selected options appropriately to the automatically selected sessions.

Miscellaneous

Force restore over existing files or database

Enable this option to let Microsoft SQL Server overwrite files it does not recognize as part of the database it is restoring. Use this option only if you receive a message from Microsoft SQL Server prompting you to use the With Replace option. This option is equivalent to using the With Replace parameter of the restore command.

Use current ASDB as original location

Enable this option if you want to use current Arcserve Backup database as original location.

Recovery Completion State

The following switches determine the condition of the database at the end of the restore job.

Leave database operational

Instructs the restore operation to roll back any uncommitted transactions. After the recovery process, the database is ready for use.

Note: If you use Automatic Selection, you do not have to choose any of the Recovery Completion State selections manually, because Arcserve Backup performs the selection of sessions and the necessary options automatically. If you do not choose Automatic Selection, you must follow Microsoft SQL Server rules regarding the restore flow. For more information, see Microsoft SQL Server documentation.

Leave database offline and able to restore differential

Instructs the restore operation not to roll back any uncommitted transactions and to leave the database in a state where it can accept additional Files-and-FileGroups, Differential, or Transaction Log restores. This is usually selected when performing manual restores.

Database Consistency Check

After restore

Enable this option to check the consistency of the database after the backup completes. To select this option, you must also choose Leave Database Operational. Selecting this option enables the following options.

Do not check indexes

Enable this option to check for consistency without checking indexes for user-defined tables.

Check only the physical consistency of the database

Enable this option to check the database for torn pages and common hardware failures. Additionally, it checks the integrity of the physical structure of the page and record headers, and the consistency between the page's object ID and index ID. This option bypasses the data validity tests normally performed in a standard database consistency check and examines only those related to physical integrity. Index checking is part of the physical integrity tests unless you specifically disable it by selecting Do not check indexes.

Agent Restore Options - Microsoft SQL Server Express Edition - Restore Database File Options

Arcserve Backup lets you specify Microsoft SQL Server Express Edition restore options and the location to restore them.

The Restore Database Files tab lets you specify the location to which the database is recovered. This tab contains the following selections:

Restore to Original Location

Lets you restore the database to its original location, overwriting the current version.

Database Move Rules

Lets you recover the database to a new drive or directory.

Agent Restore Options - Microsoft SQL Server - Restore Options

Arcserve Backup lets you specify Microsoft SQL Server restore options and the location to restore them.

The Restore Options tab lets you choose how your database is recovered. This tab contains the following selections:

Arcserve Backup Automatic Selection

Lets you automatically select all required sessions and options. This option is enabled by default for every restore job and applies selected options appropriately to the automatically selected sessions.

Restore

Database

Lets you restore the entire database.

Files and FileGroups

Lets you restore a file or file group when the database size and performance requirements make it impractical to perform a full database restore.

Partial restore

Note: This option is for only Microsoft SQL Server 2000 and SQL Server 2005.

Lets you restore part of the database to another location so that damaged or missing data can be copied back to the original database. The granularity of the partial restore operation is the database file group. The primary file and file group are always restored, along with the files that you specify and their corresponding file groups. The result is a subset of the database. File groups that are not restored are marked as offline and are not accessible.

Torn Page Repair - Online

Repairs databases in place without the need to perform a restore of the entire database. This operation is recommended when only a few pages are damaged and an immediate recovery is critical.

The database should first be taken offline by performing a Transaction Log backup with the Log Tail option. A Database Consistency Check Before Backup with the Continue if DBCC Fails option is recommended to identify damaged pages that are not recognized, and forestall the possible need of repeating the process. This option is available for all editions of Microsoft SQL Server 2005. The Torn Page Repair restore can then be performed from the latest Full or Differential backup session of that database. If the Automatic Selection option is selected, all the successive Transaction Log sessions are located, as they would for a Files-and-FileGroups restore. If a Differential session is selected, then the corresponding Full backup session will also be automatically selected. The database remains offline until the restore is complete.

Note: Microsoft recommends this only as an emergency measure. A Torn Page Repair can be used to return a damaged database to service when time is critical, but it is recommended that you migrate the database to a new disk at the earliest opportunity to forestall the risk of further errors.

Torn Page Repair - Offline

This option requires the Enterprise Edition of Microsoft SQL Server 2005. A Transaction Log backup with the Do Not Truncate option is used to obtain the latest transaction information which might need to be applied to the damaged pages. A Database Consistency Check Before Backup with the Continue If DBCC Fails option is recommended to identify any damaged pages which have not yet been encountered, and forestall the possible need to repeat the process. You can then perform the Torn Page Repair restore from the latest Full or Differential backup session of that database. If the Automatic Selection option is selected, Automatic Selection will locate all of the successive Transaction Log sessions, as they would for a Files-and-FileGroups restore. If a Differential session is selected, then the corresponding Full backup session will also be automatically selected. The database remains online during the entire process and any tables which are not affected by the damaged pages will remain accessible.

Note: In some cases, you may need to perform an additional Transaction Log Backup with the Do Not Truncate option, and restore that backup without the Automatic Selection option, to fully reactivate the repaired tables. This usually occurs if such a backup was not taken at the start of the process.

For Torn Page Repair restores, the Recovery Completion State option is restricted to the Leave Database Online option. The Database Consistency Check Before Restore option is only enabled when using the Torn Page Repair – Online option, as this is the only time the database will be online during a restore. If a Database Consistency Check was not performed before the last Transaction Log backup, this option can be used to help ensure that Microsoft SQL Server identifies any additional torn pages.

Note: Microsoft recommends this only as an emergency measure. A Torn Page Repair can be used to return a damaged database to service when time is critical, but it is recommended that you migrate the database to a new disk at the earliest opportunity to forestall the risk of further errors.

Miscellaneous

Force restore over existing files or database

Enable this option to let Microsoft SQL Server overwrite files it does not recognize as part of the database it is restoring. Use this option only if you receive a message from Microsoft SQL Server prompting you to use the With Replace option. This option is equivalent to using the With Replace parameter of the restore command.

Restricted user access after restore

If this option is selected, then a restore to Original Location will overwrite the current Arcserve Database, rather than the database that was backed up to this session. This option is used to migrate the session and log information from one Arcserve Domain to another.

Keep replication settings

Instructs the restore operation to preserve replication settings when restoring a published database to a server other than the one on which it was created. This prevents Microsoft SQL Server from resetting the replication settings when it restores a database or log backup on a warm standby server and recovers the database. Use the Keep Replication Settings option when setting up replication to work with log shipping.

You cannot select this option when restoring a backup with the Leave database non-operational, but able to restore additional transaction logs option. Use this option only with the Leave database operational, no additional transaction logs can be restored option.

Use current Arcserve Database as original location

If this option is selected, then a restore to Original Location will overwrite the current Arcserve Database, rather than the database which was backed up to this session. This option is used to migrate the session and log information from one Arcserve Domain to another.

Preserve current Arcserve Domain Memberships

If this option is selected, then the current information about Arcserve Domains, such as the Arcserve Domain name, Primary Server identity and Member Server identities, will be retrieved from the destination database before the restore begins, and written back after the restore completes, preserving this information even after the restore. This option is enabled when the “Automatic Selection”, “Leave Database Operational” and “Use current ARCserve Database as original location” options are all selected, and is selected by default when it is enabled.

In this release, Arcserve Backup retains encryption information in the Arcserve Backup database. The encrypted information can include session passwords and user profile information, all of which is tied to a Arcserve Backup domain. When using this option, the agent determines if it can associate this information with a preserved Arcserve Backup domain from the overwritten database. If the restored database and the overwritten database contain the same domains, the associations will be re-established accordingly. If the restored database and the preserved domain list do not have domains in common, the agent will behave according to the table that follows. For any restored domains that the agent cannot map to one of the preserved domains, you will need to export the keys using the DumpDB utility on one of the Arcserve Backup primary servers or stand-alone servers that is using the database.

Restored Arcserve Database	Overwritten Arcserve Database	Agent Action	Manual Follow-up
One Arcserve Backup domain	One Arcserve Backup domain	<ul style="list-style-type: none"> ■ The Arcserve Backup domain key is transferred, and session passwords and user profiles are re-associated. ■ The server information from the restored Arcserve Backup database is not retained. 	<p>You will be asked to provide the caroot password from the restored domain the first time you open the Manager Console after you restore the database. This finalizes the transfer of the Arcserve Backup domain key.</p>
One Arcserve Backup domain	Two or more Arcserve Backup domains	<ul style="list-style-type: none"> ■ The Arcserve Backup domain key is propagated. ■ The session passwords and user profiles are not re-associated. ■ The server information from the restored Arcserve Backup database is retained. ■ A dollar sign ('\$') is appended to the domain name, the primary server name, and the individual server names of the servers listed in the restored data. 	<ul style="list-style-type: none"> ■ You will be asked to provide the caroot password from the restored domain the first-time you open the Manager Console, for each Arcserve Backup domain, after you restore the database. This finalizes the migration of the Arcserve Backup domain key. ■ You must manually export and import the session passwords using the DumpDB utility on the Arcserve Backup primary or stand-alone server that is to receive the keys. <p>Note: You must execute cstop and cstart on the primary server of each Arcserve Backup domain and finalize the key migration, before importing session passwords using the DumpDB utility.</p>

Restored Arcserve Database	Overwritten Arcserve Database	Agent Action	Manual Follow-up
Two or more Arcserve Backup domains	Any number of Arcserve Backup databases	<ul style="list-style-type: none"> ■ The Arcserve Backup domain keys are not propagated. ■ The session keys and user profiles are not re-associated. ■ The server information from the restored Arcserve Backup database is retained. ■ A dollar sign ('\$') is appended to the domain name, the primary server name, and the individual server names of the servers listed in the restored data. 	<ul style="list-style-type: none"> ■ New domain keys will be created for the retained Arcserve Backup domains when the Arcserve Backup services are restarted on that domain's primary or stand-alone server. ■ You must manually export and import the session passwords using the DumpDB utility on the Arcserve Backup primary or stand-alone server that is to receive the keys. <p>Note: You must execute cstop and cstart on the primary server of each Arcserve Backup domain before importing session passwords using the DumpDB utility.</p>

Note: Before you run the DumpDB utility, you must execute cstop and cstart on all servers in each domain that uses the overwritten Arcserve Backup database.

For information about using the DumpDB utility, see the *Command Line Reference Guide*.

Log point in time restore

Stop before job mark

This option includes date and time fields in which you can set a specific date and time mark. The option recovers the database to the specified mark but does not include the transaction that contains the mark. If you do not check the After datetime check box, recovery stops at the first mark with the specified name. If you check the After datetime check box, recovery stops at the first mark with the specified name exactly at or after datetime.

Note: This option is for only Microsoft SQL Server 2000 and SQL Server 2005.

Stop at log mark

This option includes date and time fields in which you can set a specific date and time mark. The option recovers the database to the specified mark, including the transaction that contains the mark. If you do not check the After datetime check box, recovery stops at the first mark with the specified name. If you check the After datetime check box, recovery stops at the first mark with the specified name exactly at or after datetime.

Note: This option is for only Microsoft SQL Server 2000 and SQL Server 2005.

Stop at time

This option includes date and time fields in which you can enter a specific date and time. The option recovers the database to the specified date and time. This is the default option

Recovery Completion State

The following switches determine the condition of the database at the end of the restore job.

Leave database operational

Instructs the restore operation to roll back any uncommitted transactions. After the recovery process, the database is ready for use.

Note: If you use Automatic Selection, you do not have to choose any of the Recovery Completion State selections manually, because Arcserve Backup performs the selection of sessions and the necessary options automatically. If you do not choose Automatic Selection, you must follow Microsoft SQL Server rules regarding the restore flow. For more information, see Microsoft SQL Server documentation.

Leave database offline and able to restore differential

Instructs the restore operation not to roll back any uncommitted transactions and to leave the database in a state where it can accept additional Files-and-FileGroups, Differential, or Transaction Log restores. This is usually selected when performing manual restores.

Database Consistency Check

After restore

Enable this option to check the consistency of the database after the backup completes. To select this option, you must also choose Leave Database Operational. Selecting this option enables the following options.

Do not check indexes

Enable this option to check for consistency without checking indexes for user-defined tables.

Check only the physical consistency of the database

Enable this option to check the database for torn pages and common hardware failures. Additionally, it checks the integrity of the physical structure of the page and record headers, and the consistency between the page's object ID and index ID. This option bypasses the data validity tests normally performed in a standard database consistency check and examines only those related to physical integrity. Index checking is part of the physical integrity tests unless you specifically disable it by selecting Do not check indexes.

Continue Restore after Checksum Failure

Performs the restore even if consistency checking fails.

Agent Restore Options - Microsoft SQL Server - Database File Options

Arcserve Backup lets you specify Microsoft SQL Server restore options and the location to restore them.

The Database Files Options tab contains options and selections that control where you can restore your database.

Files or FileGroups

Choose the File or File Groups you want to restore from the tree.

Restore Database Files As

Restore to Original Location

Lets you restore the database to its original location. Available at the Database level. Clears any changes to the drive letters, paths, and file names. You must click the Apply button after selecting this option for the change to take effect.

Restore to Original Location Except

Available at the Database, FileGroup and Transaction Log, and File levels. Applies the requested changes to the drive letter, paths, and file names based on the location of the file when the backup was performed.

Move To Drive

Select the Move To Drive check box and enter a different drive letter in the field beside.

Move to Directory

Select the Move To Directory check box and enter a different directory path in the field beside.

Filename Pattern Change

Select the Filename Pattern Change check box, to change the filenames for the entire database, FileGroup, or Transaction Log. Enter a wildcard pattern that matches the names of the files you want to rename in the field below and enter the wildcard pattern that you want it to be renamed to in the to field.

For example, if you want to rename all the files that begin with Group as Members, enter Group* in the field and Member* in the to field.

Note: If you are using a wildcard pattern to rename files, and the pattern for the original filenames does not match with one or more of the files to which it would be applied, a yellow indicator will appear at the bottom of the dialog, and in the tree next to both the affected files and the object where the rule was applied.

Select the Rename File check box and enter a different file name, to rename a single file.

Click Apply for the changes to take effect.

Restore the Arcserve Backup Database (Different Domain)

This section describes how to restore the Arcserve Backup database and the database was backed up using an Arcserve server that resides in a different Arcserve Backup domain. You can restore the Arcserve database in the following scenarios:

- The Arcserve database is functional
- The Arcserve database is not functional and the instance hosting the Arcserve database is functional

In these scenarios you can restore the Arcserve database using the Backup Manager on the system that backed up the Arcserve database.

Important! You cannot restore the Arcserve database while there are jobs in progress. If a job tries to access the Arcserve database while the restore is in progress, the job will fail.

To restore the Arcserve Backup database that was backed up in a different Arcserve Backup domain

1. Stop all Arcserve Backup services running on the primary and the member servers in the domain using the cstop batch file.

Note: For more information, see [Stop and Start All Arcserve Backup Services](#) (see page 493).

2. Log in to the Arcserve domain containing the backup data for the database that you want to restore.

Open the Restore Manager window, click the Source tab, select the Restore by Tree method, expand the Windows Systems object, and browse to the primary server associated with the database that you want to restore.

Expand the server that you want to restore.

Based on the type of database that is running in your environment, select the following database objects:

Microsoft SQL Server 2008 Express

Expand the server object and select the following objects:

- Arcserve Backup Database object
- Microsoft SQL Server Disaster Recovery Elements

Note: If Arcserve Backup is installed in a cluster-aware environment, you must status the Microsoft SQL Server 2008 Express service in maintenance mode before you submit the restore job.

Microsoft SQL Server

Expand the server object, expand the Microsoft SQL Server object, and submit individual restore jobs for the following objects:

- System databases: [master], [msdb], and [model]
- asdb database object

3. Click the Options toolbar button.
The Global Options dialog opens.
4. Select the Operation tab, click the Disable Database Recording option, and click OK.
The database restore options are applied.

5. Click the Destination tab and select the Restore files to their original location option.

Important! If the Arcserve Backup database is a Microsoft SQL Server 2008 Express instance and Arcserve Backup is installed in a cluster-aware environment, you must place the SQL Server service in cluster maintenance mode before submitting the restore job.

Click the Submit on the toolbar to submit the restore job.

The Submit Job dialog opens.

Note: If there are jobs in progress, Arcserve Backup prompts you to restore the Arcserve database to a different location. If you cannot restore the Arcserve database to a different location, allow all jobs in progress to complete, and then restore the Arcserve database.

6. Complete the fields on the Submit Job dialog and click OK.

After the restore job is complete, complete the following tasks:

- a. Start all services on the primary and member servers in the domain using the cstart command.

Note: For more information, see [Stop and Start All Arcserve Backup Services](#) (see page 493).

- b. Using the Merge utility, merge all backup media.
- c. Perform a full backup of the Arcserve database.

Be aware of the following behaviors:

- After you restore the Arcserve Backup database, the job history for the Database Protection Job will indicate that the job is incomplete and the Activity Log will indicate that the job is in progress with an Unknown status. This behavior occurs because the data for the Activity Log and the Database Protection Job is stored in the Arcserve Backup database and the data was incomplete while backup was in progress.

In addition, the status of the Database Protection Job (Done) will be the same as it was before you submitted the restore job. This behavior occurs because the Job Queue will obtain the status of the Database Protection Job from the job scripts rather than from the Arcserve Backup database.

- The first job that you run after you recover the Arcserve Backup database appears in the Job Status Manager with the same Job ID as the restore job for the Arcserve Backup database. This behavior occurs because the Job ID assigned to the Arcserve Backup database restore job is lost after you restore the Arcserve Backup database.
- Arcserve Backup may attempt to purge data that had already been purged from the staging device when you restore the Arcserve Backup database in a disk staging environment. You will receive a warning message, however, the purge job will complete successfully.

How to Recover the Arcserve Database When the SQL Server Instance Hosting the Arcserve Database is Not Functional

A typical disaster recovery scenario consists of the following steps:

1. Reinstall Windows, if necessary.
2. Reinstall Arcserve Backup, if necessary.
3. Reinstall the Agent for Microsoft SQL Server and the Client Agent for Windows, if necessary. (The Client Agent is needed to restore Microsoft SQL Server Disaster Recovery Elements.)
4. Perform one of the following steps as appropriate:
 - If you have a Microsoft SQL Server Disaster Recovery Elements session, restore it.
 - If an offline backup exists, restore it.
 - If you do not have an offline backup or a Disaster Recovery Elements session, and you have the Microsoft SQL rebuildm.exe utility, use the utility to recreate the master and model database. For more information, see the Microsoft documentation.
 - If an offline backup or Disaster Recovery Elements backup do not exist and you do not have the Microsoft SQL rebuildm.exe utility, reinstall the Microsoft SQL Server or MSDE-based application.
5. Restore the [master] database.
6. Restart Microsoft SQL Server in normal, multi-user mode.
7. Restore the [msdb] database.
8. Restore all other databases and transaction logs, except the replication database.
9. If replication is being used, restore the replication database.

How Arcserve Database Recovery Wizard Works

Each time you run a backup job, Arcserve Backup records information in its databases about the machines, directories, and files that have been backed up, and the media that was used. This allows you to locate files whenever you need to restore them. ARCserve Database Recovery Wizard is a self-protection utility that lets you recover the Arcserve Backup database if it fails and was backed up by the Arcserve Backup domain that is using the database.

ARCserve Database Recovery Wizard interacts with the Arcserve Backup Agent for Microsoft SQL Server to facilitate the recovery of the database. You can use the wizard to recover Arcserve Backup databases that are hosted by Microsoft SQL Server and Microsoft SQL Server Express Edition.

ARCserve Database Recovery Wizard lets you recover the Arcserve Backup database from the following sources:

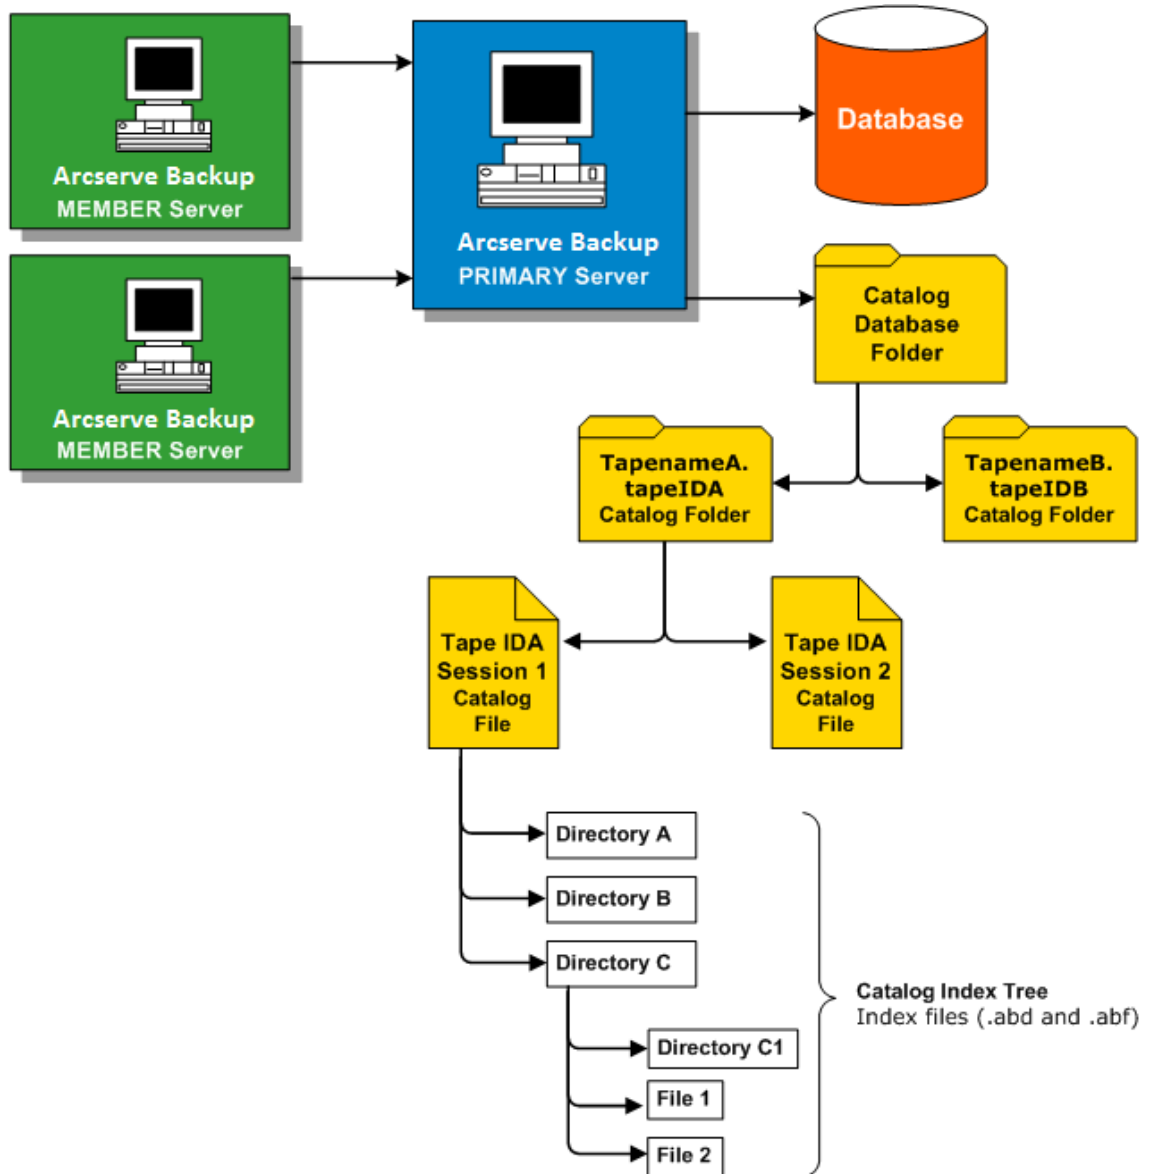
- Sessions retrieved from ASDBBackups log files. Sessions retrieved from this source can consist of full, incremental, and differential backups.
- Sessions retrieved by scanning devices connected to the backup server. Sessions retrieved from this source can consist only of full backups.

Note: ARCserve Database Recovery Wizard does not create jobs that appear in the Job Queue.

How the Catalog Database Works

As the amount of information and data that you create grows larger, the backup jobs that are used to ensure their security also grows larger. As a result, the amount of time needed to parse or analyze this information could become very time-consuming and have a negative impact on the performance and scalability of your system.

To resolve this problem, whenever Arcserve Backup performs a backup, all the job, session, and media information is loaded into a database, while a separate catalog file is also created under the catalog database folder with just the pertinent description information about each session. In addition, two corresponding catalog index files (.abd and .abf) are also generated displaying the index tree structure of the directories and files within the catalog file. These catalog index files are retained on the disk and can be used to quickly browse the content of the session in the catalog file to locate the information when needed.



More information:

[Configure the Catalog Database](#) (see page 653)

Catalog Browsing

Whenever you need to check for a directory or file to be restored, locate recovery points, or to just perform a search, instead of querying the content of the entire database, Arcserve Backup performs the query on just the catalog files with the help of the catalog index. If this catalog browsing finds the details in the catalog database folder for a specific session, it will not look in the Arcserve Backup database. However, if it does not find it, it will then attempt to look in the Arcserve Backup database again. If it still does not find the details of the session in both the catalog database folder and the Arcserve Backup database, it will then prompt you to select whether or not you want to merge the session again so that the merge process can recreate the catalog file into the catalog database folder or can regenerate the details from the tape session contents into the Arcserve Backup database.

Note: All application agent sessions except for Microsoft Exchange Server, such as SQL Server, Informix, Oracle, Microsoft SharePoint Server, Lotus Notes etc, do not support catalog browsing and the details from those sessions will be inserted into the Arcserve Backup database.

Catalog Database Pruning

Whenever you use Arcserve Backup to back up information and data, the amount of description information (catalog files and index files) stored in the catalog database folder increases. Without any controls, in time the size of the catalog database folder would increase, eventually occupying the entire backup disk, and resulting in backup errors due to insufficient free disk space.

To resolve this problem, Arcserve Backup allows you to specify a catalog database pruning threshold. The pruning threshold (or Minimum disk free space threshold) setting is accessible from the Configuration dialog of the Server Admin Manager. The selectable range of this threshold is from 1% to 99%, with the default value set to 10%, and is based on the percentage value of the detected free disk space.

Note: Arcserve Backup periodically checks the free disk space percentage on the volume where the catalog database folder is located. If the detected free space is lower than the specified percentage, a warning message will be sent to the activity log and it would automatically begin to delete catalog database files (minimum of 7 days old and starting with the oldest first) from the disk until the detected free space percentage is greater than the threshold setting.

How a Centralized Catalog Database Works

When operating in a centralized management environment (a primary server and one or more member server), Arcserve Backup centralizes all catalog files into the primary server. As a result, the catalog database files created on a member server during a backup job will be transferred (when the backup is complete) to the catalog database folder located on the associated primary server. In this way, the performance of merging and browsing catalog files that are always local to the primary server will be significantly improved and the maintenance of catalog database, such as pruning and backing up tasks will be simplified and be done only from the primary server.

The Arcserve Backup centralized catalog database helps you manage enterprise-level, multi-server environments. You can browse media information or generate reports for multiple servers at the same time. The member servers in your network update the central Arcserve Backup catalog database with media session information and details from their own database. The central catalog database is set up locally on a machine that manages the centralized catalog database. Media session information for all of the Arcserve Backup machines in your enterprise is contained there.

Configure the Catalog Database

You can configure the catalog database options to customize the location and performance of the database and the associated centralized catalog.

To configure the catalog database

1. Open the Server Admin Manager and click the Configuration toolbar button.

The Configuration dialog appears.

2. Select the Database Engine tab.

The Database Engine dialog appears displaying the catalog database options at the bottom of the dialog.

3. Complete the following fields:

- **Catalog database folder**--Lets you define the where you will store the catalog database. The catalog database folder will contain all the associated catalog files and catalog index files. You can click the ... (ellipsis) button to browse and select a different location for the catalog database folder.

By default, the catalog database folder will be stored in on the primary of stand-alone server in the following directory:

C:\Program Files\CA\ARCserve Backup\CATALOG.DB\

Note: You can only modify the catalog database folder from the Primary Server.

- **Compress catalog transfer on the following member servers**--Lets Arcserve Backup compress catalog information when the data is transferred from a member server to the primary server.

If the primary server has any associated member servers, the Compress catalog transfer on the following member servers field will be enabled, displaying the names of the member servers.

By default, this option is disabled. With this option disabled, Arcserve Backup will not compress the catalog information when it is transferred from the member server to the primary server.

- **Minimum disk free space threshold**--Lets you specify the minimum percentage of free disk space when Arcserve Backup deletes catalog files.

Default value: 10 %

Range: 1% to 99%

Note: Arcserve Backup periodically checks the free disk space percentage on the volume where the catalog database folder is located. If the detected free space is lower than the specified percentage, a warning message will be sent to the activity log and it would automatically begin to delete catalog database files (minimum of 7 days old and starting with the oldest first) from the disk until the detected free space percentage is greater than the threshold setting.

Example: If the detected free space is lower than the 10%, a warning message is sent to the activity log and it would automatically begin to delete catalog database files (minimum of 7 days old and starting with the oldest first) from the disk until the detected free space percentage is greater than 10%.

4. Click OK.

The Catalog Database options are applied.

Move the Arcserve Backup Catalog Database to a Different Location

This topic describes how to move the Arcserve Backup catalog database to a different new location. You may wish to move the catalog database to a different location when you experience any of the conditions that follow:

- The overall size of the catalog database increases significantly. For example, the size of the catalog database increased from one GB (gigabyte) to 30 GBs.
- There is a noticeable time lag when you retrieve restore data.
- More than eight hours are required to complete the Database Protection Job.
- More than four hours are required to complete the Database Pruning Job.

- The catalog database is consuming a significant amount of disk space on the C:\ drive, which affects the amount disk space required for the Windows pagefile.sys file.
- Microsoft SQL Server performance has deteriorated due to increased Arcserve Backup transactions and overhead, and you do not have a dedicated SQL Server system.
- You cannot export the Arcserve Backup database to a flat file using Microsoft SQL Server utilities due to the overall size of the database.
- You wish to modify your current Arcserve Backup database configuration such that summary information is recorded in the Arcserve Backup database and detail information is recorded in the catalog database.

Best Practices

Before you move the Arcserve Backup catalog database to a different location, consider the following best practices:

- By default, the catalog database is installed in the directory that follows:

C:\Program Files or Program Files(x86)\CA\ARCserve Backup\CATALOG.DB

If you must move the catalog database, the best practice is to move the catalog database to a location resides or communicates locally with the Arcserve Backup server.

Note: You should not move the catalog database to a remote disk that resides on a Network Attached Storage (NAS) device or a mapped network drive. These locations may require authentication to access the devices.

- The application used for the Arcserve Backup is not relevant to the location of the catalog database. However, consideration should be given for the protocols used for communication:
 - **ODBC/RPC**--The communication protocol used between member servers and the catalog database on the primary server and the communication protocol used between the catalog database on the primary server and Microsoft SQL Server.
- For configurations where the catalog database resides on a SAN device, you should consider installing an additional SCSI controller or HBA card on the Arcserve Backup server to accommodate communication between the Arcserve Backup server and the storage device.

Note: The steps that follow apply to Arcserve Backup servers using Microsoft SQL Server or Microsoft SQL Server 2008 Express Edition to host the Arcserve Backup database.

To move the Arcserve Backup catalog database to a different location

1. Complete the following tasks:
 - Ensure that all jobs, including the Database Protection Job and the Database Pruning Job, are in a Hold state.
 - Ensure that the Arcserve Backup Manager Console is closed on all servers in the Arcserve Backup domain, with the exception of the primary or stand-alone server.

2. If there are member servers in the Arcserve Backup domain, execute Cstop on all member servers to stop all Arcserve Backup services.

Note: For information about using Cstop, see [Stop and Start All Arcserve Backup Services Using Batch Files](#) (see page 493).

3. Execute Cstop on the Arcserve Backup primary or stand-alone server to stop all Arcserve Backup services.
4. After all services stop, bring the new location for the catalog database online.
5. In the new location for the catalog database, create the path.

For example:

F:\ARCserve\catalog.db

6. Copy all folders from the original location to the new location.

For example:

Original location

C:\Program Files\CA\ARCserve Backup\CATALOG.DB

New location

F:\ARCserve\catalog.db

7. After you copy the Catalog files to the new location, execute Cstart on the primary or stand-alone Arcserve Backup server to restart all Arcserve Backup services.

Note: For information about using Cstart, see [Stop and Start All Arcserve Backup Services Using Batch Files](#) (see page 493).

Allow several minutes to elapse to ensure that all Arcserve Backup services start.

8. Open the Server Admin Manager.

From the server directory tree, click the Arcserve Backup primary server or stand-alone server.

Click Configuration on the toolbar.

The Configuration dialog opens.

9. Click the Database Engine tab.

In the Catalog database folder field, specify the new path to the catalog database.

For example:

F:\ARCserve\catalog.db

Optionally, you can click the ellipsis to browse to the new location for the catalog database.

Click OK.

The Configuration dialog closes.

10. Submit a simple backup job from the primary or stand-alone server.
11. After the simple backup job is complete, submit a simple restore job from the primary or stand-alone server.
12. After the simple backup and restore jobs on the primary or stand-alone server are complete, ensure that Arcserve Backup is writing job summary and detail information to the new location for the catalog database.
13. If there are member servers in the Arcserve Backup domain, execute Cstart on one of the member servers. Allow several minutes to elapse to ensure that all Arcserve Backup services start.
14. Submit a simple backup job from a member server.
15. After the simple backup job is complete, submit a simple restore job from the member server.
16. After the simple backup and restore jobs on the member server are complete, ensure that Arcserve Backup is writing job summary and detail information to the new location for the catalog database on the primary server.
17. Execute Cstart on the remaining member servers in the Arcserve Backup domain. You should allow five minutes to elapse between Cstart executions to minimize SAN, LAN, and Arcserve Backup RPC updates.

Note: After you move the catalog database to a different location, restore job and summary and detail data from the same backup job may not reflect in the catalog database. To remedy this problem, run Repair Database Connection using the Server Configuration Wizard. For more information, see [Repair the Arcserve Database Connection on a Primary Server](#) (see page 575).

Using Microsoft SQL Server as the Arcserve Backup Database

This section contains the following topics:

[Microsoft SQL Server Database Considerations](#) (see page 658)

[Remote Database Considerations](#) (see page 659)

[Specify ODBC Communication for Remote Database Configurations](#) (see page 660)

[How to Calculate the Number of Required SQL Connections](#) (see page 661)

[How to Enable TCP/IP Communication on Microsoft SQL Server Databases](#) (see page 661)

[Database Consistency Checks](#) (see page 661)

Microsoft SQL Server Database Considerations

Review the following information if you are considering using Microsoft SQL Server for the Arcserve Backup database:

- If you are upgrading to this release and currently running Microsoft SQL Server for the Arcserve Backup database, you must continue using Microsoft SQL Server for the Arcserve Backup database.
- Arcserve Backup does not support using Microsoft SQL Server 7.0 for the Arcserve Backup database.
- By default, Arcserve Backup creates the Arcserve Backup database (ASDB) using a simple recovery model. You should retain this model for proper operation.
- Microsoft SQL Server supports local and remote communication. This capability lets you configure the Arcserve Backup database to run locally or remotely to your Arcserve Backup server.

Note: For more information, see [Remote Database Considerations](#) (see page 659).

- By default, Arcserve Backup stores information about the backed up files and directories in the Catalog Database. This behavior causes the Catalog Database to grow in size at a faster rate than the Arcserve Backup database. Given this behavior and the needs of your organization, you should plan to have a sufficient amount of free disk space to support the growth of the Catalog Database.
- For Global Dashboard, the Central Primary Server, Arcserve Backup database (ASDB) must have Microsoft SQL Server 2005 or later installed (does not support Microsoft SQL Server 2008 Express Edition or Microsoft SQL Server 2000 as its database).

Note: For a Branch Primary Server, no additional hardware or software is required beyond the minimum requirements for any Arcserve Backup primary server.

- To install Arcserve Backup with Microsoft SQL Server support, an administrative account such as the sa account, which has the right to create devices, is required for proper installation.

You should use the *sa* account when prompted for the Arcserve Backup Database (SQL) System Account during installation of Arcserve Backup with Microsoft SQL support.

- Set the database security mode to SQL security in the SQL Enterprise Manager. This applies when using SQL security as the authentication mode and the systems that you want to back up reside inside or outside the Windows domain.
- If you specify Microsoft SQL Server 2000, Microsoft SQL Server 2005, Microsoft SQL Server 2008, or Microsoft SQL Server 2012 as the Arcserve Backup database during setup, you can use Windows authentication or SQL Server authentication to communicate with the Microsoft SQL database.
- If the Microsoft SQL Server account is changed, you must make the corresponding changes using the Server Configuration Wizard.
- The Arcserve Backup Database Engine periodically polls the status of the Microsoft SQL Server database. If Microsoft SQL Server does not respond in a timely fashion, the Database Engine assumes that the Microsoft SQL Server is unavailable and shuts down (red light). To avoid this situation, you can set the registry key to an appropriately longer value to increase the wait time for Arcserve Backup Database Engine, as follows:

```
HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\Arcserve  
Backup\Base\Database\MSSQL\SQLLoginTimeout
```

- Arcserve Backup does not support local Microsoft SQL Server installations on Arcserve Backup servers in NEC CLUSTERPRO environments. In NEC CLUSTERPRO environments, you must install the Arcserve Backup database instance on a remote system.
- If the ODBC driver is configurable, the System Data Source "ASNT" under System DSN, in the ODBC Data Source Administrator should have the Client Configuration set to utilize TCP/IP communication.

Remote Database Considerations

Using a remote database provides a simple and transparent method of sharing a single database as if the database resides locally. When you use this configuration, you do not need a database on the local machine because all information is saved to the remote database. This configuration is best under the following conditions:

- There is not enough space locally for the database.
- There is no organizational requirement and you want to take advantage of the ease of management that comes with having a single location for the database.

- You require a separate server that is not a Arcserve Backup server to function as a dedicated as a Microsoft SQL Server machine.
- To protect SQL Server instances in a cluster-aware environment, you must manually install the Agent for Microsoft SQL Server on all of the cluster nodes.

Note: For information about backing up and restoring Microsoft SQL Server Databases, see the *Agent for Microsoft SQL Server Guide*.

- Use the Server Configuration Wizard to configure ODBC communication between a remote Arcserve database and the Arcserve primary or stand-alone server. This wizard lets you configure efficient communication between servers, especially when you have more than one Arcserve Backup server in your environment.
- To ensure that Arcserve Backup can communicate with the system that is hosting the Arcserve database instance, you should enable TCP/IP communication between the SQL Server database instance and the Arcserve server.

Note: For more information, see [How to Enable TCP/IP Communication on Microsoft SQL Server Databases](#) (see page 661).

Important! Microsoft SQL Server 2008 Express Edition does not support remote database communication.

Specify ODBC Communication for Remote Database Configurations

If you have another Arcserve Backup server running that uses Microsoft SQL as its database, you can redirect the local database to the remote machine. Arcserve Backup can use ODBC to connect to the Microsoft SQL server. You can direct the ODBC data source to another server if the server has SQL installed and the Arcserve Backup SQL database is properly set up. You also need to make sure the local server user is authenticated in the remote server.

To specify ODBC communication for remote database configurations

1. Open the Windows Control Panel, select Administrative Tools, Data Sources (ODBC), and System DSN.
2. Add a System Data Source labeled as follows:
Name: ASNT
Server: MachineName\InstanceName
3. Follow the on-screen instructions to test and complete the configuration.

How to Calculate the Number of Required SQL Connections

For each job that you run, you need two SQL connections. Be sure that you have set enough connections (or licenses) in your SQL server. To determine your default SQL connections, select Server and SQL server from the SQL Arcserve Manager. When you browse from the Configuration tab, you can see the user connections. Set these values to the appropriate user setting. If an error message appears, for example, “Cannot Update Record” or “Failed to Login,” you may have run out of connections. You should increase the open object to 2000.

How to Enable TCP/IP Communication on Microsoft SQL Server Databases

If you are hosting the Arcserve database instance using Microsoft SQL Server 2000 or later and the Arcserve Backup database will reside on a remote system, the installation wizard may not be able to communicate with the database on the remote system.

To help ensure that the installation wizard can communicate with the remote system, you should enable TCP/IP communication between the Arcserve Backup server and the server that will host the Arcserve Backup database before you install Arcserve Backup.

- Microsoft SQL Server 2000--To enable TCP/IP communication on Microsoft SQL Server 2000 systems, run the SQL Server Network utility and ensure that TCP/IP appears in the Enabled Protocols. If TCP/IP does not appear in the Enabled Protocols list, add TCP/IP to the list and click OK. To apply TCP/IP communication, restart all Microsoft SQL Server services.
- Microsoft SQL Server 2005 and later--To enable TCP/IP communication on Microsoft SQL Server 2005 and later systems, run the SQL Server Configuration Manager and enable TCP/IP communication for the SQL Server instance. To apply TCP/IP communication, restart all Microsoft SQL Server services.

Note: For Microsoft SQL Server 2008, you must use the SQL Server Native Client 10.0 driver.

Database Consistency Checks

When your database activity is low, we recommend that you run a database consistency check if you have a large database. Although it takes some time, it is important to determine that your SQL database is functioning well. For more information, see your Microsoft SQL guide.

Important! Be sure to monitor the log size periodically. If a log is full, the database cannot function. Although the default setting is “truncate log on checkpoint,” you should increase the log size to 50% of the database if you expect to keep a large number of records.

Specify a Arcserve Backup Database Application

The following sections describe how to configure Microsoft SQL Server and Microsoft SQL Server 2008 Express as the Arcserve Backup underlying database.

Configure Microsoft SQL Server as the Arcserve Backup Database

Using the Server Configuration Wizard, you can configure Microsoft SQL Server as the Arcserve Backup database.

Before you configure Microsoft SQL Server as the Arcserve Backup database, the following considerations apply:

- Microsoft SQL Server must be installed on the system hosting the Arcserve Backup database before you start this task.
- After you configure Arcserve Backup to use Microsoft SQL Server as the Arcserve database, the Server Configuration Wizard opens a command utility labeled `exptosql.exe` that migrates the core and detail tables from the Microsoft SQL Server 2008 Express database to the newly configured Microsoft SQL Server database.
- You can use this procedure to move the Arcserve Backup Microsoft SQL Server database to a different server.

Note: For more information about using Microsoft SQL Server as the Arcserve Backup database, see [Using Microsoft SQL Server as the Arcserve Backup Database](#) (see page 658).


To configure Microsoft SQL Server as the Arcserve Backup database

1. From the Windows Start menu, click Start, select All Programs, CA, ARCserve Backup, and click Server Configuration Wizard.

The Server Configuration Wizard opens.

2. Click the Select database option and click Next.
3. Follow the on-screen instructions to complete the configuration.

Note: If you protect data that contains Unicode-based characters from East Asian languages (for example, JIS2004) you must enable SQL collation to ensure that you can search and sort the data. To do this, click Language Support Options on the Select Database Installation Path dialog and follow the on-screen instructions to complete the configuration.



4. After the configuration is complete, the Server Configuration Wizard opens a command line window, starts exptosql.exe, and migrates the SQL Server 2008 Express core and detail tables to the new SQL Server database.

If exptosql.exe does not start, open a command line window and start exptosql.exe.

Note: By default, exptosql.exe is installed in the following directory:

C:\Program Files\CA\ARCserve Backup

Execute the following commands:

- a. **exptosql.exe core**

This is a required step. The core argument lets you migrate the core tables from the SQL Server 2008 Express database to the SQL Server database.

Important! You must execute this command immediately after the SQL Server configuration is complete.

b. **exptosql.exe detail**

This is an optional step. The detail argument lets you migrate the detail tables from the SQL Server 2008 Express database to the SQL Server database. You can execute this command, at any time, after the core migration process is complete.

Note: Depending on the size of the SQL Server 2008 Express database, the detail table migration process can require a significant amount of time to complete.

Move the Arcserve Backup Database to a Different System or Instance

Use the Server Configuration Wizard to move the Arcserve Backup database to a different system or instance.

Note: This option only applies to Microsoft SQL Server installations.

Before you move the Arcserve Backup database to a different system or instance, be aware of the following considerations:

- The Server Configuration Wizard lets you change your current Microsoft SQL Server configuration to the following types of configurations:
 - Cluster-aware
 - Remote
 - Local
- To access the new SQL Server installation, you must specify a method of authentication. You can use one of the following authentication methods:
 - Windows security
 - SQL Server security
- For remote SQL Server installations that use SQL Server authentication, you must provide the Login ID and Password for the Remote Server Administrator Account.

To move the Arcserve Backup database to a different system or instance

1. From the Windows Start menu, click Start, select All Programs, CA, ARCserve Backup, and click Server Configuration Wizard.

The Server Configuration Wizard opens.

2. Select the Select Database option and click Next.

Follow the on-screen instructions to complete the configuration.

Note: If you protect data that contains Unicode-based characters from East Asian languages (for example, JIS2004) you must enable SQL collation to ensure that you can search and sort the data. To do this, click Language Support Options on the Select Database Installation Path dialog and follow the on-screen instructions to complete the configuration.

Select database installation path
Select installation path for SQL Server

Please select an installation path for ASDB

Select Database Path:

C:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\

Overwrite DB

Warning: If you select the overwrite check-box, all data in the existing instance will be deleted and can not be recovered.

Language Support Options

< Back Next > Cancel

After the configuration is complete, you must install the Arcserve Backup database protection agent on the system hosting the SQL Server database.

3. To install the Arcserve database protection agent, do **one** of the following:
 - If the SQL Server database is installed on the Arcserve Backup primary server, open Windows Explorer and browse to the following directory:
C:\Program Files\CA\ARCserve Backup\Packages\ASDBSQLAgent
 - If the SQL server database is not installed on the Arcserve Backup primary server, open Windows Explorer and browse to the following directory:
C:\Program Files\CA\ARCserve Backup\Packages\ASDBSQLAgent
Copy the contents of the ASDBSQLAgent directory to any location on the system hosting the SQL Server database.
4. In the ASDBSQLAgent directory, double-click the following file:
SQLAgentRmtInst.exe
The Arcserve Backup Agent for SQL Setup dialog appears.
5. Complete the following fields, as required, for your installation:
 - SQL Instance Name
Specify the name of the SQL instance that you want to protect.
 - Auth Mode
Specify the authentication mode that Arcserve Backup will use to communicate with and protect the database.
If you specify SQL Authentication as the authentication mode, complete the following fields:
 - SQL SA Name
Specify the SQL system account name.
 - SQL SA Password
Specify the SQL system account password.
6. Click Install and follow the on-screen instructions to complete the installation.

Configure Microsoft SQL Server 2008 Express as the Arcserve Backup Database

Using the Server Configuration Wizard, you can configure Microsoft SQL Server 2008 Express as the Arcserve Backup database.

Before you configure Microsoft SQL Server 2008 Express as the Arcserve Backup database, be aware of the following considerations and limitations:

- To deploy Microsoft SQL Server 2008 Express in your environment, Microsoft .NET Framework 2.0 and Microsoft Data Access Components (MDAC) 2.8 Service Pack 2 must be installed on the primary server. If the Server Configuration Wizard does not detect either of these applications, the wizard installs them for you.
- Microsoft SQL Server 2008 Express does not support remote installations. You must install the Arcserve database on the Arcserve Backup primary server.
- You cannot migrate database information from a Microsoft SQL Server database installation to a Microsoft SQL Server 2008 Express database installation.

To configure Microsoft SQL Server 2008 Express as the Arcserve Backup database

1. From the Windows Start menu, click Start, select All Programs, CA, ARCserve Backup, and click Server Configuration Wizard.

The Server Configuration Wizard opens.

2. Click the **Select Database** option and then click **Next**.
3. Follow the on-screen instructions to complete the configuration.

Note: If you protect data that contains Unicode-based characters from East Asian languages (for example, JIS2004) you must enable SQL collation to ensure that you can search and sort the data. To do this, click Language Support Options on the SQL Server Express Instance dialog and follow the on-screen instructions to complete the configuration.

Arcserve Backup Logs and Reports

Arcserve Backup provides the following options for displaying logs and reports:

- [Activity Log](#) (see page 668)--Logs all Arcserve Backup activity.
- [Tape Log](#) (see page 668)--Logs all media activity (for debugging purposes only)
- [Job Log](#) (see page 668)--Logs activity related to a specific job.
- [Reports Manager](#) (see page 669)--Generates reports from the Arcserve Backup database for viewing or printing.

Activity Log Data

The Activity Log contains comprehensive information about the operations performed by Arcserve Backup. It provides an audit trail of all Arcserve Backup activity (including group activities) for every job that is run. You can scan this log every day to see if any errors have occurred. You can also use it to find out a session number in case you need to restore a specific session. The log is located on the upper right corner of the Job Status Manager.

For each job, the log includes the following:

- Time the job started and ended
- Type of job
- Average throughput of the data
- Number of directories and files processed (backed up, copied, etc.)
- Job session number and Job ID
- Result of the job
- Errors and warnings that occurred
- All child jobs listed under the parent job along with a description

The Activity Log has an organize feature which allows you to sort the log using filters, message grouping, or message post date. For more information on the Activity Log, see “Customizing Jobs.”

Tape Log

The Tape Log contains messages sent by the tape drives to Arcserve Backup. This log is not generated during normal operation. It is designed for debugging purposes only. To enable the Tape Log, use the Server Admin Configuration menu.

Note: In a cross-platform environment, the Tape Log does not display information for non-Windows servers. Only Windows server Tape Engine information is available for viewing in the GUI.

Job Log

A Job Log is generated for each job that is run by Arcserve Backup. You can specify the level of detail in the log by choosing the log options before you submit the job. For more information about the job log, see “Customizing Jobs.”

Report Manager

The Report Manager provides you with a variety of reports based on the backup activity stored in the Arcserve Backup database. You can preview a report, print to a printer or file, as well as schedule when to generate a report.

Generate Reports Using Report Manager

The Report Manager lets you generate reports about Arcserve Backup activities.

You can generate reports that Run Now from primary servers, stand-alone servers, and member servers. You can schedule reports to run at a specific time on primary servers and stand-alone servers. If you schedule a report from a member server, the report will run from the primary server, display in the Manager Console on the primary server, and will be stored in the <ARCSERVE_HOME>/Reports directory.

Report Manager Considerations

- You can view all Create now-based reports in the Report Manager window or a browser application, such as Internet Explorer.
- You must view all Schedule-based reports in the Report Manager window.

To generate reports using Report Manager

1. From the Monitor & Reports menu in the Navigation Bar on the home page, click Report.

The Report Manager opens and a collapsible tree that provides an expandable view of reports in various categories appears.

2. On the Report Categories view, select a report template from the list.

The Report template list appears on the right pane.

- Select and right-click the report that you want to generate. From the pop-up menu, specify one of the following options:

Schedule

Lets you schedule a report to run at a specific time.

When you specify this option, the Schedule Report dialog opens. On the Schedule Report dialog, follow the prompts and complete the required fields to schedule the report.

Create now

Lets you generate a report that runs now.

When you specify this option, the Create Report dialog opens. On the Create Report dialog, follow the prompts and complete the required fields to create the report now.

After Arcserve Backup creates the report, you can view the report results on the Report Content view pane as illustrated by the following graphic.

The screenshot shows two parts of the interface. The top part is a list of report files:

- Job Report_data_PRI-SER-XP_2007-07-26-11-48-22.xml
- Job Report_data_PRI-SER-XP_2007-07-26-13-27-40.xml

The bottom part shows a detailed view of the report 'Job Report_data_PRI-SER -XP_2007-07-26-13-27-40.xml'. It includes a description: 'The "Job Report" lists all of the jobs that have been run by CA ARCserve Backup. For each job type, and description.' Below this is a table with columns: ID, Type, Status, Start Time, and Owner. To the left of the table are action buttons: Print Preview, Print to Printer, Create Now, and Delete.

ID	Type	Status	Start Time	Owner
1	Backup	Finished	0/00/80 11:50 AM	caroot B
2	Backup	Finished	0/00/80 11:50 AM	caroot B
3	Backup	Finished	0/00/80 11:51 AM	caroot B
4	Backup	Finished	0/00/80 11:51 AM	caroot B
5	Prune Database	Finished	0/00/80 12:00 PM	D
6	Backup	Finished	0/00/80 1:26 PM	caroot B
7	Backup	Finished	0/00/80 1:27 PM	caroot B

Note: The Report Manager also allows you to remove reports using the delete option to delete the entire report files or delete reports based on date.

Report Manager Reports

Using Report Manager, you can generate three types of reports:

- Standard
- Custom
- Advanced

These reports are described in further detail in this section. For a summary listing of each report and type, see [Report Categories](#) (see page 672).

Standard Reports

Arcserve Backup provides several standard reports that display general backup and restore activity. The reports cover activity for job runs, media backups, and backup device errors. You can use a report filter to select the backup media you want to include in the report. Standard reports cannot be customized or scheduled to print at a specific time interval (not including the Preflight Check Report and GFS Media Prediction Report).

Note: Arcserve Backup cannot display the Preflight Check Report and the GFS Media Prediction in the Report Manager when Arcserve Backup generates the reports via a schedule.

Custom Reports

Custom reports can be modified to meet your specific needs. Although the layout is similar to standard reports, custom reports are created using templates and saved in .XML format.

Note: You can adjust the layout of a custom report by modifying the width of the report columns. Open Windows Notepad and search for the report you want to adjust. Add or change the WIDTH attribute of the FIELD tag in the report template.

Custom reports can be scheduled to run immediately, at a specified time or repeat interval, and can be scheduled from the Primary server or a Member server. You can also specify to have the generated custom report sent to you by email.

There are two types of custom reports:

- **Predefined**--Available in seven categorical types when you install Arcserve Backup. Predefined reports contain basic report data headings that you can modify to suit your specific needs using Report Writer.
- **User-created**--Using Report Writer, you can create a report without using a template as a guide. If you save your user-created report in the Arcserve Backup reports directory, the title of the report displays in the My Reports folder in Report Manager.

Note: Although you can schedule reports from a Primary server and a Member server, Arcserve Backup generates the reports from the Primary server. To ensure that user-created, custom reports run as scheduled from the Primary server, you must copy the user-created report template to the reports directory on the primary server. The default reports template directory on the Primary server is as follows:

```
C:\Program Files\CA\ARCserve Backup\templates\reports
```

Advanced Reports

Advanced reports provide you with an overview of the current data protection status in your Arcserve environment. Advanced reports are predefined, available in different types when you install Arcserve Backup, and contain report data headings that you can modify to suit your specific needs.

To run a report, you must specify the report type and the path of the file where the generated report will be saved.

Advanced reports are similar to custom reports in that they can be scheduled to run immediately or at a specified time or repeat interval. You can also elect to have the generated advanced report sent to you by email.

Report Categories

The report categories that display in Report Manager originate from an external XML schema file (categories.xml) in the Arcserve Backup home directory. You can change the display order of the report categories by editing the categories.xml file.

The following table describes the categories, and types available for standard and custom reports.

Daily Status Reports

This report category provides the status of all jobs executed within the last 24 hours, including reports that display all clients which failed backup and media written in the last 24 hours.

The following reports are available:

- Daily Job Status Report (custom)
- Daily Backup Status Report (custom)
- Daily Failed Backups Report (custom)
- Recently Written Media Report (custom)

Note: The Daily Backup Status Report contains a field called Compression Ratio. The compression ratio shows the amount of data actually written to disk after deduplication. The field is available only for sessions backed up to deduplication device groups.

Job Reports

This report category shows the status information for report jobs executed on a weekly basis. It provides reports showing all failed backups and a preflight check report that displays the status of report jobs scheduled to run at a future date.

The following reports are available:

- Job Report (standard)
- Enterprise Job Status Report (standard)
- 7 Days Job Status Report (custom)
- 7 Days Backup Status Report (custom)
- Failed Backups Report (custom)
- Preflight Check Report (standard)

Media Reports

This report category shows detailed media information about sessions backed up, including a list of media errors generated. Forecasted media schedules for GFS jobs are also available.

The following reports are available:

- Backup Media Error Report (standard)
- Session Detail Report (standard)
- Session Report (standard)
- Arcserve Backup for Laptops and Desktops Session Details Report (standard)
- GFS Media Prediction Report (standard)
- Media Usage Comparison Report (custom)
- Media Utilization Report (custom)
- Media Required for Data Recovery Report (custom)
- 7 Days Media Usage History Report (custom)
- Scratch Set Media in Device Report (standard)

Note: The Media Utilization, Session Detail and Session Reports now contain a field called Compression Ratio for Deduplication. This ratio shows the amount of data actually written to disk after deduplication. The field is only available for sessions backed up to deduplication device groups. This field is present through the Report Manager and also through Report Writer, File, Open, Arcserve Backup Home Directory, Templates, Reports.

Media Pool Reports

This report category shows detailed media pool related information including the status of media in scratch sets and GFS rotation profiles.

The following reports are available:

- Media Pool Report (standard)
- Media Pool Location Report (standard)
- GFS Rotation Profile Report (standard)
- Detailed Media Pool Report (custom)
- Media in Scratch Sets Report (custom)

Device Reports

This report category shows information about backup devices used with Arcserve Backup including the number of errors incurred during a backup on a device.

The following report is available:

- Backup Device Report (standard)

Backup Clients Reports

This report category shows backup client information including database and client agent data sizes.

The following reports are available:

- Backup Client Data Size Report (custom)
- Backup Clients and Job Associations Report (custom)
- Detailed Media Usage by Backup Clients Report (custom)

Note: The Backup Client Data Size Report now contains a field called Compression Ratio. This ratio shows the amount of data actually written to disk after deduplication. The field is only available for sessions backed up to deduplication device groups. This field is present through the Report Manager and also through Report Writer, File, Open, Arcserve Backup Home Directory, Templates, Reports.

Resource Usage History Reports

This report category shows forecasted usage information based on historical data.

The following reports are available:

- 7 Days Media Usage History Report (custom)
- Media Utilization Report (custom)
- Media Usage Comparison Report (custom)
- Backup Window and Throughput Comparison Report (custom)

Staging Reports

This report category provides you with information that you can use to analyze and manage data that was backed up to a file system device using staging.

With Staging Reports you can view status information about migration sessions, SnapLock sessions, and sessions that did not purge from staging devices. The Summary report lets you view information about a specific job or a group of jobs based upon a user-specified range of dates.

The following reports are available:

- Staging Migration Report
- Staging Purge Failed Report
- Staging SnapLock Report
- Staging Summary Report

Audit Log Reports

This report category adds audit fields to existing activity log information such as service starts and stops, or password changes. You can now identify who executed an action, from what machine, and using what application.

The following report is available:

- Audit Log Report

Statistics Reports

This report category provides an overview of the current data protection status. The reports include information about the backup and restore status. However, the output is based on the filter combinations you specify.

The following reports are available:

- Backup Attempt Success Rate: Summary Report (advanced)
- Backup Attempt Success Rate: Individual Client Report (advanced)
- Restore Attempt Success Rate Report (advanced)
- Drive Throughput Report (advanced)
- Backup Error Report (advanced)
- Clients With Most Failed Backup Attempt (advanced)
- Most Consecutive Backup Failures Report (advanced)
- Clients With Most Partial Backups Report (advanced)
- Full Backup Duration (advanced)
- Last Backup Status Report (advanced)
- Vaulting Report (advanced)

My Reports

This report category shows user-created reports that are saved in the following directory:

CAVARServe Backup\Templates\Reports

Statistics Reports

The statistics reports can accept and parse a variety of report filters; however, not all filters are required for all the reports. Based on the type of report that is being generated, only the filters required and supported by that specified report will be used.

To run any advanced report, you must specify at least the report type and the path to where the generated report will be saved.

The following report types, along with the corresponding supported filters can be generated using the statistic report category:

Backup Attempt Success Rate: Summary

This report provides information on what percentage of backup attempts that are successful and also shows percentage of incomplete and failed backup attempts.

Supported Filters:

- Start Date
- End Date
- Job Comment

Backup Attempt Success Rate: Individual Client

This report provides information on what percentage of backup attempts that are successful on a per node basis.

Supported Filters:

- Start Date
- End Date
- Job Comment

Restore Attempt Success Rate

This report provides information on what percentage of all restore attempts that are successful.

Supported Filters:

- Start Date
- End Date

Drive Throughput

This report provides information about the average throughput that is being seen on the tape drives in the system. Throughput obtained from this report can be compared against the native throughput of the drive. The output of this report can be filtered to specific drives by specifying the drive serial number.

Backup Error

This report shows the number of errors and warnings generated for the backup job for each of the backup paths during the reporting period. This helps in determining the clients with most number of errors.

Supported Filters:

- Start Date
- End Date
- Job Comment

Failed Backup Attempt

This report shows the clients with the most failed backup attempts during the reporting period.

Supported Filters

- Start Date
- End Date
- Job Comment
- Top Count (Limit output to the top 'n' clients only)

Consecutive Failed Backup Attempt

This report shows the clients with the consecutive failed backup attempts during the reporting period.

Supported Filters

- Start Date
- End Date
- Top Count (Limit output to the top 'n' clients only)

Partial Backups

This report shows the clients with the most number of partial backups. This reports help identify and restore critical file.

Supported Filters

- Start Date
- End Date
- Top Count (Limit output to the top 'n' clients only)

Full Backup Duration

This report shows the average backup time, average backup data, and average throughput for full backups of all backup paths during the reporting period.

Supported Filters

- Start Date
- End Date

Last Backup Status Report

This report shows the status of last execution of all the backup jobs in the queue. If a job is still active, it shows the current status of the running job. This report only shows the status of the job in the queue at the time the report is generated.

Supported Filters

- None

Vaulting Report

This report shows the list of tapes that will move in or out of the vault on the day of reporting.

Supported Filters

- None

Custom Report Job Scheduling

There are two ways you can schedule a custom report to run—Report Manager or Job Schedule Wizard. From Report Manager, you can schedule two of the standard reports (Preflight Check Report and GFS Media Prediction Report) and custom reports that represent Predefined or User-Created which display in the report categories.

Schedule a Custom Report Using the Report Manager

To schedule a custom report from Report Manager

1. Locate the report you want to schedule from the report list tree.
2. Click Schedule in the left panel pane next to the report description.
3. Specify the name and format type (.xml or .csv) for the report.
4. (Optional) Check the alert option box if you want to be alerted when report is sent by email and click Next.
5. Choose schedule options to run the report immediately or at a specific time and click Next.
6. Review your selection in the Job Summary page and enter a job description, if necessary.
7. Click Submit to execute the report job run.

Schedule a Custom Report Using the Job Scheduler Wizard

To schedule a custom report from the Job Scheduler wizard

1. Create a report template using the Arcserve Backup Report Writer and save it to a file.
2. Locate the Job Scheduler Wizard executable in the Arcserve Backup home directory and double click to launch it.
3. Select CAReports in the Run this program combo box and enter the report template name, the output file name where the report data will be stored, and silent mode (-s) mode as the parameters.

Note: For a full command line supported by the Report Writer, see the *Command Line Reference Guide*.

Create Custom Reports Using the Report Writer Utility

Report Writer is a Arcserve Backup utility that you can use to create custom reports. You can access Report Writer from the Quick Start, Monitor and Reports menu or the Monitor and Reports section on the Arcserve Backup home page.

Note: Reports created using Report Writer can be previewed, printed, or scheduled in Report Manager.

To create custom reports using the Report Writer utility

1. From the Quick Start menu select Monitor and Reports, and then click Report Writer.

The Arcserve Backup Report Writer window opens.

2. Select the File menu and click Open to locate the report you want if you are generating a Predefined report. Otherwise, if you are generating a User-defined report, go to Step 3.
3. Enter a name for your report in the Report Title text box. Optionally, you can enter a description of your report in the Description text box.
4. In the Available Queries table, highlight the source from which you want to gather information for your report. When you highlight a source (such as Media or a Media Pool), the Available Columns table is populated with the types of data you can collect from the selected source. For example, if you select Job in the Available Queries table, you can collect information about the Job Type, the Job Owner, the Job Start Time, and several other items.

To select an item to include in your report, highlight the item in the Available Columns table and click Add. The item then moves to the Report Columns table.

Note: You can create reports made up of information collected from multiple sources. For example, you could create a report that reports on Job Type, Tape Name, and Source Host.

5. Click Next to go to the Report Criteria screen. From this screen, you can customize your report in the following ways:
 - Set the order of the records—The records (or rows) in the columns of your report can be sorted in either ascending or descending order. By default, the records are ordered in ascending order.
 - Set the order of the columns—The column at the top of the Report Columns list appears as the first (leftmost) column in your report. To change the position of a column, highlight it in the Report Columns table and click the up or down arrow.
 - Set filters—The records for your report can be filtered for specific criteria that you define. Use the Enter Value field, with the Operators and Condition drop-down menus, to specify the criteria for each type of record (each listing in the Report Columns table) in your report. After you specify a filter, click Add Criteria to add it to the Query Criteria table.

For example, to report only on jobs with a Job ID from 150 through 250, follow these steps:

- a. Click Job, Job ID in the Report Columns table.
 - b. Set the Operators drop-down menu to “>=”, type 150 in the Enter Value field, and set the Condition drop-down menu to “and.” Then click Add Criteria.
 - c. Set the Operators drop-down menu to “<=” and type 250 in the Enter Value field. Click Add Criteria. The Query Criteria table then reflects your criteria.
6. To run your report, click Generate Report.

Note: If you are using Report Writer to generate predefined Disk Staging Reports, the Add and Remove buttons are not accessible.

Report Generation for Multiple Arcserve Backup Servers

You can generate reports for a Arcserve Backup server at any time using the -m switch with the CARports command line utility. If you want to generate reports for more than one Arcserve Backup server, it is recommended that you create and store report templates on one server, and use remote servers as data sources. The customized report templates do not have to be updated for each Arcserve Backup server. Use the -m switch for each server so that all template updates are batched as a generic job.

Note: You can use the -a switch with the CARports command line utility to enable auto-file naming to generate daily reports.

Arcserve Backup Diagnostic Utility

The Arcserve Backup Diagnostic Wizard utility is a convenient tool for gathering and packaging various Arcserve Backup and system logs, which may be necessary for troubleshooting.

The diagnostic wizard collects information about the following Arcserve Backup agents.

- Agent for Informix
- Agent for Lotus Domino
- Agent for Microsoft Exchange Server
- Agent for Microsoft SharePoint Server
- Agent for Microsoft SQL Server
- Agent for Oracle
- Agent for Sybase
- Client Agent for Windows
- Network Attached Storage Agent (NDMP NAS Option)
- Universal Agent

Note: The diagnostic utility is installed default.

Diagnostic Utility Components

The Diagnostic Utility contains two components:

- Diagnostic Wizard
- Diagnostic Report Manager

You can launch the Diagnostic Wizard from the Arcserve Backup program group. It allows you to configure what kind of report and log you want to generate.

You can run one of the two following report generation modes:

- **Express Mode**--Collects information about the local machine. Does not include advanced debugging information.
- **Advanced Mode**--Collects information about the local machine or a remote machine and generates reports with greater debugging information enabled. If you select this mode, you are prompted to rerun the relevant job so that the newly selected debug flags can be processed during the job and entered into the report.

The Diagnostic Wizard also lets you select where to place the log on your hard disk. After you complete the Diagnostic Wizard, a file is created. You can view this file from the Diagnostic Report Manager, which is also accessible from the Arcserve Backup program group.

The following sections describe the process of running and reviewing an Express mode report.

Note: To run the Diagnostic Wizard in the Advanced mode, choose the Advanced option on the Select Diagnostic Type screen, and then follow the on-screen instructions.

Configure Computers Running Windows Vista and Windows 7 Operating Systems to Communicate with the Diagnostic Wizard

By default, the Diagnostic Wizard cannot collect diagnostic information about computers running Windows Vista and Windows 7 operating systems. Therefore, you must configure computers running Windows Vista and Windows 7 operating systems to communicate with the Diagnostic Wizard.

To configure computer running Windows Vista and Windows 7 operating systems to communicate with the Diagnostic Wizard

1. Ensure that the Remote Registry service is running in the Windows Service Manager.
2. Ensure that you allow TCP port 445 to communicate through the Windows firewall on the Windows operating system.

Create Reports Using the Express Mode Diagnostic Utility

Using the Arcserve Diagnostic Wizard you can generate reports that collect diagnostic information about the local server.

Note: The Express Mode does not collect advanced debugging information about the local server.

To create reports using the Express Mode Diagnostic Utility

1. Open the Diagnostic Wizard by selecting Start, Programs, CA, ARCserve Backup, and then click Diagnostic Wizard.

The Arcserve Diagnostic Wizard opens.

2. Click Next.

The Select Diagnostic Type window opens.

From here, you can choose to collect diagnostic logs from either the local server or a remote server, as well as whether or not you want to include advanced debugging information in the report.

Choose the Express type to gather local logs without including debugging information.

3. Click Next.

Select the attributes of the machine you want to gather logs from.

4. Click Next.

Specify the location where you want to save the diagnostic information file in the Diagnostic Information File Name field and click Next.

A summary of the logs to be collected displays.

5. Click Start.

This process can be lengthy, depending on the system and the amount of information that you requested.

6. When the process is complete, click OK, and then click Finish.

After the file has been created, you may be prompted to send it to Arcserve Technical Support.

Create Reports Using the Advanced Mode Diagnostic Utility

Using the Arcserve Diagnostic Wizard you can generate reports that collect advanced debugging information about the local server.

To create reports using the Advanced Mode Diagnostic Utility

1. Open the Diagnostic Wizard by selecting Start, Programs, CA, ARCserve Backup, and then Diagnostic Wizard.

The Arcserve Diagnostic Wizard opens.

2. Click Next.

The Select Diagnostic Type window opens.

From here, you can choose to collect diagnostic logs from either the local server or a remote server, as well as advanced debugging information in the report.

Choose the Advanced type to gather local logs without including debugging information.

3. Click Next.

The Select an Arcserve Backup Job window opens.

Choose an Arcserve job.

4. Click Next.

Select the attributes of the machine you want to gather logs from.

5. Click Next.

Select the debug mode.

6. Click Next.

Specify the location where you want to save the diagnostic information file in the **Diagnostic Information File Name** field and click Next.

A summary of the logs to be collected displays.

7. Click Start.

This process can be lengthy, depending on the system and the amount of information that you requested.

8. When the process is complete, click OK, and then click Finish.

After the file has been created, you may be prompted to send it to Arcserve Technical Support.

View Reports Using the Diagnostic Report Manager

After the information you requested is collected, you can use the Diagnostic Report Manager to view it.

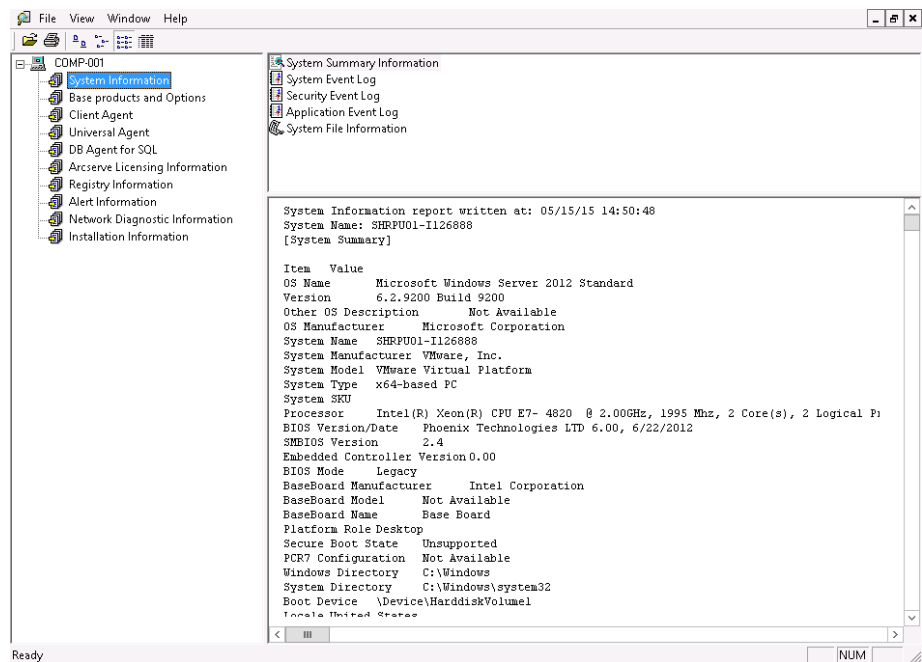
To view diagnostic reports

1. Start the Diagnostic Report Manager by selecting Start, Programs, CA, ARCserve Backup, and then select Diagnostic Report Manager from the Arcserve Backup program group.

The **Diagnostic Report Manager** opens.

2. From the File menu, choose Open.
3. Locate your saved log file and click **Open**.

A console opens that shows a list of your logs on the left-hand side. Detailed information appears in the bottom-right pane as you select items in the left-hand pane.



4. (Optional) You can also view Product Logs by selecting the Base products and Options node on the left-hand pane.

The list of all product logs appears on the right-hand pane, together with tape logs and job logs.

From the File menu on the Diagnostic Report Manager you can also export, save, and print the selected log file.

Arcserve Backup Infrastructure Visualization

Arcserve Backup Infrastructure Visualization provides a visual representation of your Arcserve Backup environment, allowing you to quickly see backup status and explore how servers, nodes, and devices are related.

Arcserve Backup Infrastructure Visualization shows each Arcserve Backup Server in a hierarchical form resembling an organization chart. Primary and Member servers are indicated at the top of the hierarchy under the Arcserve domain. The mini-map feature acts as a scaled-down version of the current view, allowing you to pan to portions of the entire graph. You can customize the Infrastructure Visualization by filtering by node name or node tier. You can also print the contents of a view. Double-clicking a specific server, node or device displays its details and permits you to access the corresponding Dashboard reports.

Example: How to Use Infrastructure Visualization

Suppose you want to see the backup status of nodes backed up by each Arcserve Backup server in your environment. To do so, open Infrastructure Visualization and switch to Nodes view, then group the nodes by subnet. The graph displays all Arcserve Backup servers with the backed up nodes grouped by subnets. At the top of each subnet group, a bar shows the last backup status of all nodes in the subnet based on a pre-defined [color scheme](#) (see page 689):

- If the entire bar is red, the backup of all the nodes in that subnet failed.
- If some portion of the bar is yellow, while the rest is green, the backup of some nodes was incomplete, while the rest were successful.

Click a node to open the details window, which shows the backup information specific to the selected node. You can check machine information such as CPU, OS and Memory, as well as check more detailed information by launching related Dashboard Reports. To launch a Dashboard Report, double-click the item in a group and then open the desired report.

Infrastructure Visualization Software Requirements

Arcserve Backup Infrastructure Visualization requires Microsoft .NET Framework 3.5 SP1.

Note: Since .NET Framework 3.5 SP1 is not supported on Itanium-based systems, Infrastructure Visualization is not supported on these systems.

Infrastructure Visualization is installed with the Arcserve Backup Server and requires no additional licenses.

To view Global Infrastructure Visualization, you must install and configure the Global Dashboard component during Arcserve Backup primary/standalone server installation. For more information about Global Dashboard installation and configuration, see the Arcserve Backup Dashboard User Guide.

Infrastructure Visualization Operations

Built-in operations like filtering, zoom controls and mini-map navigation make it easy to view the backup status of your environment.

- **Filtering**--You can filter by node name or tier, including patterns, such as PAY to find all machines whose names contain PAY. You can save filters on a per-user basis.
- **Print**--You can print out a copy of the full view, not including the mini-map, zoom slider or toolbar buttons. You can also zoom into a specific area and print it.
- **Zoom Controls**--Using a slide bar, you can enlarge and decrease the display magnification, changing the Infrastructure Visualization scale.
- **Mini-Map Navigation**--The mini-map shows a small high-level overview of the entire (current) view. For very large environments, you can zoom out the view and move the bounding box in the Mini-map view by dragging your mouse. The main display updates the view to show the portion included within the bounding box.

Some controls on the toolbar are common to all Infrastructure Visualization views, such as Refresh and Print. Other views have specific toolbars.

- **Nodes View** -- The toolbar contains Group nodes by: Subnet/Agent, Node name filter and Node tier filter controls.
- **Virtual Machine View** -- The toolbar contains Virtual Machine Type: VMware/Hyper-V, Node name filter and Node tier filter controls.
- **Device View** -- There are no controls specific to this view.

Infrastructure Visualization Color Scheme

You can determine last backup status of all nodes in a group by matching the color bar to the following color scheme key:

- Red -- indicates failed jobs
- Orange -- indicates cancelled jobs
- Yellow - indicates incomplete jobs
- Blue -- indicates jobs that have not yet been attempted
- Green -- indicates successful jobs

When the color bar is all one color, the last backup status for all nodes in the group have the same status. When status differs, the bar is shaded proportionately.

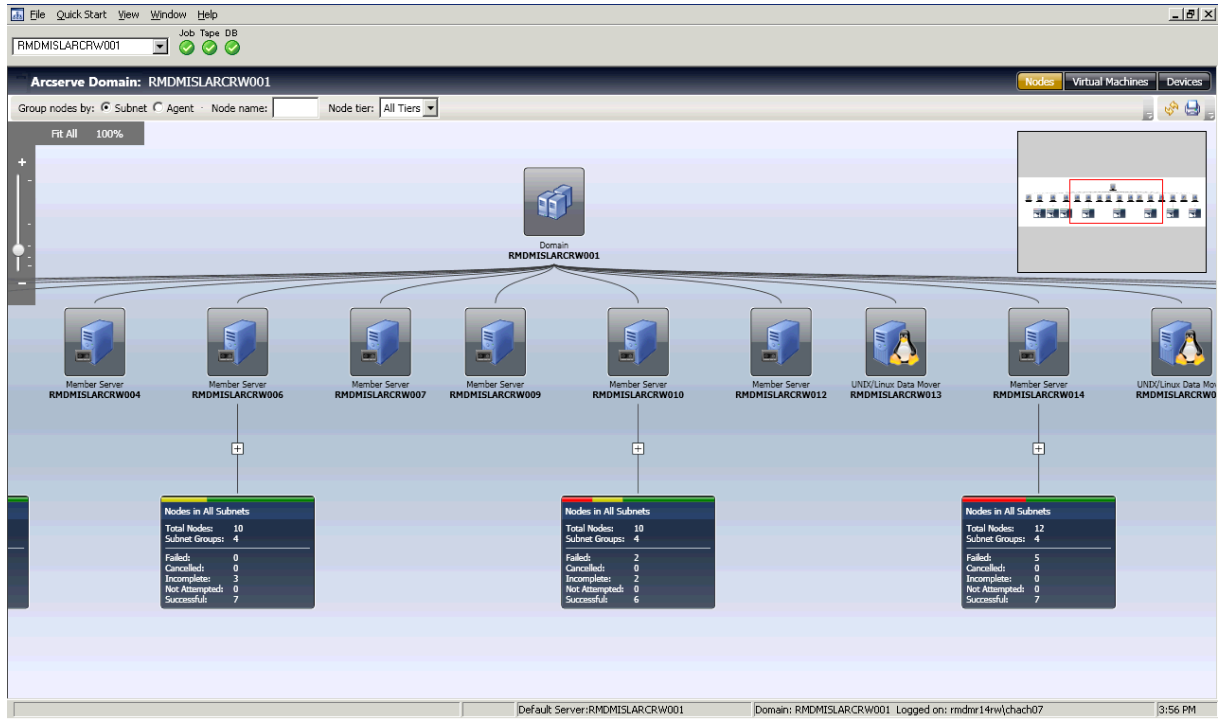
Arcserve Backup Infrastructure Visualization Views

Infrastructure Visualization is organized into the following views, showing groups of related items. Each view has a specific purpose and function. You can easily switch among views by clicking the appropriate view button at the top of your screen.

- [Nodes View](#) (see page 690).
- [Virtual Machine View](#) (see page 692).
- [Device View](#) (see page 694).

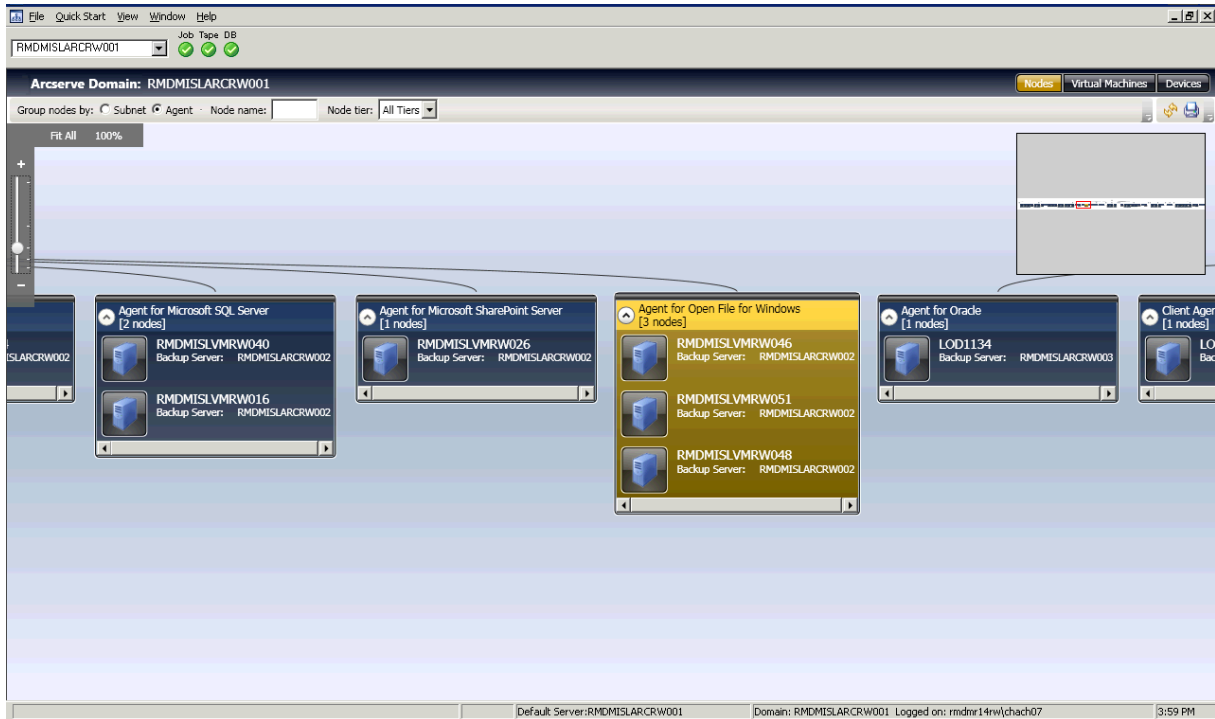
Nodes View

Nodes View is the default view and represents the backup relationship of nodes. It may be filtered by two radio buttons on the toolbar: Subnet and Agent View.



- In **Subnet** view, the nodes backed up by Arcserve Backup are displayed in subnet groupings. All servers are shown across the top and all nodes backed up by those servers are shown grouped by their subnets. Servers with nodes beneath them are displayed with a summary item and an Expansion symbol (+). Click the symbol to view subnet groups. On the summary item, there is a status bar that shows total nodes in percent format based on a pre-defined color scheme, and text details including Total Node Count, Subnet Group Count and node count for each color status.

- In **Agent** view, the nodes backed up by Arcserve Backup are grouped by installed agents. All servers are shown across the top and all nodes backed up by those servers are shown grouped by the agents installed on that server. When a node has more than one agent installed, it appears under multiple agent groups. Since last known backup status is not Agent-specific, the status bar is gray.

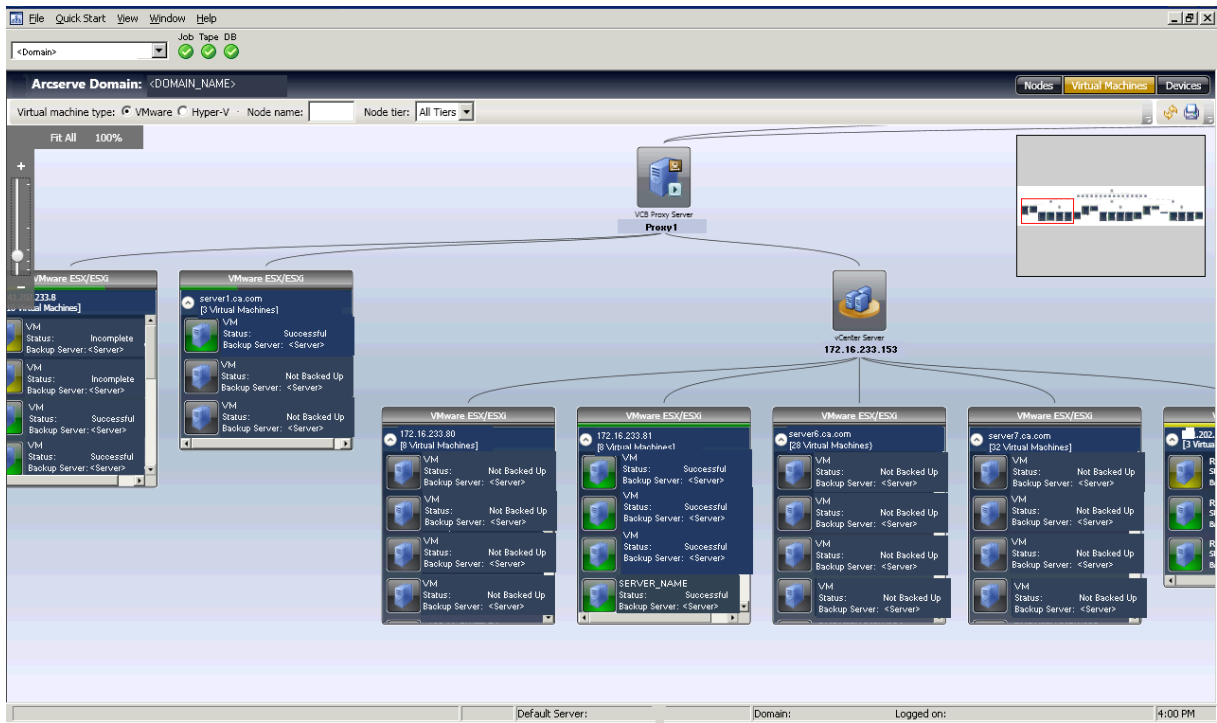


Virtual Machine View

Virtual Machine view lets you see the virtual machine environment (VMware and Hyper-V virtual machines) in the Arcserve Backup domain. In Virtual Machine view, all VMs backed up by Arcserve Backup are grouped by their VMware ESX/ESXi or Hyper-V server. You can filter the view by virtual machine as follows:


VMware

Shows the VMware ESX/ESXi Server VMs backed up using the Arcserve Backup Agent for Virtual Machines. If a VM was backed up using a standalone VMware ESX/ESXi Server, it is displayed in the hierarchy, Backup Server, VMware Proxy, VMware ESX/ESXi Server, VM. If a VM was backed up using a VMware vCenter Server, it is displayed in the hierarchy, Backup Server, VMware Proxy, VMware vCenter Server, VMware ESX/ESXi Server, VM.



Hyper-V

Shows the Microsoft Hyper-V VMs backed up using the Arcserve Backup Agent for Virtual Machines in the following hierarchy: Backup Server, Hyper-V Host Server, VM.



Backup status for each VM is represented by text and visual indicators. VMware ESX/ESXi and Hyper-V servers are displayed similar to Groups and have status bars at the top to indicate overall backup status of the VMs beneath them. Gray indicates a VM has not been backed up, which could occur when VMs are populated into the Arcserve database by the VM data population utility, but not yet backed up.

Connecting lines between a VMware proxy or Hyper-V server and a backup server indicate that at least one VM under that group was backed up.

Device View

This view lets you see the backup devices connected to the respective Arcserve Backup server in the Arcserve domain. The devices are grouped by device type.



Deduplication Devices



File System Devices



SAN Devices



Tape Devices



Virtual Tape Libraries

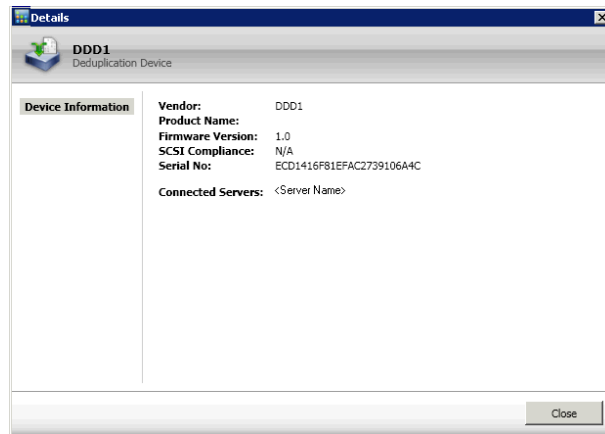


NAS Devices



Shared VTLs

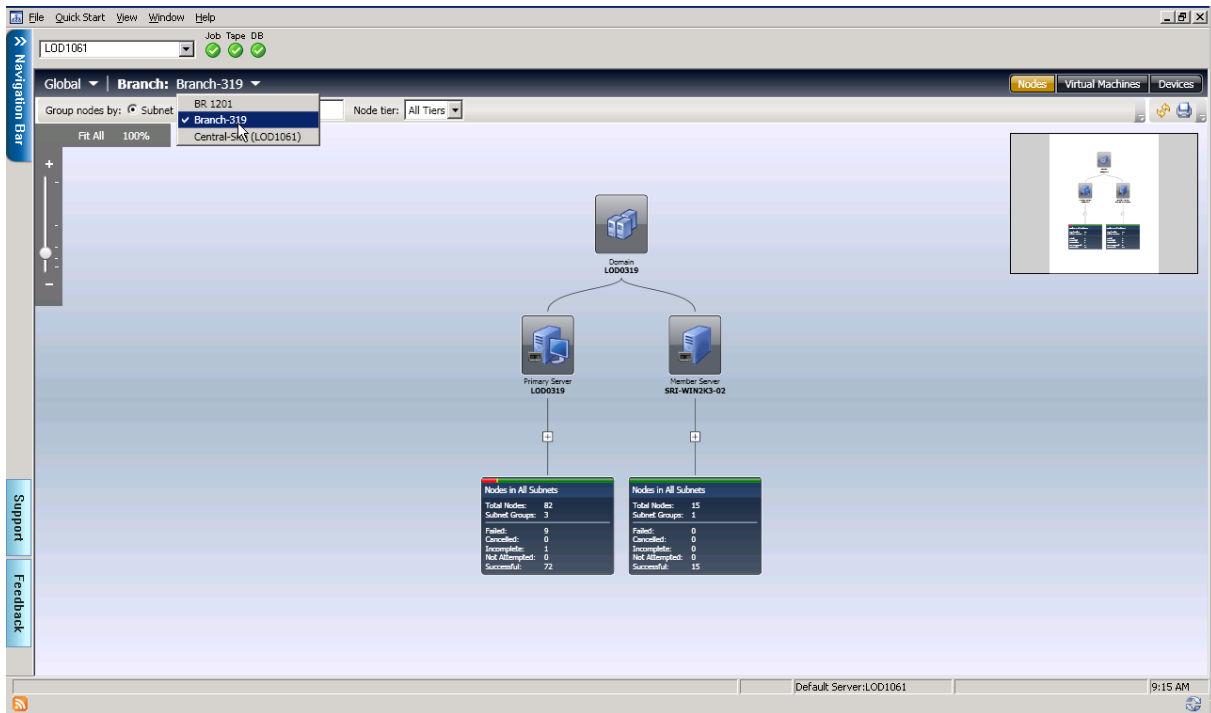
Click a device to open a Details Window that displays more information about the device. For example, for Tape Libraries, the Details Window shows the number of drives and number of slots. Each device type is indicated by icons for file system and deduplication devices, as well as for tape drives and libraries.



Note: For Arcserve Backup running in a cluster environment, the views display the information of the currently active node.

Global Infrastructure Visualization

If you have Global Dashboard installed and configured as the Central Primary server on your backup server, then Infrastructure Visualization lets you specify the display mode. In Global mode, Infrastructure Visualization lets you select an individual branch site and display visualization views for that branch. For more information about configuring the Primary server as a Central Site, see the *Arcserve Backup Dashboard Guide*.



Central Primary Server (or Central Site)

The Central Primary Server (and its associated Arcserve Backup database) is the central hub interface for storing synchronized dashboard-related information received from Branch Primary Servers. Within your Arcserve Backup environment, there can only be one primary server configured as the Central Primary Server, and a Branch Primary Server can only report to one Central Primary Server. All associated Branch Primary Servers need to be registered with this Central Primary Server to enable network communication. Communication is always one way, from a branch site to the central site.

Branch Primary Server (or Branch Site)


Any primary server (or stand-alone server) within your Arcserve Backup environment can be configured to be a Branch Primary Server. A Branch Primary Server synchronizes dashboard-related information to the designated Central Primary Server. All data is transmitted from the Branch Primary Server to the associated Central Primary Server. Within your Arcserve Backup environment, there can be multiple Branch Primary Servers, but only one Central Primary Server. In addition, a Branch Primary Server can only report to one Central Primary Server. After a primary server is configured as a Branch Primary Server and registered with the associated Central Primary Server, the corresponding dashboard data can be automatically synchronized with the Central Primary Server.

Synchronization

Data synchronization is the process of transmitting dashboard-related information from a branch site database to the central site database so that the central database contains (and reports) the same information as each of the registered branch databases. For Global Dashboard, the initial data synchronization will always be full data synchronization. All subsequent data synchronizations will be incremental. Incremental synchronization is synchronizing the data that was modified, deleted, or added since the last synchronization was performed. The synchronized data is compressed to minimize size prior to transmittal.

If the server was configured as a branch site, no Global Mode switch appears on the screen.

Local Mode shows Infrastructure Visualization views for the Central site only.



Note: To check the last update status of data from each branch site, check the Central Manager interface in the Global Dashboard window. For more information, see the *Arcserve Backup Dashboard Guide*.

If the branch site is at version r12.5 and the Central site has been upgraded to the current release, be aware of the following behavior in Global Visualization:

- Device view for the branch is blank.
- Virtual Machine view shows only VMs that were backed up.


How to View Backup Status

Infrastructure Visualization lets you see into your backup environment, showing groups of related items. In Nodes View, you can view the following:

- Group nodes by subnet or by Agent
- View backup status by node
- View backup status by agent

To launch Infrastructure Visualization, click Monitor & Reports from the Navigation Bar and then choose Infrastructure Visualization, which loads with the Nodes by Subnet view by default.

Infrastructure Visualization displays the most recent backup status only. If a particular node is backed up by more than one server, it is displayed only under the server that performed the most recent backup. The [color bar](#) (see page 689) at the top of each group provides status at a glance.



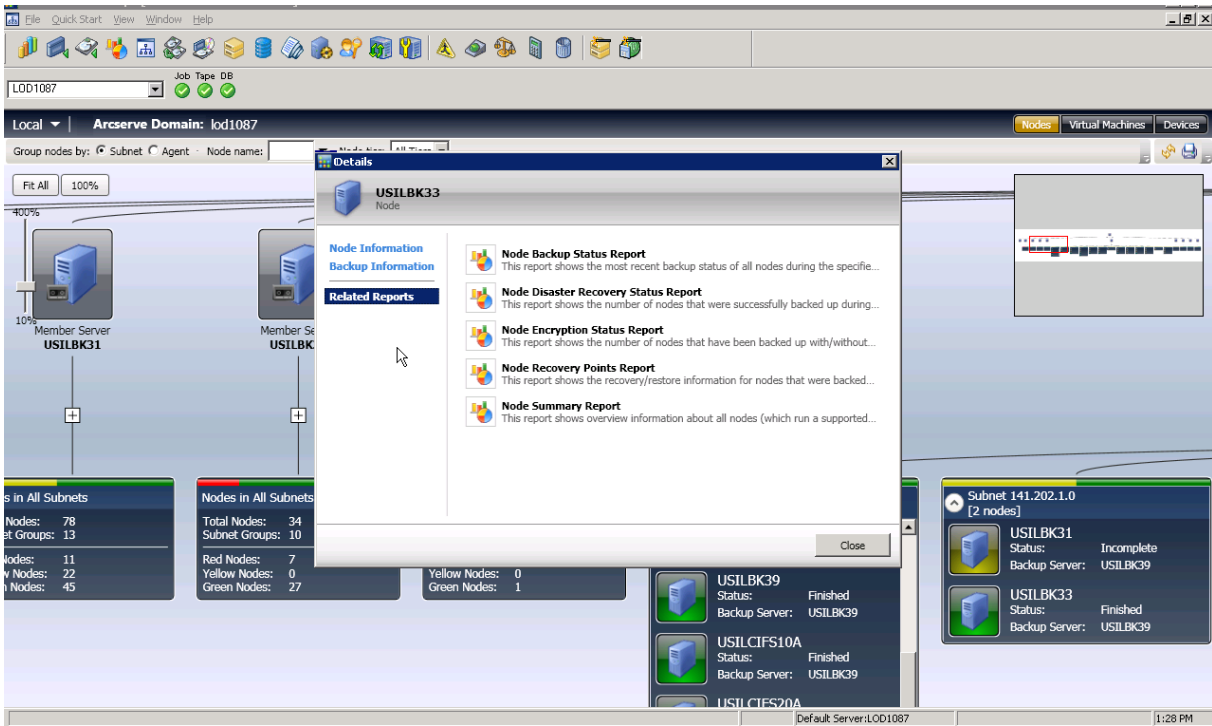
Find the Most Recent Backup Status for Backup Servers

Infrastructure Visualization helps you quickly locate information required for SRM decision-making.

To find the most recent backup status for backup servers

1. Launch the Arcserve Backup Manager and connect to the Primary Server.
2. Launch Infrastructure Visualization from the Navigation Bar.
3. Locate the subnet group for which you wish to obtain backup status. In the bar across the top of the group, you can see the failure percentage of all nodes within this subnet.

4. Click the plus sign to expand a group and view further backup details.
5. Click a server in the group to open the Details screen.



6. Click Related Reports and then choose a report item to launch. From the report, you can determine the causes of reported errors for troubleshooting purposes.

Group Nodes by Subnet or by Agent

When you launch Infrastructure Visualization, backup status is displayed in Nodes view by default. Within Nodes view, you can further control how information is displayed using two methods:

- **Group by Subnet**--Subnet Group shows all servers, all nodes backed up by those servers. By default, nodes are collapsed, but you can expand further to see the list of subnets and corresponding nodes. If a machine has multiple network cards and is part of more than one subnet, it appears multiple times.

Note: When grouping by subnet, a Data Mover node is displayed without an IP address.

- **Group by Agent**--Agent Group shows all servers, all nodes backed up by those servers according to the agents installed on each machine. If a machine has more than one agent installed, it appears multiple times.

Note: Click a node to obtain backup details, including links to Dashboard Reports.

View Backup Status by Node

Each node backed up with Arcserve Backup is shown and grouped by subnet or installed agents. If a node has multiple Network Interface Cards that are part of different subnets, that node is shown in multiple subnet groups.

If you have a UNIX/Linux Data Mover installed, it is displayed as a special backup server, showing only one node. This is because the Data Mover backs up only itself.

Note: Node-based views show the last backup status only. If a machine is backed up by two servers, it appears only under the server that performed the most recent backup.

To view backup status by node

1. Launch Infrastructure Visualization from the Navigation Bar, Monitors & Reports.
Infrastructure Visualization opens in Nodes View (default) with all Arcserve Backup servers shown across the top. The default grouping is by Subnet.
2. The group is not expanded by default. Click the + symbol to expand a group and view further backup details.

The number of nodes is shown, and the name of the backup server and backup status for each node in the subnet are displayed.

View the Backup Status for Virtual Machines

In Virtual Machine View, Infrastructure Visualization displays all VMware Proxy, VMware vCenter Server systems, VMware ESX/ESXi Host systems, and Microsoft Hyper-V systems. Virtual machines are displayed below the servers on which they reside at population or backup time. This could mean that VMware ESX/ESXi Host systems display below a VMware vCenter Server system or a VMware backup proxy system, depending on how it was populated or backed up. VMware vCenter Servers are displayed under the VMware backup proxy system.

If you enter user credentials for a specific VMware vCenter Server system, the Agent for Virtual Machines detects the corresponding ESX/ESXi server, and displays this information for each virtual machine in the Infrastructure Visualization graph.

To view the backup status for virtual machines

1. Launch Infrastructure Visualization.
2. Click Virtual Machines to change views.
3. Select the VMware or Hyper-V option to view the virtual machines of the desired type in your environment.
4. Double-click a node to obtain backup details including Dashboard Reports.

Filter Views by Node Name

The Node Name field is available in both Nodes and Virtual Machines views. Use it when you know the specific node for which you wish to view backup information. You can also search for groups of nodes with similar names using the * wildcard.

To filter views by node name

1. Launch Infrastructure Visualization.
2. Ensure Node View is active.
3. In the View-Specific toolbar, enter the name of the node you wish to view in the Node Name field. For example, PAY locates all nodes whose names contain PAY anywhere in the string.

Filter Views by Node Tier

Arcserve Backup lets you filter the Nodes or Virtual Machines Views by node tier.

To filter views by node tier

1. Launch Infrastructure Visualization.
2. Make sure Nodes view is active.
3. In the View-Specific toolbar, choose a filter from the Node Tier list:
 - High Priority
 - Medium Priority
 - Low Priority

View Devices and SANs in the Environment

Device View lets you see how backup devices are connected to the respective Arcserve Backup server in the Arcserve domain. Tape drives, libraries and devices such as file system devices and deduplication devices are shown with distinct icons that describe the type of device.

Data mover servers can also be displayed in Device View. If devices are connected to a data mover server, they display beneath the data mover server. You can connect only file system devices or SAN devices to a data mover server.

Shared devices initially appear only under the primary server. When all other Arcserve Backup servers come online, devices are displayed in the SAN.

To view devices and SANs in the environment

1. Launch Infrastructure Visualization.
2. Click Devices to change views.
Infrastructure Visualization refreshes and displays all devices connected to Backup servers.
3. (Optional) Click a device to obtain specific device details.

Dashboard Integration with Infrastructure Visualization

Within the context of a selected item in Infrastructure Visualization, you can launch Dashboard reports. For example, clicking a node launches the Details window, which provides more information about the selected node, as well as a list of the related Dashboard reports. Click a report item to open it. Reports are opened within Infrastructure Visualization, but retain the same functionality as if they are opened from Dashboard, including print, save, and email.

The reports accessible from Infrastructure Visualization are as follows:

■ Domain Reports

Note: Reports that are marked with an asterisk (*) indicate that the report is a Storage Resource Management (SRM) type of report. SRM reports let you monitor your entire storage environment at a glance and measure the status of all related resources.

- Node Summary Report *
- Volume Report *
- Disk Report *
- Network Report *
- CPU Report *
- Memory Report *
- OS Report *
- SCSI/Fiber Card Report *
- Agent Distribution Report
- License Report
- Node Tiers Report
- Node Backup Status Report
- Top Nodes with Failed Backups Reports
- Node Whose Most Recent Backup Failed Report
- Job Backup Status Report

- Backup Server Reports (primary servers, member servers, and data mover servers)
 - Backup Data Location Report
 - Job Backup Status Report
- Node and VM Node Reports
 - Node Backup Status Report
 - Node Disaster Recovery Status Report
 - Node Encryption Status Report
 - Node Recovery Points Report
 - Node Summary Report

Chapter 10: Using the Alert Manager

This section contains the following topics:

[How the Alert Manager Works](#) (see page 705)

[Alert Manager Components](#) (see page 707)

[Set Up Alerts](#) (see page 707)

[Alert Manager Configuration](#) (see page 709)

How the Alert Manager Works

Alert Manager is a notification system that sends messages to people in your organization using various methods of communication. For example, you can send alerts to the system administrator, or a hardware technician in or out of the office. You can also send alerts to groups of people in different segments of the network.

The Alert Manager does not generate its own messages. You must configure the manager with the information you want to communicate and where you want to send it. Use the Alert options in the Backup Manager or Alert configuration in Server Admin to tell Alert what information you want to communicate. Use the Alert Manager or the Alert options in the Backup Manager to tell Alert how to send information and who to send it to. For more information about how to select methods and specify recipients from within the Backup Manager, see [Backing Up Data](#) (see page 133).

The information you communicate through Alert is called an Event. Events are words or phrases that appear in the Activity Log. You can select predefined job-related events, such as “Job Completed Successfully” and “Job Incomplete.” You can also customize job-related events, such as error, warning, or notification codes. In addition, you can specify non-job related events, such as starting or stopping the Tape Engine.

You can set up alerts from the following Arcserve Backup managers and utilities:

- Backup Manager
- Restore Manager
- Media Assure & Scan Utility
- Compare Utility
- Purge Utility
- Copy Utility
- Merge Utility
- Count Utility

To select job-related events, open these managers or utilities, click the Options toolbar button, and then select the Alert tab on the Options dialog.

Job-related events can also be accessed by selecting the Utilities menu and choosing any of the Utilities menu options. To select non-job related events, in Server Admin click Config, and then the Alert tab.

After you select events and they appear in the Activity Log, Alert generates notification messages and sends them to the appropriate recipients. For more information about selecting the information you want to communicate using Alert, see [Backing Up Data](#) (see page 133) for job-related events, and [Administering the Backup Server](#) (see page 491) for non-job related events.

Alerts can be sent in the following ways:

- **Broadcasts**--Sends pop-up messages to specific computers.
- **CA Unicenter TNG Option**--Sends messages to the TNG console and WorldView repository.
- **Lotus Notes**--Sends email messages using Lotus Notes.
- **Microsoft Exchange**--Sends email messages using Microsoft Exchange.
- **Windows Event Log**--Places event information in the Event logs of local and remote computers.
- **Pager**--Sends alphanumeric pager messages.

Note: Pager options are not supported on Japanese versions of Arcserve Backup.

- **SMTP (Simple Mail Transfer Protocol)**--Sends email messages using the standard email protocol on the Internet.
- **SNMP (Simple Network Management Protocol)**--Sends messages to SNMP managers, such as HP OpenView and CA Unicenter TNG.
- **Trouble Tickets**--Sends printed documents to any print queue on your network.

Note: When defining Trouble Tickets, the values specified for the Printer Name, User Name, and Password must be less than or equal to 48 bytes.

Alert Manager Components

Alert is comprised of the following components:

- **Alert Manager**--The Alert Manager is used to configure how Alert sends its messages and to whom to send them.
- **Alert Service ([Alert Notification Server] Service)**--This service is responsible for the reception, processing, and distribution of Alert messages.
- **ALBUILD.DLL**--This .DLL acts as the channel between Alert and other applications. This file should be located in the Alert home directory.
- ***.CFG**--The application profile file is provided by an application. This *.CFG file must be present in the Windows directory so that Alert can handle messages generated by an application.

Set Up Alerts

Arcserve Backup provides event-based notification through email, pager, SNMP, broadcast, event log, or through Unicenter Network and Systems Management views. If you have Unicenter installed, you can use its Monitoring Agent to monitor the status of the Arcserve Backup processes and media, and report on the failure of backup jobs.

Example: Alert Notification

You can configure Alert to broadcast a message when a backup job finishes successfully.

To set up Alerts

1. From the Backup Manager window, click the Options toolbar button.
The Options dialog opens.
2. Click the Alert tab.
The Alert options display.

3. Click the Configure button to specify the transmission method.
The Methods & Recipients Configuration dialog appears.
4. In the Methods & Recipients Configuration dialog, click New.
The Configuration Name dialog opens.
5. Enter a name for the configuration in the Configuration Name field and then click OK.
Select the Broadcast method, and click the Add button.
The Add Broadcast Recipient dialog opens.
6. In the Group/Machine field, select your machine from the network, and click Add to add it to the Recipients field.
Or, if you know the machine name, enter the machine name into the recipient field.
Click OK and click OK again to save the configuration.
7. From the Methods & Recipients drop-down menu, select the saved configuration.
8. Select an Event from the Event drop-down menu, and click the Add button.
Now that you have set up Alert, you can proceed with your backup.
Click OK.
Click Submit on the toolbar to submit your job.
The Security and Agent Information dialog appears.
9. From the Security and Agent Information dialog, select the job you want to run.
If the user name and password do not appear, click the Security button and enter the appropriate user name and password.
Review the security information and click OK.
The Submit Job dialog screen opens.
10. Enter a description for your backup job (optional), and click OK to submit the job.
Your job, which is now active, appears on the Job Queue tab in the Job Status window. If the job is active, you can view its status by double-clicking it on the Job Queue tab to invoke the Job Properties dialog.
When the job finishes, Alert notifies you, using the specified method.

Alert Manager Configuration

Before you use the Alert notification system, you must first establish a service account. To do this, open the Alert Manager, go to the Service menu, and select Set Service Account.

Note: If the Alert Manager was previously installed with another Arcserve product, it is not reinstalled to the Arcserve Backup directory; it remains in the directory where it was previously installed.

You can send alerts using many communication mechanisms or applications. Any application that calls Alert specifies one of three event priorities—Critical, Warning, or Informational.

To view a list of the applications that call Alert, open the Alert Manager, and, in the left pane, expand Configuration and then expand Default or Arcserve Backup. You can either use the Alert default settings, which will be used by all applications that use the Alert Service, or you can enter configuration information specifically for each application. If you choose the latter, these configurations override the default Alert configurations.

The following sections describe how to configure each of the available communication mechanisms. To begin, expand Configuration, and then expand Arcserve Backup to view options discussed in the following sections.

Ports Option

The Ports option contains communication port profiles. Pagers and functions that use serial port access use these profiles. To configure, right-click Ports and select New Item. Enter the following information:

- **Port**--The name of the communications port you want the pager message to be broadcast from.
- **Data Bits**--The number of data bits, 7 or 8, which your modem uses.
- **Baud Rate**--The baud rate used by your modem.
- **Parity**--The parity setting, none, odd, or even, of your modem.
- **Stop Bits**--The number of stop bits, 1 or 2, which your modem uses.

If you want these settings to apply to any function that uses serial port access, place a check mark in the Use As Default box. When you are finished configuring port information, click OK.

Note: Pager options are not supported on Japanese versions of Arcserve Backup.

Broadcast Alerts

You can use Alert broadcasts to communicate information to specific network users or groups.

To use Broadcast Alerts, you must enable Windows Messenger services on Windows XP and Windows Server 2003 systems. The Messenger service is disabled by default on Windows XP and Windows Server 2003 systems.

Note: Windows Vista, Windows Server 2008 and later versions do not support Messenger services. As a result, Broadcast Alerts is not supported either of these platforms.

To configure broadcast options, right-click Broadcast and select New Item.

When the Broadcast Recipients page appears, enter or select all machine names in your network that you want to receive alert messages, and then click Add.

Email Notification

You can use Lotus Notes, Microsoft Exchange, or SMTP to send email notification messages to specific users.

Important! You must install Lotus Notes or Microsoft Exchange Client to set up configuration data and to send messages. See your Windows manual for instructions on how to set up your email account.

Lotus Notes

To configure Lotus Notes settings, right-click Lotus Notes and select Lotus Notes Settings. When the Lotus Notes Settings page appears, enter the following information:

- **Lotus Notes Install Path**--Enter the appropriate install path.
- **Password**--Enter your password.
- **Use Specific Account**--If you want Alert to switch to another user ID, place a check mark in this box and enter information in the following fields:
 - **ID File**--For example, joeuser.id
 - **Mail Server**--For example, NotesServer/NotesDomain
 - **Mail File**--For example, mail/joeuser.nsf

After you configure Lotus Notes Settings and right-click Lotus Notes, select New Item or Message Attributes.

If you select New Item, Alert contacts the Lotus Notes server to display the address book. Select the users to whom you want to send alerts.

If you select Message Attributes, you can attach files to the email alert. Enter a subject, click Add File to select the file you want to attach, and then click OK.

Microsoft Exchange

To configure Microsoft Exchange settings, right-click and select one of the following:

- **New Item**--Lets you select email recipients.
- **Message Attributes**--If you select this, you can attach files to the email alert. Enter a subject, click Add File to select the file you want to attach, and then click OK.
- **MS Exchange Settings**--If you select this, the Service Logon Settings dialog appears. This is the same dialog that appears when you set up a service account. Enter the domain, user name, and password you want to use with the Alert Service. Make sure the account and user you enter is an account with Login as Service rights and is also an account on the Microsoft Exchange Server. If you are running the Microsoft Exchange Client, you must also enter the name of the server and mailbox. The mailbox name is case-sensitive and should not be hidden in a folder.

Note: If you are using Microsoft Outlook, right-click your Microsoft Outlook icon and select Properties. Select Microsoft Exchange Server and click Properties to view the server and mailbox information you should enter.

Send Job Logs Via Email

In addition to sending email notification messages, you can also use Lotus Notes or Microsoft Exchange to email job logs. To do this, create a new item and select recipients. Then, in the Backup Manager, before you submit a job, click the Options icon or, from the Backup menu, select Options. When the Global Options dialog appears, click the Alert tab, place a check mark in the Attach Job Log box, and then click OK. After you submit the job, the job log is sent to the recipients you specified.

Windows Event Log Notification

You can configure the event log so that Alert puts an event for a selected server in that machine's event log.

To configure event log configurations, right-click Window Event Log and select New Item. When the Eventlog Recipients dialog appears, enter or select all machine names in your network to which you want to send Alert messages, and then click Add.

Alert Manager Pager Options

Note: Pager options are not supported on Japanese versions of Arcserve Backup.

You can use the Pager option to communicate information using alphanumeric pager messages. Before you can add pager recipients, you must configure the communication ports.

Note: For more information about configuring ports, see [Ports Option](#) (see page 709).

To set up pager configurations, right-click Pager and select New Item. When the Pager Configuration page appears, enter the following information:

- **Owner Name**--Enter the name of the pager recipient.
- **Pager Type**--Select alphanumeric pager. Numeric is not supported.
- **Pager Number**--Enter a maximum of 24 characters. If a digit, such as 9, is needed for a dial tone, you must include it in this field.

Enter a comma to indicate a one second pause. If you want a longer pause, enter a string of commas.

You can use a dash to separate digits, but it has no function. (Check your modem manual because this can vary by modem.)

- **Pager ID**--Enter up to eight digits to identify the pager that will receive the alerts.
- **Site ID**--Enter up to four digits to identify where the alert occurred. This ID is included in the message to the pager; therefore, if the number is less than four digits, use leading zeros.
- **Connection Delay**--Enter the number of seconds you want to wait before a connection is made with the pager company. This will vary with your pager company, location, time of day, telephone equipment, and telephone traffic. If the connection is not established immediately, adding a delay prevents the alert from being sent before the connection is established.
- **Message Delay**--Enter the number of seconds to wait between the time the connection is made and the time the alert message is sent.
- **Port Configuration**--Select the appropriate port configuration. See [Ports Option](#) (see page 709) for information on how to create new port profiles.

Note: When sending an alphanumeric page, consult your paging service for proper modem settings. The Alert service requires the TAP protocol for alphanumeric pages.

Pager Message Options

You can send variations of the messages in the following list to an alphanumeric pager. Substitute the bracketed words with the actual information.

- Boot Virus Detected
- Manager Detected a Virus [*virusname*] in [*path*]
- Infected File [*servername/path*] Detected
- Infected File [*path*] Accessed by user name at workstation address

Note: Pager options are not supported on Japanese versions of Arcserve Backup.

Configure SMTP Notification

You can use SMTP to send email messages to recipients using the Internet.

Follow these steps:

1. Right-click SMTP and select SMTP Settings to open the SMTP Settings screen.
2. Complete the following fields on the SMTP Settings screen:
 - **Sender Address**--The sender name that you want the alert email to display in mailboxes.
 - **Subject**--The subject that you want the alert email to display in mailboxes.
 - **SMTP Server**--The name of the SMTP server. For example, mail.yourcompany.com.

Click OK.

3. After you configure the SMTP Settings, right-click SMTP, select New Item to open the SMTP Recipients screen.
4. Complete the following fields on the SMTP Recipients screen:
 - **Address**--The email address for the recipient. For example, johnsmith@bigcompany.com.
 - **Display Name**--The name of the recipient.

Click OK.

SNMP Notification

You can use SNMP to send an SNMP trap to an SNMP manager. Examples of SNMP managers include HP OpenView, IBM NetView, and CA Unicenter TNG.

To configure SNMP settings, right-click SNMP and select New Item. When the SNMP Recipient page appears, enter the following information:

- **Manager Name**--Enter the name of the SNMP Manager.
- **Send Via**--Select one of the following options:
 - IPX--If you select this, enter the 8-byte network address of the machine where the SNMP manager is located. Next, enter the 12-byte node address of the machine where the SNMP manager is located. Use this field for Novell networks.
 - IP--If you select this, enter the IP address of the machine where the SNMP manager is located. Use this field if you are running the TCP/IP stack.

Trouble Tickets

You can use Trouble Tickets to communicate information through printed documents.

To configure Trouble Ticket settings, right-click Trouble Ticket and select New Item. When the Trouble Ticket Recipients page appears, enter the following information:

- **Company**--Enter the name of your company.
- **Location**--Enter the appropriate location information.
- **Header**--Enter the information that will appear at the top of each Trouble Ticket.

To select recipients, highlight a printer and click Add. When prompted, enter a user name and password to connect to the printer device.

In addition to using Trouble Tickets to send printed notification messages, you can also use Trouble Tickets to send job logs. To do this, create a new item and select recipients. Then, in the Backup Manager, before you submit a job, click the Options icon or, from the Backup menu, select Options. When the Global Options dialog appears, click the Alert tab, place a check mark in the Attach Job Log box, and then click OK. After you submit the job, the job log is sent to the recipients you specified.

Event Priorities

All applications calling Alert specify one of the following event priorities:

- Critical
- Warning
- Informational

Message Testing

To test any of the Alert messaging functions, from the toolbar, select Send Test Message. You should test each setting after you configure it.

To avoid unnecessary alarm, inform Alert recipients that you are performing a test.

Alert Activity Details

To review alert activity, expand the Activity group and select one of the following:

- **Alert Summary**--Displays the status of Alert.
- **Alert Event Log**--Stores every message that Alert generates. It displays the date and time a particular event occurred, the applications that sent the alert, and the application that generated the event.
- **Alert Activity Log**--Stores a historical listing of alerts.

You can view, print, or clear these logs.

Chapter 11: Managing Agents Using Central Agent Admin

This section contains the following topics:

[How the Arcserve Backup Central Agent Admin Works](#) (see page 717)

[Manage Agents](#) (see page 718)

[Configure Agents](#) (see page 719)

[Add Computers](#) (see page 721)

[Add Nodes](#) (see page 721)

[Manage Agent Logs](#) (see page 723)

[Configure SRM PKI](#) (see page 725)

[Configure SRM Exclude Paths](#) (see page 726)

[Configure Node Tiers](#) (see page 727)

How the Arcserve Backup Central Agent Admin Works

The Central Agent Admin is a central utility that manages the agent machine and lets you view agent logs and event logs, set debug level registry entries for one or more agents, and configure agent options. Using the Central Agent Admin, you can also perform basic node management tasks such as modifying agents. You can add or modify node security information without opening the Backup Manager. The Central Agent Admin also lets you perform Node Tier configuration and Agent Deployment.

The Central Agent Admin is part of the Arcserve Manager. You can open the Central Agent Admin from the Administration menu or from the Arcserve home page.

Note: The old Agent Admin will still be installed with Client Agents and be used for local configuration with same functions as before.

When you start the Central Agent Admin, it retrieves the information for all registered agent nodes and displays these agents in a tree view. The supported agents installed in that machine display when you expand the agent node. You can also view the agent properties and configuration information from the registry of the remote agent machine in the top and bottom right-hand panes.

Note: The Central Agent Admin currently supports the Client Agents, the Agent for Open Files, the Agent for SQL Server, the Agent for Microsoft Exchange Server, the Agent for Microsoft SharePoint Server, and the Agent for Oracle.

Manage Agents

Arcserve Backup Central Agent Admin lets you perform agent management tasks such as modifying agent information, configuring agents, and managing agent services.

Modify Agents

Arcserve Backup Central Agent Admin lets you add, modify, or delete agents, similar to the Backup Manager.

To modify agents

1. Open the Arcserve Backup Manager Console. From the Quick Start Menu, select Administration and click Central Agent Admin.

The Central Agent Admin window opens.

2. Click the Windows Systems object and select a remote machine.
3. Right-click the remote machine and select Modify Agent.

The Modify Agent dialog appears.

4. Enter the agent details such as host name and IP address.
5. Click OK to confirm your changes.

Configure Agent Security

Arcserve Backup Central Agent Admin lets you configure agent security similar to the Backup Manager.

To configure agent security

1. Open the Arcserve Backup Manager Console. From the Quick Start Menu, select Administration and click Central Agent Admin.

The Central Agent Admin window opens.

2. Click the Windows Systems object and select a remote machine.
3. Right-click the remote machine and select Security.

The Security dialog appears.

4. Enter the user name and password.
5. (Optional) Select or deselect one or more machines for which you want to apply or remove the same security settings.
6. Click OK to complete agent security configuration.

Start or Stop Agent Services

Using Arcserve Backup Central Agent Admin, you can start or stop the agent services.

To start or stop agent services

1. Open the Arcserve Backup Manager Console. From the Quick Start Menu, select Administration and click Central Agent Admin.
The Central Agent Admin window opens.
2. Click the Windows Systems object and select a remote machine.
3. Right-click the remote machine and select Start/Stop Services.
The Backup Agent Service Manager dialog appears.
4. Click Start Service or Stop Service to start or stop the agent services, respectively.
5. (Optional) Select Start the agent backup service as the system starts, to ensure that the service starts as soon as the system starts.

Start Agent Deployment from the Central Agent Admin

Arcserve Backup Central Agent Admin lets you deploy Arcserve Backup agents to remote systems using Agent Deployment.

To start Agent Deployment from the Central Agent Admin

1. Open the Arcserve Backup Manager Console.
From the Quick Start menu, select Administration and click Central Agent Admin.
The Central Agent Admin manager window opens.
2. Expand the Windows Systems object.
Locate the remote system.
Right-click the remote system and click Agent Deployment on the pop-up menu.
Agent Deployment starts.

Note: For more information, see Arcserve Backup Agent Deployment.

Configure Agents

Arcserve Backup Central Agent Admin lets you configure the following Arcserve Backup agents from a central location:

- Client Agent for Windows
- Agent for Microsoft SQL Server

To configure agents

1. Open the Arcserve Backup Manager Console. From the Quick Start Menu, select Administration and click Central Agent Admin.

The Central Agent Admin window opens.

2. Select an agent on the server.

Right-click an agent and select Configuration.

The Configuration dialog appears with a list of agents that you can configure.

3. Select the agent you want to configure and update the settings.
4. (Optional). Click Apply to Multiple to apply the same settings to multiple agent machines.
5. Click OK to complete the agent configuration.

To set debug level registry settings

1. Select an agent on the server.

2. Right-click an agent and select Set Debug Level on the pop-up menu.

The Set Debug Level dialog opens.

3. Set the appropriate debug level such as Normal, Detail, Debug, or Trace, and click OK.

The debug level registry settings for the agent is now complete.

Add Computers

Arcserve Backup Central Agent Admin lets you add one or more remote computers in a manner that is similar to the Backup Manager.


To add computers

1. Open the Arcserve Backup Manager Console. From the Quick Start Menu, select Administration and click Central Agent Admin.

The Central Agent Admin window opens.

2. Click the Windows Systems object and select Add Machine/Object.

The Add Agent dialog opens.



3. Complete the required fields on the Add Agent dialog and click Add.

You can now view the added computers in the left pane of the Central Agent Admin.

Add Nodes

You can use the Add, Import, and Export Nodes feature to add multiple nodes and agents into the system in either of the following ways:

To add multiple nodes and agents using the user interface

1. Open the Arcserve Backup Manager Console. From the Quick Start Menu, select Administration and click Central Agent Admin.

The Central Agent Admin window opens.

2. Right-click the Windows Systems object and select Add/Import/Export Nodes.

The Add/Import/Export Nodes dialog appears.

3. Enter the name of a node you want to add and click Add. You can also select one or more nodes from the left-pane list, and click Add or Add All.
4. (Optional) Select any node on the left-pane list and click Properties.

The Server Properties dialog appears showing the server details and list of products installed on that server. Click OK.

5. (Optional) Select any node on the right-pane list and click Security.

The Security dialog appears where you can set the user a user name and password for the node. You can also apply the same user name password to multiple nodes. Click OK.

6. Click OK.

You can now view the added nodes and agents in the Central Agent Admin.

To add multiple nodes and agents using a .csv and .txt file

1. Right-click the Windows Systems object and select Add/Import/Export Nodes.

The Add/Import/Export Nodes dialog appears.

2. Click Import and browse to a location containing .csv or .txt files.

3. Specify the name of the .csv or .txt file from the user interface.

The node and agent names are imported from the .csv or .txt file and are added into the system.

4. Click OK.

You can now view the added nodes and agents in the Central Agent Admin.

Manage Agent Logs

Arcserve Backup Central Agent Admin lets you view, export, or delete Agent logs.

Follow these steps:

1. Open the Arcserve Backup Manager Console. From the Quick Start Menu, select Administration and click Central Agent Admin.

The Central Agent Admin window opens.

2. Select an agent on the server.

The agent log file details (for example, log file name, size, type of agent, and so on), display in the top right-hand pane.

3. Right-click an agent log file and click View Log Within Specific Time Range from the pop-up menu.

The Log Retrieval Configuration dialog appears.

4. Select one of the following and click OK:

- **Retrieve full log file**--Obtains the complete log file information.

Note: Windows Server 2008 and later systems do not support viewing event logs directly from the Arcserve Backup Manager Console. To view event logs on Windows Server 2008 and later systems, you must export the event log files and then open the exported documents using a text editor such as Notepad.

- **Retrieve log file against error time**--Obtains log file information for the specified start and end times.

The agent log file opens in a text editor such as Notepad.

To export agent logs

1. Select an agent on the server.

The agent log file details (for example, log file name, size, type of agent, and so on) display in the top right-hand pane.

2. Right-click an agent log file and click Export Log To File.

The Log Retrieval Configuration dialog appears.

3. Select one of the following and click OK:

- **Retrieve full log file**--Obtains the complete log file information.
- **Retrieve log file against error time**--Obtains log file information for the specified start and end times.

The Save As dialog appears.

4. Specify a destination folder where you want to export or save the log file, and click OK.

The agent log file is exported to the specified location.

To delete agent logs

1. Select an agent on the server.

The agent log file details (for example, log file name, size, type of agents, and so on) display in the top right-hand pane.

2. Right-click the agent log file you want to delete and click Delete Selected Log.
3. Confirm whether you want to delete the agent log.

The agent log file is then deleted from the agent logs list.

Configure SRM PKI

Arcserve Backup Central Agent Admin contains a utility named SRM PKI. SRM PKI (performance key indicators) lets you monitor the performance of the agents running in your backup environment.

SRM PKI measures the following performance indicators:

- CPU usage
- Memory usage
- Disk throughput
- Network input and output

Arcserve Backup lets you enable or disable SRM PKI, specify default or custom values for the indicators, and generate alert messages when the indicators exceed values that you specified.

To configure SRM PKI

1. Open the Arcserve Backup Manager Console. From the Quick Start Menu, select Administration and click Central Agent Admin.
The Central Agent Admin window opens.
2. Right-click the Windows Systems node and select Configure SRM PKI.
The Configure SRM PKI dialog opens and the Policy tab appears.
3. For each agent, specify the options that you require:
 - **Use Default Policy**--Lets you specify the default values for each performance indicator. To specify custom values for the indicators, clear the checkmark next to Use Default Policy.
Note: You can view the status of modified threshold values for each agent by clicking the Broadcasting Status tab.
 - **Enable PKI**--Lets Arcserve Backup agents send hourly PKI values to the primary server for SRM PKI reports.
Note: For more information about SRM PKI reports, see the *Dashboard User Guide*.
 - **Enable Alert**--Lets Arcserve Backup generate alert messages in the Alert Manager when the performance of an agent exceeds your predefined PKI values.
4. Make any necessary changes for one or more listed agents in the Policy tab.

5. (Optional) Click Apply to Multiple to apply the same configuration settings to multiple agents.

When you click Apply to Multiple, the Apply to Multiple dialog opens.

To apply the same settings to multiple agents, select the individual agents, click Select All, or click Unselect All, and then click OK.

6. On the Configure SRM PKI dialog, click Apply, and then click OK.

The Configure SRM PKI dialog closes and the PKI values are applied.

Configure SRM Exclude Paths

Arcserve Backup Central Agent Admin lets you set SRM exclude paths.

To configure SRM exclude paths

1. Open the Arcserve Backup Manager Console. From the Quick Start Menu, select Administration and click Central Agent Admin.

The Central Agent Admin window opens.

2. Right-click the Windows System object and select Set SRM Exclude Path.

The Set SRM Exclude Path dialog appears.

3. Enter the SRM exclude path and click OK.
4. (Optional) You can add or delete one or more paths using Add or Delete.

Note: The Top Nodes with Most Unchanged Files dashboard report uses the SRM exclude path list to determine which files must be excluded when you generate the report. The report excludes all files in the SRM exclude paths that you specify in the Central Agent Admin.

Configure Node Tiers

You can use the Arcserve Backup Server Admin or the Central Agent Admin to change the assigned priority classifications of your Arcserve Backup nodes. These tiers are used to filter the information displayed on the Arcserve Backup Dashboard by the priority level of the monitored nodes.

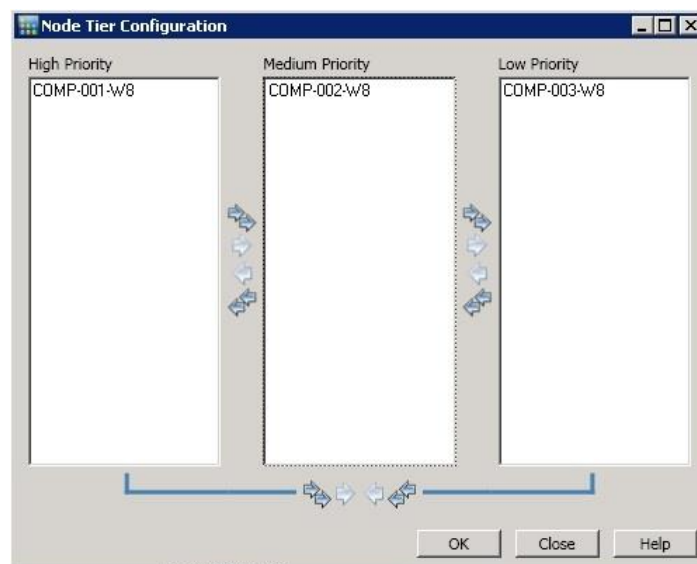
The Node Tier Configuration dialog contains three priority categories (High Priority, Medium Priority, and Low Priority), and is automatically populated when a node is added to your system and browsed. By default, a High Priority tier is configured to include all Arcserve Backup servers (Primary and Member) and any nodes with Arcserve Backup application agents installed (such as Oracle, Microsoft Exchange Server, Microsoft SQL Server, Microsoft Sharepoint Server, and so on), and a Low Priority tier is configured to include all other nodes (having file system agents installed). The Medium Priority tier is not configured to include any nodes, and is available for customized use.

The node assignments for each tier can be reconfigured and customized to meet your individual needs by using the Node Tier Configuration dialog, which is accessed from the Arcserve Backup Server Admin or from the Backup Manager (right-click Windows Systems in Source tab) or from the Central Agent Admin (right-click Windows Systems).

To configure node tiers

1. Right-click the Windows Systems object, and select Node Tier Configuration.

The Node Tier Configuration dialog opens, displaying the nodes assigned to each Tier category (High Priority, Medium Priority, Low Priority).



2. Select one or more nodes that you want to reassign to a different tier category and click on the corresponding arrow icon to move the selected nodes from one tier to another.

Note: Multiple nodes can be selected for tier assignment by using the "CTRL" or "SHIFT" key combinations.

- Click the single arrow icon to move only the selected nodes.
- Click the double arrow icon to move all the node in the tier.

3. Click OK when done.

The node tier assignments have been changed to meet your individual needs.

To view connections under the local computer

4. Select the computer from where you can open the Arcserve Backup Manager Console under the Windows Systems object and expand the computer to view the details.
5. Click Connections.

The connection details of that computer appear on the right pane.

To configure debug levels

1. Select any computer under the Windows Systems object and expand the computer to view the details.
2. Right-click the Universal Agent and select Set Debug Level from the pop-up menu.
The Configure Debug Level dialog appears.
3. Select a debug level such as Normal, Detail, Debug, or Trace, and click OK.
You have now configured the debug level for that computer.

To enable or disable SRM client

1. Select any computer under the Windows Systems object and expand the computer to view the details.
2. Right-click the Universal Agent and select Disable SRM Client to disable the SRM client. If the SRM Client is disabled, select Enable SRM Client to enable to it.
3. Click OK to confirm that you want to enable or disable the SRM client.

Chapter 12: Using Deduplication

This section contains the following topics:

[How Data Deduplication Works](#) (see page 729)

[How to Plan a Deduplication Installation](#) (see page 731)

[Deduplication Considerations](#) (see page 732)

[Create Data Deduplication Devices](#) (see page 735)

[Deduplication Device Group Configuration](#) (see page 738)

[Device Commands for Data Deduplication Devices](#) (see page 738)

[Back up Data with Deduplication](#) (see page 738)

[Recover Deduplicated Data](#) (see page 751)

[Deduplication Reports](#) (see page 757)

How Data Deduplication Works


Data deduplication is technology that allows you to fit more backups on the same physical media, retain backups for longer periods of time, and speed up data recovery. Deduplication analyzes data streams sent to be backed up, looking for duplicate "chunks." It saves only unique chunks to disk. Duplicates are tracked in special index files.

In Arcserve Backup, deduplication is an in-line process that occurs at the backup server, within a single session. To identify redundancy between the backup jobs performed on the root directories of two different computers, use [global deduplication](#) (see page 750).

During the first backup:

- Arcserve Backup scans incoming data and segments it into chunks. This process occurs in the SIS layer of the Tape Engine.
- Arcserve Backup executes a hashing algorithm that assigns a unique value to each chunk of data and saves those values to a hash file.
- Arcserve Backup compares hash values. When duplicates are found, data is written to disk only once, and a reference is added to a reference file pointing back to the storage location of the first identified instance of that data chunk.

In the diagram below, the disk space needed to backup this data stream is smaller in a deduplication backup job than in a regular backup job.



With deduplication, three files are created for every backup session:

- **Index Files (Metadata files)**
 - **Hash files**--store the markers assigned to each redundant chunk of data.
 - **Reference files**--count hashes and store the address in the data files that correspond to each hash.
- **Data files**--store the unique instances of the data you backed up.

The two index files together consume a small percentage of the total data store so the size of the drive that stores these files is not as critical as its speed. Consider a solid state disk or similar device with excellent seek times for this purpose.

During subsequent backups:

- Arcserve Backup scans incoming data and breaks it into chunks.
- Arcserve Backup executes the hashing algorithm to assign hash values.
- Arcserve Backup compares new hash values to previous values, looking for duplicates. When duplicates are found, data is not written to disk. Instead, the reference file is updated with the storage location of the original instance of the data chunk.

Note: Use Optimization for better throughputs and decreased CPU usage. With Optimization enabled, Arcserve Backup scans file attributes, looking for changes at the file header level. If no changes were made, the hashing algorithm is not executed on those files and the files are not copied to disk. The hashing algorithm runs only on files changed since the last backup. To enable Optimization, select the Allow optimization in Deduplication Backups option located on the Deduplication Group Configuration screen. Optimization is supported on Windows volumes only. It is not supported for stream-based backups, such as SQL VDI, Exchange DB level, Oracle, and VMware Image level backups.

When you must restore deduplicated data, Arcserve Backup refers to the index files to first identify and then find each chunk of data needed to reassemble the original data stream.

How to Plan a Deduplication Installation

Data deduplication happens on the Arcserve Backup Server, so it works with all Arcserve Backup Agents running in your environment. However, you must upgrade any Arcserve Backup Windows and UNIX/Linux and Mac agents to r12.5. (AS400 and Open VMS agents prior to this release need not be upgraded.)

To deduplicate data during a backup job, set up the job as usual and select a properly configured deduplication device as the backup destination, or as the staging location in a disk to disk to tape backup job. To configure deduplication devices, refer to the topic, [Deduplication Device Management](#) (see page 395). To assist you as you determine where to add deduplication device groups, consider the following:

How often does the data you back up change?

Consider deduplicating data that remains relatively stable between backups. The less data changes between backups, the greater the incidence of identifying duplicates.

How long should backup images be retained?

Consider deduplicating data that must be retained for long time periods. Deduplication fits more backups onto the same physical media.

What type of data is suitable for deduplication?

There is no limitation on data type.

How large is your data size?

Huge backup data streams are good candidates for deduplication.

What is your backup window?

Deduplication happens on the backup server, which means data is transported over the network and then deduplicated.

What are the system requirements for backup servers when performing deduplication backup jobs?

The answer to this question depends on how much data you need to back up, with approximately 110MB of data per backup stream required. The following are suggested guidelines:

For less than 500 GB, 1 CPU

For 500 GB to 2 TB, 2 CPUs

For greater than 2TB, 2 dual core CPUs

Example: How to Plan a Deduplication Installation

Suppose you backup 10 TB to a 25 TB disk, which means you can store a full backup for just one week. Using data deduplication, your first full backup might require only 8 TB of space. However, subsequent backups performed with data deduplication might require only as much as 800 GB (about 10% of its former space requirements). You would then be able to store about 20 full backups - about 5 months of backups - on the same disk.

Using this example, you can retain backup images:

- 2 weeks without deduplication
- 20 weeks with deduplication

Deduplication Considerations

Some data deduplication characteristics and considerations are as follows:

- You can specify a data deduplication device as the destination in a regular backup job.
- You can specify a data deduplication device as the staging device, the final destination device, or both. However, you cannot choose the same deduplication device for both staging and final destinations.

- You can specify different retention schedules for different jobs that all use the same deduplication device.
- You can Optimize data deduplication to improve throughput by deduplicating only the files that have changed since the last backup, except for stream-based files, such as SQL, SharePoint, Exchange, and Oracle data, which cannot be optimized. Optimization is enabled by default.
- You can create deduplication devices only on NTFS volumes.
- Deduplication groups are excluded from jobs that use * groups.
- You cannot use Encryption or Compression with deduplication devices.
- You can specify a purge policy for final destination when using a deduplication device. This is not possible using a normal FSD.
- You can specify a GFS rotation to a deduplication device where all full and incremental/differential backups are submitted to the same device, whereas GFS jobs to an FSD create daily, weekly, and monthly media.
- Due to the manner in which AS400 backup session header data is populated, the deduplication process cannot detect duplicate AS400 backup sessions and deduplicate the redundant backup sessions.
- Due to the manner in which Oracle RMAN backup session header data is populated, the deduplication process cannot detect duplicate Oracle RMAN backup sessions and deduplicate the redundant sessions. However, the Global Deduplication process lets Arcserve Backup examine and handle Oracle RMAN backup sessions. For more information, see [Global Deduplication](#) (see page 750).
- As a best practice, when backing up data on servers running Window Server 2012 related to agent backups, (for example, the Agent for Oracle and so on), enable NTFS Deduplication when you discover that the backup data compression ratios are lower than expected. When used with Arcserve deduplication, NTFS Deduplication can provide higher compression ratios of the backup data.

However, when Arcserve backup compression ratios are relatively high (for example, greater than 70%), using NTFS Deduplication with Arcserve deduplication may not increase the overall backup compression ratios.

Supported Functions Matrix

The following table shows what functions are supported with Data Deduplication.

Function	Supported	Not Supported
Compression ¹		X
Device Format	X	
Device Erase	X	

Deduplication Considerations

Function	Supported	Not Supported
Deduplication in Windows, UNIX/Linux and Mac agents prior to r12.5		X
Deduplication in AS400 and Open VMS agents prior to r12.5	X	
Deduplication in Windows, UNIX/Linux and Mac at r12.5 and later	X	
Encryption ²		X
Image Backup	X	
Migration (Copy Policy)	X	
Maximum Threshold	X	
Minimum Threshold		X
Multistreaming	X	
Multiple Concurrent Streams	X	
Multiplexing ³		X
Optimization in Deduplication	X	
Retention of Staging (Purge Policy)	X	
Scan Jobs	X	
SnapLock		X
Used by jobs using * groups		X
Used in media pools		X
Used in GFS Rotations	X	
Used as Staging Location	X	
Used as Final Destination Location	X	

¹ Compression at agent or server is not supported

² Encryption at agent or server is not supported

³ You could use multiple concurrent streams instead of multiplexing

Licensing Requirements for Deduplication

No additional license is required to perform data deduplication because functionality is built into the Arcserve Backup base product. However, you should consider the following:

- Deduplication devices can be used in disk to disk to tape or disk to tape to tape (staging) operations. However, to use the staging feature with more than two streams of backup data, you must license the [Arcserve Backup Enterprise Module](#) (see page 40).
- For deduplication, you should upgrade your Windows, UNIX, Linux, and MAC client agents to Arcserve Backup r12.5 or later.

Note: The client agents for AS400 and Open VMS do not require an upgrade.

- To back up deduplication device files, you must license the Arcserve Backup Agent for Open Files. On Windows 2003, 2008, and 2012 systems, only the license for the Agent for Open Files is required; the Agent does not need to be installed.

For more information about protecting the deduplication device itself, see [How to Back Up Deduplication Devices](#) (see page 743).

Create Data Deduplication Devices

To deduplicate data, create and choose a deduplication device group as the backup destination. When you create new deduplication devices, Arcserve Backup automatically assigns each device to a new deduplication device group.

You can create deduplication devices locally or remotely. When creating remote deduplication devices, you must manually specify security credentials by clicking the Security button on the Device Configuration dialog, as instructed in the following procedure, otherwise Arcserve Backup attempts to use the system account.

From Device Configuration, you can add one or more devices. Device Configuration verifies the validity of information specified for all devices and alerts you if a particular device failed verification.

To create data deduplication devices

1. Open the Arcserve Backup Manager Console.
From the Navigation Bar, expand Administration and click Device Configuration.
The Device Configuration screen opens.
2. Select Disk-based Devices and click Next.
The Login Server screen opens.

3. Specify the primary server name, authentication type, user name and password and then click Next.
4. Specify the server on which the deduplication device is to be created and click Next. For local servers (default), you may browse and select a path. If you wish to specify a remote server, you must have administrator rights to that server and must manually type the path.

The Disk-based Devices Configuration dialog opens.

5. Click Add to access the Deduplication Devices list.
 - Click the entry in the Device Name column to edit it or accept the default.
 - Click the entry in the Description column to edit it or accept the default.
 - Click the entry in the Data File Location column to specify a path.

Note: You can specify a path manually or browse for an existing path. To enter a location for a remote Data File, you must specify the Machine Name or IP Address, followed by the Share Name. Use the following format:

`\\MachineName\ShareName` or `\\IPAddress\ShareName`

- Click the entry in the Index File Location column to specify a path. To avoid errors, use the format specified for the Data File Location to enter a remote location, or click the arrow to browse for an existing path.
- Click the Group Name column and provide a name. If you leave this blank, a name is automatically provided that you can change in Group Configuration. This is the name you will select when submitting backup jobs that use deduplication devices.

Be aware of the following:

- The Data File Location and Index File Location fields are empty, by default. When you create the deduplication device, Arcserve Backup can create the path that you specify, if it does not exist, or you can browse for a path that exists.
- You should specify different paths for the Index File Location and Data File Location that are on NTFS volumes, and the locations should not contain data from other applications. To help you ensure the best performance, the Index File Location should reside on a disk with a fast seek time, such as a solid state disk.
- You do not need to specify user credentials when the index file location and the data file location reside on the local computer.
- If the index file and the data file reside on different remote computers, ensure that Arcserve Backup can access the remote computers using the same set of credentials. You must use this approach because Arcserve Backup lets you specify only one set of credentials when configuring data deduplication devices.

- To reduce fragmentation for the first full backup job:
 - At the beginning of the backup, the Tape Engine pre-allocates 1 GB (configurable in the registry) to the data file.
 - Before the backup is done and you reach the end of the data file, increase the size of the data file by allocating another 1 GB.
 - The last 1 GB chunk of data is rounded down so that the data file occupies the actual compressed session size after the deduplication process is complete.

Repeat this step as necessary to add more devices.

Note: Arcserve Backup supports configuring an aggregate total of 255 FSDs and DDDs (only if the number of physical devices configured is 0).

6. (Optional) If you specified remote path locations, click Security to provide login credentials.

The Security dialog opens.

Note: You must clear the Use Arcserve System Account (default) option to enable the security fields.

Complete the required fields on the Security dialog and click OK.

The Security dialog closes.

7. Click Next on the Disk-based Devices Configuration dialog to continue.

Arcserve Backup verifies the information specified for all devices in the list. If information is valid, the deduplication devices are added to the list. If any information is not valid, failed devices in the list are marked with a red Failed status. Click the corresponding Failed status to determine the cause of an individual error and resolve it. When all devices pass verification, a summary screen is displayed.

8. Click Next to return to the Welcome to Device Configuration screen or click Exit to leave Device Configuration.

Important! When you create a deduplication device, the purge policy is automatically set to four weeks. That default purge policy is inherited by every job you set up for the device. If you wish to retain backups longer than four weeks, you must adjust the purge time when you submit the backup job.

More information:

[Specify Copy and Purge Policies for Disk Staging Backups](#) (see page 213)

Deduplication Device Group Configuration

Data deduplication devices must be assigned to groups. If you do not specify your own group, a new default group is created and the deduplication device is automatically assigned to it at creation. You may not assign more than one deduplication device to the same group.

You may rename a deduplication group, remove a deduplication device from a group, or assign a deduplication device to an empty group.

You cannot convert a deduplication group to a staging group, nor can you convert a staging group to a deduplication group.

The following are some key distinctions between a staging group and a data deduplication device group:

- A staging group cannot be formatted or erased. A deduplication group can be formatted or erased.
- A staging group cannot be used as a backup destination. A deduplication group can be used as a backup destination.

Device Commands for Data Deduplication Devices

The device commands that are available for data deduplication devices are:

- **Format**--Deletes the sessions from that device and rewrites the header file with a new tape name
- **Erase**--Deletes the sessions and writes a blank header file on the device

Back up Data with Deduplication

You can back up and deduplicate data in two ways:

- **Regular backup job** -- Select a deduplication device group as the backup destination.
- **Staging backup job** -- Select a deduplication device group as the staging location, the final backup destination, or both provided you do not select the same deduplication device group.

How Regular Backup Jobs Work with Deduplication

Deduplicating data during a backup job operates much like a normal backup job, except you must select a deduplication device group as the backup destination.

- Choose Deduplication backup from the Backup Manager Start tab.
- Specify local backup options as usual except for Compression and Encryption options. Deduplication does not support Compression and Encryption. If Arcserve Backup detects an encrypted session, deduplication is skipped and the job proceeds as a normal backup job. Refer to the section, [Compression and Encryption with Deduplication](#) (see page 742) for more information.
- Select a backup source.
- Choose a deduplication device as the backup destination for a regular backup job. For more information, see [Deduplication Device Management](#). (see page 395)
- Set up a schedule, including GFS rotation, if desired. For more information, see [GFS Rotation Jobs on Deduplication Devices](#) (see page 754).
- Specify a purge policy. For more information, see [Considerations for Specifying Deduplication Device Copy and Purge Policies](#) (see page 742).

Note: For information about submitting backup jobs, see [Submit a Backup Job](#) (see page 136),

Configure Deduplication Groups to Use Staging

Regular FSD groups can be configured for staging using the Configure Staging Groups option in the device properties section of Backup Manager. This option does not apply to deduplication device groups.

Data deduplication devices can be configured for staging using only the following procedure.

To configure data deduplication device groups to use staging

1. Open the Arcserve Backup Manager Console.
From the Protection & Recovery menu in the Navigation Bar, click Backup.
The Backup Manager opens.
2. From Start tab, click Deduplication Backup and Enable Staging.
The Staging Location and Policies tabs appear in the Backup Manager.
3. Click the Staging Location tab and expand the Staging Servers object.
Browse to and select the deduplication group you wish to choose as the staging group for this backup job.

4. Click the Policies tab to specify Deduplication Staging policies.

Specify the staging policies for full, differential and incremental backups that you require for the job.

More information:

[Backup Staging Methods](#) (see page 200)

How Staging Jobs Work with Deduplication

In a disk to disk to tape operation, you may specify a deduplication device group as the staging location, the final backup destination, or both, provided the *same* deduplication device group is not selected on both tabs.

- On the Staging Location tab, select the deduplication device group, enable staging and specify a staging policy.
- On the Destination tab, select a different deduplication device group and specify a purge policy. If you do not specify a purge policy, the default value of 4 weeks is inherited from deduplication device creation for full backups, and 2 weeks for incremental/differential backups.
- On the Schedule tab, set up the rotation or GFS schedule, if desired.

For more information, see [How to Submit a Disk Staging Backup Job](#).

Back up Data with Deduplication in a Staging Backup Job

You can select to deduplicate data during the staging phase, the migration phase, or both phases of a disk staging backup job by selecting deduplication device groups on the appropriate tabs.

To back up data with deduplication in a staging backup job

1. Open the Backup Manager and click the Start tab.
From the Start tab, click Deduplication Backup and Enable Staging.
The Staging Location and Policies tabs appear in the Backup Manager.
2. Click the Staging Location tab and expand the Staging Servers object.
 - a. Browse to and select the deduplication group you want to choose as the staging group for this backup job.
 - b. Click the Policies tab to specify Deduplication Staging policies.
 - c. Specify the staging policies for full, differential and incremental backups required for your job.

3. Click the Destination tab and expand the Servers object.
 - a. Browse to and select the group you wish to use as the final destination for this backup job.

Note: You may select a regular device group or another deduplication group, but you may not select the same deduplication group you specified as the staging destination.
 - b. Click Deduplication Policy to open the Deduplication Purge Policies dialog.
 - c. Click the Full Backup tab and specify the purge policy for full backups required for the job.
 - d. Click the Differential/Incremental Backup tab and specify the purge policy for incremental and differential backups required for the job.
 - **Purge data after** -- Specify the number of weeks, days, hours, and minutes to purge the job session after the operation ends.

Note: Make sure you view the deduplication staging policy because the default deletion policy is set to four weeks. If you want to retain backups longer than four weeks, you must manually adjust the policy.
 - e. Click the Miscellaneous tab and choose the desired miscellaneous options:
 - **Purge cancelled sessions from disk** --Removes any user-cancelled sessions from the deduplication device.
 - **Purge failed sessions from disk**--Removes any sessions that fail from the deduplication device.
 - f. Click OK.
4. Click the Schedule tab and specify the schedule that you want to use for the backup job.

Note: If you choose Use Rotation Scheme and Enable GFS, the Media Pool fields are not available for deduplication device groups.
5. Click the Options button on the toolbar to open the Global Options dialog. Set up Global Options as usual.
6. Click Submit on the toolbar to submit your job, as usual.

More information:

[Submit a Backup Job](#) (see page 136)

[Global Backup Options](#) (see page 154)

[Specify Copy and Purge Policies for Disk Staging Backups](#) (see page 213)

Considerations for Specifying Deduplication Device Copy and Purge Policies

Consider the scenarios that follow when specifying deduplication copy and purge policies:

- If you use a deduplication device as the destination in a non-staging backup job, you may configure purge policies. Click Deduplication Policy. The policy is enabled by default.
 - On the Full Backup and Differential/Incremental Backup tabs, specify a purge policy, if desired. The default setting is 4 weeks for full backups and 2 weeks for incremental/differential backups.
 - On the Miscellaneous tab, choose the desired options.
 - **Purge cancelled sessions from disk** -- Use this option to delete sessions from the destination device after a backup to destination device is cancelled.
 - **Purge failed sessions from disk** -- Use this option to delete sessions from the destination device after a backup to destination device fails.

Both options help reclaim disk space as quickly as possible.

- If you use a deduplication device in a staging job as a staging device, you can specify both copy and purge policies by clicking the Policies tab. For more information, see [Specify Copy and Purge Policies for Disk Staging Backups](#) (see page 213).
- In staging jobs where you use deduplication devices for both staging and destination locations, your jobs will have two purge policies.

Purge policy is always enabled. You cannot disable purge, but you can adjust the purge schedule.

How Encryption and Compression Works with Deduplication

Compression and encryption are not supported when used with deduplication devices. However, in staging jobs, compression and encryption are not supported in the Staging phase, but are supported in the Migration Phase only when the final destination specified is a non-deduplication device. The following table lists the options available depending on the device specified.

Note: For more information about Arcserve Backup encryption and compression options, including limitations and considerations, see [Backup Manager Encryption/Compression Options](#) (see page 159).

Staging Location	Final Destination	Compression/Encryption Options Available
Non-deduplication device	Non-deduplication device	All options are available.

Staging Location	Final Destination	Compression/Encryption Options Available
Non-deduplication device	Deduplication device	Encrypt Data options not allowed Compress Data options not allowed
Deduplication device	Non-deduplication device	Encrypt Data options: "at Backup Server During Migration" Compress Data options: "at Backup Server" Note: Compression must be specified with the encryption option "at Backup Server During Migration."
Deduplication device	Deduplication device	Encrypt Data options are not allowed. Compress Data options are not allowed. You can still select them, but an error message will appear.

View Compression Results after Deduplication

You can view the compression ratios achieved after a backup job with deduplication is completed in the Activity Log. Compression is displayed as a ratio and as a percentage. This information is also stored in the Arcserve Backup database, so you may view it in Job History at the session level, job level and node level.

- From the Restore Manager, you can view session level compression ratio information.
- From the Backup, Restore or Device Managers, you can view the device/tape level compression ratios.
- From the Report Manager, you can view the session compression ratio from the Session Details and Session Reports. You can view device level or node level compression ratio from the Dashboard report.

The compression ratio is the result of the amount of actual data to be stored divided by the amount of data stored after deduplication expressed as a ratio or as a percentage.

How to Back Up Deduplication Devices

The index and data files produced during a deduplication backup job are critical to successfully restoring deduplicated data. If these files should become corrupt, Arcserve Backup will not be able to find and reassemble the data chunks needed to rebuild the original data stream, even if the deduplicated data is intact. You can back up deduplication device files, but there are some important considerations you should first understand.

- Deduplication device files are normally skipped in local backup jobs (deduplication device and Arcserve Backup are on the same machine). However, you can forcibly include them by opening Global Options and enabling "Back up deduplication device data" on the Operation tab.

- Deduplication device files are normally included in remote backup jobs (deduplication device and Arcserve Backup on different machines). In addition, the data deduplication files can reside on the same or different remote computers. However, the data and index files may not be synchronized if other backup jobs to the deduplication device are running at the same time while the deduplication device is being backed up. We therefore require licensing the Arcserve Backup Agent for Open Files and using deduplication devices on computers that support VSS. For more information, see [Back up Deduplication Device Files](#) (see page 744).
- To back up deduplication device files, you must license the Agent for Open Files.

For information about how to restore deduplication devices, see [Restore Deduplication Device Files](#) (see page 751).

Back up Deduplication Device Files

The procedure for including the deduplication device data and index files during a backup job is the same whether the device is locally or remotely connected to the Arcserve Backup server.

Note: If you are backing up data to a deduplication device, and Arcserve Backup is backing the deduplication device to another device, the data that is backed up on the deduplication device may not be complete on the other device. If you want to back up a deduplication device completely, you should back up the deduplication device while it is not in use by other jobs.

To back up deduplication device files

1. Ensure that you have licensed the Arcserve Backup Agent for Open Files so that files in use can still be backed up. If the deduplication device is connected locally to the backup server, you must issue the license on the backup server.
2. Ensure that the machine to which the deduplication device is connected supports VSS.

3. Configure backup job options as usual:
 - a. Select the deduplication data folder and index folder of the device you wish to back up. If these folders reside on different volumes as recommended, they are backed up to two different sessions.

Note: Arcserve Backup lets you protect deduplication data when the data files and the index files reside on different computers. If you are using this approach, ensure that you specify the proper source nodes on the Backup Manager, Source tab.
 - b. From the Global Options Operation tab, enable Back up deduplication device data.
 - c. From the Global Options Volume Shadow Copy Service tab, enable Use VSS and disable Revert to traditional backup if VSS fails. If you do not perform this step, then the backup job automatically enables this option when the backup job runs.
4. Save and run the backup job as usual. For more information, see [Backing up Data](#) (see page 133).

Note: The ca_backup command line utility does not support the process of backing up data that belongs to deduplication devices.

How to Replicate Deduplication Devices with Arcserve Replication

Deduplication devices can store data from numerous data sources, making protection of these devices especially critical. You can use Arcserve Replication to replicate deduplication device data, adding another layer of protection to your Arcserve Backup environment.

Replicating deduplication devices using Arcserve Replication requires you to install the Arcserve Replication Engine on both the Master and Replica servers. See the *Arcserve Replication Installation Guide* for more information.

- The local server hosting the deduplication device you wish to replicate should be designated as the Master server.
- The server running Windows Server 2003 or later should be designated as the Replica server.

Note: After you install the Arcserve Replication Engine on both the Master and Replica servers, you must create and configure a Arcserve Replication scenario for the deduplication device that you wish to protect. If a deduplication device fails, you can restore data using the VSS snapshot generated by Arcserve Replication.

Create Arcserve Replication Scenarios for Deduplication Devices

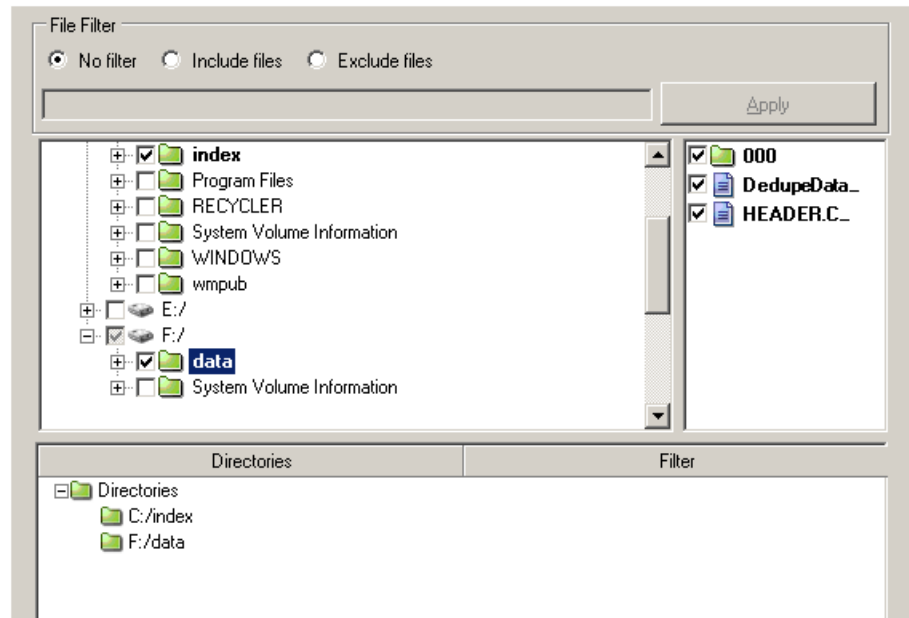
The following procedure is presented to replicate deduplication devices using Arcserve Replication scenarios, specifically, a Arcserve Replication File Server scenario. For more information, see the Arcserve Replication User Guide.

Important! The Master server is the local host for the deduplication device you wish to replicate.

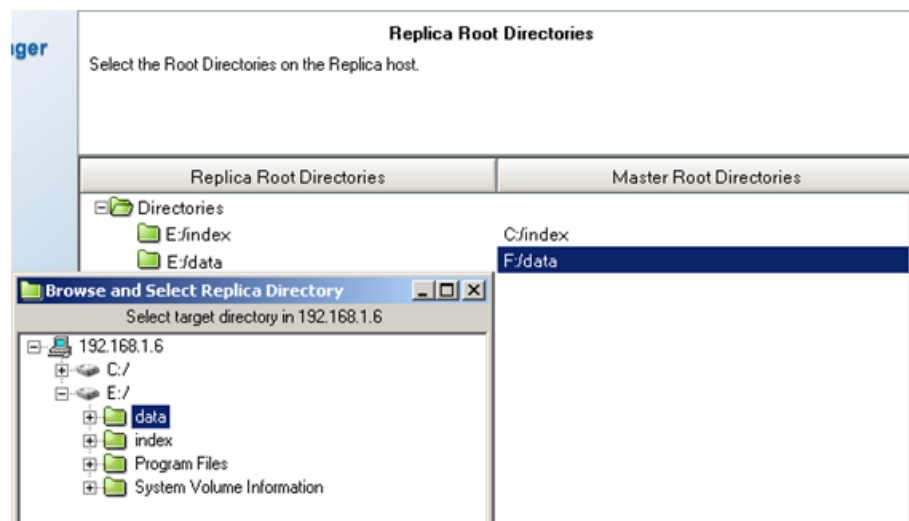
To create Arcserve Replication scenarios for deduplication devices

1. From the Arcserve Replication Manager, choose Scenario, New, or click the New Scenario button from the toolbar to launch the Scenario Creation Wizard.
2. At the Welcome screen, select Create New Scenario, select an appropriate Group, and then click Next.
3. At the Select Server and Product Type screen, select File Server, Replication and Data Recovery Scenario (DR) and Integrity Testing for Assured Recovery (AR). You must choose the Integrity Testing for Assured Recovery (AR) option to generate the VSS Snapshots used for recovering failed deduplication devices. Click Next to continue.
4. At the Master and Replica Hosts screen, provide a Scenario Name. For example, DDD. Enter the Hostname or IP address and Port Number for both Master and Replica servers. Click Next to continue.
5. Wait for Engine Verification to complete. If needed, click Install to upgrade the Engine on one or both servers and then click Next to continue.

- At the Master Root Directories screen, select the deduplication device data file folder and index file folder. Click Next to continue.



- At the Replica Root Directories screen, select the data file folder on the Replica server. Due to the size of the VSS snapshot, we recommend that you put the index and data files for the deduplication device being replicated on the same volume. Click Next to continue.



- At the Scenario Properties screen, accept the defaults and click Next to continue.
- At the Master and Replica Properties screen, accept the defaults and click Next to continue.

10. Wait for Scenario Verification to complete. If errors or warnings are listed, resolve them before continuing. Click Next to continue.
11. At the Scenario Run screen, click Finish.

You must complete the Configure Arcserve Replication Scenarios for Deduplication Devices procedure before you can run the scenario.

Configuration Considerations for Arcserve Replication Deduplication Device Scenarios

To replicate deduplication devices using Arcserve Replication, there are two configuration methods to consider:

- **Online Replication Type** -- the deduplication device is replicated to the Replica server in real-time. This may impact device performance, which you can address by configuring the scenario spool on a separate hard disk. For more information on Spool size, see the Arcserve Replication User Guide.
- **Scheduled Replication Type** -- the deduplication device is replicated to the Replica server at the time you schedule. If you specify a time during which no jobs are running, the impact to the device is minimized.

Configure Online Arcserve Replication Scenario for Deduplication Devices

Arcserve Backup lets you configure Arcserve Replication replication scenarios on deduplication devices.

To configure scenarios for online replication

1. From the Arcserve Replication Manager, select the scenario you created to replicate the deduplication device.
 - a. Click the Properties tab for this scenario.
 - b. Set the Replication, Mode property to Online.

2. From the Arcserve Replication Manager, select the Master server to which the deduplication device is locally connected.
 - a. Click the Properties tab for this server.
 - b. Set the Spool, Spool Directory property to a folder on a different hard disk than the deduplication device. This improves performance.
3. From the Arcserve Replication Manager, select the Replica server.
 - a. Click the Properties tab for this server.
 - b. Set the Scheduled Tasks, Replica Integrity Testing for Assured Recovery, Action on successful test, Create Shadow Copy (VSS) property to On.
 - c. Set its child properties as desired:
 - Number of Snapshots to keep -- 10 is used as an example. Raise or lower this value, as desired.
 - Shadow Storage Volume -- Default
 - Max Storage Size per Volume -- Unlimited
 - d. Set the Scheduled Tasks, Replica Integrity Testing for Assured Recovery, Scheduler property.
4. Save the changes.

Run the deduplication device replication scenario.

Configure Scheduled Arcserve Replication Replication Scenarios for Deduplication Devices

When you use scheduled replication for your deduplication device replication scenario, you must manually generate VSS Snapshots in order to restore data in case of device failure.

To configure scenarios for scheduled replication

1. From the Arcserve Replication Manager, select the scenario you created to replicate the deduplication device.
 - a. Click the Properties tab for this scenario.
 - b. Set the Replication, Mode property to Scheduling.
 - c. Schedule the replication time at 0:00 daily.
2. From the Arcserve Replication Manager, select the Replica server.
 - a. Click the Properties tab for this server.
 - b. Set the Scheduled Tasks, Replica Integrity Testing for Assured Recovery, Action on successful test, Create Shadow Copy (VSS) property to On.
3. Save the changes.

4. Run the deduplication device replication scenario.
5. Manually generate VSS Snapshots:
 - a. Select the Replica server for the deduplication device replication scenario you created.
 - b. Click the Replica Integrity Testing button on the Arcserve Replication Manager toolbar.
 - c. When the Replica Testing for Assured Recovery screen opens, click OK to start.

Back Up Data Using Global Deduplication

Global deduplication finds redundancies between C:\ drive backup sessions on different machines being backed up to the same deduplication device. Generally, the C:\ drive of a given machine holds operating system files, where high instances of redundancy are expected. In addition to system volumes, Global deduplication also handles Oracle RMAN sessions.

Global Deduplication happens every 6 hours but cannot happen while a backup or purge job is active, and will be interrupted if both jobs need to access the same session file.

To perform global deduplication, set up a backup job in the usual manner, ensure you have checked the Enable Global Deduplication option in Deduplication Device Group Configuration and you have select the C:\ drives of different machines specified as the backup source.

To back up data using global deduplication

1. From Deduplication Device Group Configuration, ensure the Enable Global Deduplication option is set (default setting is enabled).
2. From the Backup Manager, set up a backup job in the usual manner.
3. From the Source tab, select the Windows C:\ volume directories of different machines.
4. (Optional) If using staging, click the Staging Location tab and select a deduplication device group.
5. From the Destination tab, select a deduplication device group.
6. Complete the job setup selections as usual.
7. Run the job.

Recover Deduplicated Data

Restoring data that has been deduplicated follows the same procedure as a normal restore job. Disaster Recovery supports deduplication and also follows the same procedure as a normal disaster recovery.

The Arcserve Backup Utilities also support deduplication devices.

Note: The Purge utility is assigned a lower priority than Backup, Restore, Merge and Scan. Purge is skipped when any active backup, restore, merge, or scan jobs are active on the same deduplication device. Backup, restore, merge, and scan jobs directed to a deduplication device where a purge session is already in progress take precedence and will halt the purge session.

Restore Deduplicated Data

The process of restoring data saved to deduplication devices is similar to that of normal FSDs. You must restore from a disk, even if you have migrated data to tape as part of a staging job.

More information:

[Restoring Data](#) (see page 257)

Restore by Sessions on Deduplication Devices

If a deduplication device contains a large number of sessions, you can use the Last number of days option to filter results. When you expand the device, a progress bar appears. You can cancel progress to display a shortcut menu and then select a target session from the expanded device to restore. You can also expand the remaining sessions using the Show More option from the shortcut menu.

Restore Deduplication Device Files

As long as you have also explicitly backed up the deduplication device data and index files, you can restore deduplicated data. To restore deduplication device files, do the following:

- Browse the node on which the backed up index files are stored
- Choose the correct sessions
- Choose the index or data files you wish to restore
- Restore to an alternate location

- Remove the original deduplication device to avoid conflicts
- Create a new deduplication device configured to use the restored index and data folders
- Stop and then restart the Tape Engine to ensure that any operations taking place when the backup was initiated are invalidated
- Merge the new deduplication device so that records match the physical device.


Note: The Merge operation may be incomplete or fail if there are active sessions on the backed up Deduplication device.

Restore Deduplication Devices Using Arcserve Replication/VSS Snapshots


If a replicated deduplication device fails, you can recover the data stored on it using the VSS Snapshots.

To recover a failed deduplication device

1. From the Arcserve Replication Manager, stop the replication scenario.
2. Expose the VSS Snapshot:
 - a. Click Snapshot View and select the snapshot from which you wish to restore.
 - b. Select Expose as Drive Letter to start the restore process.



- Log on to the Replica server, open the exposed drive, and copy the exposed index and data files to an alternate location.



- From the Arcserve Backup Device Manager, remove the failed deduplication device. This device must be removed to avoid conflicts in which two devices exist with the same tapeName, randomID, and sequenceNum.
- From the Arcserve Backup Device Manager, create a new deduplication device using the index and data file paths just copied. For more information, see [Create Deduplication Devices](#) (see page 735).
- Stop and restart the Tape Engine to ensure any operations taking place during backup initiation are invalidated.
- Merge the new deduplication device so that the DB tape record is updated to match the physical deduplication tape. Use the Merge session header only option from Merge Global Option, Database menu to do so. The Merge operation may be incomplete or fail if there are active sessions on the backed up Deduplication device.

You may need to create a new Arcserve Replication scenario to replicate the new deduplication paths. To use the existing scenario, you can restore the Snapshot to previous deduplication paths, but you will need to first delete or move any files presently stored there.

Scan Jobs with Deduplication

The process of running a Scan job with deduplication is the same as a regular Scan job. If desired, click the Media Assure button to access Media Assure options, then click the Enable Media Assure option to select it.

Media Assure works with all media types but for deduplication devices where hundreds of sessions are stored, it randomly samples the sessions that meet your criteria to ensure data can be recovered.

You can scan all data (default) or only session headers, as well as set filter options. You can specify one node or multiple nodes, separating names with commas in the field provided. You cannot specify * groups for deduplication Media Assure scan jobs, but you can specify a * tape in a specific group, or you can select a specific tape. The job scans the sessions that satisfy your selection criteria and repeats periodically until deleted.

Note: Media Assure supports scanning only one group and is suspended by other jobs targeting the same media. If suspended, an error is produced, "E3708 Unable to reserve group <group name> in <minutes> minutes."

More information:

[Media Assure & Scan Utility](#) (see page 32)

Merge Jobs with Deduplication

The process of performing a Merge job with deduplication is the same as a regular Merge job. For more information, see [Merge Utility](#) (see page 30).

GFS Rotation Jobs on Deduplication Devices

Deduplication devices cannot be assigned to media pools, so consider the following when setting up a GFS or Rotation scheme on these devices:

- When you select a deduplication device as the destination device in a staging operation on a GFS or Rotation job, you will not be permitted to specify a media pool name. You will be permitted to submit the GFS or Rotation scheme without media pool.
- When you select a deduplication device as the destination device in a GFS or Rotation job in a non-staging operation, media pool is not used and media will never be overwritten. Data is written to formatted media in the deduplication device group, if one exists. If one does not exist, blank media is formatted with the current data and time.
- When you select a deduplication device as the destination device in a GFS or Rotation job in a staging operation, the behavior of the staging phase is not changed, but the migration phase will never use a media pool and never overwrite media. Data is appended to formatted media in the deduplication device group if one exists. If not, blank media is formatted with the current date and time.
- With GFS or normal rotation, with Append Media or without Append Media, backup jobs saved to deduplication devices behave in the same way.

Note: For more information about GFS rotations, see [Rotation Schemes](#) (see page 120).

Deduplication Device Purge

Purging a deduplication device is different than purging a staging FSD. When a staging FSD is purged, Arcserve Backup immediately removes the session file. But for a deduplication session purge, Arcserve Backup renames the session hash file to .hash_ToPurge and updates the reference counter. In other words, the session is merely "marked" as purged, but not actually removed because there may be other sessions that still point to the original data.

The reference counter stored in the index files is decreased. When the reference counter reaches 0, no more hashes referencing the original data exist and the data chunk is now considered to be a "hole." When Arcserve Backup finds data files with holes greater than 25%, disk space is reclaimed by a purge thread that runs every 6 hours.

Disk Fragmentation

To reduce disk fragmentation, space is always allocated for the first backup of a given session in 1 GB increments, until the session ends. If four streams are writing, each stream uses a pre-allocated chunk of disk space.

The last 1 GB chunk of data is rounded down so that the data file occupies the actual compressed session size after the deduplication process is complete. This approach helps to ensure that the disk is fragmented in 1 GB chunks.

This is done only for backups to deduplication devices and only for the first backup of one root directory on one device.

For second backup and subsequent jobs, the amount of data physically written to disk is expected to be low.

Delete Deduplication Backup Sessions

Arcserve Backup deletes deduplication backup sessions from deduplication devices based on the purge policy (retention period) that was specified when you submitted the backup job. However, deduplication backup sessions can become obsolete or redundant before the specified retention period expires. Arcserve Backup lets you delete deduplication backup sessions so that you can free disk space and remove obsolete and redundant session information from the Arcserve Backup database.

When you delete deduplication backup sessions, Arcserve Backup behaves as follows:

- When the deleted sessions are referenced by other deduplication backup sessions, the information about the sessions is removed immediately from the Arcserve Backup database. However, Arcserve Backup cannot reclaim the disk space until the deleted sessions are not referenced by other deduplication backup sessions.
- When the deleted sessions are not referenced by other deduplication backup sessions, the information about the sessions is removed from the Arcserve database and the disk space is reclaimed immediately.

Important! The deleting of deduplication backup sessions is an irreversible process. You cannot recover deleted deduplication backup sessions.

Follow these steps:

1. From the Quick Start menu on the Arcserve Backup Home Page, click Restore.
2. Click the Source tab in the Restore Manager.
From the drop-down menu, select Restore by Session.
The Sessions directory tree appears.
3. Expand Sessions to view the devices containing backup sessions.
4. Browse to the deduplication device containing the sessions that you want to delete.
Expand the deduplication device to display the deduplication backup sessions stored on the device.
5. Locate and select the sessions that you want to delete.
Right-click the sessions that you want to delete and click Delete Selected Sessions on the pop-up menu.
When prompted, click Yes to delete the selected sessions.
The sessions are deleted.

Deduplication Reports

Arcserve Backup reports have been modified to include deduplication statistics. See [Report Categories](#) (see page 672) for more information.

Chapter 13: Raw Backup and Restore of Physical Disks and Volumes

This section contains the following topics:

[Overview of Raw Backup and Restore](#) (see page 759)

[How Raw Backup Works](#) (see page 760)

[Enable Raw Backup and Restore](#) (see page 763)

[Perform Raw Backup of a Physical Disk or Volume](#) (see page 763)

[Raw Backup Restore](#) (see page 764)

Overview of Raw Backup and Restore

Arcserve Backup lets you back up and restore physical disks and physical volumes that may or may not have a file system. For example, this capability lets you back up the following items:

- Oracle database snapshots
- Non-Windows file system partitions
- Nameless partitions or volumes, that is, volumes without a drive letter
- Any snapshot mounted as a disk to the client system

Licensing Requirements for Raw Backup of Physical Disks and Volumes

You must have the following licenses to use the raw backup of physical disks and volumes feature:

- License for the Arcserve Backup Client Agent for each server on which you want to use the feature
- License for the Arcserve Backup Enterprise Module for each server from which you want to back up the client agents

Example: Licensing requirements for raw backup of physical disks and volumes

If you want to use the raw backup of physical disks and volumes feature on servers A, B, and C, you must install a license for the Arcserve Backup Client Agent on all three servers. In addition, you must register one license for the Enterprise Module on the backup server from which you back up these three servers.

How Raw Backup Works

To perform a raw backup of physical disks and volumes, Arcserve Backup obtains exclusive access to the device, which allows Arcserve Backup to capture a consistent backup image. Arcserve Backup reads the data sequentially, by blocks, and then copies the image to the staging device or the Arcserve Backup server.

Supported Functions

You can use the following Arcserve Backup functions with raw backup and restore:

- Backup estimation
- Compression
- Encryption

Limitations on Performing Raw Backup and Restore Operations

Consider the following limitations when performing a raw backup of physical disks and volumes:

- Arcserve Backup does not support performing raw incremental and differential backups. If you submit such a backup job, Arcserve Backup will automatically change it to a full backup job.
- Arcserve Backup does not execute backups using Volume Shadow Copy Service (VSS) snapshot technology.
- Arcserve Backup does not support the backup and restore of physical disks or physical volumes for clusters. Therefore, these devices will not be displayed under the cluster virtual node in the Source tab of the Backup Manager.

- Arcserve Backup does not support the backup and restore of physical disks or physical volumes for removable media. Therefore, these devices will not be displayed in the Source tab of the Backup Manager.
- Backup jobs may fail when Arcserve Backup cannot obtain exclusive access to the device.
- When a dynamic physical disk is restored to another physical disk, Arcserve Backup does not copy the partition information to the destination physical disk. Therefore, the volumes and partitions of the source physical disk are not shown on the destination physical disk after the restore. In other words, you can restore dynamic disks only to their original location. In addition, if you have multiple dynamic disks, then you must restore all of these dynamic disks to their original locations to reinstate the original volume partitioning.

The reason for restoring dynamic disks to their original location is:

The partition table on a dynamic disk does not contain an entry for each volume on the disk because the volume information is stored in the dynamic disk database. Each dynamic disk of a system contains a replica of this dynamic disk database. The location of the database is determined by the partition style of the disk.


- On Master Boot Record (MBR) disks, the database is contained in the last 1 megabyte (MB) of the disk.
- On Globally Unique Identifier Partition Table (GPT) disks, the database is contained in a 1 MB reserved (hidden) partition known as the Logical Disk Manager (LDM) Metadata partition.

As a result, unless this database is written to the disk during the restore, the partition information cannot be restored. For a system with multiple dynamic disks, all the disks have to be restored to their original locations because each disk contains a copy of the database, and database copies must be identical to restore the original partition information.

- The physical volumes corresponding to system or boot volumes and the disks on which these physical volumes reside are not displayed in the Source tab of the Backup Manager.
- The Filter option is not available for raw backup and restore of physical disks and volumes.

Naming Convention of Physical Disks and Volumes

After you enable raw backup and restore, you can view the physical disks and volumes that are connected to the agent in the Backup Manager under the Source tab. The following illustration shows a section of the Source tab that has displayed the physical volumes:



PhysicalDisk_<Disk id>:

Indicates a physical disk. <Disk id> represents the identification tag of the disk. Arcserve Backup and Windows Disk Management display the <Disk id> in the same way.

PhysicalVolume_<Disk id>_<Volume id>

Indicates a physical volume. <Volume id> is a drive letter or the GUID. The GUID is a hexadecimal number that displays only when the volume is an unnamed volume (no drive letter association). For example, PhysicalVolume_Disk 2_E: means that the volume E: resides on Disk 2.

Enable Raw Backup and Restore

The capability to perform raw backup and restore of physical disks and volumes is disabled by default. You must enable the option for each agent.

To enable raw backup and restore of a physical disk or volume

1. From the Windows Start menu, click Start, Programs, CA, ARCserve Backup, and Backup Agent Admin.

The ARCserve Backup Agent Admin window appears.

2. Click Options, Configuration.

The Configuration window appears.

3. Click Enable physical disk/volume backup and restore.
4. Click OK.

The raw backup and restore feature is enabled for the agent.

Perform Raw Backup of a Physical Disk or Volume

You can perform a raw backup of a physical disk or volume using one of the following approaches:

- **Regular backup**--Lets you specify a regular device group as the backup destination.
- **Deduplication backup**--Lets you specify a deduplication device group as the backup destination.
- **Staging backup**--Lets you specify regular device group or a deduplication device group as the staging location, the final destination media, or both.

Note: You should not specify the same deduplication device group for the staging location and the final destination media for deduplication backups.

To perform raw backups of a physical disk or volume

1. Open the Arcserve Backup Manager Console.
2. From the Navigation Bar, click Quick Start, and then click Backup.
The Backup Manager opens.
3. Specify the type of backup.
4. (Optional) Select the Enable Staging check box.
5. Click the Source tab and locate the physical disks or volumes that you want to back up.
6. Click the Schedule tab and define the schedule for the backup job.

7. (Optional) Click the Staging Location tab and specify a staging device group.

Note: The Staging Location tab displays only when you click the Enable Staging check box on the Start tab.

8. (Optional) Click the Policies tab and specify the Tape staging policies and Copy policies for the job.

Note: The Policies tab displays only when you click the Enable Staging check box on the Start tab.

9. Click the Destination tab and select the device group to store the backup data.

10. Click Submit.

The Submit Job dialog opens.

11. Complete the required fields on the Submit Job dialog and click OK.

The job is submitted.

Note: Arcserve Backup verifies that you have a valid license for the Enterprise Module on the server where you want to run the backup job. If Arcserve Backup detects the required licenses, the backup job is submitted. If Arcserve Backup does not detect the required licenses, the backup job is not submitted.

Entire Node Backup

When you select an entire node to back up, Arcserve Backup does not perform a raw backup of the physical volumes corresponding to Windows File System volumes. The Client Agent for Windows backs up these volumes in the traditional manner. In addition, Arcserve Backup does not perform raw backup on the physical disks that host any of the File System volumes, to avoid backing up duplicate data.

Raw Backup Restore

You can restore raw backups of physical disks and volumes using the following approaches:

- Restore the backup data to an alternate location as a binary file.
- Restore the backup data to its original location.
- Restore the backup data to another physical disk or volume.

Restore to an Alternate Location as a File

Arcserve Backup lets you restore the raw backups of physical disks or volumes to an alternate location as a binary file. You can choose to overwrite all files, rename files, skip existing files, or overwrite with newer files.

To restore a raw backup to an alternate location as a binary file

1. Open the Arcserve Backup Manager Console.
2. From the Navigation Bar, click Quick Start, and then click Restore.
The Restore Manager opens.
3. Click the Source tab.
4. Browse to and select the source physical disk or volume.
5. Click the Destination tab and select a folder on a physical disk or volume.

Note: The restore destination must not be the same location as the source directory that you specified on the Source tab.

6. Click the Schedule tab and define the schedule for the restore job.
7. Click Submit on the toolbar to submit the job.

The restore job is submitted and Arcserve Backup restores the raw backup data to the alternate locations as a binary file.

Restore to the Original Location

Arcserve Backup lets you restore the raw backup of a physical disk or volume to the same location from where you backed it up.

To restore a raw backup to the original location

1. Open the Arcserve Backup Manager Console.
2. From the Navigation Bar, click Quick Start, and then click Restore.
The Restore Manager opens.
3. Click the Source tab.
4. Browse to and select the source physical disk or volume.
5. Click the Destination tab.
Ensure that Restore files to their original location(s) is selected.
6. Click the Schedule tab and define the schedule for the restore job.

7. Click Submit on the toolbar to submit the job.

The restore job is submitted and Arcserve Backup restores the raw backup data to the original location.

Restore to Another Physical Disk or Volume

You can restore the raw backup of a physical disk or volume to another physical disk or volume. You can restore the raw backup of a physical disk to another physical disk and not to a physical volume; similarly, you can restore from a physical volume to another physical volume and not to a physical disk.

To restore a raw backup to another physical disk or volume

1. Open the Arcserve Backup Manager Console.
2. From the Navigation Bar, click Quick Start, and then click Restore.
The Restore Manager opens.
3. Click the Source tab and select the source physical disk or volume.
4. Click the Destination tab and select a physical disk or volume that is different from the one you selected as the source.
5. Click the Schedule tab and select when you want the restore process to start.
6. Click Submit on the toolbar.

The restore starts or is saved to be executed at the scheduled time.

Note: Before writing to the target device, Arcserve Backup compares the size of the target device with the size of the device in the backup session you have selected to restore. If the size of the target device is smaller, Arcserve Backup fails the restore job.

Chapter 14: Backing Up and Recovering D2D/UDP Data

This section contains the following topics:

[Specify D2D/UDP Backup Options](#) (see page 767)

[How Arcserve Backup Processes Encrypted D2D/UDP Backup Sessions](#) (see page 768)

[Submit Backup Jobs of D2D/UDP Nodes](#) (see page 768)

[Submit Backup Jobs of D2D/UDP Nodes through Proxy Servers](#) (see page 770)

[Submit Static Backup Jobs Containing D2D/UDP Data](#) (see page 772)

[Submit Dynamic Backup Jobs Containing D2D/UDP Data](#) (see page 774)

[Recover D2D/UDP Data at File Level Granularity](#) (see page 775)

[Recover D2D/UDP Data at Application Level Granularity](#) (see page 777)

[Recover Files and Folders from D2D/UDP Recovery Points](#) (see page 783)

[Recover D2D/UDP Data from Raw Sessions](#) (see page 785)

[Configure Arcserve Backup to Generate D2D/UDP Catalog Files](#) (see page 787)

Specify D2D/UDP Backup Options

From the Backup Manager, you can open the D2D/UDP Home Page so that you can specify D2D/UDP backup options. The backup options let you specify behaviors such as:

- The backup source and the backup destination
- The schedule for each type of backup
- The advanced settings for backup jobs
- Pre and post backup operations

Follow these steps:

1. Open the Arcserve Backup Manager Console.

From the Quick Start menu, click Backup.

The Backup Manager window opens.

2. Click the Source tab.

The source tree appears.

3. Expand the D2D/UDP Servers object.

Right-click the D2D/UDP server that you want to configure and click Modify D2D/UDP Backup Settings on the pop-up menu.

The D2D/UDP Backup Settings dialog opens.

4. Specify the backup settings that you require for your D2D/UDP backups.

Note: For more information about specifying D2D/UDP backup settings, click Help on the Backup Settings dialog.

How Arcserve Backup Processes Encrypted D2D/UDP Backup Sessions

D2D/UDP lets you protect sensitive data using data encryption. The data is protected using an encryption password that you specify when the backup is submitted. To recover the D2D/UDP data, you specify the password when you submit the restore.

To back up D2D/UDP sessions to Arcserve Backup media, you add the D2D/UDP servers to the Backup Manager Source directory tree before you can submit backups. The process of adding the D2D/UDP servers to the Source tree requires you to specify the D2D/UDP computer name and credentials (user name and password). Arcserve Backup uses the D2D/UDP credentials to retrieve the D2D/UDP encryption password, to decrypt the data, and to back up the data to Arcserve Backup media. As a result, Arcserve Backup stores the D2D/UDP backup sessions on Arcserve Backup media in decrypted format.

To recover D2D/UDP data from Arcserve Backup media, a password is not required. If you want to encrypt the D2D/UDP data on Arcserve Backup media, you can specify Arcserve Backup encryption options when you submit the job. For more information about encryption options, see Backup Manager Encryption/Compression Options in the *Administration Guide*.

Submit Backup Jobs of D2D/UDP Nodes

Arcserve Backup lets you back up D2D/UDP sessions to Arcserve Backup media. The Arcserve Backup media can be file systems devices, tape media, NAS devices, and devices that reside on a SAN. The process of backing up D2D/UDP data is identical to the steps that are required to back up files, folders, nodes, servers, and so on.

This backup method lets Arcserve Backup synthesize full backup sessions from the full and incremental D2D/UDP backup sessions. The synthesized backup sessions can be used to recover D2D/UDP data at file level, folder level, and application level granularity, and to perform full system recoveries of the D2D/UDP servers, such as bare metal recovery.

Follow these steps:

1. From the Backup Manager, select the [Start](#) (see page 139), [Source](#) (see page 140), [Destination](#), and [Schedule](#) (see page 152) tabs to specify the options that you require for the job.

Note: To back up the D2D/UDP sessions using multistreaming, on the Destination tab, click the Multistreaming option and then specify the maximum number of streams that you want to allow for the job.

2. Click the Options toolbar button to specify the backup options that are required for the job. For more information, see [Global Backup Options](#) (see page 154).
3. Click the Submit on the toolbar to submit the job.

The Security and Agent Information dialog opens.

4. On the Security and Agent Information dialog, edit or confirm the security and agent information for your job, and click OK.
5. When the Submit Job dialog opens, select Run Now to run the job immediately, or select Run On and select a date and time when you want the job to run.

Note: For more information about the Run Now option, see Job Queue Tab.

6. Enter a Job Name for your job.
7. If you selected multiple sources to back up and want to set the priority in which the job sessions initiate, click Source Priority. Use the Top, Up, Down, and Bottom buttons to change the order in which the jobs are processed. When you finish setting priorities, click OK.
8. To save the job as a Arcserve Backup job script, click the Save Job button.
9. To save the job template, click the Save Template button.
10. To preflight check the job, click the Preflight Check button. If the preflight check failed, click the Cancel button to modify the job settings.
11. On the Submit Job dialog, click OK.

The job is submitted.

Submit Backup Jobs of D2D/UDP Nodes through Proxy Servers

Arcserve Backup lets you back up D2D/UDP sessions to Arcserve Backup media through a proxy server. When Arcserve Host-Based VM Backup is installed on the proxy server, you can submit the backup job through the Host-Based VM Backup proxy server. The Arcserve Backup media can be file systems devices, tape media, NAS devices, and devices that reside on a SAN.

When you backup D2D/UDP servers through a proxy server, Arcserve Backup retrieves all of the D2D/UDP backup sessions for all of the D2D/UDP servers that are protected by the proxy server. When Arcserve Host-Based VM Backup is installed on the proxy server, Arcserve Backup retrieves all of the D2D/UDP backup sessions for all of the D2D/UDP virtual machines that are protected by the proxy server.

This backup method lets Arcserve Backup synthesize full backup sessions from the full and incremental D2D/UDP backup sessions. The synthesized backup sessions can be used to recover D2D/UDP data at file level, folder level, and application level granularity, and to perform full system recoveries of the D2D/UDP servers, such as bare metal recovery.


The process of backing up D2D/UDP through a proxy server is almost identical to the steps that are required to back up files, folders, nodes, servers, and so on.

When backing up D2D/UDP nodes on a D2D/UDP server that failed through a proxy server; the D2D/UDP server is still supported and backed up as long as there was at least one backup running when the D2D/UDP server failed. The details from the D2D/UDP destination are cached and used for subsequent backups when the D2D/UDP server fails.

Follow these steps:

1. From the Backup Manager, click the [Start](#) (see page 139) tab and specify the type of backup job that you want to submit.
2. Click the [Source](#) (see page 140) tab to display the D2D/UDP nodes and proxy servers.

Specify the proxy servers containing the nodes that you want to back up as illustrated by the following screen:



3. Click the Destination tab to specify the location where you want to store the backup data.

Note: To back up the D2D/UDP sessions using multistreaming, on the Destination tab, click the Multistreaming option and then specify the maximum number of streams that you want to allow for the job.

4. Click the [Schedule](#) (see page 152) tab to specify when and the frequency that you want the job to run.
5. Click the Options toolbar button to specify the backup options that are required for the job. For more information, see [Global Backup Options](#) (see page 154).
6. Click the Submit on the toolbar to submit the job.
The Security and Agent Information dialog opens.
7. On the Security and Agent Information dialog, edit or confirm the security and agent information for your job, and click OK.
8. When the Submit Job dialog opens, select Run Now to run the job immediately, or select Run On and select a date and time when you want the job to run.
Note: For more information about the Run Now option, see Job Queue Tab.
9. Enter a Job Name for your job.
10. If you selected multiple sources to back up and want to set the priority in which the job sessions initiate, click Source Priority. Use the Top, Up, Down, and Bottom buttons to change the order in which the jobs are processed. When you finish setting priorities, click OK.
11. To save the job as a Arcserve Backup job script, click the Save Job button.
12. To save the job template, click the Save Template button.
13. To preflight check the job, click the Preflight Check button. If the preflight check failed, click the Cancel button to modify the job settings.
14. On the Submit Job dialog, click OK.

The job is submitted.

Submit Static Backup Jobs Containing D2D/UDP Data

The Enable Static Backup option lets you submit scheduled backups of source groups and computers and maintain a static set of source volumes. With static backups, you can submit staging and deduplication backups of D2D/UDP data.

You can submit static backup jobs only to D2D/UDP recovery point servers that appear below the D2D/UDP recovery point Servers object, which includes D2D/UDP Agent-based servers and D2D/UDP Agentless VM nodes.

Static backups affect only the immediate subordinate objects of the source group or the computer specified when the job was submitted. Subsequent backups will not include objects and volumes that were added to the source group or the computer after you submitted the backup. Arcserve Backup backs up the original source volumes dynamically. Subsequent job executions will include changes to only the files and folders contained in the original source volumes.

Note: For more information, see [Static Job Packaging](#) (see page 305).

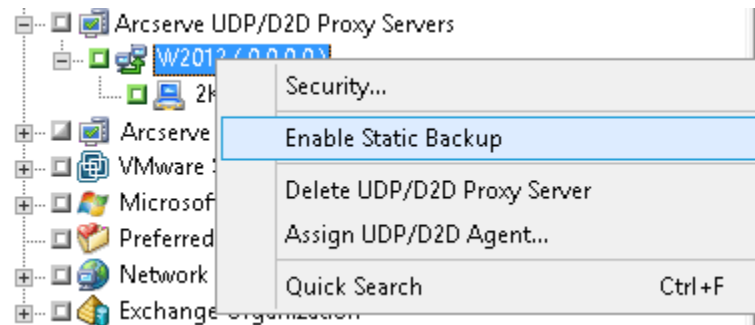
Follow these steps:

1. Open the Backup Manager and click the Start tab.
The backup types appear.
2. Select Normal backup and then click the Source tab.
The Source directory tree appears.


3. Select Classic View from the drop-down list.

Expand the D2D/UDP Recovery Point Servers object and browse to the D2D/UDP Recovery Point Servers that you want to back up.

- a. Click the check box next to the D2D/UDP Recovery Point Server name.
- b. Right-click the D2D/UDP Recovery Point Server and click Enable Static Backup on the pop-up menu.



The Static Backup option is applied to the computer, as illustrated by the following screen:



Note: The Static Backup option remains applied to the specified computer until you disable this option. You can repeat this step to disable the Enable Static Backup option.

4. Click the Schedule tab and specify the schedule that you want to use for the backup job.

Note: For more information, see [Rotation Schemes](#) (see page 307) and [Custom Schedules](#) (see page 318).

5. Click the Destination tab.

The Destination groups directory tree appears.

6. Specify the group that you want to use to store the backup data.

The storage group is applied.

7. Click Options on the toolbar.

The Options dialog opens.

8. Specify the options that you require for the job.

Note: For more information, see [Global Backup Options](#) (see page 154).

Click OK.

The backup options are applied.

9. (Optional) Expand the contents of the computer.
Select a drive or volume contained in the computer.
Repeat the previous two steps to apply backup options to the specified drive or volume.
(Optional) Repeat this step for all drives or volumes in the computer.
10. Click Submit on the toolbar.
The Submit Job dialog opens.
11. Complete the required fields on the Submit Job dialog and click OK.
The job is submitted.

Submit Dynamic Backup Jobs Containing D2D/UDP Data

Arcserve Backup lets you back up D2D/UDP data dynamically. Dynamic packaging means that the content of your selection is determined when the job runs.

For example, if you choose to back up a source group or a server, and the nodes or volumes listed on that source group or server change between the time you scheduled the job and the time the job runs, the nodes and volumes changed from the time the job actually runs are backed up.

Follow these steps:

1. Open the Backup Manager and click the Start tab.
The backup types appear.
2. Select Normal backup and then click the Source tab.
The Source directory tree appears.
3. Select Classic View from the drop-down list.
Expand the D2D/UDP Recovery Point Servers object and browse proxy server containing the D2D/UDP Agent-based or agentless nodes that you want to back up.
Click the check box next to the D2D/UDP Recovery Point Server containing the D2D/UDP nodes that you want to back up.
4. Click the Schedule tab and specify the schedule that you want to use for the backup job.
Note: For more information, see [Rotation Schemes](#) (see page 307) and [Custom Schedules](#) (see page 318).
5. Click the Destination tab.
The Destination groups directory tree appears.

6. Specify the group that you want to use to store the backup data.
The storage group is applied.
7. Click Options on the toolbar.
The Options dialog opens.
8. Specify the options that you require for the job.
Note: For more information, see [Global Backup Options](#) (see page 154).
Click OK.
The backup options are applied.
9. (Optional) Expand the contents of the computer.
Select a drive or volume contained in the computer.
Repeat the previous two steps to apply backup options to the specified drive or volume.
(Optional) Repeat this step for all drives or volumes in the computer.
10. Click Submit on the toolbar.
The Submit Job dialog opens.
11. Complete the required fields on the Submit Job dialog and click OK.
The job is submitted.

Recover D2D/UDP Data at File Level Granularity

Arcserve Backup lets you recover D2D/UDP backup data at file level granularity from Arcserve Backup media. Use these steps when you want to recover individual files, folders, directories, and so on from Arcserve Backup media to D2D/UDP servers.

The steps listed below describe how to perform the following recoveries:

- Recover D2D/UDP data that was backed up via D2D/UDP servers at file level granularity.
- Recover D2D/UDP virtual machine data that was backed up via D2D/UDP Central Host-Based VM Backup proxy servers at file level granularity.

Note: The following steps describe how to submit restore jobs using the Restore by Session method. The process of restoring D2D/UDP data and D2D/UDP virtual machine data using the Restore by Tree method is identical to that of all other types of data. For more information about the Restore by Tree method, see "Restoring Data."

Follow these steps:

1. Open the Arcserve Backup Manager Console.

From the Quick Start menu, click Restore.

The Restore Manager window opens.


2. From the drop-down list, click Restore by Session.

The backup sessions appear in the Session tree.

3. From the Session tree, expand Sessions and locate the D2D/UDP server that you want to recover.

Expand the server that you want to recover and locate the session from which you want to recover data at file level granularity.

Next, click the check boxes next to the files, folders, or directories that you want to recover, as illustrated by the following screen:



4. Click the Destination tab.

The Destination options appear.

5. Verify that the Restore files to their original locations option is specified and then click the Schedule tab.

The Schedule options appear.

6. Specify the Schedule options that you require for the job.

7. Click Options on the toolbar to specify the global options that you require for the job.

Note: For more information, see [Global Restore Options](#) (see page 271).

8. Click Submit on the toolbar to submit the job.

The Restore Media dialog opens to display a list of the media that is required to perform the operation.

9. If the required media resides on multiple Arcserve Backup servers, click the drop-down list and specify the backup server that you want to execute the job and then click OK.

The Session User Name and Password dialog opens.

10. On the Session User Name and Password dialog, edit or confirm the user name and password that is required to access the sessions and click OK.

The Submit Job dialog opens.

11. Complete the required fields on the Submit Job dialog and click OK.

The job is submitted.

Recover D2D/UDP Data at Application Level Granularity

Arcserve Backup lets you recover Microsoft SQL Server and Microsoft Exchange Server database objects that were backed up from D2D/UDP to Arcserve Backup media. You can recover the database objects at application level granularity.

The steps listed below describe how to perform the following recoveries:

- Recover D2D/UDP nodes that were backed up via D2D/UDP Recovery Point Servers at application level granularity.

Be aware of the following:

- Arcserve Backup lets you recover Microsoft Exchange Server data only at the mailbox level. To recover data at message level granularity, recover the data to a temporary location and then use D2D/UDP to recover at a Granular Recovery Technology (GRT) level. For more details, see the D2D/UDP user guide.
- The following steps describe how to submit restore jobs using the Restore by Session method. The process of restoring D2D/UDP data nodes using the Restore by Tree method is identical to that of all other types of data. For more information about the Restore by Tree method, see "Restoring Data."


Follow these steps:

1. Verify that one of the following conditions exist on the D2D/UDP destination server:
 - You are logged in to the D2D/UDP destination server.
 - The D2D/UDP service is running on the D2D/UDP destination server.
2. Open the Arcserve Backup Manager Console.
From the Quick Start menu, click Restore.
The Restore Manager window opens.
3. From the drop-down list, click Restore by Session.
The backup sessions appear in the Session tree.

- From the Session tree, expand Sessions and locate the D2D/UDP node that you want to recover.

Expand the server that you want to recover and locate the session from which you want to recover Microsoft database data at application level granularity.

Next, click the check boxes next to the objects that you want to recover, as illustrated by the following screen:



- (Optional) To specify options that affect only this restore job, right-click selected objects, select Local Options on the pop-up menu and then do either or both of the following options:

Microsoft SQL Server databases--**Complete the following fields on the D2D/UDP SQL Server Options dialog:**

- **Restore files to their original location**--This option lets you restore data to the original location from where the backup image was captured.
- **Dump files only**--This option lets Arcserve Backup dump the selected Microsoft SQL database files to the specified folder. In addition, you can specify or browse to the folder location where the dump file is restored to.
- **Restore to alternative location**--This option lets you restore data to an alternative location.

Because the backups can be copied to network locations, multiple SQL Server instances can be used. Multiple database restore can be performed (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 browse to the alternate location where the database is restored to.

Microsoft Exchange Server databases--**Complete the following fields on the D2D/UDP Exchange Options dialog:**

- **Restore to Original Location**--This option lets you restore data to the original location from where the backup image was captured.

Note: When restoring the logs folder to the original location, the files that are in the logs folder will be skipped.

- **Dump files only**--This option lets you restore only the dump files. For this option, Arcserve Backup restores the Microsoft Exchange database file to a specified folder, and will not bring it online after recovery is complete. You can then move this file to a different server and mount it to an Exchange server to view the data that is contained in the file.

Note: When a Recovery Mailbox Database exists and you restore it with the 'Dump file only' option, it fails.

Replay log on database--This option lets you specify when the database files are dumped to the destination folder, you can replay and apply all Exchange transaction log files and commit them to the database file.

Note: This option is not applicable for Microsoft Exchange Server 2003 databases.

- **Restore to Recovery Storage Group (Exchange 2007) / Restore to Recovery Database (Exchange 2010/2013)**--This option lets you restore the database to a Recovery Storage Group (RSG) or Recovery Database.

An RSG is a storage group that can be used for recovery purposes. You can restore Microsoft Exchange Server Mailbox Databases from a backup in a Recovery Storage Group and then recover and extract data from it, without affecting the production database that is being accessed by end users.

If a single storage group or database (except a public folder database) from the same storage group are selected to restore, the default restore destination is Restore to Recovery Storage Group (or Restore to Recovery Database).

If multiple storage groups or databases from multiple storage groups are selected to restore, Exchange can be restored only to the original location or restore with Dump file only option. The default restore destination is "Restore to original location".

Before restoring an Exchange 2007 database to a Recovery Storage Group, you must create a Recovery Storage Group and Mailbox Database with the same name.

Example:

If you want to restore MailboxDatabase1 from the First Storage Group to a Recovery Storage Group, you need to create a Recovery Storage Group and add the database "MailboxDatabase1" to the Recovery Storage Group.

Note: This option is not applicable for Microsoft Exchange Server 2003.

- **Dismount Database**--This option lets you dismount the database before you restore and mounts the database after restored.

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 D2D/UDP, these two options are controlled by this "Dismount the database before restore and mount the database after restore" option. With this option, D2D/UDP 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 mounts the database after the restore process is completed. In addition, if checked, this option also allows 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.

Click OK.

6. Click the Destination tab.

The Destination options appear.

7. From the Destination tab, do one of the following:

- Click the check box next to Restore files to their original locations. With this option specified, Arcserve Backup restores the data to the location where the backup was taken.
- Clear the check box next to Restore files to their original locations. With this option specified, you must specify an alternative location. In addition, Arcserve Backup restores the data to the alternative location using the local restore options specified for the database.

Example: Machine-A contains the following database:

C:\Program Files\SQL Server\example.dat

With the Restore files to their original locations option specified, Arcserve Backup restores the database to original path on Machine-A. If you clear the check box next to Restore files to their original locations, specify Machine-B as the alternative location, and the local restore option Restore to Original Location is specified for the database, Arcserve Backup restores the database to the original path (C:\Program Files\SQL Server\example.dat) on Machine-B.

Click the Schedule tab.

The Schedule options appear.

- Specify the Schedule options that you require for the job.
- Click Options on the toolbar to specify the global options that you require for the job.

Note: For more information, see [Global Restore Options](#) (see page 271).

- Click Submit on the toolbar to submit the job.

The Restore Media dialog opens to display a list of the media that is required to perform the operation.

- If the required media resides on multiple Arcserve Backup servers, click the drop-down list and specify the backup server that you want to execute the job and click OK.

The Session User Name and Password dialog opens.

- On the Session User Name and Password dialog, edit or confirm the user name and password that is required to access the sessions and click OK.

The Submit Job dialog opens.

- Complete the required fields on the Submit Job dialog and click OK.

The job is submitted.

Recover Files and Folders from D2D/UDP Recovery Points

D2D/UDP creates point-in-time snapshot images of your backups each time a successful backup is complete. Arcserve Backup lets you open D2D/UDP so that you can find files and folders and submit D2D/UDP restore jobs.

Follow these steps:

- Open the Arcserve Backup Manager Console.

From the Quick Start menu, click Backup.

The Backup Manager window opens.

- Click the Source tab.

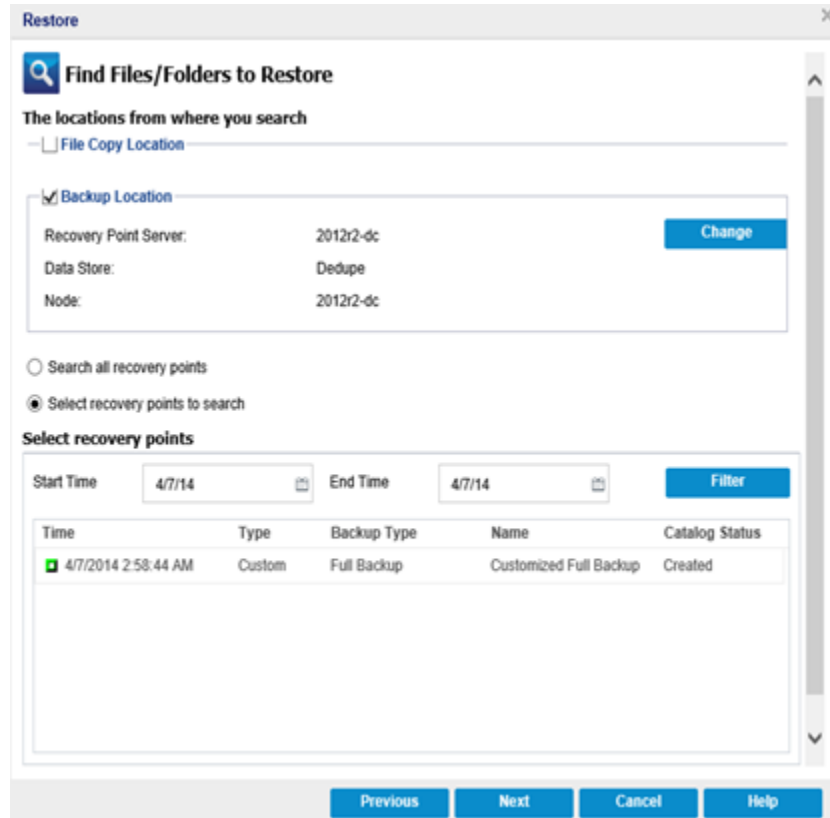
The source tree appears.

- Expand the D2D/UDP Servers object and locate the D2D/UDP server that you want to restore.

Note: If the D2D/UDP server is assigned to a proxy server, expand the D2D/UDP Proxy Servers object to locate the server.

Right-click the D2D/UDP server that you want to recover, point to D2D Restore, and click Find Files/Folders to Restore on the pop-up menu.

The Find Files/Folders to Restore dialog opens.



- To submit a restore, follow the prompts and complete the required fields on the Find Files/Folders to Restore dialog.

Note: For more information about how to submit the restore, click Help on the Find Files/Folders to Restore dialog.

Recover D2D/UDP Data from Raw Sessions

Arcserve Backup lets you back up D2D/UDP recovery points and stores the data on Arcserve Backup media. The D2D/UDP Recovery Points are stored on Arcserve Backup media as raw sessions.

Arcserve Backup lets you recover Arcserve UDP data at file, folder, and application level granularity with respect to Microsoft SQL Server and Microsoft Exchange Server applications. However, if you want to recover individual email messages from the mailbox from Exchange GRT enabled backups, recover the raw sessions to the D2D/UDP server, and then use D2D/UDP to recover the individual email messages from the Exchange GRT sessions from the D2D/UDP server. To recover individual email messages, open the D2D/UDP home page and browse the recovered raw session to locate the objects that you want to recover.

The raw sessions can also be used to perform bare metal recovery (BMR) of D2D/UDP servers. Using the raw sessions, the BMR process is a two-phased approach:

1. Recover the raw session to a shared folder, network file share, or a device that the failed server can access while completing the BMR process.
2. Boot the server that you want to recover using the D2D/UDP BMR media and then browse to the location where you recovered the raw session. Then follow the on-screen steps to complete the BMR process.

Note: For information about recovering files and folders and performing BMR on D2D/UDP servers, see the D2D/UDP documentation.

Be aware of the following:

- Arcserve Backup recovers raw sessions using the same directory structure as the D2D/UDP backup destination.

Example: E:\restore is the alternative location.

- Arcserve Backup recovers raw sessions using the following pattern:
E:\restore\VSTORE\\
- Arcserve Backup recovers the catalog files for raw sessions using the following pattern:
E:\restore\CATALOG\\
- Arcserve Backup lets you restore raw sessions only to alternative locations; not to original locations.

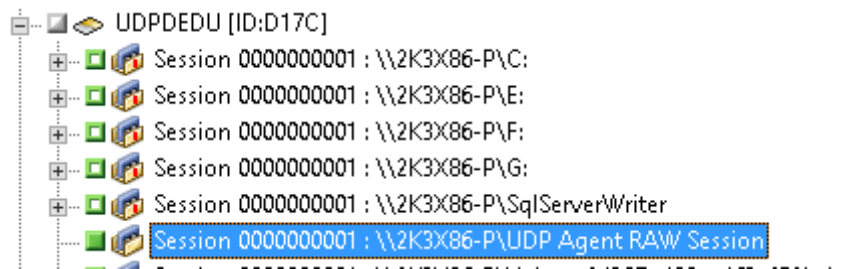
The steps listed below describe how to perform the following recoveries:

- Recover D2D/UDP nodes from raw sessions.

Note: The following steps describe how to submit restore jobs using the Restore by Session method. The process of restoring D2D/UDP Agent-based or agentless nodes using the Restore by Tree method is identical to that of all other types of data. For more information about the Restore by Tree method, see "Restoring Data."

Follow these steps:

1. Open the Arcserve Backup Manager Console.
From the Quick Start menu, click Restore.
The Restore Manager window opens.
2. From the drop-down list, click Restore by Session.
The sessions appear in the Session tree.
3. Locate the D2D/UDP server containing the raw session that you want to recover.



- Click the check box next to the raw session.
Click the Destination tab.
The Destination options appear.
4. On the Destination tab, clear the check box next to Restore files to their original location(s), and specify the alternative location where you want to recover the data.
Click the Schedule tab.
The Schedule options appear.
 5. Specify the Schedule options that you require for the job.
 6. Click Options on the toolbar to specify the global options that you require for the job.
Note: For more information, see [Global Restore Options](#) (see page 271).
 7. Click Submit on the toolbar to submit the job.

The Restore Media dialog opens to display a list of the media that is required to perform the operation.

8. If the required media resides on multiple Arcserve Backup servers, click the drop-down list and specify the backup server that you want to execute the job and then click OK.

The Session User Name and Password dialog opens.

9. On the Session User Name and Password dialog, edit or confirm the user name and password that is required to access the sessions and click OK.

The Submit Job dialog opens.

10. Complete the required fields on the Submit Job dialog and click OK.

The job is submitted.

After the restore job completes, use the sessions that you restored to an alternative location to recover the D2D/UDP data.

Configure Arcserve Backup to Generate D2D/UDP Catalog Files

When Arcserve Backup backs up D2D/UDP backup sessions, the backup process does not generate catalog files for the D2D/UDP sessions. The process of generating catalog files is a time consuming task that can increase the time that is required to back up the D2D/UDP backup sessions (backup window). To help reduce the backup window, (by default) Arcserve Backup does not generate catalog files.

Arcserve Backup requires catalog files to recover D2D/UDP backup sessions at file level, folder level, and application level granularity. When you need to recover D2D/UDP backup sessions at file level, folder level, and application level granularity, you can configure Arcserve Backup to generate the catalog files for the D2D/UDP backup sessions during the backup operation.

The following steps describe how to configure Arcserve Backup to generate and to not generate catalog files during backup operations.

Follow these steps:

1. Log in to the D2D/UDP Recovery Point Server and open Windows Registry Editor.

Important! Log in to the D2D/UDP Recovery Point Server when you back up the node using the D2D/UDP Recovery Point Servers object in the Backup Manager Source tree.

Note: Using the D2D/UDP Recovery Point Server to perform a backup on a D2D/UDP nodes, supports generating catalog files, only if the D2D/UDP Recovery Point Server is available. If the D2D/UDP Recovery Point Server is not available, then generating the catalog file fails.

2. Browse to the following key:

HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA ARCServe
Backup\ClientAgent\Parameters\GenerateCatalog

3. Define the key value of GenerateCatalog of REG_DWORD type as one of the following:
 - Generate catalogs: **1**
 - (Default) Do not generate catalogs: **0**
4. Close Windows Registry Editor.

Chapter 15: Backing Up and Recovering Arcserve UDP Nodes from the Arcserve UDP Recovery Point Server

This section contains the following topics:

[Submit Backup Jobs of Arcserve UDP Nodes through the Arcserve UDP Recovery Point Server](#) (see page 789)

[Submit Static Backup Jobs Containing Arcserve UDP Data](#) (see page 791)

[Submit Dynamic Backup Jobs Containing Arcserve UDP Data](#) (see page 793)

[Recover Arcserve UDP Data at File Level Granularity](#) (see page 795)

[Recover Arcserve UDP Data at Application Level Granularity](#) (see page 797)

[Recover Arcserve UDP Data from Raw Sessions](#) (see page 802)

[Configure Arcserve Backup to Generate Arcserve UDP Catalog Files](#) (see page 805)

Submit Backup Jobs of Arcserve UDP Nodes through the Arcserve UDP Recovery Point Server

Arcserve Backup lets you back up Arcserve Unified Data Protection (UDP) sessions that are protected under the RPS data stores to Arcserve Backup media through a proxy server. The Arcserve Backup media can be file systems devices, tape media, NAS devices, and devices that reside on a SAN.

When you back up Arcserve UDP servers through an Arcserve UDP Recovery Point Server (RPS), Arcserve Backup retrieves all of the Arcserve UDP backup sessions for all of the Arcserve UDP servers that are protected by the Arcserve UDP Recovery Point Server data store.

This backup method lets Arcserve Backup synthesize full backup sessions from the full and incremental Arcserve UDP backup sessions. The synthesized backup sessions can be used to recover Arcserve UDP data at file level, folder level, and application level granularity, and to perform full system recoveries of the Arcserve UDP servers, such as bare metal recovery.

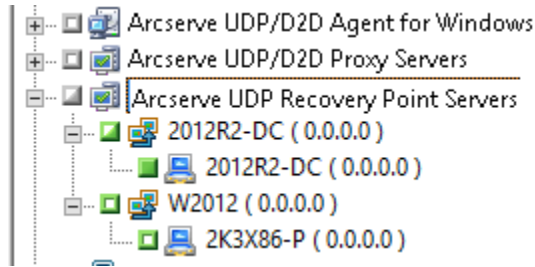
The process of backing up an Arcserve UDP server through an Arcserve UDP Recovery Point Server is almost identical to the steps that are required to back up files, folders, nodes, servers, and so on.

Backing up Arcserve UDP nodes through an Arcserve UDP Recovery Point Server can also back up Arcserve UDP nodes or Arcserve UDP services that are down and not working.

Follow these steps:

1. From the Backup Manager, click the [Start](#) (see page 139) tab and specify the type of backup job that you want to submit.
2. Click the [Source](#) (see page 140) tab to display the Arcserve UDP nodes and proxy servers.

Specify the proxy servers containing the nodes that you want to back up as illustrated by the following screen:




3. Click the Destination tab to specify the location where you want to store the backup data.

Note: To back up the Arcserve UDP sessions using multistreaming, on the Destination tab, click the Multistreaming option and then specify the maximum number of streams that you want to allow for the job.

4. Click the [Schedule](#) (see page 152) tab to specify when and the frequency that you want the job to run.
5. Click the Options toolbar button to specify the backup options that are required for the job. For more information, see [Global Backup Options](#) (see page 154).
6. Click the Submit on the toolbar to submit the job.

Note: If the Arcserve UDP sessions is protected by Arcserve UDP session password, then the following message appears to let you enable the Arcserve Backup encryption.



The Security and Agent Information dialog opens.

7. On the Security and Agent Information dialog, edit or confirm the security and agent information for your job, and click OK.

8. When the Submit Job dialog opens, select Run Now to run the job immediately, or select Run On and select a date and time when you want the job to run.
Note: For more information about the Run Now option, see Job Queue Tab.
9. Enter a Job Name for your job.
10. If you selected multiple sources to back up and want to set the priority in which the job sessions initiate, click Source Priority. Use the Top, Up, Down, and Bottom buttons to change the order in which the jobs are processed. When you finish setting priorities, click OK.
11. To save the job as a Arcserve Backup job script, click the Save Job button.
12. To save the job template, click the Save Template button.
13. To preflight check the job, click the Preflight Check button. If the preflight check failed, click the Cancel button to modify the job settings.
14. On the Submit Job dialog, click OK.
The job is submitted.

Submit Static Backup Jobs Containing Arcserve UDP Data

The Enable Static Backup option lets you submit scheduled backups of source groups and computers and maintain a static set of source volumes. With static backups, you can submit staging and deduplication backups of Arcserve UDP data.

You can submit static backup jobs only to Arcserve UDP Recovery Point Servers that appear below the Arcserve UDP Recovery Point Servers object, which includes Arcserve UDP Agent-based servers and Arcserve UDP agentless VM nodes.

Static backups affect only the immediate subordinate objects of the source group or the computer specified when the job was submitted. Subsequent backups will not include objects and volumes that were added to the source group or the computer after you submitted the backup. Arcserve Backup backs up the original source volumes dynamically. Subsequent job executions will include changes to only the files and folders contained in the original source volumes.

Note: For more information, see [Static Job Packaging](#) (see page 305).

Follow these steps:

1. Open the Backup Manager and click the Start tab.

The backup types appear.


2. Select Normal backup and then click the Source tab.

The Source directory tree appears.


3. Select Classic View from the drop-down list.

Expand the Arcserve UDP Recovery Point Servers object and browse to the Arcserve UDP Recovery Point Servers that you want to back up.

- a. Click the check box next to the Arcserve UDP Recovery Point Server name.
- b. Right-click the Arcserve UDP Recovery Point Server and click Enable Static Backup on the pop-up menu.



The Static Backup option is applied to the computer, as illustrated by the following screen:



Note: The Static Backup option remains applied to the specified computer until you disable this option. You can repeat this step to disable the Enable Static Backup option.

4. Click the Schedule tab and specify the schedule that you want to use for the backup job.

Note: For more information, see [Rotation Schemes](#) (see page 307) and [Custom Schedules](#) (see page 318).

5. Click the Destination tab.

The Destination groups directory tree appears.

6. Specify the group that you want to use to store the backup data.

The storage group is applied.

7. Click Options on the toolbar.

The Options dialog opens.

- Specify the options that you require for the job.

Note: For more information, see [Global Backup Options](#) (see page 154).

Click OK.

The backup options are applied.

- (Optional) Expand the contents of the computer.

Select a drive or volume contained in the computer.

Repeat the previous two steps to apply backup options to the specified drive or volume.

(Optional) Repeat this step for all drives or volumes in the computer.

- Click Submit on the toolbar.

The Submit Job dialog opens.

- Complete the required fields on the Submit Job dialog and click OK.

The job is submitted.

Submit Dynamic Backup Jobs Containing Arcserve UDP Data

Arcserve Backup lets you back up Arcserve UDP data dynamically. Dynamic packaging means that the content of your selection is determined when the job runs.

For example, if you choose to back up a source group or a server, and the nodes or volumes listed on that source group or server change between the time you scheduled the job and the time the job runs, the nodes and volumes changed from the time the job actually runs are backed up.

Follow these steps:

- Open the Backup Manager and click the Start tab.

The backup types appear.

- Select Normal backup and then click the Source tab.

The Source directory tree appears.

- Select Classic View from the drop-down list.

Expand the Arcserve UDP Recovery Point Servers object and browse proxy server containing the Arcserve UDP Agent-based or agentless nodes that you want to back up.

Click the check box next to the Arcserve UDP Recovery Point Server containing the Arcserve UDP nodes that you want to back up.

4. Click the Schedule tab and specify the schedule that you want to use for the backup job.

Note: For more information, see [Rotation Schemes](#) (see page 307) and [Custom Schedules](#) (see page 318).

5. Click the Destination tab.

The Destination groups directory tree appears.

6. Specify the group that you want to use to store the backup data.

The storage group is applied.

7. Click Options on the toolbar.

The Options dialog opens.

8. Specify the options that you require for the job.

Note: For more information, see [Global Backup Options](#) (see page 154).

Click OK.

The backup options are applied.

9. (Optional) Expand the contents of the computer.

Select a drive or volume contained in the computer.

Repeat the previous two steps to apply backup options to the specified drive or volume.

(Optional) Repeat this step for all drives or volumes in the computer.

10. Click Submit on the toolbar.

The Submit Job dialog opens.

11. Complete the required fields on the Submit Job dialog and click OK.

The job is submitted.

Recover Arcserve UDP Data at File Level Granularity

Arcserve Backup lets you recover Arcserve UDP backup data at file level granularity from Arcserve Backup media. Use these steps when you want to recover individual files, folders, directories, and so on from Arcserve Backup media to Arcserve UDP nodes.

The steps listed below describe how to perform the following recoveries:

- Recover Arcserve UDP data that was backed up via Arcserve UDP Recovery Point Servers at file level granularity.

Note: The following steps describe how to submit restore jobs using the Restore by Session method. The process of restoring Arcserve UDP data nodes using the Restore by Tree method is identical to that of all other types of data. For more information about the Restore by Tree method, see "Restoring Data."

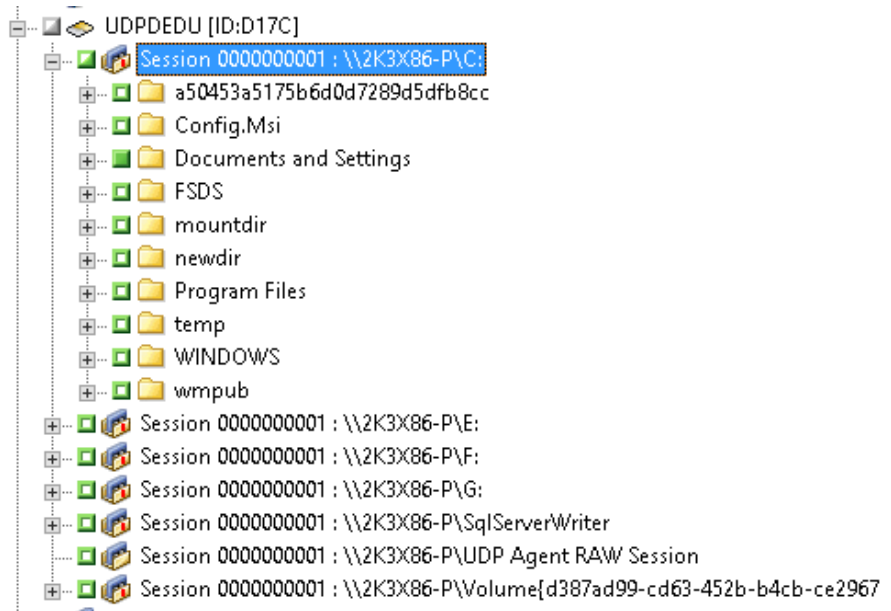
Follow these steps:

1. Open the Arcserve Backup Manager Console.
From the Quick Start menu, click Restore.
The Restore Manager window opens.
2. From the drop-down list, click Restore by Session.
The backup sessions appear in the Session tree.

3. From the Session tree, expand Sessions and locate the Arcserve UDP nodes that you want to recover.

Expand the nodes that you want to recover and locate the session from which you want to recover data at file level granularity.

Next, click the check boxes next to the files, folders, or directories that you want to recover, as illustrated by the following screen:



4. Click the Destination tab.
The Destination options appear.
5. Verify that the Restore files to their original locations option is specified and then click the Schedule tab.
The Schedule options appear.
6. Specify the Schedule options that you require for the job.
7. Click Options on the toolbar to specify the global options that you require for the job.

Note: For more information, see [Global Restore Options](#) (see page 271).

8. Click Submit on the toolbar to submit the job.
The Restore Media dialog opens to display a list of the media that is required to perform the operation.
9. If the required media resides on multiple Arcserve Backup servers, click the drop-down list and specify the backup server that you want to execute the job and then click OK.

The Session User Name and Password dialog opens.

10. On the Session User Name and Password dialog, edit or confirm the user name and password that is required to access the sessions and click OK.

If the Arcserve UDP session is protected by a session password then the Arcserve UDP session password is required. (The Arcserve UDP session password is the password that was given in the Arcserve UDP Recovery Point Server plan for the respective Arcserve UDP sessions backed up to tape).

The Submit Job dialog opens.

11. Complete the required fields on the Submit Job dialog and click OK.

The job is submitted.

Recover Arcserve UDP Data at Application Level Granularity

Arcserve Backup lets you recover Microsoft SQL Server and Microsoft Exchange Server database objects that were backed up from Arcserve UDP to Arcserve Backup media. You can recover the database objects at application level granularity.

The steps listed below describe how to perform the following recoveries:

- Recover Arcserve UDP nodes that was backed up via Arcserve UDP Recovery Point Servers at application level granularity.

Be aware of the following:

- Arcserve Backup lets you recover Microsoft Exchange Server data only at the mailbox level. To recover data at message level granularity, recover the data to a temporary location and then use Arcserve UDP to recover at a Granular Recovery Technology (GRT) level. For more details, see the Arcserve UDP user guide.
- The following steps describe how to submit restore jobs using the Restore by Session method. The process of restoring Arcserve UDP node data using the Restore by Tree method is identical to that of all other types of data. For more information about the Restore by Tree method, see "Restoring Data."

Follow these steps:

1. Verify that one of the following conditions exist on the Arcserve UDP destination server:

- You are logged in to the Arcserve UDP destination server.
- The Arcserve UDP service is running on the Arcserve UDP destination server.

2. Open the Arcserve Backup Manager Console.

From the Quick Start menu, click Restore.

The Restore Manager window opens.


3. From the drop-down list, click Restore by Session.

The backup sessions appear in the Session tree.

4. From the Session tree, expand Sessions and locate the Arcserve UDP node that you want to recover.

Expand the server that you want to recover and locate the session from which you want to recover Microsoft database data at application level granularity.

Next, click the check boxes next to the objects that you want to recover, as illustrated by the following screen:



5. (Optional) To specify options that affect only this restore job, right-click selected objects, select Local Options on the pop-up menu and then do either or both of the following options:

Microsoft SQL Server databases--**Complete the following fields on the Arcserve UDP SQL Server Options dialog:**

- **Restore files to their original location**--This option lets you restore data to the original location from where the backup image was captured.
- **Dump files only**--This option lets Arcserve Backup dump the selected Microsoft SQL database files to the specified folder. In addition, you can specify or browse to the folder location where the dump file is restored to.
- **Restore to alternative location**--This option lets you restore data to an alternative location.

Because the backups can be copied to network locations, multiple SQL Server instances can be used. Multiple database restore can be performed (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 browse to the alternate location where the database is restored to.

Microsoft Exchange Server databases--**Complete the following fields on the Arcserve UDP Exchange Options dialog:**

- **Restore to Original Location**--This option lets you restore data to the original location from where the backup image was captured.

Note: When restoring the logs folder to the original location, the files that are in the logs folder will be skipped.

- **Dump files only**--This option lets you restore only the dump files. For this option, Arcserve Backup restores the Microsoft Exchange database file to a specified folder, and will not bring it online after recovery is complete. You can then move this file to a different server and mount it to an Exchange server to view the data that is contained in the file.

Note: When a Recovery Mailbox Database exists and you restore it with the 'Dump file only' option, it fails.

Replay log on database--This option lets you specify when the database files are dumped to the destination folder, you can replay and apply all Exchange transaction log files and commit them to the database file.

Note: This option is not applicable for Microsoft Exchange Server 2003 databases.

- **Restore to Recovery Storage Group (Exchange 2007) / Restore to Recovery Database (Exchange 2010/2013)**--This option lets you restore the database to a Recovery Storage Group (RSG) or Recovery Database.

An RSG is a storage group that can be used for recovery purposes. You can restore Microsoft Exchange Server Mailbox Databases from a backup in a Recovery Storage Group and then recover and extract data from it, without affecting the production database that is being accessed by end users.

If a single storage group or database (except a public folder database) from the same storage group are selected to restore, the default restore destination is Restore to Recovery Storage Group (or Restore to Recovery Database).

If multiple storage groups or databases from multiple storage groups are selected to restore, Exchange can be restored only to the original location or restore with Dump file only option. The default restore destination is "Restore to original location".

Before restoring an Exchange 2007 database to a Recovery Storage Group, you must create a Recovery Storage Group and Mailbox Database with the same name.

Example:

If you want to restore MailboxDatabase1 from the First Storage Group to a Recovery Storage Group, you need to create a Recovery Storage Group and add the database "MailboxDatabase1" to the Recovery Storage Group.

Note: This option is not applicable for Microsoft Exchange Server 2003.

- **Dismount Database**--This option lets you dismount the database before you restore and mounts the database after restored.

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, these two options are controlled by this "Dismount the database before restore and mount the database after restore" option. With this option, Arcserve UDP 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 mounts the database after the restore process is completed. In addition, if checked, this option also allows 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.

Click OK.

6. Click the Destination tab.

The Destination options appear.

7. From the Destination tab, do one of the following:

- Click the check box next to Restore files to their original locations. With this option specified, Arcserve Backup restores the data to the location where the backup was taken.
- Clear the check box next to Restore files to their original locations. With this option specified, you must specify an alternative location. In addition, Arcserve Backup restores the data to the alternative location using the local restore options specified for the database.

Example: Machine-A contains the following database:

C:\Program Files\SQL Server\example.dat

With the Restore files to their original locations option specified, Arcserve Backup restores the database to original path on Machine-A. If you clear the check box next to Restore files to their original locations, specify Machine-B as the alternative location, and the local restore option Restore to Original Location is specified for the database, Arcserve Backup restores the database to the original path (C:\Program Files\SQL Server\example.dat) on Machine-B.

Click the Schedule tab.

The Schedule options appear.

- Specify the Schedule options that you require for the job.
- Click Options on the toolbar to specify the global options that you require for the job.

Note: For more information, see [Global Restore Options](#) (see page 271).

- Click Submit on the toolbar to submit the job.

The Restore Media dialog opens to display a list of the media that is required to perform the operation.

- If the required media resides on multiple Arcserve Backup servers, click the drop-down list and specify the backup server that you want to execute the job and click OK.

The Session User Name and Password dialog opens.

- On the Session User Name and Password dialog, edit or confirm the user name and password that is required to access the sessions and click OK.

The Submit Job dialog opens.

- Complete the required fields on the Submit Job dialog and click OK.

The job is submitted.

Recover Arcserve UDP Data from Raw Sessions

Arcserve Backup lets you back up Arcserve UDP recovery points and store the data on Arcserve Backup media. The Arcserve UDP Recovery Points are stored on Arcserve Backup media as raw sessions.

Arcserve Backup lets you recover Arcserve UDP data at file, folder, and application level granularity with respect to Microsoft SQL Server and Microsoft Exchange Server applications. However, if you want to recover individual email messages from the mailbox from Exchange GRT enabled backups, recover the raw sessions to the Arcserve UDP server, and then use Arcserve UDP to recover the individual email messages from the Exchange GRT sessions from the Arcserve UDP server. To recover individual email messages, open the Arcserve UDP home page and browse the recovered raw session to locate the objects that you want to recover.

The raw sessions can also be used to perform bare metal recovery (BMR) of Arcserve UDP servers. Using the raw sessions, the BMR process is a two-phased approach:

1. Recover the raw session to a shared folder, network file share, or a device that the failed server can access while completing the BMR process.
2. Boot the server that you want to recover using the Arcserve UDP BMR media and then browse to the location where you recovered the raw session. Then follow the on-screen steps to complete the BMR process.

Note: For information about recovering files and folders and performing BMR on Arcserve UDP servers, see the Arcserve UDP documentation.

Be aware of the following:

- Arcserve Backup recovers raw sessions using the same directory structure as the Arcserve UDP backup destination.

Example: E:\restore is the alternative location.

- Arcserve Backup recovers raw sessions using the following pattern:

E:\restore\VSTORE\\

- Arcserve Backup recovers the catalog files for raw sessions using the following pattern:

E:\restore\CATALOG\\

- Arcserve Backup lets you restore raw sessions only to alternative locations; not to original locations.

The steps listed below describe how to perform the following recoveries:


- Recover Arcserve UDP nodes from raw sessions.

Note: The following steps describe how to submit restore jobs using the Restore by Session method. The process of restoring Arcserve UDP Agent-based and agentless nodes using the Restore by Tree method is identical to that of all other types of data. For more information about the Restore by Tree method, see "Restoring Data."

Follow these steps:

1. Open the Arcserve Backup Manager Console.
From the Quick Start menu, click Restore.
The Restore Manager window opens.
2. From the drop-down list, click Restore by Session.
The sessions appear in the Session tree.

3. Locate the Arcserve UDP server containing the raw session that you want to recover.



Click the check box next to the raw session.

Click the Destination tab.

The Destination options appear.

4. On the Destination tab, clear the check box next to Restore files to their original location(s), and specify the alternative location where you want to recover the data.

Click the Schedule tab.

The Schedule options appear.

5. Specify the Schedule options that you require for the job.
6. Click Options on the toolbar to specify the global options that you require for the job.

Note: For more information, see [Global Restore Options](#) (see page 271).

7. Click Submit on the toolbar to submit the job.

The Restore Media dialog opens to display a list of the media that is required to perform the operation.

8. If the required media resides on multiple Arcserve Backup servers, click the drop-down list and specify the backup server that you want to execute the job and then click OK.

The Session User Name and Password dialog opens.

9. On the Session User Name and Password dialog, edit or confirm the user name and password that is required to access the sessions and click OK.

The Submit Job dialog opens.

10. Complete the required fields on the Submit Job dialog and click OK.

The job is submitted.

After the restore job completes, use the sessions that you restored to an alternative location to recover the Arcserve UDP data. For information about how to recover individual Arcserve UDP files and folders, and how to perform BMR operations on Arcserve UDP servers, see the Arcserve UDP documentation.

Note: For information about recovering files and folders and performing BMR on Arcserve UDP servers, see the Arcserve UDP documentation.

Configure Arcserve Backup to Generate Arcserve UDP Catalog Files

When Arcserve Backup backs up Arcserve UDP backup sessions, the backup process does not generate catalog files for the Arcserve UDP sessions. The process of generating catalog files is a time consuming task that can increase the time that is required to back up the Arcserve UDP backup sessions (backup window). To help reduce the backup window, (by default) Arcserve Backup does not generate catalog files.

Note: The File System Catalog Generation process is optional and disabled by default on Arcserve UDP 5.0 and later releases.

Arcserve Backup requires catalog files to recover Arcserve UDP backup sessions at file level, folder level, and application level granularity. When you need to recover Arcserve UDP backup sessions at file level, folder level, and application level granularity, you can configure Arcserve Backup to generate the catalog files for the Arcserve UDP backup sessions during the backup operation.

The following steps describe how to configure Arcserve Backup to generate and to not generate catalog files during backup operations.

Follow these steps:

1. Log in to the Arcserve UDP Recovery Point Server and open Windows Registry Editor.

Important! Log in to the Arcserve UDP Recovery Point Server when you back up the node using the Arcserve UDP Recovery Point Servers object in the Backup Manager Source tree.

Note: Using the Arcserve UDP Recovery Point Server to perform a backup on a Arcserve UDP node, supports generating catalog files, only if the Arcserve UDP Recovery Point Server is available. If the Arcserve UDP Recovery Point Server is not available, then generating the catalog file fails.

2. Browse to the following key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA ARCServe  
Backup\ClientAgent\Parameters\GenerateCatalog
```

3. Define the key value of GenerateCatalog of REG_DWORD type as one of the following:
 - Generate catalogs: **1**
 - (Default) Do not generate catalogs: **0**
4. Close Windows Registry Editor.

Chapter 16: Using Synthetic Full Backup

This section contains the following topics:

[Overview of Synthetic Full Backup](#) (see page 807)

[How Synthetic Full Backup Works](#) (see page 808)

[Submit a Synthetic Full Backup Job](#) (see page 811)

[Specify the Synthetic Schedule for a Synthetic Full Backup](#) (see page 813)

[Submit a D2D2T Synthetic Full Backup Job Using a Rotation Scheme](#) (see page 813)

[Submit a D2D2T Synthetic Full Backup Job Using a GFS Scheme](#) (see page 819)

[Submit a D2D2T Synthetic Full Backup Job Using a Custom Schedule](#) (see page 825)

[Run a Real Full Backup on the Next Full Backup Day](#) (see page 830)

[Submit an On-Demand Synthetic Full Backup](#) (see page 831)

[Scan Synthetic Full Backup Session Only](#) (see page 831)

[How the Purge Policy and a DB Pruning Job Works for Synthetic Full Backup](#) (see page 832)

[How Makeup Jobs Work for Synthetic Full Backup](#) (see page 832)

[Point In Time Restore](#) (see page 833)

[Synthetic Full Backup Considerations](#) (see page 836)

[Working with Synthetic Full Backup](#) (see page 837)

Overview of Synthetic Full Backup

Synthetic full backup (SFB) lets you synthesize a previous full session and its subsequent incremental sessions to a full session.

Using the synthetic full backup option you can:

- Enable synthetic full backup and select the synthetic schedule.
- Synthesize incremental sessions to a new synthetic full backup session.
- Generate a real full backup based on the synthetic schedule for data assurance.
- Run a real full backup on the next full backup day without waiting for the scheduled full backup to run.
- Submit on demand synthetic full backups without waiting for the scheduled data synthetic job to run.
- Allow scans of synthetic full backup sessions for data assurance.
- Migrate synthetic full backup sessions to tape, cloud, and disk devices using D2D2T.
- Perform point-in-time restores.
- Use the disaster recovery processes and point-in-time restores to recover data. Arcserve Backup lets you perform a disaster recovery using a synthetic full backup session without referring to the previous full or incremental sessions.

SFB is supported on the following applications and platforms:

- Client Agent - All platforms that Arcserve Backup for Windows Client Agent supports. SFB is only applicable to r16 or higher Windows Client Agents (not UNIX/Linux).
- Server - All platforms that Arcserve Backup server supports. Synthetic full backup is a 'job type' that you can submit on any Arcserve Backup r16 servers, primary server, member server, or standalone server.
- **Note:** To perform synthetic full backup jobs, you must install and license Arcserve Backup Agent for Open Files on the computers that you want to back up.

How Synthetic Full Backup Works

You can use Arcserve Backup to perform a synthetic full backup. A synthetic full backup lets you synthesize a previous full session and its subsequent incremental sessions to a full session for Windows Client Agents.

Synthetic full backup lets you perform the following tasks:

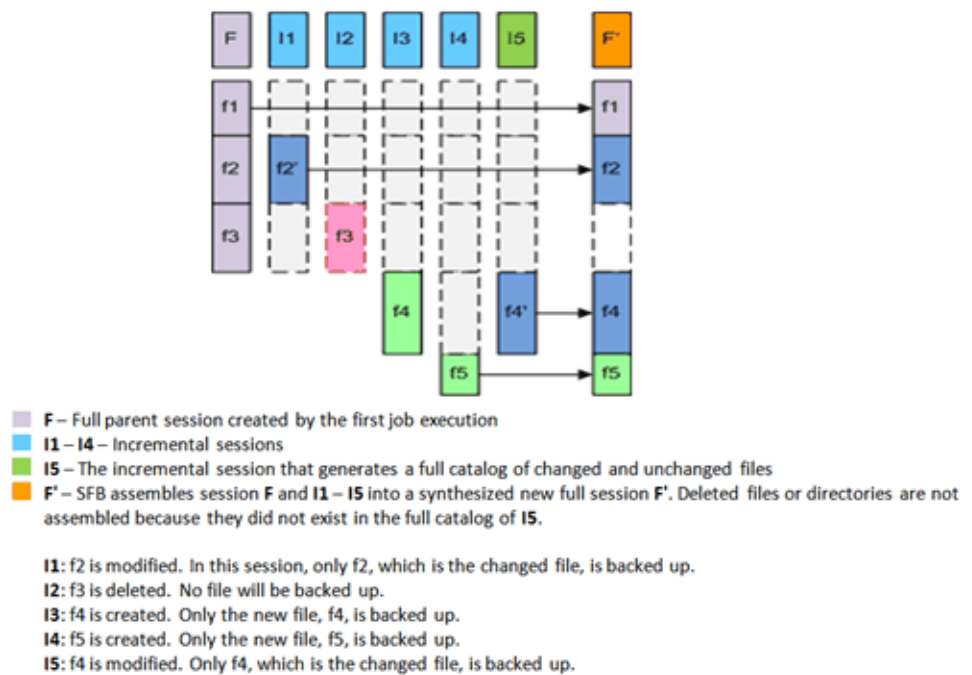
- Synthesize incremental sessions to a new synthetic full backup session.
- Generate real full backups that are based on the synthetic schedule for data assurance.
- Run a real full backup on the next full backup day without waiting for the scheduled full backup to run.
- Submit on demand synthetic full backups without waiting for the scheduled data synthetic job to run.
- Allow scans of synthetic full backup sessions for data assurance.
- Migrate synthetic full backup sessions to tape, cloud, and disk devices using D2D2T.
- Perform point-in-time restores.
- Use the disaster recovery processes and point-in-time restores to recover data. Arcserve Backup lets you perform a disaster recovery using a synthetic full backup session without referring to the previous full or incremental sessions.
- Enable synthetic full backup and select the synthetic schedule.

Note: To include objects in the backup source that resides on servers, you must first enter the servers into the Client Database. The Arcserve Backup software must be running to add servers to the database.

You can perform synthetic full backups on the following applications and platforms:

- Client Agent - All platforms that Arcserve Backup for Windows Client Agent supports. SFB is only applicable to r16 or higher Windows Client Agents (not UNIX/Linux).
- Server - All platforms that Arcserve Backup server supports. Synthetic full backup is a 'job type' that you can submit on any Arcserve Backup r16 servers, primary server, member server, or standalone server.

The following diagram illustrates how Arcserve Backup creates synthetic full backup sessions:



To complete job configuration, make the appropriate selections on the Source, Schedule, Policies, and Destination tabs, and then click Options on the toolbar to set additional backup job properties. If you selected a staging job, you must also specify a staging location. When done, click Submit to submit your job.

Be aware of the following limitations:

- SFB is supported only on disk staging and deduplication devices. You can migrate SFB sessions to tape, cloud, and disk devices with D2D2T.
- SFB supports only Arcserve Backup r16 (or higher) Windows client agents. However, if there are other agents in the backup source in the same job (such as database agents or pre-r16 Windows client agents), on the day of the synthetic full backup, Arcserve Backup uses the original scheduled backup method for the other agents. Arcserve Backup synthesizes only Arcserve Backup r16 (or higher) Windows client agents.

- Detailed information about synthetic full backups can be viewed in the Job Backup Status Report, Node Backup Status Report, and Total Protection Size Reports on the Dashboard.
- To perform synthetic full backup jobs, you install and license Arcserve Backup Agent for Open Files on the computers that you want to back up.
- The following behavior is expected when you modify a synthetic full backup job:
 - When you change the values of the encryption or compression options that are applied to the job, a real full backup job runs on the next day a full backup job is scheduled to run.
 - When you change the values of the synthetic backup job schedule or the purge policies, the purge time for all related sessions is updated. A real full backup job will not run.
 - When you add nodes or disks to the synthetic backup job, a real full backup job runs on only the added nodes and disks. A real full backup job does not run for the job.
- The following behavior is expected when you change the backup destination for a device (group or tape):
 - When the destination device for the backup job is to a data deduplication device, a real full backup job runs on the next day a full backup job is scheduled to run.
 - When the destination device for the backup is a file system device or a data deduplication device with migration to tape and you change the staging location, a real full backup job runs on the next day a full backup job is scheduled to run.
 - When the destination device for the backup is a file system device or a data deduplication device with migration to tape and you change the destination tape, the job runs based on the schedule for the job. (A real full backup does not run the next day that the job is scheduled to run.)

Submit a Synthetic Full Backup Job

The Enable Synthetic Full Backup option lets you submit scheduled backups of source groups and computers. With synthetic full backups, you can submit Normal backups with staging and deduplication backups that consist of your first (parent) full backup and all subsequent incremental backups combined into one session.

Note: Synthetic full backup only supports the Windows Client Agent at r16 or higher. Synthetic full backup is not supported for Windows Client Agents at a version prior to r16, Database Agents, or application agents. If you submit a synthetic full backup job to back up r16 or higher Windows Client Agents, Windows Client Agents prior to r16, Database Agents, or application agents together in one job, it only synthesizes the r16 or higher Windows Client Agent. For the agents and application agents that do not support synthetic full backup, the traditional backup method will be used instead.

Example: How Synthetic Full Backup Works

A synthetic full backup (SFB) lets you synthesize a previous full session and its subsequent incremental sessions to a full session for Windows Client Agents.

Note: To perform synthetic full backup jobs, you must install and license Arcserve Backup Agent for Open Files on the computers that you want to back up. If you do not install and license the agent, Arcserve Backup converts synthetic full backups to regular backups, and produces a warning message in the Activity Log.

- Submit a SFB enabled backup job.
- When the job launches, the first execution will be a real full backup and subsequent backups will be incremental backups. The full session will be the first parent full session, and subsequent incremental backups execute according to the incremental backup policy you have defined.
- When the last incremental backup job executes, it queries Client Agent for all changed and unchanged files. After the last incremental backup, the Data Synthetic job is created.
- The Data Synthetic job assembles the last full session (parent session) and subsequent incremental sessions into one SFB session.
- The new SFB session becomes the parent full session for the next SFB job (unless a real full backup session is executed after the SFB session).

To submit a synthetic full backup job

1. Open the Backup Manager and click the Start tab.
2. Select either Normal backup or Deduplication backup, and Enable Synthetic Full Backup.

Note: When you select Normal backup and Enable Synthetic Full Backup, Enable Staging is selected by default; if you disable the Enable Staging option, you cannot select the Enable Synthetic Full Backup option.

3. Select the [Source](#) (see page 140), [Schedule](#) (see page 152), [Policies](#) (see page 813), and Destination tabs to specify the options that you require for the job. If you selected a staging job, you must also specify staging location.
4. Click the Options toolbar button to specify global options that you require for the job. For more information, see [Global Backup Options](#) (see page 154).
5. Click the Submit toolbar button to submit your job.

The Security and Agent Information dialog opens.

6. Edit or confirm the security and agent information for your job, and click OK.
7. When the Submit Job dialog opens, select Run Now to run the job immediately, or select Run On and select a date and time when you want the job to run.

Note: For more information about the Run Now option, see Job Queue Tab.

8. Enter a description for your job.
9. If you selected multiple sources to back up and want to set the priority in which the job sessions initiate, click Source Priority. Use the Top, Up, Down, and Bottom buttons to change the order in which the jobs are processed. When you finish setting priorities, click OK.
10. To save the job as a Arcserve Backup job script, click the Save Job button.
11. To save the job template, click the Save Template button.
12. To preflight check the job, click the Preflight Check button. If the preflight check failed, click the Cancel button to modify the job settings.
13. On the Submit Job dialog, click OK.

The job is submitted.

Note: For information about how to use disk staging (D2D2T) and tape staging (D2T2T) to manage your backup operations, see How Backup to Disk to Tape Works.

Specify the Synthetic Schedule for a Synthetic Full Backup

Arcserve Backup lets you specify the synthetic schedule for synthetic full backups. The synthetic schedule lets you define when to execute synthetic full backup jobs and options for specifying the schedule for real full backup jobs.

To specify the synthetic schedule for a synthetic full backup

1. Open the Backup Manager and click the Start tab.
On the Start tab, select either Normal backup or Deduplication backup, and Enable Staging and Enable Synthetic Full Backup.
2. Click the Schedule tab.
3. Specify the following synthetic schedule and options as required for the job:

Synthetic after backup job

Lets you specify the amount of time in hours and minutes that must elapse before Arcserve Backup executes a synthetic full backup (not a real full backup) after the last incremental backup job.

Synthetic Schedule

Lets you specify the synthetic schedule. You specify the number of weeks or months that elapse before a synthetic full backup is executed, and the day of the week to execute the synthetic full backup. Based on the schedule type, you can specify the day of the week and the week of the month to execute the synthetic full backup.

Do full backup

Lets you specify whether to execute a real full backup (not a synthetic full backup), and the frequency of this backup.

Note: Specifying a real full backup changes a backup method (incremental/differential) to a real full backup.

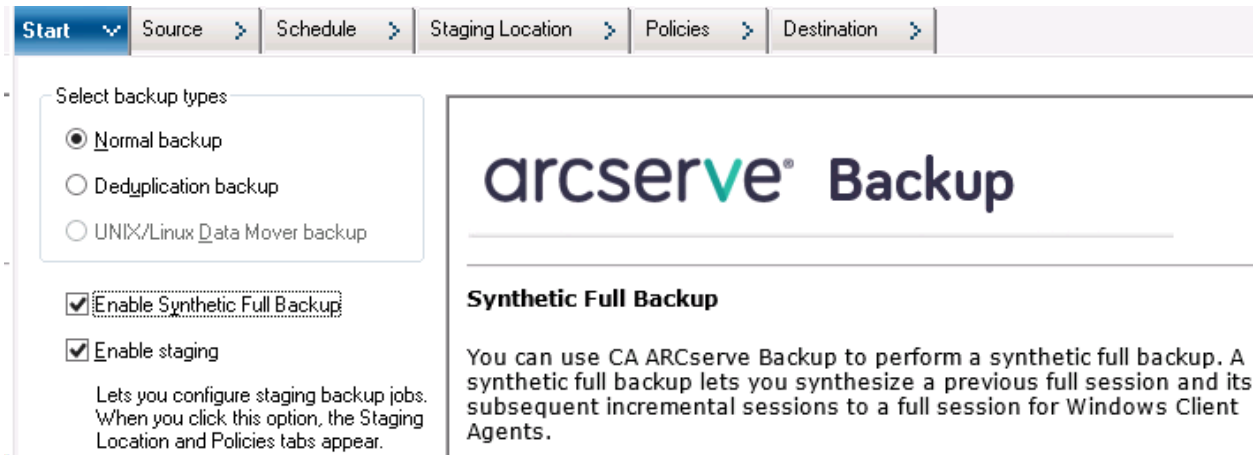
Submit a D2D2T Synthetic Full Backup Job Using a Rotation Scheme

The Backup Manager provides you with scheduling options and rotation schemes to help you establish an automatic backup strategy. You can specify a D2D2T synthetic full backup job using a rotation scheme that suits your environment.

Note: To perform synthetic full backup jobs, you must install and license Arcserve Backup Agent for Open Files on the computers that you want to back up. If you do not install and license the agent, Arcserve Backup converts synthetic full backups to regular backups, and produces a warning message in the Activity Log.

To submit a D2D2T synthetic full backup job using a rotation scheme

1. Open the Backup Manager and click the Start tab.
2. Select Normal backup and Enable Synthetic Full Backup. The Enable Staging option is selected by default.




Be aware of the following:

- When you select Normal backup and Enable Synthetic Full Backup, Enable Staging is selected by default. This is because synthetic full backup can only be enabled for Normal Backup with the Enable Staging option. You can migrate the data to Tape, Cloud, and Disk devices.
- When you select Deduplication backup and Enable Synthetic Full Backup, Enable Staging is not selected by default. This is because synthetic full backup supports submitting to a deduplication device without the Enable Staging option. However, you can also submit a synthetic full backup job to a deduplication device using the Enable Staging option.
- Synthetic full backup does not support Unix/Linux Data Mover backup.

3. Click the Source tab and select the backup source.

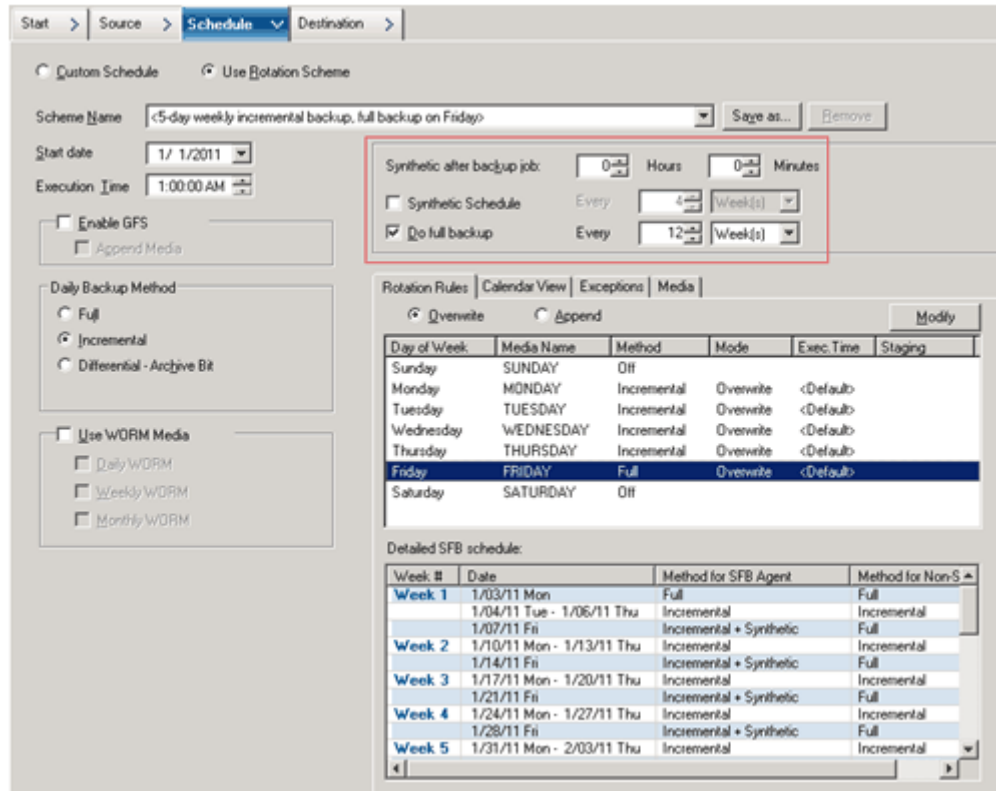
Note: Synthetic full backup only supports r16 or higher Windows Client Agents. You can also select other agents (such as pre-r16 Windows Client Agents, database agents, or application agents) along with r16 or higher Windows Client Agents to submit a mixed job. In this mixed job, it will not change the backup method of those agents that do not support synthetic full backup.

- Click the Schedule tab, select the Use Rotation Scheme option, and select the <5-day weekly incremental backup, full backup on Friday> scheme.



Note: No matter what daily backup method you select, it executes an incremental backup for file system agents except on the first day and on the days that match the do full backup schedule. On the days that match the do full backup schedule, it executes a full backup for the file system agents.

5. Set the Synthetic Schedule.



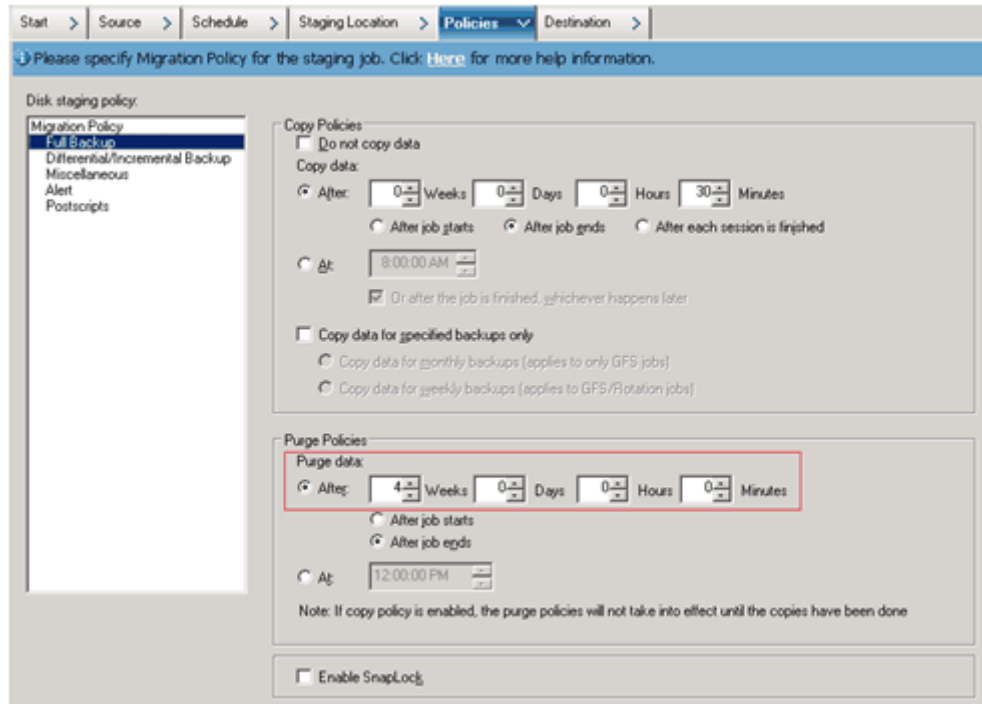
Be aware of the following:

- **Synthetic after backup job:** After the last incremental backup job completes, a data synthetic job will be generated as a run-on job. To avoid multiple jobs competing for CPU/Disk resources, you can select when to launch the data synthetic job after the last incremental backup job is complete.
- **Synthetic Schedule:** You can set the schedule for how long to run a data synthetic job to generate a synthetic full backup session. If this option is not selected, the synthetic schedule will be set as every one (1) week, which is the default setting. In this scenario, a data synthetic job will run every Friday. You can also select the Synthetic Schedule check box to set the Synthetic Schedule as every *n* weeks.
- **Do Full Backup:** You can set the schedule for how long to run a traditional full backup for data assurance.

Ensure the Do Full Backup Schedule is longer than the Synthetic Schedule when you enable the Do Full Backup option. Otherwise, the data synthetic job may not have a chance to run. You can check the Detailed SFB Schedule to know when the data synthetic job will run, when the full backup job will run, and when the incremental backup job will run.

6. Click the Staging Location tab and select the staging location for your job.

7. Click the Policies tab and set the Migration Policy.



8. Click the Destination tab, select tape as your destination, and specify the Media pool name.

9. Click the Options toolbar button to specify global options that you require for the job. For more information, see [Global Backup Options](#) (see page 154).

10. Click the Submit toolbar button to submit your job.

The Security and Agent Information dialog opens.

11. Edit or confirm the security and agent information for your job, and click OK.

The Submit Job dialog opens.

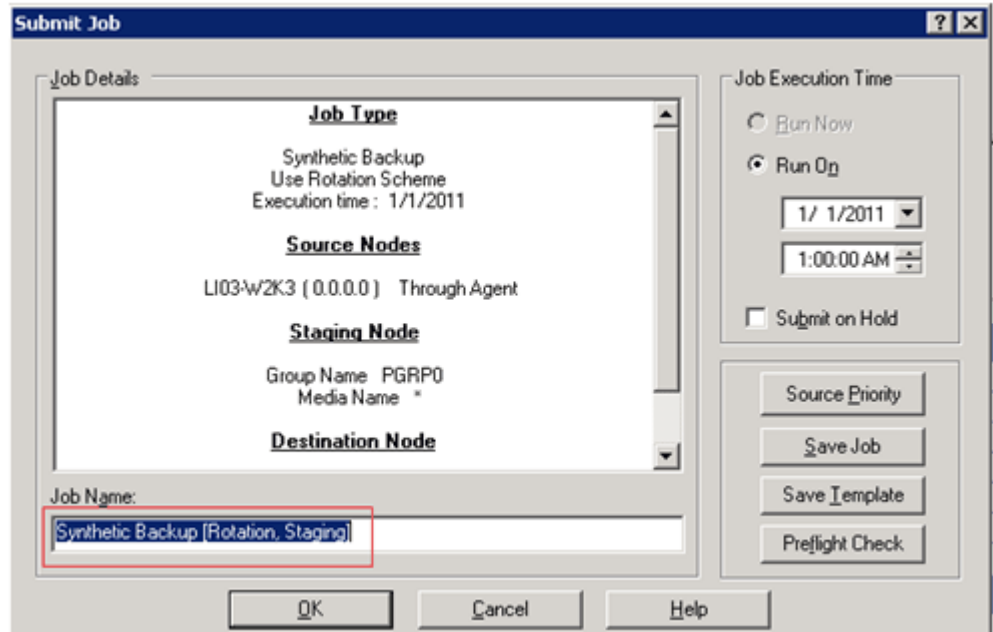
12. Specify the options that you require for the job:

Run Now

Specifies to run the job immediately.

Run On

Lets you specify a date and time when you want the job to run, and enter a description for your job. The job name for the Synthetic Full Backup will start with *Synthetic Backup*:



Source Priority

Lets you specify the priority in which job sessions initiate if you selected multiple sources to back up. Use the Top, Up, Down, and Bottom buttons to change the order in which the jobs are processed.

Save Job

Saves the job as a Arcserve Backup job script.

Save Template

Saves the job template.

Preflight Check

Specifies to preflight check the job. If the preflight check fails, click the Cancel button to modify the job setting.

- Click OK when you are done specifying your settings.

The job is submitted to the Job Queue.

Job Name	Backup Se...	Job No.	Job ID	Status	Execution Time	Job Type	Last Result
Database protection job	LI03-W...	2		HOLD	12/23/2010 11:00 AM	Backup (Ro...	
Database pruning job	LI03-W...	1	95	READY	1/08/2011 12:00 PM	DB Pruning	Finished
Synthetic Backup [Rotation, Staging]	LI03-W...	13	56	READY	1/10/2011 1:00 AM	Synthetic B...	Finished
Data Migration Job	LI03-W...	14	59	DONE	1/07/2011 7:27 PM	Migration	Finished
Data Migration Status	LI03-W...						
Data Synthetic Job	LI03-W...	8	58	DONE	1/07/2011 7:27 PM	Synthetic B...	Finished

Note: Similar to a migration job, after the last incremental backup completes, a new data synthetic job will be created and added into the job queue under the master job. If the data synthetic job fails, a makeup job will be created and added into the job queue.

Submit a D2D2T Synthetic Full Backup Job Using a GFS Scheme

The Grandfather-Father-Son (GFS) rotation strategy lets you maintain backups on a daily, weekly, and monthly basis. You can specify a D2D2T synthetic full backup job using this strategy.

Note: To perform synthetic full backup jobs, you must install and license Arcserve Backup Agent for Open Files on the computers that you want to back up. If you do not install and license the agent, Arcserve Backup converts synthetic full backups to regular backups, and produces a warning message in the Activity Log.

To submit a D2D2T synthetic full backup job using a GFS scheme

- Open the Backup Manager and click the Start tab.
- Select Normal backup and Enable Synthetic Full Backup. The Enable Staging option is selected by default.

Start | Source | Schedule | Staging Location | Policies | Destination

Select backup types

- Normal backup
- Deduplication backup
- UNIX/Linux Data Mover backup

Enable Synthetic Full Backup

Enable staging

Lets you configure staging backup jobs. When you click this option, the Staging Location and Policies tabs appear.

arcserve Backup

Synthetic Full Backup

You can use CA ARCserve Backup to perform a synthetic full backup. A synthetic full backup lets you synthesize a previous full session and its subsequent incremental sessions to a full session for Windows Client Agents.

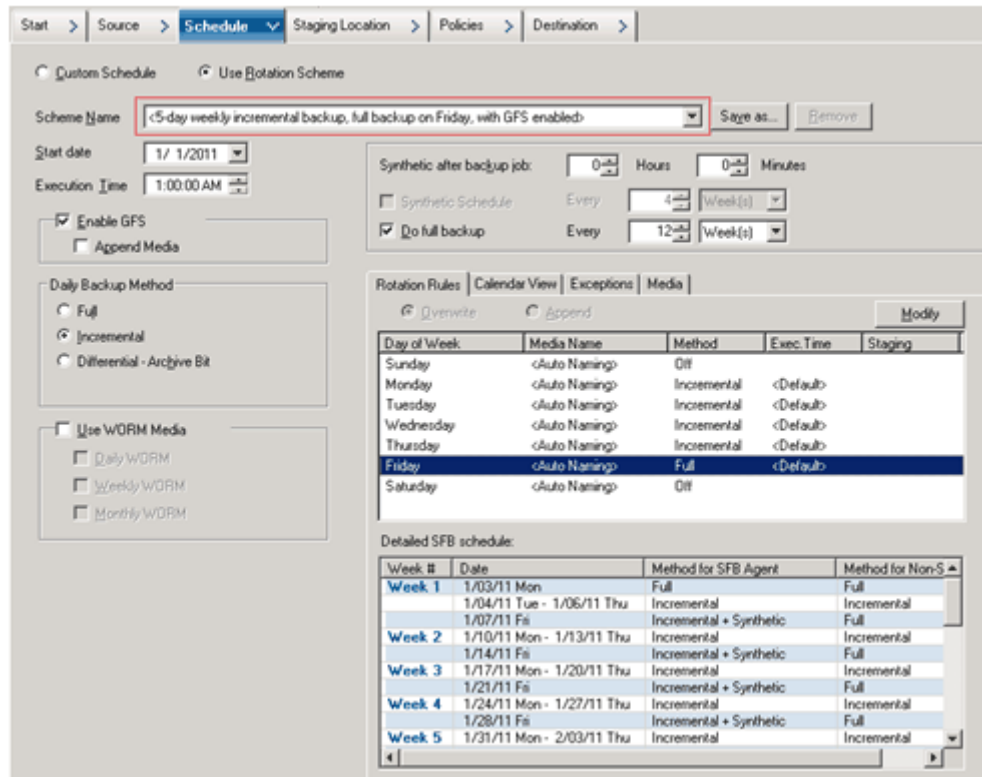
Be aware of the following:

- When you select Normal backup and Enable Synthetic Full Backup, Enable Staging is selected by default. This is because synthetic full backup can only be enabled for Normal backup with the Enable Staging option. You can migrate the data to Tape, Cloud, and Disk devices.
- When you select Deduplication backup and Enable Synthetic Full Backup, Enable Staging is not selected by default. This is because synthetic full backup supports submitting to a deduplication device without the Enable Staging option. However, you can also submit a synthetic full backup job to a deduplication device using the Enable Staging option.
- Synthetic full backup does not support Unix/Linux Data Mover backup.

3. Click the Source tab and select the backup source.

Note: Synthetic full backup only supports r16 or higher Windows Client Agents. However, you can also select other agents (such as pre-r16 Windows Client Agents, database agents, or application agents) along with r16 or higher Windows Client Agents to submit a mixed job. In this mixed job, it will not change the backup method of those agents that do not support synthetic full backup.

4. Click the Schedule tab, select the Use Rotation Scheme option, and select the <5-day weekly incremental backup, full backup on Friday, with GFS enabled> scheme:



Note: No matter what daily backup method you select, it executes an incremental backup for file system agents except on the first day and on the days that match the do full backup schedule. On the days that match the do full backup schedule, it executes a full backup for the file system agents.

5. Set the Synthetic schedule.

The screenshot shows the Backup Exec console configuration for a synthetic full backup job. The 'Synthetic after backup job' section is highlighted with a red box, showing the following settings:

- Synthetic after backup job: 0 Hours, 0 Minutes
- Synthetic Schedule: Every 4 Week(s)
- Do full backup: Every 12 Week(s)

The 'Rotation Rules' table shows the backup schedule for each day of the week:

Day of Week	Media Name	Method	Exec. Time	Staging
Sunday	<Auto Naming>	Off		
Monday	<Auto Naming>	Incremental	<Default>	
Tuesday	<Auto Naming>	Incremental	<Default>	
Wednesday	<Auto Naming>	Incremental	<Default>	
Thursday	<Auto Naming>	Incremental	<Default>	
Friday	<Auto Naming>	Full	<Default>	
Saturday	<Auto Naming>	Off		

The 'Detailed SFB schedule' table shows the sequence of backup jobs over five weeks:

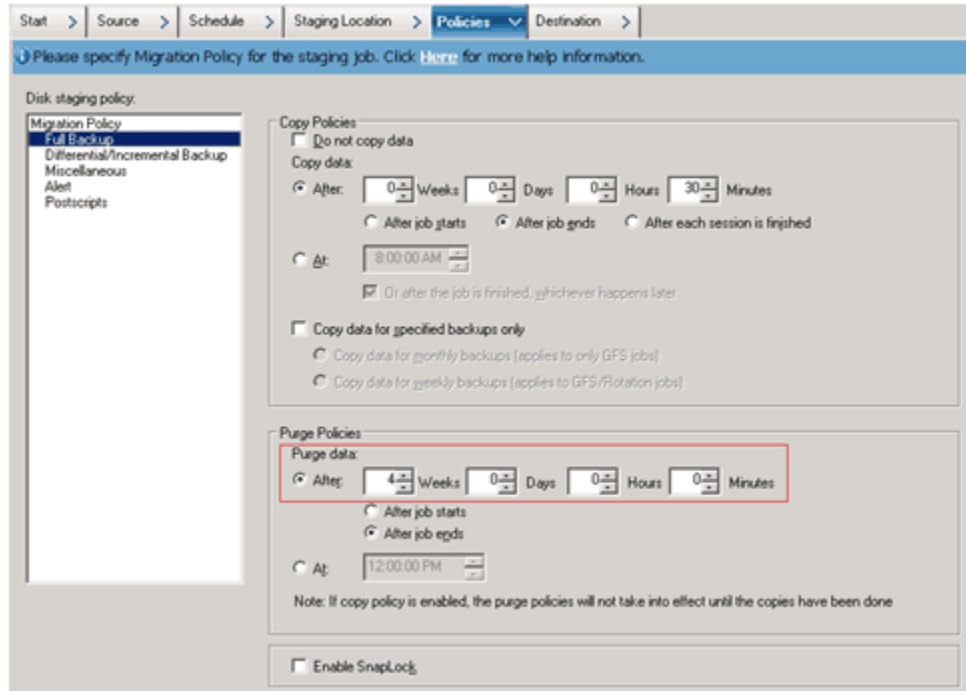
Week #	Date	Method for SFB Agent	Method for Non-S
Week 1	1/03/11 Mon	Full	Full
	1/04/11 Tue - 1/06/11 Thu	Incremental	Incremental
	1/07/11 Fri	Incremental + Synthetic	Full
Week 2	1/10/11 Mon - 1/13/11 Thu	Incremental	Incremental
	1/14/11 Fri	Incremental + Synthetic	Full
Week 3	1/17/11 Mon - 1/20/11 Thu	Incremental	Incremental
	1/21/11 Fri	Incremental + Synthetic	Full
Week 4	1/24/11 Mon - 1/27/11 Thu	Incremental	Incremental
	1/28/11 Fri	Incremental + Synthetic	Full
Week 5	1/31/11 Mon - 2/03/11 Thu	Incremental	Incremental

Be aware of the following:

- **Synthetic after backup job:** After the last incremental backup job completes, a data synthetic job will be generated as a run-on job. To avoid multiple jobs competing for CPU/Disk resources, you can select when to launch the data synthetic job after the last incremental backup job is complete.
- **Synthetic Schedule:** This option is not applicable when Enable GFS is selected. This means you cannot specify a Synthetic Schedule as n weeks for a GFS job, and the Synthetic Schedule will be set as every one (1) week. In this scenario, the data synthetic job will run every Friday.
- **Do Full Backup:** You can set the schedule for how long to run a traditional full backup for data assurance.

Ensure that the Do Full Backup Schedule is longer than the Synthetic Schedule when you enable the Do Full Backup option. You can check the Detailed SFB Schedule to know when the synthetic job will run, when the full backup job will run, and when the incremental backup job will run.

6. Click the Staging Location tab and select the staging location for your job.
7. Click the Policies tab and set the Migration Policy.



8. Click the Destination tab, select tape as your destination, and specify the Media pool name.
9. Click the Options toolbar button to specify global options that you require for the job. For more information, see [Global Backup Options](#) (see page 154).
10. Click the Submit toolbar button to submit your job.
The Security and Agent Information dialog opens.
11. Edit or confirm the security and agent information for your job, and click OK.
The Submit Job dialog opens.

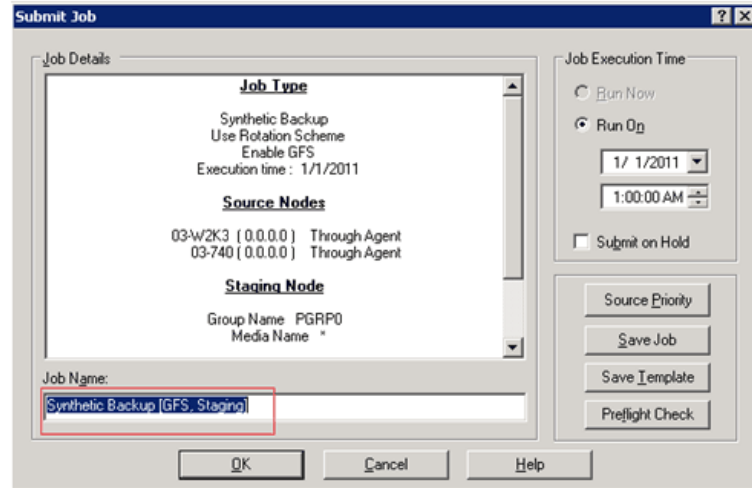
12. Specify the options that you require for the job:

Run Now

Specifies to run the job immediately.

Run On

Lets you specify a date and time when you want the job to run, and enter a description for your job. The job name for the Synthetic Full Backup will start with *Synthetic Backup*:



Source Priority

Lets you specify the priority in which job sessions initiate if you selected multiple sources to back up. Use the Top, Up, Down, and Bottom buttons to change the order in which the jobs are processed.

Save Job

Saves the job as a Arcserve Backup job script.

Save Template


Saves the job template.

Preflight Check

Specifies to preflight check the job. If the preflight check fails, click the Cancel button to modify the job setting.

13. Click OK when you are done specifying your settings.

The job is submitted to the Job Queue.



Job Name	Backup Se...	Job No.	Job ID	Status	Job Type	Last Result
Database protection job	L103-W...	2		HOLD	Backup (Ro...	
Database pruning job	L103-W...	1	82	READY	DB Pruning	Finished
Synthetic Backup [GFS, Staging]	L103-W...	10	67	HOLD	Synthetic B...	Cancelled
Data Migration Status	L103-W...					
Synthetic Backup [GFS, Staging]	L103-W...	16	83	READY	Synthetic B...	Finished
Data Migration Job	L103-W...	26	85	DONE	Migration	Finished
Data Migration Status	L103-W...					
Data Synthetic Job	L103-W...	27	86	DONE	Synthetic B...	Finished
Synthetic Backup [Rotation, Staging]	L103-W...	13	56	HOLD	Synthetic B...	Finished

Note: Similar to a migration job, after the last incremental backup completes, a new data synthetic job will be created and added into the job queue under the master job. If the data synthetic job fails, a makeup job will be created and added into the job queue.

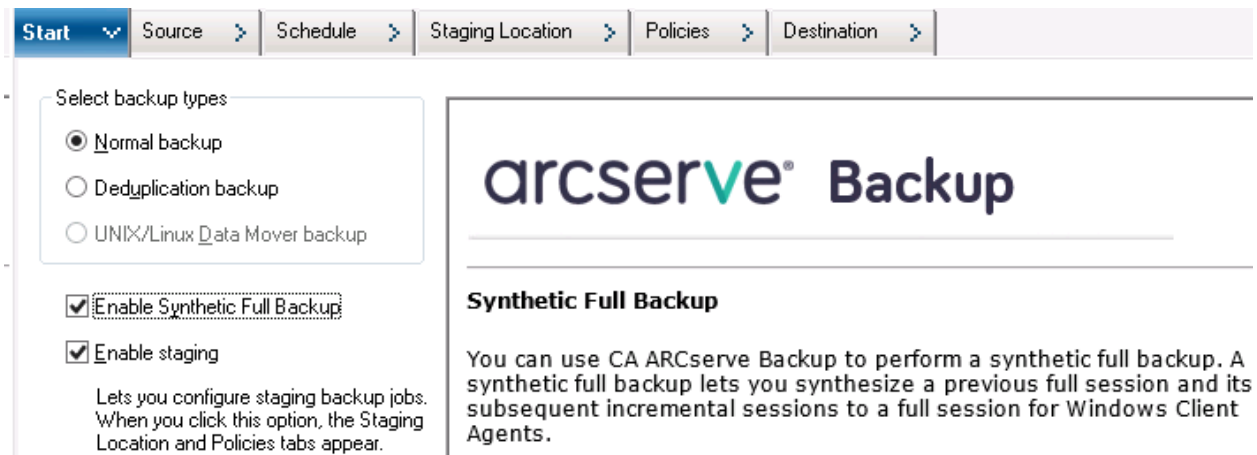
Submit a D2D2T Synthetic Full Backup Job Using a Custom Schedule

The Backup Manager provides you with scheduling options and rotation schemes to help you establish an automatic backup strategy. You can specify a D2D2T synthetic full backup job using a custom backup schedule that suits your environment.

Note: To perform synthetic full backup jobs, you must install and license Arcserve Backup Agent for Open Files on the computers that you want to back up. If you do not install and license the agent, Arcserve Backup converts synthetic full backups to regular backups, and produces a warning message in the Activity Log.

To submit a D2D2T synthetic full backup job using a custom schedule

1. Open the Backup Manager and click the Start tab.
2. Select Normal backup and Enable Synthetic Full Backup. The Enable Staging option is selected by default.



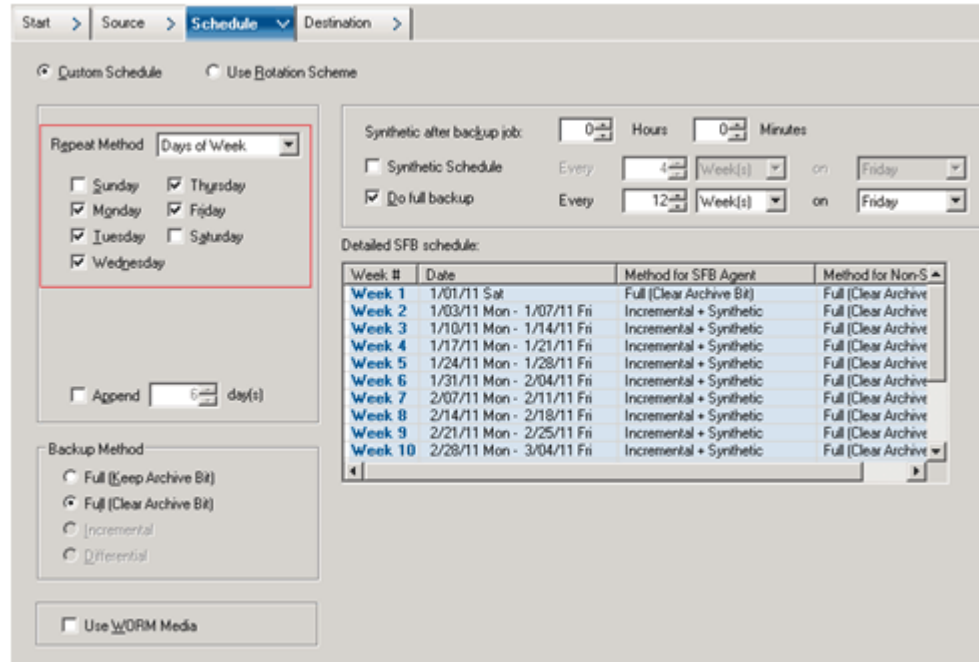
Be aware of the following:

- When you select Normal backup and Enable Synthetic Full Backup, Enable Staging is selected by default. This is because synthetic full backup can only be enabled for Normal backup with the Enable Staging option. You can migrate the data to Tape, Cloud, and Disk devices.
- When you select Deduplication backup and Enable Synthetic Full Backup, Enable Staging is not selected by default. This is because synthetic full backup supports submitting to a deduplication device without the Enable Staging option. However, you can also submit a synthetic full backup job to a deduplication device using the Enable Staging option.
- Synthetic full backup does not support Unix/Linux Data Mover backup.

- Click the Source tab and select the backup source.

Note: Synthetic full backup only supports r16 or higher Windows Client Agents. You can also select other agents (such as pre-r16 Windows Client Agents, database agents, or application agents) along with r16 or higher Windows Client Agents to submit a mixed job. In this mixed job, it will not change the backup method of those agents that do not support synthetic full backup.

- Click the Schedule tab, select the Custom Schedule option, and specify the Repeat Method.



Note: No matter what daily backup method you select, it executes an incremental backup for the file system agents except on the first day and on the days that match the do full backup schedule. On the days that match the do full backup schedule, it executes a full backup for the file system agents.

5. Set the Synthetic Schedule.

Start > Source > **Schedule** > Destination >

Custom Schedule Use Rotation Scheme

Repeat Method: Days of Week

Sunday Thursday
 Monday Friday
 Tuesday Saturday
 Wednesday

Append 6 day(s)

Backup Method:
 Full (Keep Archive Bit)
 Full (Clear Archive Bit)
 Incremental
 Differential

Use WORM Media

Synthetic after backup job: 0 Hours 0 Minutes

Synthetic Schedule Every 4 Week(s) on Friday
 Do full backup Every 12 Week(s) on Friday

Detailed SFB schedule:

Week #	Date	Method for SFB Agent	Method for Non-S
Week 1	1/01/11 Sat	Full (Clear Archive Bit)	Full (Clear Archive Bit)
Week 2	1/03/11 Mon - 1/07/11 Fri	Incremental + Synthetic	Full (Clear Archive Bit)
Week 3	1/10/11 Mon - 1/14/11 Fri	Incremental + Synthetic	Full (Clear Archive Bit)
Week 4	1/17/11 Mon - 1/21/11 Fri	Incremental + Synthetic	Full (Clear Archive Bit)
Week 5	1/24/11 Mon - 1/28/11 Fri	Incremental + Synthetic	Full (Clear Archive Bit)
Week 6	1/31/11 Mon - 2/04/11 Fri	Incremental + Synthetic	Full (Clear Archive Bit)
Week 7	2/07/11 Mon - 2/11/11 Fri	Incremental + Synthetic	Full (Clear Archive Bit)
Week 8	2/14/11 Mon - 2/18/11 Fri	Incremental + Synthetic	Full (Clear Archive Bit)
Week 9	2/21/11 Mon - 2/25/11 Fri	Incremental + Synthetic	Full (Clear Archive Bit)
Week 10	2/28/11 Mon - 3/04/11 Fri	Incremental + Synthetic	Full (Clear Archive Bit)

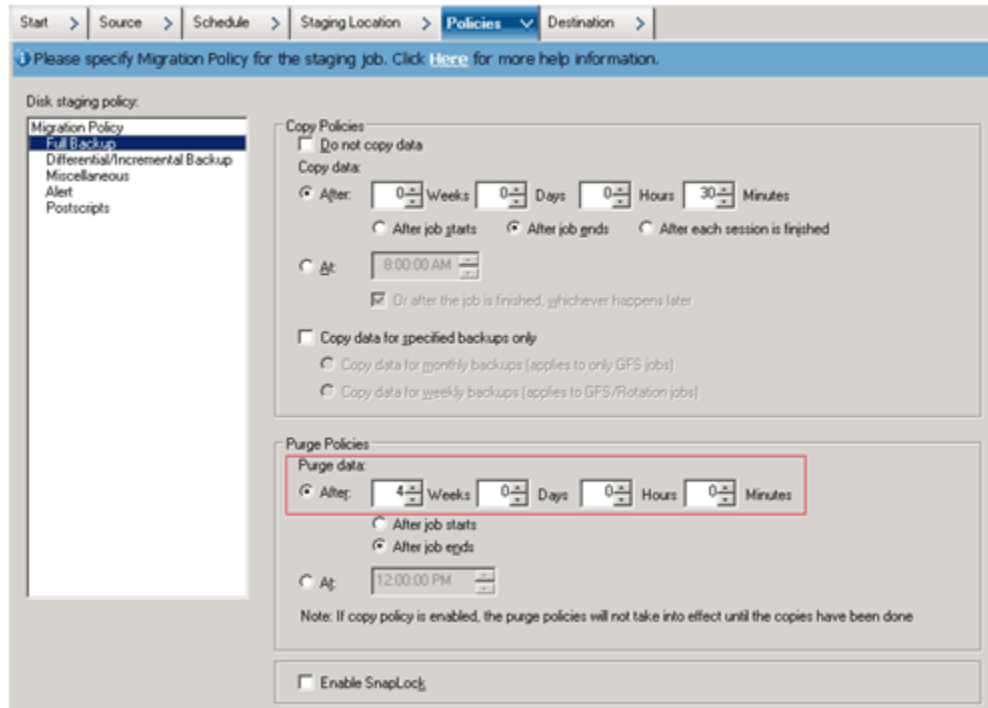
Be aware of the following:

- **Synthetic after backup job:** After the last incremental backup job completes, a data synthetic job will be generated as a run-on job. To avoid multiple jobs competing for CPU/Disk resources, you can select when to launch the data synthetic job after the last incremental backup job is complete.
- **Synthetic Schedule:** You can set the Synthetic Schedule for how long to run a synthetic job to generate a synthetic full backup session. If this option is not selected, the Synthetic Schedule will be set as every one (1) day, which is the default setting. In this scenario, the data synthetic job will run after the incremental backup is completed on every backup day. You can also select the Synthetic Schedule check box to set the Synthetic Schedule as *n* weeks.
- **Do Full Backup:** You can set the schedule for how long to run a traditional full backup for data assurance.

Ensure that the Do Full Backup Schedule is longer than the Synthetic Schedule when you enable Do Full Backup. Otherwise, the data synthetic job may not have a chance to run. You can check the detailed SFB Schedule to know when the synthetic job will run, when the full backup job will run, and when the incremental backup job will run.

6. Click the Staging Location tab and select the staging location for your job.

7. Click the Policies tab and set the Migration Policy.



8. Click the Destination tab, select tape as your destination, and specify the Media pool name.
9. Click the Options toolbar button to specify global options that you require for the job. For more information, see [Global Backup Options](#) (see page 154).
10. Click the Submit toolbar button to submit your job.
The Security and Agent Information dialog opens.
11. Edit or confirm the security and agent information for your job, and click OK.
The Submit Job dialog opens.

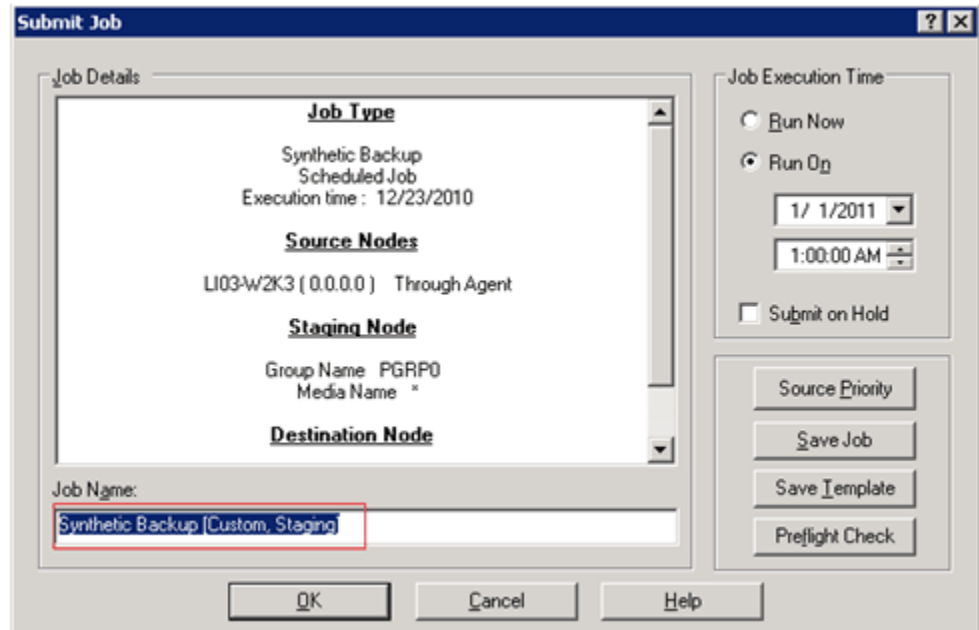
12. Specify the options that you require for the job:

Run Now

Specifies to run the job immediately.

Run On

Lets you specify a date and time when you want the job to run, and enter a description for your job. The job name for the Synthetic Full Backup will start with *Synthetic Backup*:



Source Priority

Lets you specify the priority in which job sessions initiate if you selected multiple sources to back up. Use the Top, Up, Down, and Bottom buttons to change the order in which the jobs are processed.

Save Job

Saves the job as a Arcserve Backup job script.

Save Template


Saves the job template.

Preflight Check

Specifies to preflight check the job. If the preflight check fails, click the Cancel button to modify the job setting.

13. Click OK when you are done specifying your settings.

The job is submitted to the Job Queue.



Job Name	Backup Se...	Job No.	Job ID	Status	Job Type	Last Result
Database protection job	LI03-W...	2		HOLD	Backup (Ro...	
Database pruning job	LI03-W...	1	91	READY	DB Pruning	Finished
Synthetic Backup [Custom, Staging]	LI03-W...	15	103	READY	Synthetic B...	Finished
Data Migration Job	LI03-W...	25	105	DONE	Migration	Finished
Data Migration Status	LI03-W...					
Data Synthetic Job	LI03-W...	26	106	DONE	Synthetic B...	Finished
Synthetic Backup [GFS, Staging]	LI03-W...	10	67	HOLD	Synthetic B...	Cancelled
Data Migration Status	LI03-W...					

Note: Similar to a migration job, after the last incremental backup completes, a new data synthetic job will be created and added into the job queue under the master job. If the data synthetic job fails, a makeup job will be created and added into the job queue.

Run a Real Full Backup on the Next Full Backup Day

You can run a real full backup to generate a full backup session even though it is not a scheduled real full backup day.

To run a real full backup on the next full backup day

1. Open the Job Status Manager and click the Job Queue tab.
2. Right-click the synthetic job that you want to convert and select the option Run real full on next full day.

A real full backup will be executed on the next full backup day. The next full backup day is the day that will execute a full backup based on the original schedule, not the synthetic schedule.

For example, on a Monday, you submit a “5-day weekly incremental backup, full backup on Friday” rotation SFB job with a synthetic schedule every 4 weeks, and you right-click the job and select the option “Run real full on next full day” on the first Thursday, Arcserve Backup will run a real full backup on the first Friday.

Similarly, if you modify a SFB job, it will also run a real full backup on the next full backup day.

Submit an On-Demand Synthetic Full Backup

When you submit a synthetic full backup job with the Enable Point in Time Restore option selected, you can submit an on-demand synthetic full backup job to synthesize incremental sessions and their full session to a new full session even though it is not a scheduled synthetic full backup day.

To submit an on-demand synthetic full backup

1. Open the Job Status Manager and click the Job Queue tab.
2. Right-click the synthetic job that you want to run immediately and select Synthesize Now.

The synthetic job is submitted.

Scan Synthetic Full Backup Session Only

You can submit a separate media assure job to only scan all of the Synthetic Full Backup sessions.

To submit a media assure job to scan SFB sessions

1. Open Media Assure and Scan Manager.
2. Select the tape that contains the Synthetic Full Backup sessions and click Media Assure.
3. Select the Enable Media Assure and Scan Synthetic Full Backup session only options, and click OK.
4. Click Submit.

The media assure job is submitted. The job scans all of the Synthetic Full Backup sessions on the tape that were not scanned previously.

How the Purge Policy and a DB Pruning Job Works for Synthetic Full Backup

Tape data, ASDB information, and the Catalog file of sessions pending synthesization cannot be purged by the job engine and a DB pruning job. It will use the following algorithm to get a minimum data retention time for the session that is pending synthesization:

$\text{MAX}(\text{MIN}(\text{Synthetic Cycle Time} * 2, \text{Real Full Cycle Time}), \text{Session Retain Time in Purge Policy})$

This ensures there is always at least one full backup (regardless if it is a real full backup or synthetic full backup) session and all subsequent incremental sessions. Because Arcserve Backup will not prune catalog files for these sessions, the CATALOG.DB folder may use disk space, the size of which depends on your backup sources and purge policy.

Arcserve Backup treats one full backup (regardless if it is a real full backup or synthetic full backup) session and all subsequent incremental sessions as a unique session set that will be used for the next synthetic job. Therefore, the full backup session and all subsequent incremental sessions will have the same purge policy even if you have set a different purge policy for full backups and incremental backups. Arcserve Backup calculates the purge time for the full session only and applies the same purge time for all subsequent incremental sessions.

You can determine the purge time for synthetic related sessions by viewing the master job log in the activity log.

How Makeup Jobs Work for Synthetic Full Backup

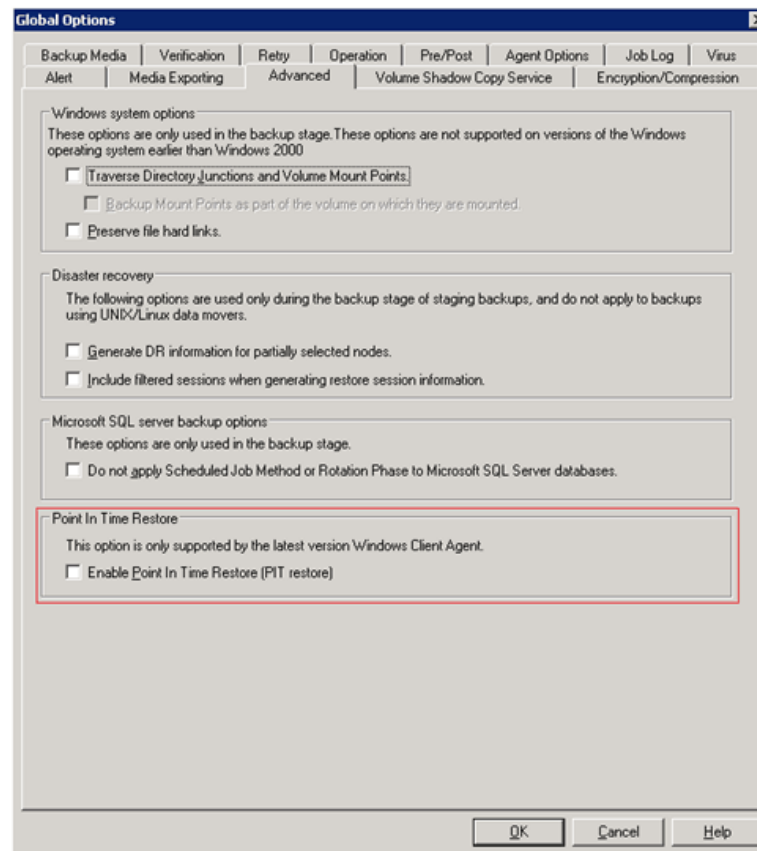
On the synthetic full backup day, if some incremental sessions finished but other incremental sessions failed, a normal makeup job for the failed client agent is submitted. A data synthetic job for synthesizing all finished sessions is also submitted. If the makeup job for the failed client agent is finished, it creates the incremental sessions for the failed client agent, and then submits another data synthetic job for synthesizing the failed client agent.

Point In Time Restore

With synthetic full backup, you can enable Point in Time (PIT) restore.


To enable Point in Time restore

1. Select the Enable Synthetic Full Backup option on Start tab.
2. Click the Options button then click the Advanced panel.
3. Select the Enable Point in Time Restore (PIT restore) option.




Note: When the Enable Point in Time Restore (PIT restore) option is enabled, no matter what daily backup method you select, it executes a Point in time incremental backup for file system agents except on the first day and on the days that match the Do full backup schedule. On the days that match the Do full backup schedule, it executes a full backup for file system agents.

The Backup method in the Activity log and the Session method in the Restore by Session window will be changed from “Incremental” to “Point in time incremental.” However, the Backup method on the Backup GUI is still “Incremental.”



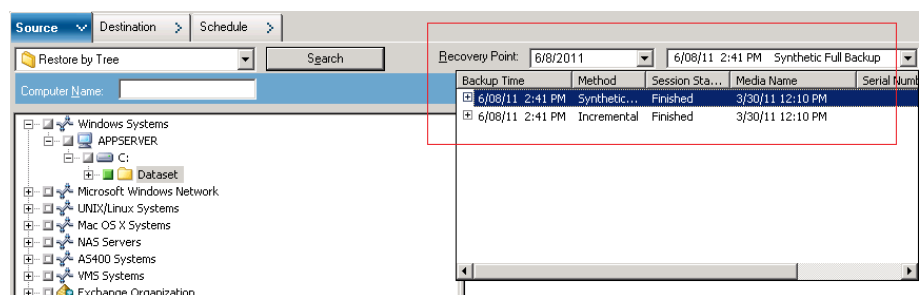
For the Point in time incremental backup method, the Client Agent will send unchanged catalog information on the normal incremental day. Then, *mergcat.exe* will index both the changed catalog file and the unchanged catalog file to a full file system structure. After it is indexed, you can do a Point in Time restore job. However, this takes more disk space and CPU usage.

When the Enable Point in Time Restore (PIT restore) option is disabled, the daily backup method is Incremental. This is the default setting for a SFB job.



For the Incremental backup method, the Client Agent will not send an unchanged catalog file on the normal incremental day. The Client Agent will only send unchanged catalog information on the synthetic day. *Mergcat.exe* has no chance to index the unchanged catalog file. This will save disk space and CPU usage.

To submit a Point in Time restore job, you select a node from the source tree and select a recovery point date, and the related session versions are listed in the Recovery Point drop-down list. After you select a recovery point, the related session information is updated in the Tree view and in the List view.



Synthetic Full Backup Considerations

When you submit a synthetic full backup job, keep in mind the following points:

- Synthetic full backup only supports r16 Windows Client Agents. It does not support Windows Client Agents before r16, database agents, or application agents. However, you can still submit a synthetic full backup job to back up mixed sources (including r16 Windows Client Agents, pre-r16 Windows Client Agents, database agents, or application agents). For the agents not supported by synthetic full backup, it uses the daily backup method you had previously specified.
- Synthetic full backup only supports disk-staging devices and deduplication devices. You can migrate synthetic full backup sessions to tape, cloud, and disk devices with D2D2T.
- Ensure that you define the *Do full backup* option on the Backup Manager Schedule tab for data assurance when you submit a synthetic full backup job.
- When you modify an existing synthetic full backup job, a traditional full backup runs the next day that a full backup is scheduled to run. To disable this behavior, modify the following registry key:
 1. Open the following key on the backup server:
`HKEY_CURRENT_USER\Software\ComputerAssociates\CA ARCserve Backup\Base\ASMgr`
 2. Specify the following Value Name:
`(DWORD) SFBCheckedRealFull`
 3. Specify the following Value:
`1`
The value 1 disables this behavior.

Note: A traditional full backup job always runs when you make specific changes to the job. Specifically, you change the backup source, the synthetic schedule, the disk staging location, or the Encryption/Compression option. This approach helps to ensure that you can perform synthetic backups on the days synthetic backups are scheduled to run when you make these specific to the job.
- When you submit a synthetic full backup job, the Arcserve Backup user interface displays your 90-day detailed synthetic schedule for your reference. However, the detailed synthetic schedule that displays in the user interface may not be consistent with your actual schedule when you modify existing synthetic full backup jobs.
- When you schedule a synthetic full backup job, the Detailed SFB Schedule field displays the current date as the default start date. However, when you select the Run On date in the Submit Job screen, the synthetic full backup job is scheduled based on the day you specify.

- With synthetic full backup, you can schedule a synthetic full backup using a timeframe that is longer than a week. Using this setting, the default retention time (six days) of the media pool for a rotation scheme cannot prevent overwriting necessary data. Therefore, when you submit a D2D2T synthetic full backup job with a rotation scheme and a synthetic schedule of more than a week, verify that the retention time of the media pool is longer than the synthetic schedule, which will require more media.

Note: To perform synthetic full backup jobs, you must install and license Arcserve Backup Agent for Open Files on the computers that you want to back up.

Working with Synthetic Full Backup

The following information addresses common questions about how Synthetic Full Backup works.

What types of agents does Synthetic Full Backup support?

Synthetic Full Backup only supports r16 or higher Windows Client Agents. It does not support Windows Client Agents prior to r16, database agents, or application agents.

Note: To perform synthetic full backup jobs, you must install and license Arcserve Backup Agent for Open Files on the computers that you want to back up.

What backup methods can be submitted with the Synthetic Full Backup option?

You can submit three types of backup with Synthetic Full Backup:

- Normal backup with Enable staging
- Deduplication backup with Enable staging
- Deduplication backup

What types of devices does Synthetic Full Backup support?

Synthetic Full Backup supports Disk-staging devices and Deduplication devices. However, you can migrate the Synthetic Full Backup session to Tape, Cloud, and FSD with D2D2T.

Does a Synthetic Full Backup job let you back up the supported and unsupported agents?

Yes. Although Synthetic Full Backup only supports r16 or higher Windows Client Agents, you still can submit a Synthetic Full Backup job to back up the mixed sources. The mixed sources include the r16 or higher Windows Client Agent, Windows Client Agents prior to r16, database agents, or application agents. When you submit a Synthetic Full Backup job that contains mixed sources, it only synthesizes the r16 or higher Windows Client Agents. For the agents and application agents that do not support Synthetic Full Backup, the traditional backup method will be used instead (the backup method for those unsupported agents does not change).

What is the daily backup method for the mixed Synthetic Full Backup job for one week?

Assume you submit a Synthetic Full Backup with the *5-day weekly incremental backup, full backup on Friday* scheme to back up an r16 Windows Client Agent, an r15 Windows Client Agent, and a SQL 2008 agent together on Monday.

Here is the daily backup method for the agent for one week:

	Monday	Tuesday	Wednesday	Thursday	Friday
r16 Windows Client Agent	Full	Incremental	Incremental	Incremental	Incremental + Synthetic Full Backup
r15 Windows Client Agent	Full	Incremental	Incremental	Incremental	Full
SQL 2008 Agent	Full	Incremental	Incremental	Incremental	Full

When I select the GFS scheme, the synthetic schedule is disabled. When will the data synthetic job launch?

For the Synthetic Full Backup job with GFS scheme, the data synthetic job will be launched on the full backup day. For example, if you choose the *5-day incremental backup, Full backup on Friday with GFS enabled* scheme, the data synthetic job will be launched on every Friday. If you choose the *7-day incremental backup, Full backup on Sunday with GFS enabled* scheme, the data synthetic job will be launched on every Sunday.

What is the purge policy of a Synthetic Full Backup job?

If you submit a Synthetic Full Backup job with enable staging, there are two purge policies; one is for the migration policy, and the other is for the synthetic policy. For staging a Synthetic Full Backup job, the sessions related to synthetic on a staging device are purged according to the Synthetic Purge Policy. The sessions not related to synthetic on a staging device are purged according to the Migration Purge Policy.

For a Normal Deduplication synthetic job, the sessions related to synthetic on a Deduplication device are purged according to the Synthetic Purge Policy. The sessions not related to synthetic on a Deduplication device are purged according to purge configuration in the Deduplication Policy.

Note: You can calculate the Synthetic Purge Policy with the following formula:

$\text{MAX}(\text{MIN}(\text{Synthetic Cycle Time} * 2, \text{Real Full Cycle Time}), \text{Session Retain Time in Purge Policy})$

When is a real full backup job run?

A real full backup is launched in any of the following conditions:

- On the scheduled real full backup day
- If a synthetic full backup job was modified, a real full backup is launched on the next scheduled full backup day
- If the last data synthetic job failed on the synthetic day, a real full backup is launched on the next synthetic day
- If there are no parent full backup sessions found when the data synthetic job launches on the synthetic day, it will convert the data synthetic job to a real full backup instead.

Chapter 17: Using Arcserve Backup in a Storage Area Network

This section contains the following topics:

[How to License the Storage Area Network \(SAN\) Option](#) (see page 841)

[The SAN Environment](#) (see page 842)

[Install the SAN Option](#) (see page 845)

[Using the SAN Option](#) (see page 848)

How to License the Storage Area Network (SAN) Option

To successfully license Arcserve Backup Storage Area Network (SAN) Option, you must fulfill the following installation requirements:

- You must install and license the option to perform backup operations to libraries that are shared on a SAN.
- You must install the option on the Arcserve Backup Primary server.
- You must issue all licenses on the primary server.
- Ensure that you have a sufficient number of Storage Area Network (SAN) Option licenses to support your environment.

The Central Management Option is a prerequisite component for the Storage Area Network (SAN) Option.

The Storage Area Network (SAN) Option is a count-based license. You must issue one license for all Arcserve servers that share a library with another Arcserve server.

Examples: How to License the Storage Area Network (SAN) Option

The following examples describe how count-based licensing works with the Storage Area Network (SAN) Option:

- Your environment consists of a primary server and three member servers. The primary server and the three member servers share a multiple drive library on a SAN. This configuration requires you to issue four Storage Area Network Option (SAN) licenses issued on the primary server. All servers in the Arcserve domain are sharing a library.
- Your environment consists of a primary server and three member servers. Two member servers share a multiple drive library and the third member server is configured with a locally attached, multiple drive library. This configuration requires you to issue four Tape Library Option licenses and three Storage Area Network (SAN) Option licenses on the primary server. All servers in the Arcserve domain have access to a multiple drive library; however, three Arcserve servers are sharing a library.

The SAN Environment

The servers within a SAN group include one primary SAN server and one or more SAN-attached member servers. The primary SAN server is the most important server because no tasks can occur until it initializes the tape libraries on the SAN.

The primary SAN server is responsible for the following actions:

- Initializes, governs, and maintains a shared device on the SAN
- Coordinates the use of shared library resources among the servers in the SAN, preventing conflict if two servers try to allocate a device or media at the same time

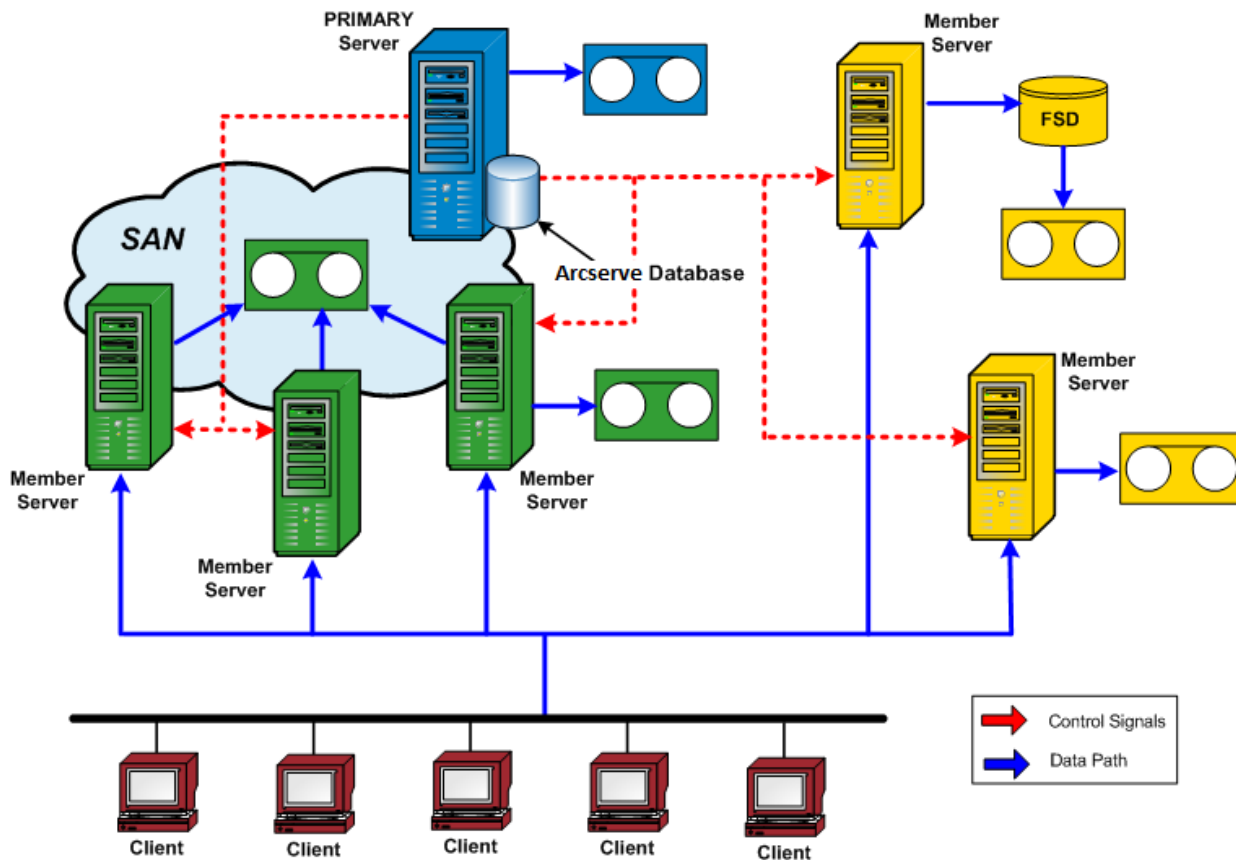
You can designate any Arcserve Backup SAN server as the primary SAN server. However, because the primary SAN server is responsible for managing and initializing the shared SAN, you should use your most reliable server as the primary SAN server.

How Arcserve Backup Works in a SAN

During Arcserve Backup installation, you create your SAN domain with the specified primary SAN server and associated member SAN servers. In a SAN environment, all servers are divided into one of three groups; a primary SAN server, SAN-attached member server, or non-SAN member server.

Within a domain, there can be only one primary server and each SAN-attached member server can report to only one primary SAN server. In addition, a primary SAN server can belong to only one domain.

The following diagram shows how a SAN is configured with a primary SAN server that has Arcserve Backup and the SAN option installed:



When a job is ready to run, the option reserves the device and media. Once reserved, that device and media are no longer available to any other jobs across all SAN servers.

The SAN Option enables Arcserve Backup servers to share one or more tape libraries by creating a virtual ring. Any backup or restore jobs on a server that has the option installed run as local jobs. As the backup progresses, Arcserve Backup sends data over the SAN hardware to the tape libraries for storage instead of over the LAN cabling. This provides greater speed, reduces network traffic, and maximizes backup and restore throughput.

Server Management in a SAN

Using the SAN Option does not change the way you manage Arcserve Backup servers. You continue to connect to each SAN Option server through the Backup or Restore Manager to schedule backup and restore jobs and to manage your Arcserve Backup database.

Backup Plans

You should plan a backup strategy that is appropriate for your SAN configuration. You should consider the impact of multiple Arcserve Backup servers sharing a single device. For example, if your backup device contains two tape drives and there are five option servers sharing the media libraries, you should not schedule five backup jobs to begin simultaneously. This would force the option to determine which two of the five jobs to begin first. Instead, you should carefully consider and schedule the start times for jobs to meet your backup strategy and allow you to control the schedule sequence.

Note: Each scheduled backup job waits in the queue until a tape drive is available to perform the backup.

Benefits of Using the Option

The SAN Option provides the following benefits:

- Decreased Costs--Allows servers to share one or more tape libraries.
- Improved Backup and Restore Speed--Eliminates the need for remote backups through your Local Area Network (LAN).
- Efficiency--Centralizes the backup of hardware and media.
- Flexibility--Optimizes the flexibility by redirecting or reconfiguring in case of a device failure.

Terminology

The following terms are commonly used in a SAN environment:

- Storage Area Network (SAN)--A high-speed network designed for sharing attached tape libraries.
- SAN server group--A group of Arcserve Backup servers that can share a set of tape libraries on a Storage Area Network.
- Primary SAN server--The Arcserve Backup server that initializes the shared tape libraries and is responsible for controlling usage and detecting any changes in status of these devices.

- SAN-attached Member server--Servers in a SAN that are assigned to the primary server to use the shared tape libraries.
- Shared device--A device on a SAN used by a SAN server group.

Install the SAN Option

This section explains how to install and configure the SAN option on your primary and member-attached SAN servers from one central location.

Operating System Compatibility

For information about operating system compatibility, see the Arcserve Backup for Windows Readme.

Installation Prerequisites

Before you install the SAN Option, verify the following prerequisites:

Note: Arcserve Backup supports libraries configured with one drive. If your library has more than one drive, you must license the Arcserve Backup Tape Library Option to enable multi-drive capabilities.

- Your system requirements meet the minimum requirements needed to install the option. For more information about installation requirements, see the *Implementation Guide*.
- Your system meets the minimum hardware and software requirements needed to install Arcserve Backup and the Arcserve Backup Tape Library Option (if necessary).
- You have installed all appropriate SAN hardware device drivers for adapters to access the devices attached to the Fibre Channel adapter.

Note: For information about the SAN hardware and drivers, see the Arcserve Backup for Windows Certified Device List. You can access the Certified Device List from the Arcserve Backup home page.

- You have Arcserve Backup and the Central Management Option installed on the computer on which you want to install the option. In addition, if you have a multi-drive library you must also have the Tape Library Option installed. This computer can be either a local computer or a remote computer.

Note: If these applications are not already installed, you must install them when you install the SAN Option.

- These options (Central Management Option, SAN Option, and Tape Library Option) are all installed on the Primary server only.
- There is license count for the Primary server and each SAN member server. There is one Central Management Option license for the entire SAN, and one SAN and Tape Library Option license for each server in the SAN (Primary server and all associated SAN member servers.)

- You have made a note of the default installation path.
- You have administrator privileges to install software on the computers on which you want to install the option.

Note: Contact your Arcserve Backup administrator to obtain the proper privileges and information if you do not have them.
- You know the user names, passwords and IP addresses for the primary SAN server and SAN-attached member servers.
- You have all SAN hardware and related device drivers installed.
- The Windows backup server identifies all of the appropriate SAN devices, including the medium changer and tape drives.
- All SAN servers in your storage area network can communicate with one another by pinging each server by name or by pinging their IP address with the display server name switch.

IP connectivity and name resolution among all servers participating in the SAN is essential. To ensure you have IP connectivity and name resolution, you may need to update the IP host file on each server so that the name and IP address of each server is present in the IP host file of all other servers. The IP host file on each server resides in `\windows\system32\drivers\etc`.

SAN Option Installation

The SAN Option follows the standard installation procedure for the system components, agents, and options of the Arcserve Backup. Start all installation sessions by running setup.exe.

- You can install the Arcserve Backup base product, agents, and options all in one session
- You can install the Arcserve Backup base product first, and then install the agents and options separately later.
- You can install the Storage Area Network (SAN) Option on the primary SAN server (only).

For more information on installation, see the *Implementation Guide*.

The installation process is very flexible and allows you to decide whether to install the different system components, agents, and options of Arcserve Backup in one installation session or in multiple installation sessions. The preferred method is to install all of the components in one installation session. You can, however, install each component sequentially in individual sessions or install selected components in one session and other components in individual sessions later.

Before starting the installation process, decide which system components, agents, and options of Arcserve Backup you want to install during this session. Then, gather the prerequisite information for each of the agents and options you want to install. You can find this information in each agent and option guide. Select the combination of installation sessions that best meets your needs.

For example, to install the Arcserve Backup server, the Tape Library Option, and the Agent for SQL Server, you can use any of the following combinations of installation sessions:

- Install the server, the option, and the agent in the same installation session.
- Use three separate installation sessions; one session to install the server, a second session to install the option, and a third session to install the agent.
- Use two separate installation sessions. When using two separate sessions, you can group the components in the following ways: Install the server in one session and the option and agent in a separate session; install the server and option in one session and the agent in a separate session; or install the server and agent in one session and the option in a separate session.

Uninstall the Storage Area Network Option

The Storage Area Network (SAN) Option is a primary server and stand-alone server based installation. You must use the Server Admin to uninstall server based options. For more information, see [Install and Uninstall Arcserve Backup Server Based Options](#) (see page 590).

Using the SAN Option

This section provides the information you need to use the SAN Option. Specifically, it explains how you can use the option to perform the following tasks:

- Create shared device groups
- Backup and restore data
- Manage devices
- Manage media
- Monitor job status
- Control when jobs run
- Create reports and logs
- Use virtual libraries

Storage Area Network (SAN) Configuration

The Storage Area Network (SAN) Configuration utility lets you configure the relationships between the Primary and Distributed Arcserve Backup servers. After all the servers have the SAN Option installed, you can go to any server to run the SAN Configuration utility.

The executable for the SAN Configuration utility is stored in the Arcserve Backup installation directory on the backup server. For example:

```
C:\Program Files\CA\ARCserve Backup\ELOConfig.exe
```

To start the SAN Configuration utility, access the above directory and then double-click the file named ELOConfig.exe.

Create Shared Device Groups

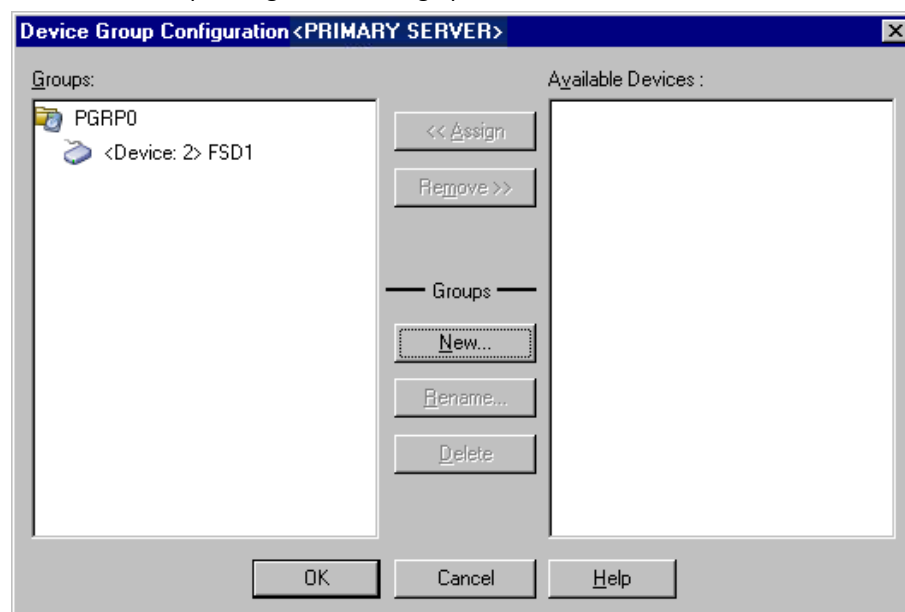
Creating shared device groups is the key to the flexibility and efficiency of Arcserve Backup.

Note: Shared device groups can only be modified, created, or deleted from the Primary server.

To create shared device groups

1. From the Device Manager window, select Configure Groups from the properties pane.

The Device Group Configuration dialog opens.



2. Click New.

The New Group dialog opens.



3. Select the type of device group in the Type field, and enter a name for the device group in the Name field. Click OK.

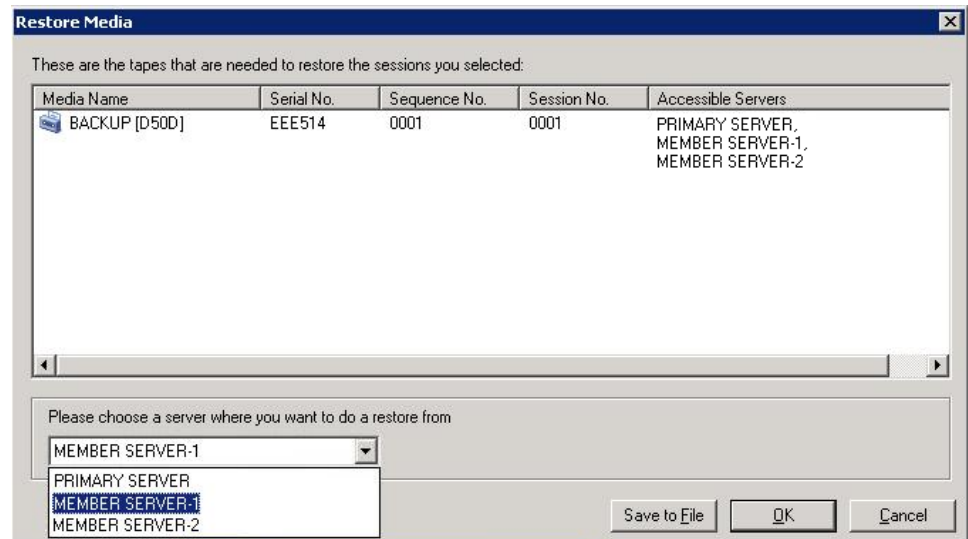
The new device group is displayed in the Device Group Configuration dialog.

Data Backup and Restore in a SAN Environment

You must use the Backup Manager or Restore Manager to configure and submit backup or restore jobs in your SAN environment. These backup and restore jobs run locally on the server where the SAN Option is installed. The data is transferred over the SAN hardware to the library instead of over the LAN cabling. This speeds up the job processing and reduces Ethernet traffic. If you use the Media view, information on backups performed by all SAN servers is available from the database.

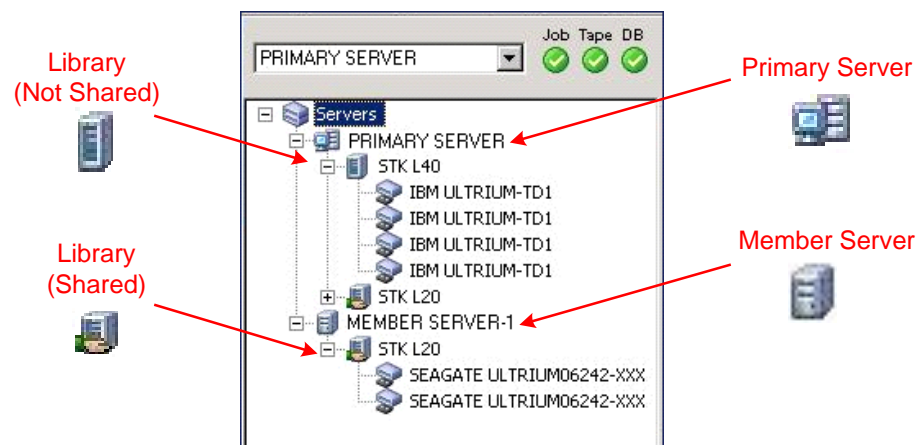
The option also provides a large number of backup and restore options, filters, and scheduling features for your jobs. For more information, see "Backing Up Data" and "Restoring Data."

When submitting a restore job of data found on a tape inside a SAN-attached Library you can use the drop down menu in the Restore Media pop up which lists all the SAN servers that the tape is accessible from. The server you select here will be the server on which the Restore operation will run and it does NOT have to be the server that actually did the Backup.



Device Management

Use the Device Manager to display information about the storage devices connected to your SAN server group, the media in these devices, and the status of these devices. Through the Device Manager, you can view all of the shared devices connected to your SAN server group.



Consider the following when managing devices:

- Each Arcserve Backup SAN Option installed server in the SAN server group displays the same view of the SAN attached devices.
- If you change the device configuration on the primary server (reconfiguring a library as a RAID, or add more drives to the library for example), you must stop the tape engine service on all SAN servers (primary and all member), then start the primary server Tape Engine first. After the primary server Tape Engine is running, you can then start each of the SAN attached member servers to see the new configuration correctly.

For more information about managing devices, see [Device Manager](#) (see page 395).

Media Management

Consider the following when managing media:

- Because SAN servers share media, be careful when selecting media from a scratch set. Scheduled jobs can be affected if the media is unavailable.
- Only one administrator of media pools should supervise the SAN server group.
- A tape in a save set cannot be destroyed, formatted, or erased, unless it is moved to a scratch set.

For more information about media pool management, see [Media Pool Manager](#) (see page 454).

Media Pools

A media pool is a collection of media that is managed as a set and shared in the SAN. Each media pool is assigned a name, and the media is organized by serial numbers. Manage media pools from the Arcserve Backup Media Pool Manager window, which you access by clicking the Media Pool Manager icon.

For more information about media pools, see [How Media Pools Work](#) (see page 445).

Control of Job Runtime

Arcserve Backup can determine if a device is being used by a job, even if it is on another SAN-attached server. It can then wait until the device is free before starting another job.

When there are many jobs waiting in the queue for the same device, there is no way to determine which job will run next. If priority is important, configure the start times based on how long you think the previous job will take. By carefully arranging the start times, you should have no more than one job ready to begin at a time.

Note: For more information about scheduling, managing jobs, and the Job Status Manager, see [Tasks You Can Perform Using the Job Status Manager](#) (see page 319).

Reports and Logs

Arcserve Backup provides the following options for displaying logs and reports:

- Activity Log--Contains comprehensive information about the operations performed by Arcserve Backup. It provides an audit trail of all backup activity, including every job that is run, and also displays the session number if you need to restore a session..

The activity log can be viewed from the Job Status Manager.

- Tape Log (TAPE.LOG)--Contains all Tape Engine-related messages.
- Job Log--Tracks activity related to a specific job.
- Report Manager--Generates reports from the Arcserve Backup database for viewing or printing. You can open the Report Manager from the Quick Access menu and view reports such as the Job report, Backup Media Error report, Session report, Backup Device report, and the Media Pool report.

For more information about reporting, see [Arcserve Backup Logs and Reports](#) (see page 667).

Arcserve Virtual Libraries

The Tape Library Option is used in conjunction with the Arcserve virtual libraries to provide you with a versatile tool for addressing a wide range of storage requirements. The Arcserve virtual libraries work seamlessly over the Tape Library Option, allowing you to configure physical libraries into smaller virtual (logical) libraries. These virtualized libraries can share the same robotics and import/export slots, which in turn allows drives and storage slots to be grouped together.

When you use Arcserve virtual libraries, be aware of the following restrictions:

- You can configure Arcserve virtual libraries for shared Tape Libraries on the primary server only.
- If the device configuration on the primary server changes (for example, you reconfigure a library into multiple Arcserve virtual libraries), you must stop the tape engine service on all SAN servers (primary and all member), then start the primary server Tape Engine first. After the primary server Tape Engine is running, you can then start each of the SAN attached member servers to see the new configuration correctly

For more information about Arcserve virtual libraries, see [Virtual Library Configuration Option](#) (see page 380).

Appendix A: Troubleshooting

This section provides troubleshooting information to help you identify and resolve problems that you may encounter when using Arcserve Backup.

This section contains the following topics:

[Log in Problems](#) (see page 855)

[Authentication Problems](#) (see page 863)

[Backup and Restore Problems](#) (see page 872)

[Media Problems](#) (see page 883)

[SAN Configuration Problems](#) (see page 890)

[Cluster-based Backup and Restore Problems](#) (see page 895)

[Miscellaneous Problems](#) (see page 898)

Log in Problems

This section contains the following topics:

[Unable to Log In After Changing the caroot Password](#) (see page 855)

[Makeup Jobs Created When the Media is Full](#) (see page 857)

[Unable to Log In to Arcserve Backup After Changing the Computer Name](#) (see page 857)

[Arcserve Backup Cannot Communicate after Changing the IP Address of a Arcserve Backup Server](#) (see page 858)

Unable to Log In After Changing the caroot Password

Valid on Windows platforms.

Symptom:

I changed the password for the caroot account. Why is it saying invalid password when I try to log in to Arcserve Backup?

Solution:

Your password did not change at the time of setup. There are various reasons for this; your machine name may have extended characters or you may have a machine name in a language other than English. If so, run the following debugging authentication commands (replace AB_MACHINE with your machine name) so that you can send the logs to Arcserve Customer Support for investigation:

Note: The caroot password can consist of any combination of alphanumeric and special characters, but may not exceed 15 bytes. A password totaling 15 bytes equates to approximately 7 to 15 characters.

1. Ping the machine by name. For example:

```
ping.exe AB_MACHINE
```

where AB_MACHINE is your machine. If this does not work, resolve the name to an IP address by changing the etc/hosts file or on the DNS.

2. Enter the following command

```
ipconfig /all > ipconfig.log
```

3. Enter the following command to tell Arcserve Customer Support if the portmapper is running on your machine:

```
netstat -na >netstat.log
```

4. Enter the following command to let Arcserve Customer Support know which Arcserve Backup services have registered with the rpc server running on the client machine:

```
rpcinfo.exe -p AB_MACHINE >rpcinfo.log
```

where AB_MACHINE is your machine.

5. Enter the following command:

```
rpcinfo.exe -t AB_MACHINE 395648 1 > caauthd.txt
```

where AB_MACHINE is your machine.

Note: Using '>' to a file does not show the results on the screen.

6. Set up the following registry key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA  
ARCserve Backup\Base\LogARCserve[DWORD]DebugLogs =1.
```

This creates the rpc.log file in the Arcserve Backup home directory under \log.

Makeup Jobs Created When the Media is Full

Valid on Windows 64-bit operating systems.

Symptom

While performing a backup to tape operation using the Client Agent for Windows, Arcserve detects a media full condition. You must replace the media within 20 minutes of the detection. If you replace the media after 20 minutes of the detection elapses, the following events occur:

- Error E3392 (Backup server TCP reconnection timeout) is recorded in the Activity Log.
- The job completes successfully with a completion status of Failed.
- Arcserve creates a Makeup Job.

Solution

The remedies for this problem are as follows:

- Although the job completed successfully, Error E3392 caused the job to appear to fail. The Makeup Job was created because Arcserve detected a failed job. In conclusion, you can safely delete the Makeup Job.
- You can increase the time-out waiting period value by modifying the following registry keys:

HKEY_LOCAL_MACHINE\Software\ComputerAssociates\CA ARCserve
Backup\ClientAgent\Parameters\SendTimeOut

Default: 1200 (seconds)

HKEY_LOCAL_MACHINE\Software\ComputerAssociates\CA ARCserve
Backup\ClientAgent\Parameters\ReceiveTimeOut

Default: 1200 (seconds)

Example: To increase the waiting period to 60 minutes, modify the above DWORD values to 3600.

Unable to Log In to Arcserve Backup After Changing the Computer Name

Valid on Windows platforms.

Symptom:

I changed the name of a machine that has Arcserve Backup installed and rebooted it. Why can't I log in to the Arcserve Backup Manager Console anymore?

Solution:

The computer name is a name that your computer uses to identify itself in a network or a domain. In a centralized management environment, an Arcserve domain can consist of a primary server and one or more member servers, or a stand-alone server. Arcserve Backup uses the computer names of the primary server and the member servers to establish communication between the servers.

For more information about how to process computer name changes in an Arcserve domain, see [How to Process Computer Name Changes in an Arcserve Domain](#) (see page 542).

Arcserve Backup Cannot Communicate after Changing the IP Address of a Arcserve Backup Server

Valid on Windows platforms.

Overview

There are several reasons why the IP address would be changed on a Arcserve Backup server. Some of the common reasons are as follows:

- The network interface card (NIC) was replaced in a Arcserve Backup server. When the computer rejoins the network, the IP address is different from that of the previous network card.
- The Arcserve Backup server communicates with a DHCP server to obtain an IP address and the Arcserve Backup server connects to a different DHCP server.

Symptoms

After you change the IP address on a primary server, a stand-alone server, and a system hosting the Arcserve Backup Manager Console, Arcserve Backup demonstrates the following behavior:

- On a member server and a system hosting the Arcserve Backup Manager Console, the value of the Domain is "None" in the Default Server and Security fields on the Manager Console.
- Error messages appear when you click the Backup link in the Navigation Bar on the member server. For example, a pop-up message appears "Connecting to primary server."

- Pinging the primary server from the Command Line on a member server, returns the message "Request timed out."
- Equivalence errors may occur when you execute a task using a Arcserve Backup command line utility on the primary server or a stand-alone server. For example, one of the following messages may appear:

Ntuser not validated in authentication server on *Hostname*.
Do you want to create equivalence (default : y)?

After you change the IP address on a member server, Arcserve Backup demonstrates the following behavior:

- On a member server, the value of the Domain is "None" in the Default Server and Security section on the Manager Console.
- From the primary server or a system hosting the Arcserve Backup Manager Console, you cannot log in to the member server.
- Pinging a member server from the Command Line on the primary server, returns the message "Request timed out."
- Equivalence errors may occur when you execute a task using a Arcserve Backup command line utility on a member server. For example, one of the following messages may appear:

Ntuser not validated in authentication server on *Hostname*.
Do you want to create equivalence (default : y)?

Solutions

To remedy the communication problems, use the procedure that corresponds with the type of Arcserve Backup server where the IP address was changed.

IP Address Changed on the Primary Server or Stand-alone Server

Important! After you change the IP address of a primary server or a stand-alone server, basic backup jobs and restore jobs from the server itself can complete successfully. Additionally, the host names of the member servers should display in the Arcserve Backup managers on the primary server. However, to ensure that you can successfully complete backup jobs and restore jobs on member servers, you must complete the modification described in Step 1 on the member server before completing any other task.

1. If the Arcserve Backup server is a primary server, stop and restart the Arcserve Backup services using the following commands:
 - cstop
 - cstart
2. If there are member servers in your Arcserve Backup domain, open the Windows Command Line on the member server.

Execute the ipconfig command using the /flushdns switch. For example:

```
c:\documents and settings\windows user name>ipconfig /flushdns
```

Note: You must repeat this step on all member servers in your Arcserve Backup domain.

3. Create equivalence on the primary server or stand-alone server using the ca_auth command. The syntax for this task is as follows:

```
ca_auth [-cahost HOST-NAME] -equiv add ntuser HOST-NAME ARCserveBackupUser [caroot_username] [caroot_password]
```

Note: For more information about using the ca_auth command, see the *Command Line Reference Guide*.

IP Address Changed on a Member Server

1. Create equivalence on the member server using the ca_auth command. The syntax for this task is as follows:

```
ca_auth [-cahost HOST-NAME] -equiv add ntuser HOST-NAME ARCserveBackupUser [caroot_username] [caroot_password]
```

Note: For more information about using the ca_auth command, see the *Command Line Reference Guide*.

2. Open the Windows Command Line on the primary server.

Execute the ipconfig command using the /flushdns switch. For example:

```
c:\documents and settings\windows user name>ipconfig /flushdns
```

IP Address Changed on a Server that is Hosting the Arcserve Manager Console

No action is required when the IP address is changed on a server that is hosting the Arcserve Backup Manager Console. You can continue to manage other Arcserve Backup servers without experiencing communication problems.

IP Address Changed on an Agent System

Select one of the following corrective actions:

- Open the Windows Command Line on the Arcserve Backup server that is backing up the agent system.

Execute the ipconfig command using the /flushdns switch. For example:

```
c:\documents and settings\windows user name>ipconfig /flushdns
```

- If you added the agent system to the primary server, a member server, or a stand-alone server by referencing the IP address rather than the agent system's host name, you can log in the Arcserve Backup server and manually change the IP address of the agent system. To do this, complete the following steps:
 1. Log in to the Arcserve Backup server, open the Backup Manager, and expand the Windows system object.
 2. Right-click the agent system and select Modify Agent from the pop-up menu.
The Agent Option dialog opens.
 3. Specify the new IP address in the IP address field and click OK.
The new IP address is applied to the agent system.

IP Address Changed on a System Using a Static IP Address

Select one of the following corrective actions:

- When you use a static IP address, you can register the new static IP address on the DNS server. Based on the type of server (for example, a primary server, member server, and so on), use the ipconfig command task described in the previous sections to refresh their local DNS client.

This action lets you resolve the cache for establishing a new relationship between the host name and IP address.


- If you do not register the new static IP address with the DNS server, you must modify the Hosts file on the servers to reflect this change.

To remedy this scenario, do the following:

1. From Windows Explorer, open the following file using a text editing application such as Notepad:

C:\WINDOWS\system32\drivers\etc\hosts

2. Specify the static IP address and host name of the system as illustrated by the following screen:



```
hosts - Notepad
File Edit Format View Help
# Copyright (c) 1993-1999 Microsoft Corp.
#
# This is a sample HOSTS file used by Microsoft TCP/IP for Windows.
#
# This file contains the mappings of IP addresses to host names. Each
# entry should be kept on an individual line. The IP address should
# be placed in the first column followed by the corresponding host name.
# The IP address and the host name should be separated by at least one
# space.
#
# Additionally, comments (such as these) may be inserted on individual
# lines or following the machine name denoted by a '#' symbol.
#
# For example:
#
#       102.54.94.97       rhino.acme.com           # source server
#       38.25.63.10      x.acme.com              # x client host
127.0.0.1        localhost
172.31.255.255  comp-001
```

3. Close the file and save the changes.

Note: When you use this solution, you must modify the Hosts file when you change the IP address, and delete the specified information when you revert to a dynamic IP address.

Authentication Problems

This section contains the following topics:

[Authentication Security Settings](#) (see page 863)

[Restricted Users Cannot Access the Activity Log and the Audit Log](#) (see page 865)

[Authentication Errors Occur When Stopping and Starting the CAportmapper Service](#) (see page 868)

[Credential Problems with Cloud Connection](#) (see page 868)

[Unable to Open Report Writer](#) (see page 869)

[Unable to Create Reports from Report Manager](#) (see page 870)

[Updating Passwords Using ca_jobsecmgr Does Not Complete Successfully](#) (see page 871)

Authentication Security Settings

The following section provides guidance to help you address authentication and security-related issues when using Arcserve Backup. Because symptoms of security-related issues vary widely, this section includes possible resolutions only.

Possible Resolutions

The following list of resolutions can help you address security-related issues:

- Ensure that Arcserve Backup has properly authenticated the caroot account. Use the Server Configuration Wizard to perform this authentication. Select the Password for Backup Server Logon and Administration option to set the caroot account and password.
- Ensure that the Arcserve Backup folder is shared with:
 - Administrator--Full Control
 - Arcserve Backup System Account--Full Control
 - Backup Operators--Change and Read
- If you are having general problems understanding what rights your backup account needs to perform storage functions in your environment, consider the following information.

If you are backing up only your local Arcserve Backup server, the Arcserve Backup System account configured at installation has sufficient rights (Administrator and Backup Operator).

If you are backing up remote data within your domain (through the Client Agent for Windows or through the network facility of Arcserve Backup), your backup account requires additional rights. The following is a general outline of common permissions necessary for a powerful backup account. You can tailor your backup account to match your needs, and some rights may not be required in your environment.

Note: Security requirements for storage-related functions are dependent upon the resources accessed. Windows security rules and requirements should be considered at all times.

The backup account should have the following Group Rights:

- Administrator
- Backup Operator

Note: A user in the Backup Operator Group does not have rights to access the Arcserve Backup database. As a result, member servers are not visible to the user in the Backup Manager.

- Domain Administrator

The backup account should have the following Advanced Rights:

- Act as part of Operating System
 - Log on Locally
 - Log on as a service
- When prompted by Arcserve Backup to enter security within a domain, always use domain\username as the context.
 - If you have established a connection between two computers with one login/password session, Session Credential Conflicts can occur if you attempt to establish a second connection with the same login/password session. Consider any existing sessions you may have and how these may affect Arcserve Backup ability to access a resource.
 - The security entered in Arcserve Backup jobs is static and does not update dynamically if the Windows security account information changes at the operating system level. If you change the account information packaged in your Arcserve Backup jobs, you must modify the jobs and repackage them with the proper security information.
 - You must back up remote Registry and System State information through the Arcserve Backup Client Agent for Windows.
 - If you manually stopped and restarted the CA Remote Procedure Call service (CASportmap) without using the cstop and cstart command, the service cannot communicate with its port assignments properly. This can prevent a user account with caroot equivalence from logging in to the Arcserve Backup domain.

To remedy the inability to log in to the Arcserve Backup domain, run the cstop command and then run the cstart command. This enables the service to communicate properly and lets the user account with caroot equivalence log in to the Arcserve Backup domain.

Restricted Users Cannot Access the Activity Log and the Audit Log

Valid on Windows Server 2003, Windows Vista, and Windows Server 2008 systems.

Symptom

When you log in to Arcserve Backup using Windows authentication and a Windows account with restricted privileges (for example, Backup Operator and Remote Desktop User), you cannot access the Arcserve Backup Activity Log and Audit Log.

Note: This behavior does not occur when you configure Arcserve Backup to authenticate with the Arcserve Backup database using SQL Server authentication.

Solution

To remedy this behavior, grant all Windows accounts that require access to the Activity Log and the Audit Log the privilege to connect to SQL Server using Microsoft SQL Server authentication.

To grant privileges on Microsoft SQL Server 2000

1. Allow Microsoft SQL Server to communicate by adding the application to the Windows firewall exceptions list.
2. Allow the restricted Windows accounts to communicate as members of the SQL instance group or the sysadmin group.

To grant privileges on Microsoft SQL Server 2008 Express Edition and Microsoft SQL Server 2005 or later

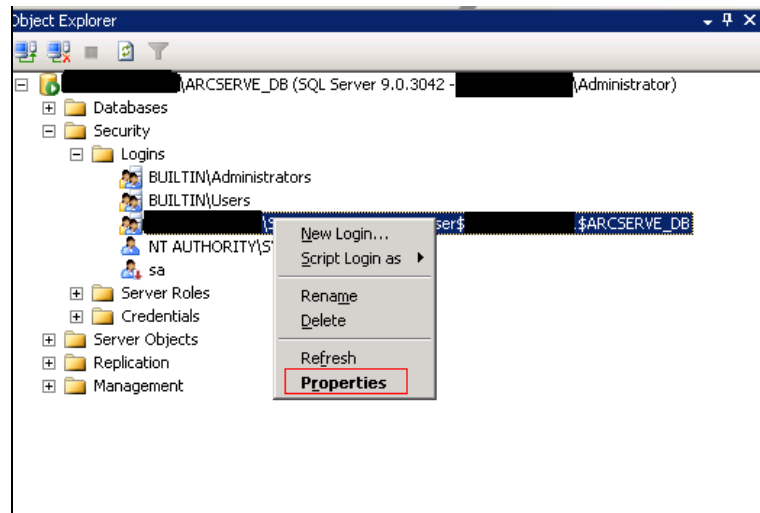
1. Allow Microsoft SQL Server to communicate by adding the application to the Windows firewall exceptions list.
2. Allow the restricted Windows accounts to communicate as members of the SQL instance group or the sysadmin group

3. Add the restricted Windows accounts into Microsoft SQL Server by doing the following:

- a. Open Microsoft SQL Server Management Tool.


Open Object Explorer.

Expand the Arcserve Backup server, Security, and Logins as illustrated by the following screen:



- b. Right-click the Arcserve Backup instance and click Properties on the pop-up menu.

The Login Properties dialog opens as illustrated by the following screen.



- c. Click User Mapping.
- d. In the Users mapped to this logon field, select the databases that you want to map by click the Map check box.
- e. In the Database role membership for field, select the roles that you want to apply to this user for the selected database and click OK.

Authentication Errors Occur When Stopping and Starting the CAportmapper Service

Valid on Windows platforms.

Symptom

Authentication errors occur that prevent you from opening the Manager Console after you stop and restart the CAportmapper service.

Solution

This condition only occurs under the following sequence of events:

- All Arcserve Backup services are running.
- You stop the CAportmapper service using either the Net Stop command or by stopping the service from the Windows Computer Management console.
- You restart the CAportmapper service.

Important! You must stop and start the CAportmapper Service using the `cstop` or `cstart` command. These commands let you stop and start all Arcserve Backup services sequentially, based on their dependencies with other Arcserve Backup services.

More information:

[Stop and Start All Arcserve Backup Services Using Batch Files](#) (see page 493)

Credential Problems with Cloud Connection

Valid on Windows Server 2003, Windows Server 2008, and Windows Server 2012 systems.

Symptom

Log in to the Device Manager to create a cloud connection. After entering your credentials, the connection passes whether the credentials entered were right or wrong. You later access the cloud but fail to connect because the credentials you entered can possibly be invalid.

Solution

If the error encountered is a proxy error, then you need to check if the proxy parameters are correct. For example: IP address or username and password.

If the error encountered is a network error, then you need to check if the vendor site is available.

Unable to Open Report Writer

Valid on Windows platforms.

Symptom

When you open Report Writer from the Manager Console, the following message appears:

You must be logged in to Windows using an account with administrative privileges to open Report Writer. For more information, see the Arcserve Backup documentation.

Solution

This problem occurs when you log in to the Arcserve Backup server using a domain user account that is not a member of an administrative group. If you apply caroot equivalence to the account, you can open the manager console, however, you cannot open Report Writer.

The solution to this problem is to allow the domain user account to join an administrative group. However, if you do not want to allow the user account to join an administrative group, complete the following alternative solution:

1. Log in to the Arcserve Backup server using a Windows administrative account.
2. Open Windows Registry Editor and browse to the following key:
HKEY_LOCAL_MACHINE\Software\Microsoft\Internet Explorer
3. Edit the permissions for this key by doing the following:
 - a. Right-click Internet Explorer and click Permissions on the pop-up menu.
The Permissions for Internet Explorer dialog opens.
 - b. Add the user account and allow Full Control for this key.
 - c. Close all dialogs and then log off.
4. Log in to the Arcserve Backup server using the domain user account.
Open the manager console.
Open Report Writer.
Report Writer opens.

Unable to Create Reports from Report Manager

Valid on Windows platforms.

Symptom

When you try to create reports from Report Manager, the following message appears:

Arcserve Backup cannot write to the specified directory or cannot communicate with Arcserve Backup engines.

Solution

This problem can occur under the following conditions:

- The user does not have write access to the specified directory.
To remedy this condition, verify that the user has write permissions to the specified directory.
- The path specified is not correct.
To remedy this condition, verify that the path specified is correct.
- Report Manager cannot communicate with the Arcserve Tape Engine, the Arcserve Database Engine, or both.
To remedy this condition, verify that the Tape Engine service and the Database Engine service are running on the backup server.

Updating Passwords Using ca_jobsecmgr Does Not Complete Successfully

Valid on Windows platforms.

Symptom

The ca_jobsecmgr command provides a means to specify whether you want your changes (user name and password) to be applied to a specific Arcserve Backup server or to all backup servers in the current Arcserve Backup domain. Updating passwords using ca_jobsecmgr -s may not complete the password update process.

Solution

Complete the following steps to correct this behavior.

1. Log in to the Arcserve Backup server.
2. Open the Backup Manager and click the Source tab to display the source directory tree.
3. From the Backup menu, click Source Context Menu and then click Security to open the Security dialog.
4. Complete the required information on the Security dialog and click OK.
5. (Optional) For the jobs that are not scheduled, resubmit the failed jobs.

Backup and Restore Problems

This section contains the following topics:

[Jobs Do Not Start on Schedule](#) (see page 872)

[Cannot Back Up Open Files](#) (see page 873)

[Restore Job Fails on Citrix Server](#) (see page 874)

[Local Restore Data Backed Up with Compression and/or Encryption Failed](#) (see page 875)

[Arcserve Backup Does Not Restore Data Based on File Access Time](#) (see page 875)

[GUI Freezes in Active Directory Restore Mode](#) (see page 876)

[Scheduled Backup Jobs Fail After Changing the Login Credentials for Agent Computers](#) (see page 876)

[Running a Full Backup \(Clear Archive Bit\) or Incremental Backup on a Source Previously Backed Up in a SFB Job Causes Data Loss](#) (see page 877)

[Unable to Run an On-Demand Synthetic Full Backup](#) (see page 878)

[Restore Job Activity Log Errors](#) (see page 879)

[Arcserve D2D Backup Jobs Fail When Submitted from the Backup Manager](#) (see page 880)

[Arcserve Backup Does Not Back up Files that were Moved or Renamed](#) (see page 880)

[Arcserve Backup Cannot Perform a Hyper-V Backup Job That Contains Data on SMB Shares](#) (see page 882)

[Arcserve Backup Cannot Perform a Hyper-V Writer or SQL VSS Writer Back Up Job That Contains Data on SMB Shares](#) (see page 882)

[Memory Usage Increases When Restoring Data From a Remote FSD Located on a Windows 2008 SP2 and Windows 2008 R2 Server](#) (see page 882)

[Restore Job Fails on MAC Agents](#) (see page 883)

Jobs Do Not Start on Schedule

Valid on Windows, UNIX, and Linux platforms.

Symptom

Scheduled jobs do not start on schedule.

This problem is most likely to occur when you have multiple Arcserve Backup servers in a centralized management environment and the Arcserve Backup primary server, the member servers, or the Arcserve Console reside in different time zones.

Solution

To remedy this problem, synchronize the system time on the primary server with the system time on all member servers in the Arcserve Backup domain.

Use the Windows Time Service to complete this task.

Note: For information about how to synchronize the time using the Windows Time Services, see the Windows Help and Support.

Cannot Back Up Open Files

The following section provides guidance to help you address issues related to open files when using Arcserve Backup.

Possible Problems and Resolutions

If a particular resource you are backing up is locked or in use by the operating system, you may receive the following errors. These errors may be preceded by error code W3404.

Note: The Arcserve Backup Agent for Open Files reconciles many common open file errors. If you are not using this Agent, you should consider doing so. We also recommend that you perform remote backups using the Arcserve Backup Client Agent for Windows.

MS Error Code	Cause and Resolution
SHARING VIOLATION	<p>Cause: File sharing violation error. Another process (such as an application service) was using a target file when Arcserve Backup ran a Backup Job.</p> <p>Resolution: Stop all services and applications using the target file and restart the Backup.</p>
ACCESS DENIED	<p>Cause: A target file for the backup job was not accessible, or another process (such as an application service) was using a target file when Arcserve Backup ran the backup job.</p> <p>Resolution: Ensure that your user account has sufficient rights to access the target file, and stop all services and applications using the target file before restarting the backup job.</p>
FILE NOT FOUND	<p>Cause: A target file has been deleted or moved between the submission and the execution of a Backup Job.</p> <p>Resolution: Modify and repackage Job and retry.</p>
PATH NOT FOUND	<p>Cause: A target file path has been deleted or changed between the submission and the execution of a Backup Job.</p> <p>Resolution: Modify and repackage Job and retry.</p>

MS Error Code	Cause and Resolution
BAD NET PATH	<p>Cause: A Backup Job is submitted to a Remote Machine and a target network path was not detected because of a missing path or network protocol delay.</p> <p>Resolution: Confirm your network environment and retry the Backup Job.</p>

Restore Job Fails on Citrix Server

Valid on Windows platforms.

Symptom

A restore job fails on a server running Citrix. The list that follows describes the environment on the Citrix server:

- Citrix 4.0
- Microsoft SQL Server (hosting the Citrix database instance)
- Client Agent for Windows
- Agent for Microsoft SQL Server

Solution

When you restore a Citrix 4.0 server, the job may fail because the Microsoft SQL Server instance hosting the Citrix database instance will not start after the restore job is complete.

To remedy this problem, restart the Citrix database instance manually.

Local Restore Data Backed Up with Compression and/or Encryption Failed

Valid on Windows platforms.

Symptom

During the backup, which is using compression and/or encryption, the file is increasing in size. When you perform a local restore, the job is incomplete with Error E3453 - Unable to write stream data.

Solution

1. Create a DWORD value called RestoreDCENDataByWriteFile under the following registry key and set it to 1.
HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\Base\Task\Restore
2. Select the files being restored that failed, set the destination, and submit a new restore job.

Arcserve Backup Does Not Restore Data Based on File Access Time

Valid on Windows platforms.

Symptom

Arcserve Backup does not restore data when filtered based on the last file access time.

Solution

Using the Restore Manager and the ca_restore command line utility, you can recover files based on the last time they were accessed. However, Arcserve Backup does not store the last file access time in the backup records. As a result, Arcserve Backup cannot restore data based on the last file access time.

The remedy for this problem is to modify the registry key listed below on the agent computer and then submit the backups.

HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA ARCServe
Backup\ClientAgent\Parameters\FileAccessTimeFlag

Value: 1

GUI Freezes in Active Directory Restore Mode

Valid on Windows platforms.

Symptom

The Arcserve Backup manager freezes when Windows is started in the Active Directory Restore Mode and you cannot restore the Active Directory when Windows is in the Active Directory Restore mode.

Solution

This behavior occurs because Microsoft SQL Server Express Edition and Microsoft SQL Server cannot function when you start Windows in the Active Directory Restore mode. To remedy this problem, submit the Active Directory restore job using Windows Normal Mode and execute the Active Directory restore job after restarting Windows using the Active Directory Restore mode.

Scheduled Backup Jobs Fail After Changing the Login Credentials for Agent Computers

Valid on all operating systems.

Symptom

Scheduled backup jobs fail after you change the login credentials (User name, Password, or both) for agent computers.

Solution

This is expected behavior.

Arcserve Backup stores information about the login credentials for agent computers in the Arcserve Backup database. As a scheduled backup job runs, Arcserve Backup retrieves the login credentials from the database, which allows Arcserve Backup to log in to the agent and process the scheduled backup job. If you change the login credentials for an agent computer without updating the database, Arcserve Backup cannot log in to the agent to process the job, which causes the job to fail.

As a best practice, you should update the Arcserve Backup database with the new login credentials immediately after you change the login credentials on your agent computers.

To update the Arcserve Backup database with the new login credentials, do the following:

1. Open the Backup Manager and click the Source tab.

The backup source directory tree appears.

Note: From the Views menu, you can select either Classic View or Group View.

2. Locate the target agent computer.

Right-click the agent computer and click Security on the pop-up menu.

The Security dialog opens.

3. Specify the new User name, Password, or both that is required to log in to the agent computer and click OK.

The Security dialog closes and the Arcserve Backup database is updated with the current login credentials for the agent computer.

Running a Full Backup (Clear Archive Bit) or Incremental Backup on a Source Previously Backed Up in a SFB Job Causes Data Loss

Valid on Windows platforms.

Symptom

Running a Full backup (Clear archive bit) or Incremental Backup to back up a source that was backed up in a synthetic full backup job causes data loss.

This occurs when you submit a backup job with Synthetic backup enabled and select folders and files as the backup source. If the data changes in that backup source and you submit a Normal full backup (Clear archive bit) or Incremental Backup on this source prior to the next scheduled synthetic full backup job, the changed data is missed in the next point-in-time incremental backup because the archive bit is cleared for the changed files in the folder. This ultimately causes the changed files to be missed in the next synthetic full backup session. Arcserve Backup provides one warning message: *W29893 Missed file XXX, the file may be backed up by another job*. Arcserve Backup will do a real full backup on the next synthetic day.

Solution

If you want to submit a Normal full backup job to back up the same source as the synthetic full backup job, it is recommended that you use a Normal full (Keep archive bit) backup job.

Unable to Run an On-Demand Synthetic Full Backup

Valid on Windows platforms.

Symptom

When you submit an on-demand synthetic full backup job, the job does not execute and a message dialog opens indicating the reason why the synthetic full backup was not executed.

You may receive one of the following four messages:

1. Cannot submit synthetic job since this job has been modified

The synthetic full backup job was modified and no real full backup and incremental job is done.

2. Unable to submit the data synthetic job, or there are no sessions that need to be synthesized

The synthetic full backup job cannot synthesize full sessions because there are no suitable sessions that can be synthesized. This can occur due to two possible scenarios:

- Sessions were already synthesized before this job was submitted, so there is no need to synthesize them again.
- A previous incremental session was not backed up successfully.

3. Data synthetic job submitted successfully

4. Data synthetic job submitted successfully
n sessions were skipped that failed

Even though there were some sessions that were synthesized to full sessions, some of the sessions were either skipped or failed. This can occur due to two possible scenarios:

- Sessions were already synthesized before this job was submitted, so there is no need to synthesize them again.
- A previous incremental session was not backed up successfully.

Solution

1. No action is required for this message. It will do a real full backup on the next synthetic day.
2. If you have any sessions that failed, verify the previous incremental session and make any necessary adjustments to ensure a successful backup.
3. No action is required for this message.
4. If you have any sessions that failed, verify the previous incremental session and make any necessary adjustments to ensure a successful backup.

Restore Job Activity Log Errors

Symptom

If there are any unreachable cloud devices when you submit a restore job, the restore job completes successfully but there may be the following errors in the Restore Job Activity Log:

E6008 [server name] [date time]

Device communication failure [Cloud Device [cloud device name]: [Cloud library internal error.]]

E6112 [server name] [date time]

Failed to read the tape's header. (Serial Number=[])

Solution

When you perform a restore job, the task will send the tape name, random ID, and sequence number to the Tape Engine in which to find the group the device is located. Once found, the Tape Engine will enumerate all the groups and open the devices in the groups, read the tape header, and then compare the read tape name, random ID, and sequence number with the inputs. If they are all equal, the group is found. When the restore job attempts to open the cloud device and read the tape header, then one of the two errors described previously are returned. This is because the Tape Engine is not aware if it requires the cloud device or not. This does not affect the actual restore job, which completes successfully.

Arcserve D2D Backup Jobs Fail When Submitted from the Backup Manager

Valid on Windows Platforms.

Symptom

When you submit Arcserve D2D backups from the Backup Manager, the jobs fail when the source data includes Arcserve D2D sessions that were backed up previously by Arcserve Backup. Error message AW0813 appears in the Activity Log.

Solution

This is expected behavior. If a backup job consists of Arcserve D2D sessions that were backed up previously, Arcserve Backup fails the job and reports error message AW0813 in the Activity Log. This behavior is designed to help prevent backing up Arcserve D2D data redundantly.

If you want to modify this behavior to allow Arcserve Backup to back up Arcserve D2D sessions that were backed up previously, do the following:

1. From the Arcserve D2D server (node) that you are backing up, open Windows Registry Editor.
2. Open the following key:

HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA ARCserve
Backup\ClientAgent\Parameters\AllowRedundantD2DBackups

3. Change the DWORD value of AllowRedundantD2DBackups to 1.

Note: If the above-described Registry key is not present on the node that you are backing up, you must create the key.

4. Close Windows Registry Editor.

Arcserve Backup Does Not Back up Files that were Moved or Renamed

Valid on Windows platforms.

Symptom

When performing file system backups, Arcserve Backup does not back up files that were moved or renamed on a node after the last backup completed. Arcserve Backup does not generate error or warning messages in the Activity Log.

Solution

This is expected behavior. The Client Agent for Windows lets you back up files only if they are modified after the last backup completed. The Client Agent for Windows cannot identify files that were moved or renamed as files that were modified.

To help prevent this behavior from occurring, you must install and license Arcserve Backup Agent for Open Files, and verify that the CA Universal Agent Service is running in your backup environment before you rename or move files on the node. The CA Universal Agent service allows Arcserve Backup to detect files and directories that were changed or moved on the agent node. If the Universal Agent Service was not running when the files were moved, renamed, or both, ensure you have installed and licensed Arcserve Backup Agent for Open Files, and submit a full backup of the node to help ensure that Arcserve Backup backs up all of the files.

When you move or rename a file or folder, the Archive bit for the object is not updated. As a result, Arcserve Backup does not back up the file or folder the next time an incremental backup job runs. To help ensure that the Universal Agent detects and backs up the moved and renamed files, the Universal Agent enumerates all fixed volumes that are present on the node at specific intervals. The Universal Agent then monitors each volume for moved or renamed files and folders. If a file or folder is detected, the archive bit for the object is updated.

You can enable and disable this function of CA Universal Agent Service. The function is enabled by default. To modify the default behavior, do the following:

1. Open Windows Registry Editor.
2. Browse to the following registry key.
`HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA ARCserve Backup\UniversalClientAgent\Parameters\VolChangeMgr`
3. Change the value that you require for VolChangeMgr.
Default value--1 Enabled
Optional value--0 Disabled
4. Close Windows Registry Editor.
5. Restart CA Universal Agent Service.

The default time interval for volume enumeration is one hour. To modify the default behavior, do the following:

1. Open Windows Registry Editor.
2. Browse to the following registry key.
`HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA ARCserve Backup\UniversalClientAgent\Parameters\RefreshVolEnum`
3. Specify the value that you require for ResfreshVolEnum.
Default value--3600000 (1 hour in milliseconds)
Note: One minute equals 60000 milliseconds.
4. Close Windows Registry Editor.
5. Restart CA Universal Agent Service.

Arcserve Backup Cannot Perform a Hyper-V Backup Job That Contains Data on SMB Shares

Valid on Windows platforms.

Symptom

When you submit Hyper-V backup jobs where the disks or configuration files of the virtual machine are on remote SMB (Server Message Block) shares, the backup job fails.

Solution

This behavior occurs because the File Server VSS Agent Service is not enabled on the file server. You can enable the File Server VSS Agent Service by logging on to the Server Manager and selecting the following menu options: Server Roles, File and Storage Services, File and iSCSI Services, and File Server VSS Agent Service.

Arcserve Backup Cannot Perform a Hyper-V Writer or SQL VSS Writer Back Up Job That Contains Data on SMB Shares

Valid on Windows platforms.

Symptom

When you submit a Hyper-V or SQL VSS Writer backup job where the application has data on remote SMB (Server Message Block) shares, the backup job fails.

Solution

This behavior occurs because the File Server VSS Agent Service is not enabled on the file server. You can enable the File Server VSS Agent Service by logging on to the Server Manager and selecting the following menu options: Server Roles, File and Storage Services, File and iSCSI Services, and File Server VSS Agent Service.

Memory Usage Increases When Restoring Data From a Remote FSD Located on a Windows 2008 SP2 and Windows 2008 R2 Server

Valid on all operating systems.

Symptom

The system response becomes very slow because of the increase in memory usage. The reason why this occurs is because the behavior of the Windows operating system as the system file cache consumes most of the physical RAM.

Solution

To resolve this issue, follow the work-around in the knowledge base article from the Microsoft link: <http://support.microsoft.com/kb/976618/en-us>.

Restore Job Fails on MAC Agents

Valid only on Macintosh operating systems.

Symptom

A restore job fails on a MAC agent when you attempt to restore from multiple MAC sessions as part of a single restore job.

Solution

To resolve this issue, go to the "/etc/sysctl.conf" file and modify the parameter "kern.sysv.shmall" to "kern.sysv.shmall=4096". Then restart the computer.

Media Problems

This section contains the following topics:

[Tape Errors Occur When Backing Up or Restoring Data](#) (see page 883)

[Arcserve Backup Cannot Detect RSM Controlled Devices on x64 Platforms](#) (see page 885)

[Arcserve Backup Does Not Detect a Cleaning Tape](#) (see page 886)

[Hardware Does Not Function as Expected](#) (see page 887)

[Autoloaders and Changers Appear Offline](#) (see page 888)

[Catalog Database Log Files Consume a Large Amount of Disk Space](#) (see page 888)

[Unrecognized Vaults Appear in the Media Management Administrator](#) (see page 890)

Tape Errors Occur When Backing Up or Restoring Data

The following section provides guidance to help you address issues related to tape errors when using Arcserve Backup.

Possible Problems

If you receive an error that suggests that there is something wrong with one of your tapes, you should take corrective action as soon as possible to ensure the security of your data. Before replacing your tape, however, you should make certain that it is the tape that is causing the problem, and not another part of your system. Try these steps to rule out the possibility that the problem is being caused by something other than the tape:

- Check the history of the Activity Log for the task that caused the error. Although you may get a media error, it may only be the consequence of an earlier error.

For example, during a backup job, you may receive a SCSI port error. After receiving this error, you may get errors that indicate a problem with the tape, or even with the drive, but it is possible that these errors are only a consequence of the problems with the SCSI port. Therefore, you should check the Activity Log for all the messages and errors you received prior to receiving the error that indicated a problem with your tape. In this way, you can determine whether there is actually a problem with your tape, or if the tape error was the consequence of another problem.
- Monitor the library robot. If the robot is not functioning properly, you may get tape errors. Make sure that the robot can move tapes in and out of the drives.
- Rule out the possibility of a mechanical problem with the drive. To do so, try one of these options:
 - Clean the drive, then perform the same task again.
 - If the tape still causes errors after the drive has been cleaned, move the tape to a drive that you know is in good working order and try the same task again. If you get the same error, then it is likely that there is a problem with the tape.

Note: If your drives are inside a library, and you want to try your tape in a different drive, the problem drive must be offline. If Arcserve Backup did not automatically set the drive to an offline status when it detected the media error, right-click the library and select Offline from the pop-up menu.
 - Try the same task on the same drive, but with a different tape. If you get the same error, then it is likely that the tape is fine, but that there is a problem with the drive or some other system component.

Possible Resolutions

After you have determined that there is a problem with the tape--part of the tape is unreadable, the tape is physically damaged in some way, and so on--you should replace the tape as soon as possible. Before you do that, you will need to back up the data on the bad tape to a reliable tape. You have two options at this point:

- Copy the data to a new tape
- Create a new backup tape

Copy Data to a New Tape

If you can read data from the tape, follow the steps below. If no data can be read, see [Create New Backup Tapes](#) (see page 885) for steps about creating a new backup tape.

To copy data to a new tape

1. Try moving the tape to a drive that you know is in good working order. You can also try cleaning the drive.
2. Use the Tapedcopy utility to copy the data from the bad tape to the new tape.

Note: If the bad tape was part of a library, export the tape from the library so that it does not get used again.

Create New Backup Tapes

If you are unable to read any data from the bad tape, follow these steps to create a new backup tape.

To create new backup tapes

1. Remove the bad tape. If the bad tape is part of a library, export it.
2. Insert a new tape and resubmit the backup job.

Arcserve Backup Cannot Detect RSM Controlled Devices on x64 Platforms

Valid on Windows Server 2003 x64 platforms.

Symptom

Arcserve Backup is installed on a Windows Server 2003 x64 system. From the Device Manager, Arcserve Backup cannot detect devices controlled by the Removable Storage Manager (RSM), and you cannot enable and disable RSM controlled devices.

Solution

The Arcserve Backup Manager Console is designed using x86 architecture. The Manager Console cannot detect RSM controlled devices on all Windows 64-bit platforms. To remedy this limitation, you can enable and disable RSM controlled devices using the RSM Computer Management utility.

Arcserve Backup Does Not Detect a Cleaning Tape

Valid on Windows platforms.

Symptom

There is a cleaning tape in the library but Arcserve Backup does not detect a cleaning tape.

Solution

To remedy this problem, Arcserve Backup must detect the location (slot) of the cleaning tape. There are two methods that you can use to let Arcserve Backup detect the location (slot) of the cleaning tape.

- **Method 1**--Let Arcserve Backup discover the location (slot) of the cleaning tape. To do this, complete the following steps:

1. Insert a cleaning tape into any available slot in your library.
2. Open the Device Manager window, right-click the library and select Inventory from the pop-up menu.

Arcserve Backup inventories the media in the slots. After the inventory process is complete, Arcserve Backup detects the presence of a cleaning tape. The slot where the cleaning tape resides becomes the cleaning slot.

- **Method 2**--Manually specify the location (slot) of the cleaning tape. To do this, complete the following steps:

1. Open the Device Manager window, right-click the library and select Properties from the pop-up menu.

The Library Properties dialog opens.

2. Click the Cleaning tab.

The Cleaning options appear.

3. From the Available Slots lists, click an available slot and then click the Add button.

The available slot moves to the Clean Slots list.

4. Click OK.
5. Insert the cleaning tape into the slot specified.

Hardware Does Not Function as Expected

The following sections provide guidance to help you address hardware-related issues when using Arcserve Backup.

Possible Problems

If you are having hardware-related issues with Arcserve Backup, you may experience the following symptoms:

- E6300 Windows NT SCSI Port Errors in the Arcserve Backup Activity Log.
- Slots not showing the proper status or not updating properly.
- Devices not listed properly in Arcserve Backup Device screen.
- Critical hardware errors in the Arcserve Backup Activity Log.
- Inability to properly configure your tape device.
- Hardware-related inconsistencies in day-to-day Arcserve Backup functions.

Possible Resolutions

The following list of resolutions can help you address hardware-related issues:

- Ensure that the operating system is properly recognizing your devices. If the operating system is having a problem seeing the devices, Arcserve Backup may not function properly.
- Ensure that the latest Device Patch is installed for Arcserve Backup.
- Check the Arcserve Backup Certified Device list to ensure compatibility with your device's firmware.
- Ensure that the proper SCSI drivers are loaded for your SCSI adapter.
- Try using different tapes to ensure that the errors are not media-related.
- Check physical connections and SCSI cabling. Errors can occur because of physical problems, such as a bent SCSI pin.
- If you are running Arcserve Backup on a Windows platform, run Device Configuration by selecting Device Configuration from the Configuration menu. Choose Enable/Disable Devices (For RSM). If you see your devices listed in the Available Devices window, ensure that the check box is selected. Doing so gives Arcserve Backup full control over your device, and does not allow the Windows Removable Storage Manager service to interfere.
- Check if third-party device monitoring or controlling services are running and, if necessary, disable these services, as they may be conflicting with the Arcserve Backup ability to control the device.

- If you are using Library Quick Initialization, ensure that you disable this option when you are troubleshooting hardware and devices. You can then apply the Library Quick Initialization option after troubleshooting has been completed.

Note: The Library Quick Initialization option can be found on the General tab on the Library Properties dialog.

Autoloaders and Changers Appear Offline

Valid on all Windows platforms. Affects stand-alone libraries and changers.

Symptom

In Device Manager, the device appears offline. Jobs associated with the device fail. The Tape Engine is running.

Solution

When the Tape Engine detects problems with devices, Arcserve Backup statuses the device as offline and generates a message, similar to the following, in the Activity Log:

```
[Library Failure: Manual intervention required [Device:5][Omega REV LOADER]]
```

To remedy this problem, do the following:

1. Stop the Tape Engine.
2. Disconnect the device and correct the problem with the device.

Note: For information about troubleshooting the device, see the device manufacturer's documentation.

3. Connect the device to Arcserve Backup.
4. Start the Tape Engine.

If the device is working properly, Arcserve Backup will detect the device and status the device as online.

Catalog Database Log Files Consume a Large Amount of Disk Space

Valid on Windows Server 2003, Windows Server 2008, and Windows Server 2012 systems.

Symptom

Arcserve Backup is generating a high number of catalog database log files that are consuming a large amount of disk space on the Arcserve Backup server.

Solution

Arcserve Backup generates catalog database log files that can be used for debugging purposes. The log files are stored in the following directory on the Arcserve Backup server:

```
C:\Program Files\CA\ARCserve Backup\LOG
```

By default, Arcserve Backup generates up to 3 log files, and the overall size of each log file can be up to 300 MB. However, Arcserve Backup lets you customize the behavior of the log files by creating registry keys that control the number and size of the log files.

If you created the required optional keys and specified high values for the number and size of the log files, the log files can consume a large amount of disk space on the backup server.

To control the number and size of the catalog database log files, do the following:

1. Open Windows Registry Editor.
2. Open the following key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA ARCServe Backup\Base\Database
```

3. Modify the following string values:

- **DebugLogFileSize**

- Range: 1 MB to 1024 MB (1 GB)
- Recommended value: 300 MB

Note: With this value, Arcserve Backup generates a new log file after the catalog database logging activity causes the current log file to exceed the specified value.

- **LogFileNum**

- Range: 1 to 1023
- Recommended value: 3

Be aware of the following:

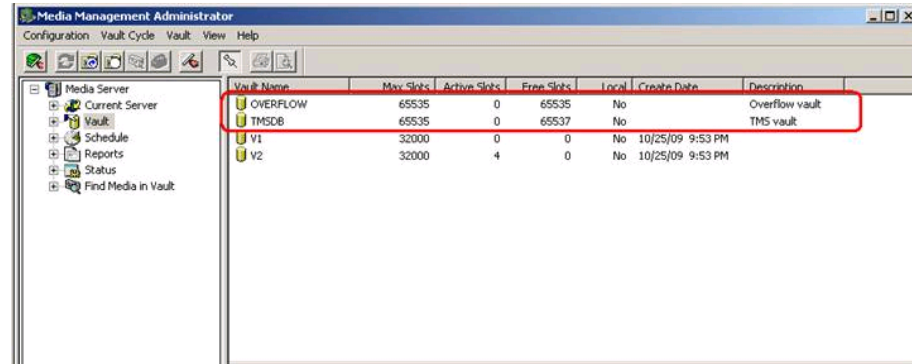
- With this value, Arcserve Backup retains the number of log files specified.
- If catalog database logging activity causes the number of log files to exceed the specified limit, Arcserve Backup deletes the oldest log files until the number of log files is equal to the specified limit.

Unrecognized Vaults Appear in the Media Management Administrator

Valid on Windows backup servers managing data mover servers.

Symptom

Two unrecognized vaults named TMSDB and OVERFLOW appear in the Media Management Administrator (MM Admin). You did not configure either vault. The following diagram illustrates the vaults:



Solution

This is expected behavior. Arcserve Backup creates two default vaults without media associations when you install UNIX and Linux Data Mover, register the data mover servers with the primary server, and migrate MM Admin data from your previous Arcserve Backup installation. The unrecognized vaults will not affect your backup and storage operations.

SAN Configuration Problems

This section provides the information you need to troubleshoot Arcserve Backup for Windows SAN installations.

Devices are Not Shared

Valid on Windows platforms.

Symptom

Devices attached to the Arcserve (SAN) Primary server are not marked as shared or slots could not be shown in the Arcserve Backup GUI.

The mechanism that is used to determine whether a device is “shared” is now dynamic. The Arcserve Backup Member server is responsible for detecting “shared” devices and reporting these “shared” devices to the Arcserve Backup (SAN) Primary server. Therefore, if the Arcserve Backup Tape Engine service is not running on any of the SAN-attached Member servers, then no “sharing” is occurring and the devices will not show as being “shared”.

Solution

Make sure the Arcserve Backup Tape Engine service has been started on at least one of the SAN-attached Member servers. If necessary, start the Tape Engine service on one or all of the SAN-attached Member servers.

Devices are Not Shared and the Tape Engine is Running

Valid on Windows platforms.

Symptom

The Tape Engine service is up and running on all SAN-attached servers in the Arcserve Backup Domain, but the devices are NOT marked as "shared" in the Arcserve Backup GUI. The SAN-attached devices are not detected by all the SAN attached servers correctly.

Solution

Analysis of your SAN zoning maybe necessary to make sure that all servers participating in the SAN can see all devices that are “shared” through the SAN. To do this you need to check the following:

- Check on each server that the “shared” devices can be seen by the Operating System, by checking the Windows Device Manager.
 - If the “shared” devices cannot be seen by Windows, then double check your SAN zoning to make sure this server is included. If it is included, reboot the server to get Windows to discover the devices. When you get Windows to see the devices, then you can restart the Tape Engine on that machine.
 - If the server having problems seeing the devices is the Primary server, then you need to restart the Arcserve Backup Tape Engine service on this server and then on all SAN-attached Member servers in that domain.

- Check if the “shared” devices can be seen by Arcserve Backup through the Arcserve Backup Device Manager, looking under each server.
 - If “shared” devices cannot be seen by Arcserve Backup, but Windows does see them, then you need to restart the Tape Engine service.
 - If the server having problems seeing the devices is the Primary server, then you need to restart the Arcserve Backup Tape Engine service on this server and then all SAN-attached Member servers in that domain.

Shared Devices Appear as Unavailable or Offline

Valid on Windows platforms.

Symptom

The shared devices on SAN-attached Member servers are marked as unavailable or offline.

This could be caused by the order in which the Arcserve Backup Tape Engine services were started in the domain.

- If the Arcserve Backup Member server's Tape Engine service was started before the Primary server's Tape Engine has finished initializing, the Member server will wait for the Primary server for a period of time, but eventually will start without being able to "share" the devices.
- If the Arcserve Backup Member server's Tape Engine service was started before the Primary server's Tape Engine service was started, the Member server will wait for the Primary server, for a period of time, but eventually will start, without being able to "share" the devices.

Solution

Check the Arcserve Backup Activity log to see when, and in what order the services might have started. Make sure all devices on the Arcserve Backup Primary server are initialized, then just restarting the Tape Engine service on the Member server(s) should be enough to resolve this.

Important! When starting the Tape Engine service in a Arcserve Backup Domain that is part of a SAN, it is important to always start the Primary server's service first and let it fully initialize before starting the Tape Engine service on any of the Member servers.

Shared IBM Devices Appear as Unavailable or Offline

Valid on Windows platforms.

Symptom

The shared IBM tape devices on SAN-attached Member servers are marked as unavailable or offline.

If you have IBM Ultrium/LTO tape drives, in some cases the IBM LTO Tape driver (Windows Server 2003) will issue a SCSI Reserve command to the tape devices when the Arcserve Backup Primary server opens the devices. This is not a problem for the Primary server, but when the SAN-attached Member server tries to access these tape devices it will fail and the devices will not be usable from that server.

Solution

In the Arcserve Backup Activity log for the Member server, you may see SCSI Port errors when the Arcserve Backup Tape Engine service is started and tries to access these drives.

You can configure the IBM tape driver to not issue the SCSI Reserve command by performing the following registry procedure.

1. In the registry editor, access the registry key:

`\\KEY_LOCAL_MACHINE\\SYSTEM\\CurrentControlSet\\Services\\`

2. If the IBMtape driver is installed, search under the services key for a key name that is equal to the file name of the IBMtape driver.

For example, if the file name of the IBMtape driver `ibmtp2k3.sys`, the key name is `ibmtp2k3`.

Select the key, add a DWORD value named "DisableReserveUponOpen," and set it to 1.

3. Exit registry editor and reboot the server.

Backup Jobs Fail

Valid on Windows Platforms.

Symptom

- The Arcserve Backup SAN License for the server you are trying to run the backup on has expired.
- Each SAN-attached server must have a SAN License. The licenses are all applied to the Primary server in the corresponding Arcserve Backup Domain.
- The Arcserve Backup TLO License for the server you are trying to run the backup on has expired.
- Each SAN-attached server sharing a multiple drive library must have a TLO License. The licenses are all applied to the Primary server in the corresponding Arcserve Backup Domain.
- The Arcserve Backup Primary server Tape Engine service is no longer available.

Solution

1. Check the Arcserve Backup Activity log for any SAN license errors.
If necessary add the applicable SAN license.
2. Check the Arcserve Backup Activity log for any TLO license errors.
If necessary add the applicable TLO license.
3. Check the Arcserve Backup Primary server Tape Engine service status
 - a. In the Arcserve Backup GUI, access the Server Admin screen and check the status of the Tape Engine service.
 - b. In the Arcserve Backup Activity log, check for a Primary server Tape Engine stop event.
 - c. In the Windows system Event Viewer, check for a Tape Engine service stop event or an exception.

If the problem is related to the Primary Tape Engine service being down, then restart and try again.

If the problem persists, contact Arcserve support for online technical assistance and a complete list of locations, primary service hours, and telephone numbers.

Cluster-based Backup and Restore Problems

This section contains the following topics:

- [Prevent Job Failures](#) (see page 895).
- [Back Up MSCS Nodes on Remote Machines](#) (see page 896).
- [Back Up Arcserve Backup Database in a Cluster Environment](#) (see page 897).
- [Job Failure: Media Not Mounted](#) (see page 897).

Prevent Job Failures

Valid on Windows platforms.

Symptom:

How can I stop Arcserve Backup services in a cluster node without failover occurring?

Solution:

When a Arcserve Backup server is configured as cluster-aware, all critical Arcserve Backup services will be monitored by the cluster application (MSCS or NEC CLUSTERPRO/ExpressCluster). If some service fails, the cluster application will try to restart it or trigger a failover if the restart attempt fails. This means that you can no longer stop a service by using the Arcserve Backup Server Admin. If you attempt to stop a Arcserve Backup service, you will get a pop-up message indicating that it is not permitted.

To shut down any Arcserve Backup services for maintenance or configuration changes when you do not want Arcserve Backup to fail over to another node, perform the following procedure:

- For MSCS clusters, see [Stop MSCS From Monitoring Arcserve Backup Services on Windows Server 2003 Systems](#) (see page 917).
- For NEC CLUSTERPRO/ExpressCluster clusters, see [Stop HA Service Monitoring by NEC CLUSTERPRO/ExpressCluster](#) (see page 930).

Back Up MSCS Nodes on Remote Machines

Valid on Windows platforms.

Symptom:

How can I reliably back up MSCS Nodes with Arcserve Backup installed on remote machines?

For information about recovering clusters, see the *Arcserve Backup Disaster Recovery Option Guide*.

Solution:

The Arcserve Backup Windows File System Agent must be installed on each node of the cluster.

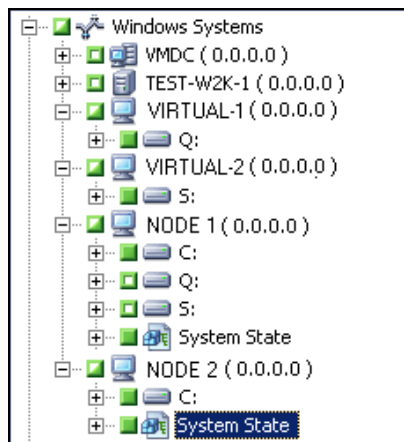
The challenge is to back up the shared disk reliably even if cluster shared disks fail over from one node to another. This can be done as follows:

1. Back up each of the nodes with their private disks and system state, using the hostname when submitting the backup jobs.

Note: Because shared disks can move from one node to another and there is no reliable way of predicting which node will own the shared disks during backup, **do not** back up shared disks using the machine hostname.

2. Back up the shared disks, using the cluster virtual name when submitting the backup job. If the shared disks fail over from one node to another, the cluster virtual node name fails over with it, so that Arcserve Backup always backs up the cluster shared disks. To ensure this, set up the cluster dependencies so that the cluster name and cluster shared disks fail over at the same time.

Note: To provide disaster protection for your cluster nodes, you must perform a full backup of each node.



Back Up Arcserve Backup Database in a Cluster Environment

Valid on Windows platforms.

Symptom:

How do I effectively backup the Arcserve Backup database (ASDB) in a cluster environment? (so that it can be recovered using the recoverdb operation)

Solution:

To ensure the backed up ASDB session can be used by the recoverdb operation you must backup the ASDB through the network name which you set during the setup phase.

For example:

1. For MS SQL Server 2008 Express, you must use the virtual name which the Arcserve Backup is deployed on.
2. For MS SQL Server 2005 cluster, you must use the virtual name which the SQL Server Cluster is deployed on. (In this case, make sure that you set the correct virtual name of the SQL Server cluster when installing Arcserve Backup. To find out the SQL Server Cluster virtual name, refer to the SQL Server Cluster document.

Job Failure: Media Not Mounted

Valid on Windows platforms.

Symptom:

When my jobs fail over from one cluster node to another, I receive messages such as "Please mount media XYZ, 1234." How do I resolve this problem?

Solution:

If you select a backup media on the Destination tab of the Backup Manager when submitting a backup job, the job backs up only to that specific media. If the backup device is not shared among the cluster nodes, the specific media is not available after failover. As a result, the backup operation fails. To resolve this problem, select Destination at Group Level in the Backup Manager when submitting backup jobs.

This problem does not occur if you are backing up to a shared device.

Miscellaneous Problems

This section contains the following topics:

[Discovery Service Does Not Function Properly](#) (see page 898)

[Arcserve Backup Servers and Agent Servers Cannot Communicate with Each Other](#) (see page 898)

[SRM PKI Alert is Enabled by Default](#) (see page 900)

[Job Queue Log Files Consume a Large Amount of Disk Space](#) (see page 902)

[Unable to View all Job Logs for Rotation Backups after the Database is Pruned](#) (see page 904)

Discovery Service Does Not Function Properly

The following section provides guidance to help you address issues related to discovery service problems when using Arcserve Backup.

Symptom

You may experience a problem in discovering Arcserve Backup applications on a specific machine. It is possible that the machine to be discovered is not located on the same subnet as the machine the discovery service is running (the default setting for Discovery Service is the local subnet).

Solution

Choose the Subnet sweep option in the Configuration window and restart the Discovery Service. Or, you can add the specific subnet or machine name (IP address) and restart the Discovery Service.

Arcserve Backup Servers and Agent Servers Cannot Communicate with Each Other

Valid on Windows Server 2008 R2 and Windows Server 2012 systems.

Symptom

Arcserve Backup servers, Arcserve Backup agent servers, or both, may not be able to communicate with each other after you change the Windows Server 2008 R2 firewall connection settings.

Solution

To ensure that Arcserve Backup primary servers, member servers, and stand-alone servers running Windows Server 2008 R2 can browse, back up, and restore data that resides on Arcserve Backup agent servers running Windows Server 2008 R2, you must allow the Arcserve Backup applications on the backup servers and the agent servers to communicate using one of the following Windows network location types:

- Windows domain
- Windows private network
- Windows public network

For information about how to allow applications communicate using the above Windows network location types, see the Windows Server 2008 R2 documentation.

SRM PKI Alert is Enabled by Default

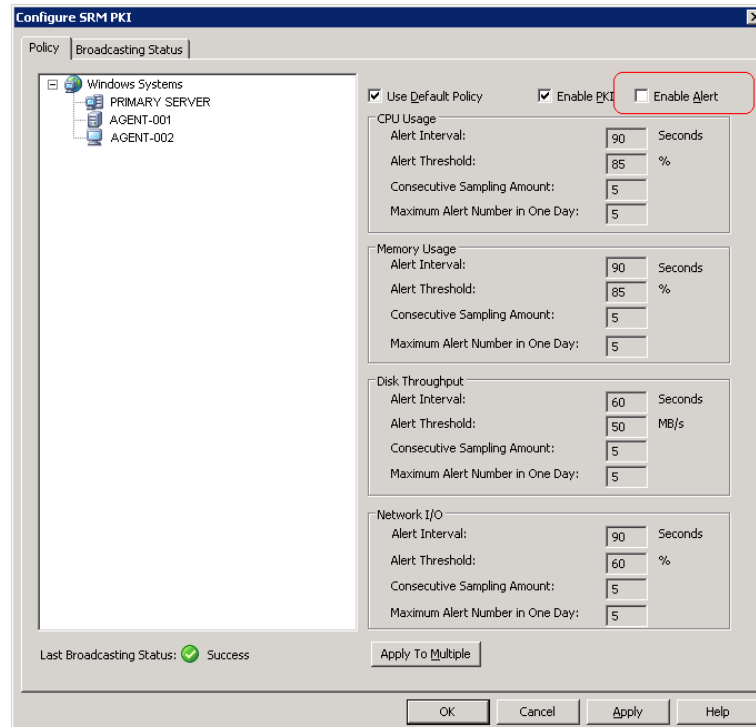
Valid on Windows platforms.

Symptom

Arcserve Backup includes an option named SRM PKI (performance key indicators) that lets you monitor the performance of agents running in your backup environment. Using the Central Agent Admin you can configure Arcserve Backup to generate alert messages when an agent's performance falls below your predefined performance key indicators.

Note: The alert messages appear in the Alert Manager based on how you configure the Alert Manager. For more information, see "Using the Alert Manager."

After you perform a new installation of Arcserve Backup, the Enable Alert option is disabled, by default, as illustrated by the following dialog:



However, when you upgrade from a previous version of Arcserve Backup, the Enable Alert option may be enabled, by default. Based on specific scenarios, the following solution describes the corrective actions required to disable the Enable Alert option.

Solutions:

To disable the Enable Alert option, do one of the following:

Solution 1:

You want to disable the Enable Alert option for all the agents that exist currently in your backup environment

1. Open the Central Agent Admin

Right-click the Windows Systems object and click Configure SRM PKI on the pop-up menu.

The Configure SRM PKI dialog opens.

2. From the agent tree (left pane) click an agent.

Clear the checkmark next to Enable Alert

Click Apply to Multiple.

The Apply to Multiple dialog opens.

3. On the Apply to Multiple dialog, select the individual agents, click Select All, or click Unselect All, and then click OK.

The Enable Alert option is disabled.

Solution 2:

You want to disable the Enable Alert option for all the agents that exist currently in your backup environment and the newly added agents

Note: The following steps describe the corrective actions for Arcserve Backup servers using Microsoft SQL Server 2005 for the Arcserve Backup database.

1. Open the Arcserve Backup database instance using the Microsoft SQL Server Management Console.
2. After you click Connect, open asdb, tables, and dbo.tbl_wcf_pkiAlertCft.
3. In the validalert field, change the value from 1 (enable alert) to 0 (disable alert).
4. Open the Windows Command Line.

Change the directory to the Microsoft SQL Server installation directory.

For example:

```
C:\Program Files\Microsoft SQL Server\90\Tools\Binn
```

5. Change the value of the validalert field by executing the following command:

```
SQLCMD.exe -S <Server_Name>\<ARCserve_Instance_Name> -d asdb
```

Look up the value of validalert by executing the following commands:

```
SELECT validalert FROM tbl_wcf_pkiAlertCfg;  
go
```

Change the value of validalert by executing the following commands:

```
UPDATE tbl_wcf_pkiAlertCfg SET validalert=0 where validalert=1;  
go
```

The Enable Alert option is disabled.

Job Queue Log Files Consume a Large Amount of Disk Space

Valid on Windows platforms.

Symptom

The Job Queue log files are consuming a large amount of disk space on the Arcserve Backup server.

Solution

Arcserve Backup stores debugging information about the Job Queue in one or more log files named JobQueue.log. The log files are stored in the following directory on the Arcserve Backup server:

```
%HOME%\LOG
```

Example:

```
C:\Program Files\CA\ARCserve Backup\LOG\JobQueue.log
```

As the log files reach a specified size (for example, 300 MB), Arcserve Backup renames the log file and creates a new Job Queue log file. By default, Arcserve Backup deletes the log files 31 days after they were generated.

Based on the following conditions, the Job Queue log files can consume a large amount of free disk space on the backup server:

- Quantity of agents and nodes you are protecting
- Quantity of scheduled jobs

Arcserve Backup lets you control the size and quantity of the Job Queue log files by doing the following:

1. From the Arcserve Backup server, open Windows Registry Editor.
2. To limit the size of the log files that Arcserve Backup generates, open the following registry key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA ARCServe  
Backup\Base\QueueSystem\DebugFileSize
```

Modify the value as required:

- **Default:** 100000000 bytes (approximately 100 MB)
 - **Range:** 10000000 bytes (approximately 10 MB) to 1000000000 bytes (approximately 1 GB)
 - **Recommended value:** 100000000 bytes (approximately 100 MB)
3. To limit the quantity of log files that Arcserve Backup retains, open the following registry key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA ARCServe  
Backup\Base\Task\Common\JobQueueMaxFiles
```

Modify the value as required:

- **Default:** 10
 - **Range:** 5 to 30
 - **Recommended value:** 10
4. (Optional) To decrease the level of debugging detail, create the following registry key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA ARCServe  
Backup\Base\QueueSystem\DebugLevel
```

The values for this key are as follows:

- **Range:** 1 to 5
- **Recommended value:** 3

Note: A higher value increases the level of debugging detail.

Specify a value of 1 or 2.

Be aware of the following:

- **Values too low**--Arcserve Backup may not be able to provide you with a sufficient level of information to debug problems.
- **Values too high**--The log files may consume too much free disk space on the backup server.

Unable to View all Job Logs for Rotation Backups after the Database is Pruned

Valid on Windows platforms.

Symptom

When you want to view all job logs relating to rotation and repeating jobs, you can view only job logs for the most recent execution for the job.

Solution

When Arcserve Backup database pruning jobs run, the pruning job deletes all job logs for a particular job with the exception of the most recent log file. This behavior prevents you from having the capability to view all of the log files relating to repeating jobs and jobs defined with rotation schedules. To configure database pruning jobs to retain all of the job logs associated with repeating and rotation jobs, complete the following steps:

1. From the backup server, open the following key using the Windows Registry Editor:

x86 Operating Systems:

HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA ARCServe Backup\Base\Database

x64 Operating Systems:

HKEY_LOCAL_MACHINE\SOFTWAREWow6432Node\ComputerAssociates\CA ARCServe Backup\Base\Database

2. Create the following keyword:

DoNotPruneRecentJobLogs

3. Define the DWORD (32-bit) value for the keyword as 1.

Be aware of the following:

- After you create this keyword, Arcserve Backup prunes the job logs based on the value that you define for the Prune activity logs older than option for the Database Engine. For more information, see [Database Engine Configuration](#) (see page 514).
- This behavior applies to only job logs. Arcserve Backup does not prune the Activity Log for recent executions of jobs.
- You can view the job logs in the Job Queue, the Job History, and the Database Manager.

Appendix B: Using Arcserve Backup in a Cluster-aware Environment

The following topics provide an overview of Arcserve Backup cluster support, which allows you to back up and recover data in a cluster environment. In addition, these topics also provide information about configuring Arcserve Backup as a cluster-aware backup server with high availability capabilities.

Installation of Arcserve Backup in a cluster environment with job failover capability is supported for the following cluster platforms:

- Microsoft Cluster Server (MSCS) in x86/x64
- NEC CLUSTERPRO/ExpressCluster X 2.1 for Windows
- NEC CLUSTERPRO/ExpressCluster X 3.0 for Windows
- NEC CLUSTERPRO/ExpressCluster X 3.1 for Windows

This section contains the following topics:

[Install Arcserve Backup into a Cluster-aware Environment](#) (see page 905)

[Cluster Overview](#) (see page 905)

[Protecting Your Cluster with Arcserve Backup](#) (see page 913)

Install Arcserve Backup into a Cluster-aware Environment

For information about how to install and upgrade Arcserve Backup into a cluster-aware environment, see the *Implementation Guide*.

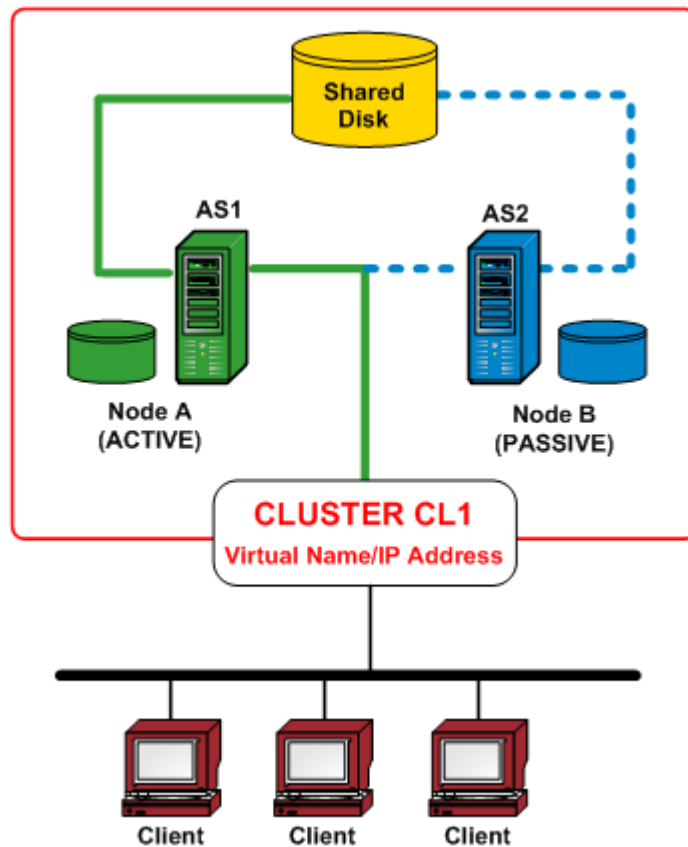
Cluster Overview

A computer cluster is a group of connected computers that work together closely so that in many respects they can be viewed as though they are a single computer. Clusters can be categorized in two types: HA (High Availability) and High Performance. Within HA clusters, there are two working modes: active/active or active/passive. Currently, Arcserve Backup can only be deployed in an active/passive HA mode.

The primary function of a cluster occurs when one server (or node) in a cluster fails or is taken offline. In a cluster environment, the other node in the cluster will then take over the failed server's operations. Arcserve Managers using server resources experience little or no interruption of their work because the resource functions move transparently from the active node to the failover node.


The servers within a cluster environment are not only connected physically by cables, but also programmatically through clustering software. This connection allows clustered servers to take advantage of features (such as fault tolerance and load balancing) that are unavailable to stand-alone server nodes. Clustered servers can also share disk drives that contain important information, such as a clustered database.

For example, assume node A and node B form a clustered Arcserve Backup HA server. The Arcserve Backup cluster server will work only in the "active/passive" mode, and as a result, only one Arcserve Backup instance is running at the same time. In this environment, Arcserve Managers could connect to the Arcserve Backup server AS1 or Arcserve Backup server AS2 without knowing which node is active and currently hosting their server. The virtual server name and IP address ensures that the server location is transparent to Arcserve Backup applications. To the Arcserve Manager, it appears that the Arcserve Backup server is running on a virtual server called CL1.



When one of the software or hardware resources fails or is shut down, a failover occurs. Resources (for example: applications, disks, or an IP address) migrate from the failed active node to the passive node. The passive node takes over the Arcserve Backup server resource group and now provides service.

If node A fails, node B automatically assumes the role of the active node. To an Arcserve Manager, it is exactly as if node A were turned off and immediately turned back on again. The location of the active node (A or B) in the Cluster (CL1) is transparent to Arcserve Backup.



How Failover Works

Failover is the process of having cluster resources migrate (or transfer) from an unavailable node to an available node. Failover is automatically initiated when a failure is detected on one of the cluster nodes. The cluster monitors resources to determine when a failure has occurred and then takes action to recover from the failure by moving the clustered resource(s) to another node in the cluster.

In a Arcserve Backup HA cluster environment, Arcserve Backup is installed in each cluster node, but only one instance will be running. In this cluster, the active node will automatically take control of the backup resources and is referred to as the backup server. Other instances of Arcserve Backup that are hosted in a passive node are referred as the standby (or failover) server and the cluster system will only activate one of them in case of failover. If the active node fails, then all backup resources will migrate to a passive node, which then becomes the new active node. The new active node begins to function as the backup server, and continues the original backup operations and maintains all previous job scheduling and media management services.

Arcserve Backup provides the following types of failover protection:

- **Planned Failovers**--Planned failovers occur when it is necessary to perform maintenance on the active node within a cluster and you want Arcserve Backup to migrate the cluster resources from that active node to a passive node within the cluster. Examples for planned failovers are system maintenance, disaster recovery tests, and training. A planned failover can only be executed when no jobs are running and no other Arcserve Backup-related services (such as media operations, reporting, etc.) are occurring on both the primary and member backup servers.
- **Unplanned Failovers**--Unplanned failover can occur because of hardware or software failures. When the active node in a cluster fails, jobs are dispersed from the failed server and critical data (such job information) is saved into a shared disk. When failover occurs, the cluster system will move the shared disk into a passive node and activate the Arcserve Backup instance in that node. After the Arcserve Backup services are resumed in the failover server, any failed jobs from the previous server are rerun in a new active cluster node. If checkpoint information was created by the job before failover occurred, the restarted job will resume from the checkpoint.

Resource Group

A cluster resource is any physical or logical component that can be physically shared between multiple cluster nodes, but can only be hosted (owned) by one active node at a time. The virtual IP address, virtual computer name, shared disk, and even the applications are considered cluster resources. A cluster system allows you to categorize these resources as a "group" for a specific functionality purposes. These resource groups can be treated as a "container" of resources. A cluster resource group is a logical unit for application deployment, which means that a cluster-aware application must be installed into a group and bind itself with the resources that are associated with that group. The resource group is the minimum unit for failover purposes.

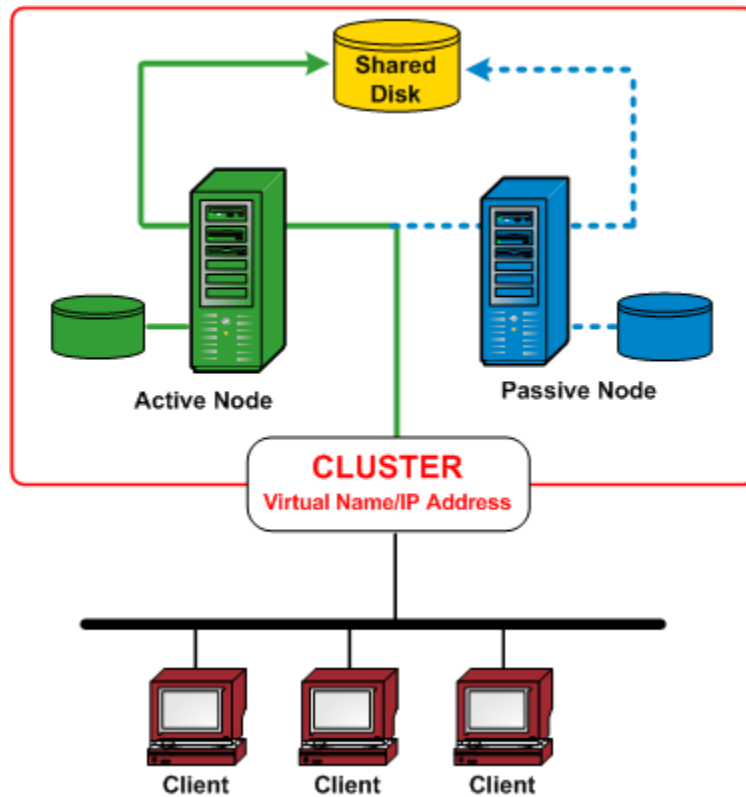
Virtual Name and Virtual IP Address

The virtual server name is independent of the name of the physical server on which the virtual server runs and can migrate from server to server. In a cluster environment, the active node will always use the cluster virtual name and IP address to provide service instead of the physical hostname and IP address. Through the use of clusters, virtual servers are created so that when another server takes its place, the services can still be available. The virtual name and IP address are linked with Arcserve Backup. Similarly, other cluster-aware applications (SQL/Exchange Cluster) often create a dedicated virtual name and IP address for high-availability purposes during their installation.

Unlike a physical server, a virtual server is not associated with a specific computer, and it can fail over from one server to another. If the server that is hosting the virtual server fails, clients can still access its resources by using the same virtual server name, but they will be redirected to a different server in the cluster.


Shared Disks

Shared disks provide shared location for cluster-aware applications to save data. Shared disks allow cluster-aware applications, which might run on different nodes because of failover, to gain access to a logical volume in a consistent way, as if they are local at each of the nodes. Each virtual shared disk corresponds to a logical volume that is actually local at one of the nodes, which is called the server or primary node. Each node within the cluster must have access to a shared disk to operate within the cluster. The cluster system is configured so that only the active node can access the shared disk at any time.



Mirrored Disks

Mirrored disks provide a shared location for cluster-aware applications to save data. Mirrored disks (applicable to NEC clusters only) are separate disk devices that are physically attached to their host separately, but work like one single device logically. Mirrored disks contain an exact duplicate of the disk that it mirrors. Data is stored twice by writing to both the local disk and its remote mirrored disk. If a disk fails, data does not have to be rebuilt and can be easily recovered by copying it from the mirrored disk to the replacement disk. It is recommended that mirrored disks reside on different devices so that a single-point disk failure cannot damage both copies of the data. The main disadvantage of mirrored disks is that the effective storage capacity is only half of the total disk capacity because all data gets written twice. The cluster system is configured so that only the active node can access the mirrored volume and sync-up the data between different physical disks.



Quorum Disks

In addition to the resource groups created for each clustered application, a cluster always has a resource group to represent the quorum of the cluster. This resource group, by default named Cluster Group, is created when the cluster is created. In a shared disk quorum, the disk containing the quorum resource is called the quorum disk, and it must be a member of the default Cluster Group. A quorum disk is used to store cluster configuration database checkpoints and log files which help manage the cluster as well as maintain consistency. The quorum resource is used to decide which nodes of the cluster are supposed to form the cluster. Because the cluster configuration is kept on a quorum disk resource, all nodes in the cluster must be able to access and communicate with the node that owns it.

Note: Quorum Disks apply to only Microsoft Cluster Server (MSCS) environments.

Arcserve Backup HA Server to Support of Job Failover

Clustered Arcserve Backup servers provide service through virtual name and support backup job failover capability. When the active Arcserve Backup server in a cluster fails, these backup jobs are dispersed from the failed server to other Arcserve Backup servers in the cluster. After Arcserve Backup services are resumed in another cluster node, any failed jobs from the previous server are rerun in a new cluster node.

Arcserve Backup HA server supports two types of failovers; planned failovers and unplanned failovers.

■ Planned Failovers

Planned failovers occur when it is necessary to perform maintenance on the active node within a cluster and you want Arcserve Backup to migrate the cluster resources from that active node to a passive node within the cluster. Examples for planned failovers are system maintenance, disaster recovery tests, and training.

When a planned failover occurs, Arcserve Backup recovers in another node with all scheduled jobs retained.

■ Unplanned Failovers

Unplanned failover can occur because of hardware or software failures. When unplanned failover occurs, Arcserve Backup recovers in another node, picks up the failed job from the Arcserve Backup job queue, and resumes the job from the point where it failed. If a failover occurs, the job resume is based on a checkpoint mechanism as follows:

- For a local backup job, the job will resume at the volume level after a failover.
For example, if a backup job involves two volumes: C and D, and a failover occurs when the backup for volume C was finished and the backup for volume D was ongoing. After the failover, the backup job will restart and skip the backup for volume C and continue to backup volume D.
- For a remote backup job, the job will resume at the host level.
For example, if a backup job involves Host1 and Host2, and a failover occurs when the backup for Host1 is finished, but the backup for Host2 is not. After the failover, the backup job will restart and skip the backup for Host1 and continue to backup Host2 (in this case, the backup for Host2 does not skip any volumes that might be backed up before failover).

Jobs running on other backup servers instead of the HA server of the domain will rarely be impacted by failover. For example, when the Primary server is HA server, and it fails over, the jobs running on member servers are not impacted except in one situation. If you are using a HA primary server, the jobs that run on the member servers may fail when an unplanned failover occurs on the HA primary server. (The failures only occur when the jobs on the member servers are finishing when the failover occurs)


Note: If you are using Arcserve Backup Agents to backup the active node of the cluster or the virtual node, and an unplanned failover occurs (the active node is down), the job would become incomplete. To ensure that these nodes could be backed up after failover, you should configure the jobs to create makeup jobs.

Protecting Your Cluster with Arcserve Backup

For mission-critical applications deployed into a cluster environment, the data is the most valuable investment and protection of this data is essential. A cluster environment always involves multiple physical nodes, a virtual name/IP address, and cluster-specific applications, all of which bring additional complexities for the backup and restore application. To address these complexities, Arcserve Backup provides multiple backup and restore capabilities for servers operating in a cluster environment.

Note: Arcserve Backup supports cluster environments for Microsoft Cluster Server (MSCS) and NEC Cluster Server (CLUSTERPRO/ExpressCluster).

The following diagram shows a typical active/passive cluster environment. The active node in this cluster is associated with two names and IP addresses: one for the physical name of the machine and the other for the virtual name created by the cluster itself or the cluster-aware application. The passive node is associated with only one name, the physical name of the machine. To fully protect the cluster, you need to install the Arcserve Backup agent into both of these physical nodes. In each of these instances, depending on the protected target, Arcserve Backup will be deployed to protect your cluster and back up your data using either the physical node or the virtual node.



■ **Protect data using physical node**

To protect the system state of each cluster node and the local application data, you need to schedule the backup job based on the physical name/IP address of the machine. For the active node (1), you can back up all attached disks, including the local disks as well as shared disks. For a passive node (2), you can only back up the local disks. However, it is not a best practice to back up a shared disk based only on physical name. In a cluster environment, the role of each node (active and passive) may change dynamically due to a failover condition. If you specify the physical name of the failed node, the backup will fail and the data that is located on the shared disk will not be backed up.

- **Protect data using virtual node**

In a cluster-aware application (SQL Server cluster or MS Exchange cluster), all data is saved on a shared disk to provide HA capability. To back up this data, the Arcserve Backup agent (installed on each physical node) will archive the data into the shared disk via the virtual name and IP address of the cluster (3). Under normal conditions, Arcserve Backup will back up data from the shared disk using the virtual name and IP address of the cluster as the source rather than the physical name and IP address of the active node. The advantage of doing this is if the active node fails or is shut down, the failover mechanism of cluster will make the passive node become the new active node, and Arcserve Backup will automatically continue to perform backups from the shared disk. As a result, you can schedule rotation backup jobs to protect your data located in the shared disk regardless which cluster node is active.

Note: To back up any application specified data (for instance, a SQL Server database), you should deploy the corresponding Arcserve Backup agent and perform the backup using the virtual name associated with this cluster-aware application.

MSCS Protection

Microsoft Cluster Server (MSCS) software provides a clustering technology that keeps server-based applications highly available, regardless of individual component failures. For MSCS, there are two basic types of targets that need to be protected by backup: cluster self-protection, in which the cluster itself is protected (metadata and configuration information) and clustered-application protection.

How Arcserve Backup Integrates with MSCS

Arcserve Backup is a fault-tolerant application, capable of supporting failover in cluster environments. Arcserve Backup protects cluster nodes by backing up and restoring cluster-specific resources such as shared disks, quorum resources, disk signatures, and cluster registry hives. Microsoft Cluster Server (MSCS) allows multiple Windows based servers to connect with one another so that they appear, to network clients, to be a single, highly available system.

With the MSCS support provided by Arcserve Backup, you can:

- Back up and restore MSCS nodes
- Run on, and take advantage of, MSCS high availability features such as:
 - Job failover from one Arcserve Backup node in a cluster to another node
 - High availability through automatic failover of Arcserve Backup services from one node in a cluster to another node
 - Install Arcserve Backup to an Active/Passive cluster as your SAN primary server and allow continuation of distributed server backups upon failover
 - Manageability through standard cluster management tools
- Provide disaster protection for MSCS nodes using the Arcserve Backup Disaster Recovery Option. For more information, see the *Disaster Recovery Option Guide*.
- Backup and restore applications, such as MS SQL Server and MS Exchange Server, installed on MSCS clusters using the Arcserve Backup agents. For more information on available agents, see the *Implementation Guide*.

MSCS Cluster Self-Protection

For MSCS, all cluster configuration information resides in a Cluster Database. The Cluster Database is located in the Windows registry on each cluster node and contains information about all physical and logical elements in a cluster, including cluster objects, their properties, and cluster configuration data. The Cluster Database contains the cluster state data that is replicated among nodes to ensure that all nodes in the cluster have a consistent configuration. The Cluster Database registry is located in %WINDIR%\CLUSTER\CLUSDB.

The Cluster Database is part of the Windows System State. When the System State is selected for backup, the Cluster Database is automatically included in this backup. Therefore, the Cluster Database is included in the system state backup only if the node is a part of a cluster and the cluster service is running on that node. If the cluster service is not running, the Cluster Database is not backed up.

To protect a cluster node itself and reduce the potential risk caused by accidental node failure, you should back up the following data using the physical name of the nodes:

- all data on the local disks contained in the Windows boot/system partitions
- system state data

During restore operations, you first need to determine the severity of the problem. If you cannot boot the node at all, see *Recovering Clusters in the Arcserve Backup Disaster Recovery Option Guide*. If you can boot the operating system and only the cluster database is damaged, you will not be able to selectively restore the Cluster Database as a single entity, it must be restored as part of a System State session restore.


Note: To back up and restore the Cluster Database, it is sufficient to back up and restore the Windows System State with the cluster service running. When a cluster node is in the Directory Service Restore mode, the Logon properties of the Cluster Service User Account must be set as Administrator to ensure that Arcserve Backup can be accessed while in Windows Safe Mode.

MSCS Cluster Application Protection

Arcserve Backup offers fast and intelligent backup and restore operations of applications such as Microsoft SQL Server and Microsoft Exchange Server installed on a cluster. For the most current list of available agents, see the readme file or access the Arcserve web site. For information on backing up and restoring applications installed on a cluster, see the corresponding Arcserve Backup agent guide. For example, for information on backing up and restoring a Microsoft SQL Server, see the Agent for Microsoft SQL Server Guide.

Stop MSCS From Monitoring Arcserve Backup Services on Windows Server 2003 Systems

When a Arcserve Backup server is configured as cluster-aware, all critical Arcserve Backup services will be monitored by MSCS. If some service fails, MSCS will try to restart it or trigger a failover if the restart attempt fails. This means that you can no longer stop a service by using the Arcserve Backup Server Administrator. If you attempt to stop a Arcserve Backup service, the following message opens:



However, in some situations, you may want to stop Arcserve Backup services. For example, you need to stop various services so that you can perform hardware maintenance.

Note: For information about how to perform this task on Windows Server 2008 systems, see [Stop MSCS From Monitoring Arcserve Backup Services on Windows Server 2008 Systems](#) (see page 919).

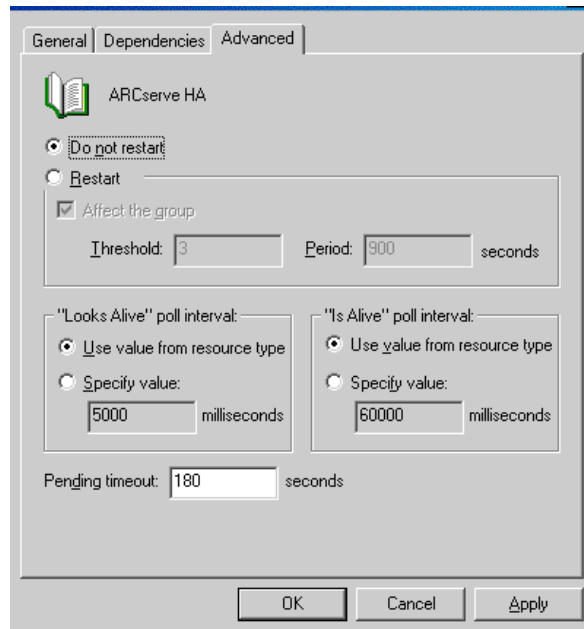
To stop MSCS from monitoring Arcserve Backup services on Windows Server 2003 systems

1. Open Microsoft Cluster Administrator.

Note: Cluster Administrator is a utility provided by Microsoft and is installed on servers that have MSCS installed. From the Cluster Administrator, you perform most of the configuration and management tasks associated with clusters.

2. Select the group that the Arcserve server is deployed in, and locate the applicable Arcserve resource. Right-click the Arcserve resource and click Properties on the pop-up menu.

The ARCserve HA Properties dialog appears.



3. Click the Advanced tab.

From the Advanced tab, click the Do not restart option, and then click OK.

With this option specified, MSCS will not try to restart the Arcserve Backup services automatically or initiate the failover process.

Note: All Arcserve Backup services are controlled by the Arcserve HA resource. However, the Tape Engine service and ASDB service are also controlled by additional resources. See the following table to identify the resources that need to be changed for each Arcserve Backup service. For each of the applicable resources, you need to set the Advanced property to Do not restart.


Service Name	Controlling Resources
Tape Engine	Arcserve HA Arcserve Registry

Service Name	Controlling Resources
ASDB (only for SQL2008 Express)	Arcserve ASDB Arcserve HA Arcserve Registry
Others (DB Engine, Job Engine, ...)	Arcserve HA

- Using the Windows Service manager, stop the applicable Arcserve Backup service to allow you to perform the necessary maintenance.
- When you have completed the maintenance, restore all settings.

Stop MSCS From Monitoring Arcserve Backup Services on Windows Server 2008 Systems

When a Arcserve Backup server is configured as cluster-aware, all critical Arcserve Backup services will be monitored by MSCS. If some service fails, MSCS will try to restart it or trigger a failover if the restart attempt fails. This means that you can no longer stop a service by using the Arcserve Backup Server Administrator. If you attempt to stop a Arcserve Backup service, the following message opens:



In some situations, you may want to stop Arcserve Backup services. For example, you may need to stop various services so that you can perform hardware maintenance.

Note: For information about how to perform this task on Windows Server 2003 systems, see [Stop MSCS From Monitoring Arcserve Backup Services on Windows Server 2003 Systems](#) (see page 917).

To stop MSCS from monitoring Arcserve Backup services on Windows Server 2008 systems

- Access the Failover Cluster Manager.

The Cluster Administrator dialog appears.

Note: Cluster Administrator is a utility provided by Microsoft and is installed on servers that have MSCS installed. From the Cluster Administrator, you perform most of the configuration and management tasks associated with clusters.

- Select the Arcserve Backup server where the Arcserve service is deployed in, and locate the applicable Arcserve resource. Right-click the ARCserve service or application and click Properties on the pop-up menu.

The ARCserve HA Properties dialog opens.

3. Click the Policies tab.

Click the If resource fails, do not restart option, and then click OK.

With this option specified, MSCS will not try to restart the Arcserve Backup services automatically or initiate the failover process.


Note: All Arcserve Backup services are controlled by the Arcserve HA resource. However, the Tape Engine service and ASDB service are also controlled by additional resources. See the following table to identify the resources that need to be changed for each Arcserve Backup service. For each of the applicable resources, you need to set the Advanced property to Do not restart.

Service Name	Controlling Resources
Tape Engine	Arcserve HA Arcserve Registry
ASDB (only for SQL2008 Express)	Arcserve ASDB Arcserve HA Arcserve Registry
Others (DB Engine, Job Engine, ...)	Arcserve HA

4. Using the Windows Service manager, stop the applicable Arcserve Backup service to allow you to perform the necessary maintenance.
5. When you have completed the maintenance, restore all settings.

Stop MSCS From Monitoring Arcserve Backup Services on Windows Server 2012 Systems

When a Arcserve Backup server is configured as cluster-aware, all critical Arcserve Backup services will be monitored by MSCS. If some service fails, MSCS will try to restart it or trigger a failover if the restart attempt fails. This means that you can no longer stop a service by using the Arcserve Backup Server Administrator. If you attempt to stop a Arcserve Backup service, the following message opens:



In some situations, you may want to stop Arcserve Backup services. For example, you may need to stop various services so that you can perform hardware maintenance.

Note: For information about how to perform this task on Windows Server 2003 systems, see [Stop MSCS From Monitoring Arcserve Backup Services on Windows Server 2003 Systems](#) (see page 917).

Follow these steps:

1. Access the Failover Cluster Manager.

The Cluster Administrator dialog appears.

Note: Cluster Administrator is a utility provided by Microsoft and is installed on servers that have MSCS installed. From the Cluster Administrator, you perform most of the configuration and management tasks associated with clusters.

2. Select the Arcserve Backup server where the Arcserve cluster role is deployed in, and locate the applicable Arcserve resource. Right-click the ARCserve HA resource and click Properties on the pop-up menu.

The ARCserve HA Properties dialog opens.

3. Click the Policies tab.

Click the If resource fails, do not restart option, and then click OK.

With this option specified, MSCS will not try to restart the Arcserve Backup services automatically or initiate the failover process.

Note: All Arcserve Backup services are controlled by the Arcserve HA resource. However, the Tape Engine service and ASDB service are also controlled by additional resources. See the following table to identify the resources that need to be changed for each Arcserve Backup service. For each of the applicable resources, you need to set the Advanced property to Do not restart.

Service Name	Controlling Resources
Tape Engine	Arcserve HA Arcserve Registry
ASDB (only for SQL2008 Express)	Arcserve ASDB Arcserve HA Arcserve Registry
Others (DB Engine, Job Engine, ...)	Arcserve HA

4. Using the Windows Service manager, stop the applicable Arcserve Backup service to allow you to perform the necessary maintenance.
5. When you have completed the maintenance, restore all settings.

Rebuild Cluster Resources Manually

In most cases, the installation process automatically creates the necessary HA cluster resources without user input. However, there may be cases where you need to create these cluster resources manually.

Before you manually create the new resources, you should stop and delete all existing cluster resources from the group where Arcserve Backup is deployed. For more information on deleting cluster resources, see [Delete Cluster Resources](#) (see page 923).

To rebuild cluster resources manually

1. Open a command console and change current directory to %bab_home% (where, %bab_home% represents the actual Arcserve Backup install path).
2. Run the "babha.exe -postsetup" utility to set up new Arcserve cluster resources.

When a cluster-aware installation is successfully finished, a Post Setup pop-up screen appears with an option to create HA resources.

3. Select the Create HA resources for MSCS option and click OK to create the new cluster resources.

Note: You should specify this option only after you completed the Arcserve Backup installation on the last node in the cluster.

The following new Arcserve cluster resources are created:

- **Windows Server 2003**--On Windows Server 2003 systems, the following cluster resources are created:
 - Arcserve HA
 - Arcserve ASDB
 - Arcserve Registry
 - Arcserve Share
- **Windows Server 2008**--On Windows Server 2008 systems, the following cluster resources are created:
 - Arcserve HA
 - Arcserve ASDB
 - Arcserve Registry
- **Windows Server 2012**--On Windows Server 2012 systems, the following cluster resources are created:
 - Arcserve HA
 - Arcserve ASDB
 - Arcserve Registry

Delete Arcserve Backup Cluster Resources

Before you can create new cluster resources, you must delete all existing cluster resources from the group where Arcserve Backup is deployed.

The available MSCS cluster resources are as follows:

- Windows Server 2003 Systems
 - Arcserve High Availability
 - Arcserve database
 - Arcserve Registry
 - Arcserve Share
- Windows Server 2008 Systems
 - Arcserve High Availability
 - Arcserve database
 - Arcserve Registry
- Windows Server 2012 Systems
 - Arcserve High Availability
 - Arcserve database
 - Arcserve Registry

Follow these steps:

1. Do one of the following:

Windows Server 2003 systems:

- Open the Cluster Administrator.
- Right-click the group name and click Take offline on the pop-up menu.

Windows Server 2008 systems:

- Open the Failover Cluster Manager.
- Right-click the service or application name and click Take this service or application offline.

The state of the Arcserve cluster resources is changed from Online to Offline.

Windows Server 2012 systems:

- Open the Failover Cluster Manager.
- Right-click the cluster role and stop role.

The state of the Arcserve cluster resources is changed from Online to Offline.

2. Based on the Windows Server system you use, perform one of the following options:
 - **For Windows Server 2003 or 2008 systems:**

Select the Arcserve Group that the Arcserve server is deployed in.
 - **For Windows Server 2012 systems:**

Select the Arcserve cluster role that the Arcserve server is deployed in.

Locate the corresponding Arcserve cluster resources.

Right-click each Arcserve cluster resource and click Delete on the pop-up menu.

The selected Arcserve cluster resources are deleted.

Manage Arcserve Backup Cluster Servers in an MSCS Cluster

The Server Configuration Wizard lets you perform various management tasks to specify how Arcserve Backup servers function in a cluster environment. In a cluster environment, these management tasks can only be made on the active node and must also be made for all nodes within the cluster. These management tasks include the following:

- Changing the database
- Promoting a member server to a primary server
- Demoting a primary server to a member server.

To manage the Arcserve Backup Cluster Servers in an MSCS cluster

1. Delete all cluster resources. For more information, see [Delete Arcserve Backup Cluster Resources](#) (see page 923).

All Arcserve Backup cluster resources are deleted.
2. From the Arcserve Backup home directory, run the cstart.bat utility to start all Arcserve Backup services.

3. From the Start menu, access the Server Configuration Wizard to run the ARCServeCfg.exe utility for the active node and make the necessary change. Do not check the "Last Cluster Node" checkbox on the last screen of the Server Configuration Wizard.

- For more information about changing the database, see [Specify a Arcserve Backup Database Application](#) (see page 662).

Note: Local SQL Server is not supported when NEC CLUSTERPRO/ExpressCluster is used to make Arcserve Backup highly available.

- For more information about promoting a member server to a primary server, see [Promote a Member Server to a Primary Server](#) (see page 565).
- For more information about demoting a primary server to a member server, see [Demote a Primary Server to a Member Server](#) (see page 568).

Note: When this utility is run on the first node within a cluster, it will run in the normal mode.

The first "active" cluster node is configured for the new property and a new arcservecfg.ICF configuration file is created.

4. From the Arcserve home directory, run the cstop.bat utility to stop all Arcserve Backup services.

All Arcserve Backup services are stopped.

5. Do one of the following:

Windows Server 2003 systems:

- Open the Cluster Administrator.
- Right-click the group name and select Move Group on the pop-up menu.

Windows Server 2008 systems:

- Open the Failover Cluster Manager.
- Right-click the service or application name and click Move this service or application to another node on the pop-up menu.

Windows Server 2012 systems:

- Open the Failover Cluster Manager.
- Right-click the ARCServe cluster role name and click Move this role to another node on the pop-up menu.

The status of the original node will be changed to "passive" and the status of the next node within the cluster will be changed to "active".

6. From the Start menu, access the Server Configuration Wizard to run the ARCServeCfg.exe utility for the new active node and make the necessary change.

Note: When this utility is run again on any subsequent nodes in the same cluster, it will detect the existence of the arcservecfg.ICF configuration file and automatically run the utility in the cluster mode.

The next "active" cluster node is configured for the new property.

7. Repeat steps 5 and 6 for all remaining nodes in the cluster. When you are performing this configuration procedure on the last node in the cluster, check the "Last Node" checkbox on the last screen of the Server Configuration Wizard.

All nodes in the cluster are configured for the new property.

8. Do one of the following:

Windows Server 2003 systems:

- Open the Cluster Administrator.
- Right-click the group name and select Move Group on the pop-up menu .

Windows Server 2008 systems:

- Open the Failover Cluster Manager.
- Right-click the service or application name and click Move this service or application to another node on the pop-up menu.

Windows Server 2012 systems:

- Open the Failover Cluster Manager.
- Right-click the Arcserve cluster role name and click Move this role to another node on the pop-up menu.

The status of the last node will be changed to "passive" and the status of the original node within the cluster will be changed back to "active".

9. Create all Arcserve Backup cluster resources manually. For more information, see [Rebuild Cluster Resources Manually](#) (see page 922).

The new Arcserve cluster resources are created.

10. Do one of the following:

- **Windows Server 2003 systems**--From the Cluster Administrator, right-click the group and click Bring Online on the pop-up menu.
- **Windows Server 2008 systems**--From the Failover Cluster Manager, right-click the service or application name and click Bring this service or application online on the pop-up menu.
- **Windows Server 2012 systems**--From the Failover Cluster Manager, right-click the ARCServe cluster role name and click Start role on the pop-up menu.

The state of the new Arcserve cluster resources is changed from Offline to Online.

Change the Arcserve Backup Domain in a MSCS Cluster

In a MSCS cluster environment, you can move a member server to a different Arcserve Backup domain. Changes to the domain in a cluster environment can only be made on an active node and must be changed for all nodes within the cluster.

To change the Arcserve Backup domain in a MSCS cluster

1. Delete all cluster resources. For more information, see [Delete Arcserve Backup Cluster Resources](#) (see page 923).
All Arcserve Backup cluster resources are deleted.
2. From the Arcserve Backup home directory, run the cstart.bat utility to start all Arcserve Backup services.
All Arcserve Backup services are started.
3. From the Start menu, access the Server Configuration Wizard to run the ARCServeCfg.exe utility for the active node and specify the new Arcserve Backup domain. For more information about changing a domain, see [Move a Member Server to a Different Arcserve Backup Domain](#) (see page 572).
The first "active" cluster node is configured for the new domain.
4. From the Arcserve home directory, run the cstop.bat utility to stop all Arcserve Backup services.
All Arcserve Backup services are stopped.
5. From the Cluster Administrator, right-click on the group name and from the pop-up menu, select Move Group to change the active node.
The status of the original node will be changed to "passive" and the status of the next node within the cluster will be changed to "active".
6. From the Arcserve Backup home directory, run the cstart.bat utility to start all Arcserve Backup services.
All Arcserve Backup services are started.
7. From the Arcserve home directory, run the cstop.bat utility to stop all Arcserve Backup services.
All Arcserve Backup services are stopped.
8. Repeat steps 5 through 7 for all remaining nodes in the cluster.
All nodes in the cluster have been changed to the new domain.
9. From the Cluster Administrator, right-click on the group name and from the pop-up menu, select Move Group to change the active node back to the original node.
The status of the last node will be changed to "passive" and the status of the original node within the cluster will be changed back to "active".

10. Create all Arcserve Backup cluster resources manually. For more information, see [Rebuild Cluster Resources Manually](#) (see page 922).

Note: You should create Cluster resources based on new Arcserve database type. The new Arcserve cluster resources are created.

11. From the Cluster Administrator, right-click on the group name and from the pop-up menu, select Bring Online.

The state of the new Arcserve cluster resources is changed from Offline to Online.

NEC CLUSTERPRO/ExpressCluster Protection

NEC CLUSTERPRO/ExpressCluster is a high-availability clustering solution that provides fast recovery and high reliability to maximize critical applications and data availability. NEC clusters offer integrated application and data protection that enables fast, easy recovery and continuity of critical systems.

NEC clusters allow multiple Windows-based servers to connect with one another so that they appear to network clients to be a single, highly available system. Arcserve Backup supports NEC CLUSTERPRO/ExpressCluster X 2.1 for Windows, NEC CLUSTERPRO/ExpressCluster X 3.0 for Windows and NEC CLUSTERPRO/ExpressCluster X 3.1 for Windows.

For NEC clusters, there are two basic types of targets that need to be protected by backup: cluster self-protection, in which the cluster itself is protected (metadata and configuration information) and clustered-application protection.

How Arcserve Backup Integrates with NEC CLUSTERPRO

Arcserve Backup is a fault-tolerant application, capable of handling failover and providing backup and restore capabilities for data residing in cluster environments.

NEC CLUSTERPRO/ExpressCluster allows multiple Windows-based servers to connect with one another so that they appear to network clients to be a single, highly available system. Arcserve Backup supports NEC CLUSTERPRO/ExpressCluster X 2.1 for Windows, NEC CLUSTERPRO/ExpressCluster X 3.0 for Windows and NEC CLUSTERPRO/ExpressCluster X 3.1 for Windows. Similar to MSCS, we need protect the cluster itself and these clustered applications.

Arcserve Backup support for NEC CLUSTERPRO/ExpressCluster offers the following advantages:

- Ability to run on NEC CLUSTERPRO/ExpressCluster and take advantage of high availability features such as:
 - Automatic failover of Arcserve Backup services from one node in a cluster to another node.
 - Ability to fail jobs over from one Arcserve Backup node in a cluster to another node when Arcserve Backup failover occurs.
 - Ability to restart jobs after failover.
 - Ability to install Arcserve Backup on an Active/Passive cluster as the SAN primary server to allow the continuation of distributed server backup operations after failover.
 - Ability to use NEC cluster management tools.
- Data backup and restore functionality for NEC cluster nodes.
- Disaster protection for NEC CLUSTERPRO/ExpressCluster nodes through the Disaster Recovery Option. For more information, see the *Disaster Recovery Option Guide*.

NEC Cluster Server Self-Protection

For NEC CLUSTERPRO/ExpressCluster, all cluster configuration information resides in a file system as regular files.

To protect a cluster node itself and reduce the potential risk caused by accidental node failure, you should back up the following data using the physical name of the nodes:

- all data on the local disks contained in the Windows boot/system partitions
- system state data

During restore operations, you first need to determine the severity of the problem. If you cannot boot the node at all, see Recovering NEC Clusters in the *Arcserve Backup Disaster Recovery Option Guide*. If you can boot the operating system and only the NEC cluster files are damaged, refer to the applicable NEC CLUSTERPRO/ExpressCluster document to manually restore these configuration files that are related to NEC clusters.


NEC CLUSTERPRO/ExpressCluster Application Protection

For NEC CLUSTERPRO/ExpressCluster, there are few applications that are native cluster-aware. Native cluster-aware refers to some applications that are aware they will run in cluster environment to support HA and take special considerations in design.

For NEC clusters, few applications are designed to be cluster-aware and many do not recognize the NEC virtual name/IP address. However, for some of the more popular applications, NEC provides specific documentation for configuring these applications as “cluster-aware” and deploying them so that you can perform backup and restore jobs. Refer to NEC website for a list of all supported applications and detailed information on how to configure them as cluster-aware. If you have deployed one of these supported applications, refer to the corresponding NEC documentation for details about how to perform a backup and restore with cluster support.

Stop HA Service Monitoring by NEC CLUSTERPRO/ExpressCluster

When a Arcserve Backup server is configured as cluster-aware, all critical Arcserve Backup services will be monitored by NEC CLUSTERPRO/ExpressCluster. If some service fails, NEC CLUSTERPRO/ExpressCluster will try to restart it or trigger a failover if the restart attempt fails. This means that you can no longer stop a service by using the Arcserve Backup Server Administrator. If you attempt to stop a Arcserve Backup service, you will see a pop-up message:



However, in some situations, you may want to stop some Arcserve Backup service. For example, you may want to stop the Tape Engine so that you can perform hardware maintenance.


Note: This section contains graphics that correspond with NEC CLUSTERPRO/ExpressCluster version 8.0. If you are running a more recent version of NEC CLUSTERPRO/ExpressCluster, see your NEC CLUSTERPRO/ExpressCluster documentation.

To stop NEC CLUSTERPRO/ExpressCluster from monitoring Arcserve Backup services

1. Access the Task Manager.

The Task Manager window appears.

Note: You can only stop monitoring services from the active node. If you attempt to perform this task on a passive node, the Application/Service Name list on the Task Manager will be empty.



2. Locate and select the applicable Arcserve service. Right-click on the service and from the pop-up menu, select Stop Monitoring. A confirmation screen appears asking you to confirm or cancel your request to stop monitoring the selected service. Click OK.

The selected Arcserve Backup service is no longer being monitored by NEC CLUSTERPRO/ExpressCluster.

Change the Arcserve Backup Domain in NEC CLUSTERPRO/ExpressCluster

In a NEC CLUSTERPRO/ExpressCluster cluster environment, you can move a member server to a different Arcserve Backup domain. Changes to the domain in a cluster environment can only be made on an active node and must be changed for all nodes within the cluster.

To change the Arcserve Backup domain in a NEC cluster

1. Stop the cluster group. For more information, see [Stop NEC Cluster Groups](#) (see page 934).

Note: You must stop the group to edit the group properties.

2. Remove the registry sync and edit the start.bat and stop.bat scripts to disable Arcserve Backup scripts added during installation. For more information, see [Disable Arcserve Backup in NEC Cluster Scripts](#) (see page 935).

3. From the Arcserve Backup home directory, run the cstart.bat utility to start all Arcserve Backup services.

All Arcserve Backup services are started.

4. From the Start menu, access the Server Configuration Wizard to run the ARCserveCfg.exe utility for the active node and specify the new Arcserve Backup domain. For more information about changing a domain, see [Move a Member Server to a Different Arcserve Backup Domain](#) (see page 572).

The first "active" cluster node is configured for the new domain.

5. From the Arcserve home directory, run the cstop.bat utility to stop all Arcserve Backup services.

All Arcserve Backup services are stopped.

6. From the Cluster Manager, right-click on the group name and from the pop-up menu, select Move Group to change the active node.

The status of the original node will be changed to offline (passive) and the status of the next node within the cluster will be changed back to online (active).

7. From the Arcserve Backup home directory, run the cstart.bat utility to start all Arcserve Backup services.

All Arcserve Backup services are started.

8. From the Arcserve home directory, run the cstop.bat utility to stop all Arcserve Backup services.

All Arcserve Backup services are stopped.

9. Repeat steps 6 through 8 for all remaining nodes in the cluster.

All nodes in the cluster have been changed to the new domain.

10. From the Cluster Manager, right-click on the group name and from the pop-up menu, select Move Group to change the active node back to the original node.

The status of the last node will be changed to offline (passive) and the status of the original node within the cluster will be changed back to online (active).

11. Rebuild the NEC Cluster Scripts and Registry Sync. For more information, see [Enable Arcserve Backup in NEC Cluster Scripts](#) (see page 938).

The new NEC HA scripts are created and the registry is synchronized.

12. Start the cluster group.

Manage Arcserve Backup Cluster Servers in NEC CLUSTERPRO/ExpressCluster

The Server Configuration Wizard lets you perform various management tasks to specify how Arcserve Backup servers function in a cluster environment. In a cluster environment, these management tasks can only be made on the active node and must also be made for all nodes within the cluster. These management tasks include the following:

- Changing the database
- Promoting a member server to a primary server
- Demoting a primary server to a member server.

To manage the Arcserve Backup Cluster Servers in NEC CLUSTERPRO/ExpressCluster

1. Stop the failover group that the Arcserve Backup is deployed on and then suspend the cluster group.

For more information, see [Stop NEC Cluster Groups](#) (see page 934).

Note: You must stop the group to edit the group properties.

2. Remove the registry sync and edit the start.bat and stop.bat scripts to disable Arcserve Backup scripts added during installation. For more information, see [Disable Arcserve Backup in NEC Cluster Scripts](#) (see page 935).

Resume the cluster and then start the failover group.

3. From the Start menu, access the Server Configuration Wizard to run the ARCserveCfg.exe utility for the active node and make the necessary change. Do not check the "Last Cluster Node" checkbox on the last screen of the Server Configuration Wizard.

- For more information about changing the database, see [Specify a Arcserve Backup Database Application](#) (see page 662).

Note: Local SQL Server is not supported when NEC CLUSTERPRO/ExpressCluster is used to make Arcserve Backup highly available.

- For more information about promoting a member server to a primary server, see [Promote a Member Server to a Primary Server](#) (see page 565).
- For more information about demoting a primary server to a member server, see [Demote a Primary Server to a Member Server](#) (see page 568).

Note: When this utility is run on the first node within a cluster, it will run in the normal mode.

The first "active" cluster node is configured for the new property and a new arcservecfg.ICF configuration file is created.

4. From the Start menu, access the Server Configuration Wizard to run the ARCServeCfg.exe utility for the new active node and make the necessary change.

Note: When this utility is run again on any subsequent nodes in the same cluster, it will detect the existence of the arcservecfg.ICF configuration file and automatically run the utility in the cluster mode.

The next "active" cluster node is configured for the new property.

5. Repeat steps 4 and 5 for all remaining nodes in the cluster. When you are performing this configuration procedure on the last node in the cluster, check the "Last Node" checkbox on the last screen of the Server Configuration Wizard.

All nodes in the cluster are configured for the new property.

6. From the Cluster Manager, right-click on the group name and from the pop-up menu, select Move Group to change the active node back to the original node.

The status of the last node will be changed to offline (passive) and the status of the original node within the cluster will be changed back to online (active).

7. Stop the failover group that the Arcserve Backup server is deployed on and then suspend cluster services.

Rebuild the NEC Cluster Scripts and Registry Sync. For more information, see [Enable Arcserve Backup in NEC Cluster Scripts](#) (see page 938).

The new NEC HA scripts are created and the registry is synchronized.

8. Resume the cluster and then start the failover group.

Stop NEC Cluster Groups

If you need to edit the group properties (for example to edit the start.bat or stop.bat files, or remove or add registry sync) you must first stop the group. In addition, if you need to remove Arcserve Backup from NEC CLUSTERPRO/ExpressCluster, you must also stop the group.

Note: This section contains graphics that correspond with NEC CLUSTERPRO/ExpressCluster version 8.0. If you are running a more recent version of NEC CLUSTERPRO/ExpressCluster, see your NEC CLUSTERPRO/ExpressCluster documentation.

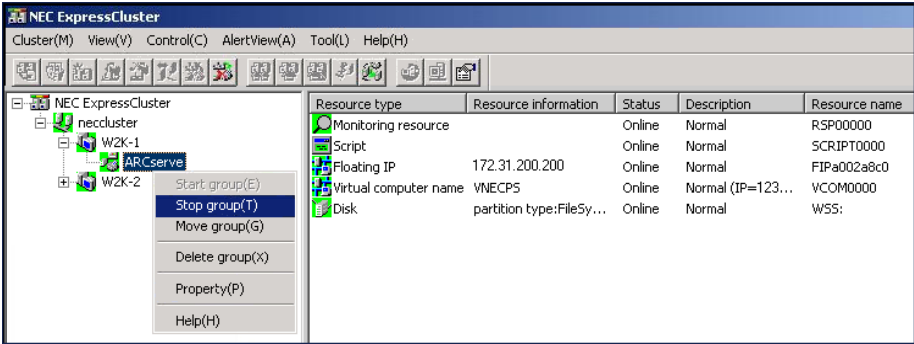
To stop the NEC cluster group

1. Access the Cluster Manager.

The Cluster Manager window appears.

- 2. From the tree listing, right-click on the Arcserve group, and from the pop-up menu select Stop group.

A confirmation pop-up screen appears.



- 3. Click OK.

The selected group is stopped.

Disable Arcserve Backup in NEC Cluster Scripts

Cluster scripts and registry keys are inserted during the NEC post-setup process. When upgrading from a previous release, the cluster scripts need to be disabled and the registry key need to be deleted.

Note: This section contains graphics that correspond with NEC CLUSTERPRO/ExpressCluster version 8.0. If you are running a more recent version of NEC CLUSTERPRO/ExpressCluster, see your NEC CLUSTERPRO/ExpressCluster documentation.

To disable Arcserve Backup in NEC cluster scripts


- 1. Access the Cluster Manager.

The Cluster Manager window appears.

Note: Cluster Manager is a utility provided by NEC and is installed on servers that have NEC CLUSTERPRO/ExpressCluster installed. From the Cluster Manager, you perform most of the configuration and management tasks associated with clusters.

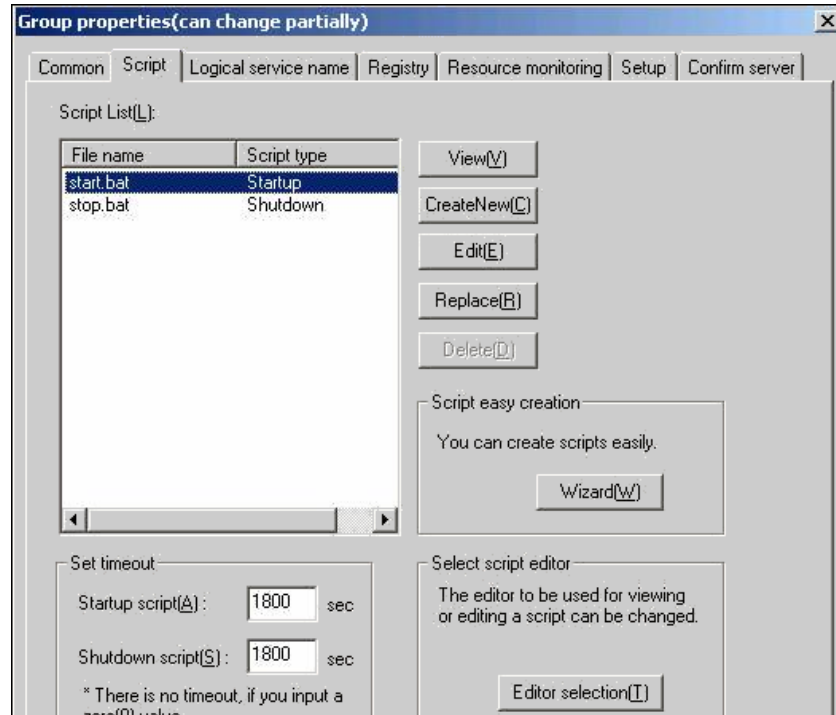
2. Select the NEC Group that the Arcserve server is deployed in, and locate the corresponding Arcserve cluster resources. Right-click on each Arcserve cluster resource and from the pop-up menu, select Property.

The Group property dialog appears.



3. Select the Reference and Change option. When the Group properties dialog opens, select the Script tab.

The Script tab dialog appears.



4. From the Script list, select start.bat and click Edit. When the start.bat script appears, locate the REM SET process script (two locations) and set the value to zero as follows:

```
SET process=0
```

Note: In the start.bat file, the REM SET process script is located after NORMAL and after FAILOVER.

The start.bat script is modified.

5. From the Script list, select stop.bat and click Edit. When the stop.bat script appears, locate the REM SET process script (two places) and set the value to zero as follows:


```
SET process=0
```

Note: In the stop.bat file, the REM SET process script is located after NORMAL and after FAILOVER.

The stop.bat script is modified.

- From the Group properties dialog, select the Registry tab.

The Registry dialog appears.



- From the Registry key list, select the existing registry key and click Delete.

The Registry key is deleted.

Enable Arcserve Backup in NEC Cluster Scripts

Cluster scripts and registry keys are inserted during the NEC post-setup process. During the upgrade process, the cluster scripts are disabled and the registry key is deleted. When the upgrade is finished, these cluster scripts need to be enabled and registry keys need to be rebuilt.

Note: This section contains graphics that correspond with NEC CLUSTERPRO/ExpressCluster version 8.0. If you are running a more recent version of NEC CLUSTERPRO/ExpressCluster, see your NEC CLUSTERPRO/ExpressCluster documentation.

To enable Arcserve Backup in NEC cluster scripts

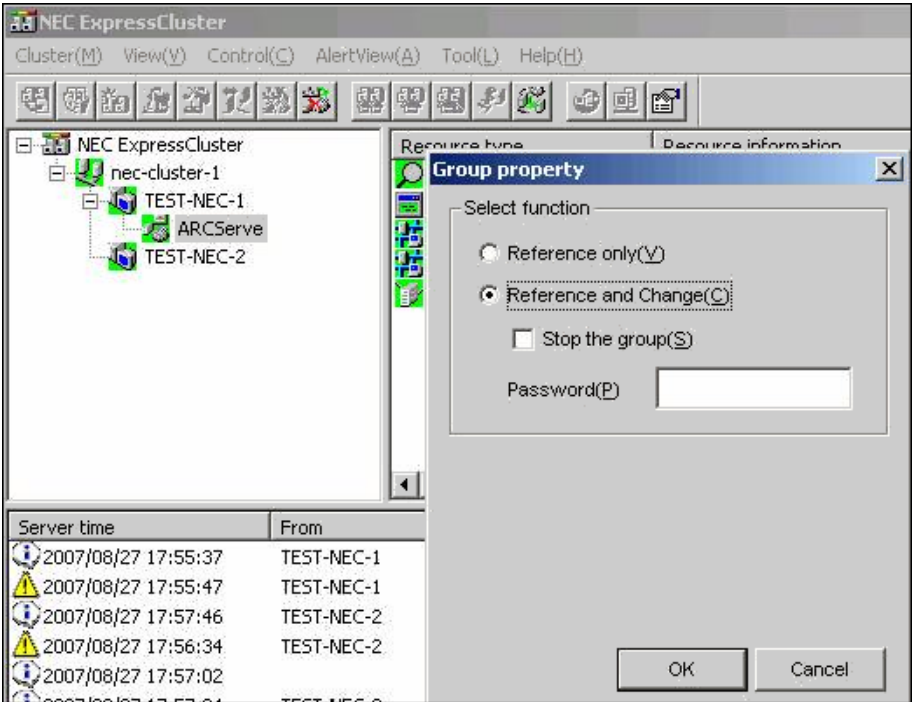
- Access Cluster Manager.

The Cluster Manager dialog appears.

Note: Cluster Manager is a utility provided by NEC and is installed on servers that have NEC CLUSTERPRO/ExpressCluster installed. From the Cluster Manager, you perform most of the configuration and management tasks associated with clusters.

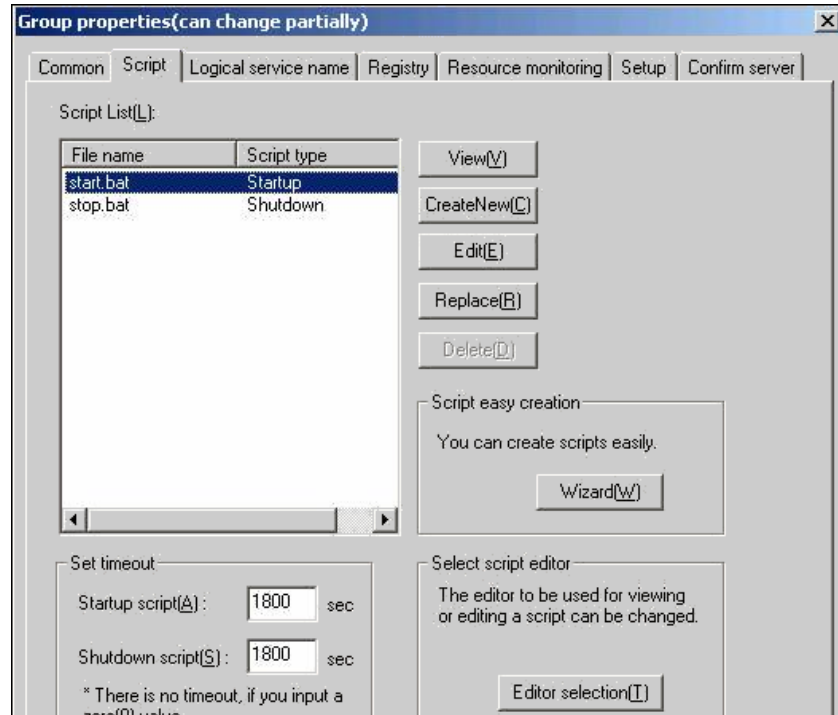
- 2. Select the NEC Group that the Arcserve server is deployed in, and locate the corresponding Arcserve cluster resources. Right-click on each Arcserve cluster resource and from the pop-up menu, select Property.

The Group property dialog appears.



3. Select the Reference and Change option. When the Group properties dialog opens, select the Script tab.

The Script tab dialog appears.



4. From the Script list, select start.bat and click Edit. When the start.bat script appears, locate the REM SET process script (two places) and set the value to 1 as follows:

```
SET process=1
```

Note: In the start.bat file, the REM SET process script is located after NORMAL and after FAILOVER.

The start.bat script is modified.

5. From the Script list, select stop.bat and click Edit. When the stop.bat script appears, locate the REM SET process script (two places) and set the value to 1 as follows:


```
SET process=1
```

Note: In the stop.bat file, the REM SET process script is located after NORMAL and after FAILOVER.

The stop.bat script is modified.

6. From the Group properties dialog, select the Registry tab. When the Registry dialog opens, click Add.

The Add/Change registry key dialog appears.



7. Add the Registry key that corresponds with your computer's architecture:

- **x86 Platforms:**

HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\CA_ARCserve Backup\Base

- **x64 Platforms:**

HKEY_LOCAL_MACHINE\SOFTWAREWow6432Node\ComputerAssociates\CA_ARCserve Backup\Base

Click OK.

The Registry key is added to the Registry key list on the Group Properties dialog.

NetApp NAS Cluster Protection

The NetApp cluster is a configuration of nodes, each running the Data ONTAP operating system. This section describes the Arcserve Backup NDMP NAS Option support for NetApp Cluster running the Data ONTAP 8.2.

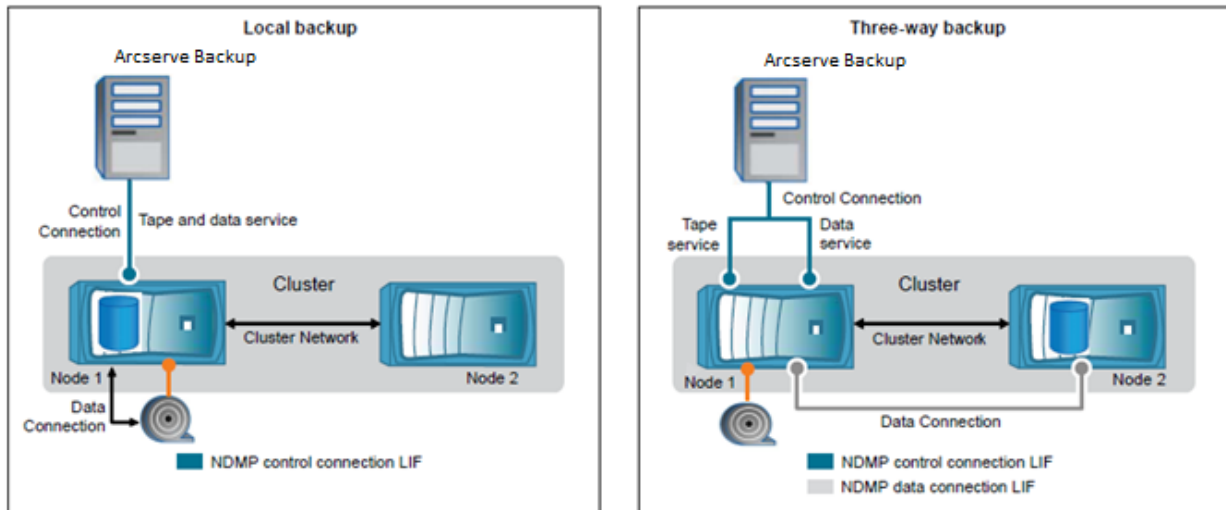
NDMP Version 4 Cluster Aware Backup Extension Interface

The NDMP Cluster Aware Backup Extension (CAB) interface leverages the extensibility of NDMP v4 to define the mechanism and protocol for enabling efficient backups of the file systems on a NetApp cluster. Refer to the [NDMP specification](#) for further details.

Arcserve Backup r16.5 SP1 update RO72520 supports the CAB extension.

NetApp Cluster Backup and Restore and Supported Modes

The following diagram describes the local backup and the three way backup of the NetApp NAS Cluster.



There are two NDMP modes available with NetApp Cluster running Data ONTAP 8.2:

- **NDMP node-scope-mode:**

A cluster is in a mixed version state where Data ONTAP 8.2 and earlier versions in the Data ONTAP 8.x release family are running on nodes, NDMP follows the node-scoped behavior. This NDMP behavior continues even after upgrading to Data ONTAP 8.2.

In the node-scoped NDMP mode, you can perform tape backup and restore operations at a node level.

You can manage NDMP at a node level by using NDMP options and commands. It is required that you use NDMP specific credentials to access a storage system to perform tape backup and restore operations.

Refer to the following link to view the NetApp documentation on commands for managing node-scoped NDMP mode.

https://library.netapp.com/ecm/ecm_download_file/ECMP1196817

- **Vserver aware NDMP mode:**

In a newly installed cluster where all nodes are running Data ONTAP 8.2, NDMP is in the Vserver aware mode by default.

If you want to perform tape backup and restore operations in the Vserver aware mode, it is recommended that you disable the node-scoped NDMP mode. From the [NetApp documentation](#), use the command to disable node-scoped NDMP mode.

Administrators can back up and restore all volumes hosted across different nodes in a cluster of a Vserver.

Configuring Arcserve Backup

You can configure Arcserve Backup to establish an NDMP control connection on any of the LIF (Logical Interface) types in a cluster. In Vserver aware NDMP mode, you can determine the availability of volumes and tape devices for backup and restore operations depending upon these LIF types and the status of the Cluster Aware Backup (CAB) extension.

The following tables show the availability of volumes and tape devices for NDMP control Connection LIF types in both Node-scoped mode and Vserver mode when configured from the Arcserve Backup Manager.

As a backup administrator, you can select any of the following LIFs, which Arcserve Backup supports, to back up a NetApp NAS cluster. Administrators can configure these LIFs to submit backup and restores.

Node-scoped mode		
NDMP Control Connection Type	Volumes available for backup or restore	Tape devices available for backup or restore
Node-management LIF	All volumes hosted by a node.	Tape devices connected to the node hosting node management LIF.
Vserver Mode with CAB extension supported		
NDMP Control Connection Type	Volumes available for backup or restore	Tape devices available for backup or restore
Node-management LIF	All volumes hosted by a node.	Tape devices connected to the node hosting node management LIF.
Data LIF	All volumes that belong to the Vserver that hosts the data LIF.	None
Cluster-management LIF	All volumes in the cluster.	All tape devices in the cluster.

Limitations

The NetApp Clustered Data ONTAP® 8.2 introduces the scalability limit for NDMP sessions that are based on the system memory size of the storage system. For more information, click on the following link for the Clustered Data ONTAP® 8.2 documentation.

<https://library.netapp.com/ecmdocs/ECMP1196874/html/GUID-479784EB-7101-4F16-9580-37FFF7C6D040.html>

Arcserve Backup can support only the maximum of 7 NAS devices per NAS node when the device is backed up from using the Node management LIF. A NAS device can be a changer or a tape drive.

In the Vserver aware NDMP mode, Arcserve Backup can support only the maximum of 7 NAS devices per cluster when the device is backed up from using the Cluster management LIF. A NAS device can be a changer or a tape drive.

For more details, you can refer to the following NetApp Clustered Data ONTAP® 8.2 Commands: Manual Page Reference documentation for increasing the NDMP session count: https://library.netapp.com/ecm/ecm_download_file/ECMP1196817

Appendix C: Using JIS2004 Unicode Characters with Arcserve Backup

This section contains the following topics:

[Introduction to JIS2004 Unicode Characters](#) (see page 947)

[Configuration Requirements for JIS2004 Unicode](#) (see page 947)

[Platforms Supporting JIS2004 Unicode Characters](#) (see page 948)

[Tasks You Can Perform Using JIS2004 Unicode Characters with Arcserve Backup](#) (see page 948)

[Arcserve Backup Applications Supporting JIS2004 Unicode Characters](#) (see page 949)

[Limitations of Using JIS2004 Unicode Characters with Arcserve Backup](#) (see page 950)

Introduction to JIS2004 Unicode Characters

Various Windows platforms support the capability to process data and display text and symbols using Unicode characters. Arcserve Backup supports the capability to display JIS2004 Unicode characters that relate to the following Windows system attributes:

- On client agent systems, Arcserve Backup displays folder names, file names, and registry strings using JIS2004 Unicode characters.
- On systems hosting application agents, Arcserve Backup displays database names, table names, and database instance names using JIS2004 Unicode characters.

Configuration Requirements for JIS2004 Unicode

The following configuration requirements apply when you require support for JIS2004 Unicode characters in your Arcserve Backup environment:

- All servers in a Arcserve Backup domain (primary and member servers) must be running this release of Arcserve Backup and have the same language packs installed.
- Arcserve Backup r16, r15, r12.5, r12 SP1, and Arcserve Backup r12 agents and the Arcserve Backup r16.5 base product cannot co-exist on the same computer. However, the Arcserve Backup r16.5 base product can co-exist with these agents in the same network.
- To back up and restore Microsoft Exchange Server and Microsoft SharePoint Server data with support for JIS2004 Unicode characters, you must enable the Arcserve Catalog database.

- All Arcserve Backup servers sharing a single Microsoft SQL Server database should be upgraded to the same version of Arcserve Backup.
- All Arcserve Backup domains sharing a single Microsoft SQL Server database must specify the same SQL Server collation setting. You can specify the SQL Server collation setting from a primary and stand-alone server using the Server Configuration Wizard.

Platforms Supporting JIS2004 Unicode Characters

The following Windows operating systems support JIS2004 Unicode characters:

- Windows Server 2008, Japanese version.
- Windows Server 2008 R2, Japanese version.
- Windows Vista, Japanese version.
- Windows Server 2003, Japanese version, with the Japanese fonts patch.
Note: For more information, see the Microsoft website.
- Windows XP Japanese version, with the Japanese fonts patch.
Note: For more information, see the Microsoft website.
- Windows 7, Japanese version.
- Windows 8, Japanese version.
- Windows Server 2012, Japanese version.

Tasks You Can Perform Using JIS2004 Unicode Characters with Arcserve Backup

You can perform the following tasks when you run Arcserve Backup on JIS2004 Unicode character-based operating systems:

- Browse and view system and volume information in the Arcserve Backup managers, view logs, and generate reports without displaying unrecognized text.
- Back up systems hosting Arcserve Backup agents.
- Back up files, folders, databases, tables, instances, and Microsoft Exchange messages.
- Restore data by tree, session, query, and media.
- View job details and Activity Log data in the Job Status Manager.
- Specify local and global filters using JIS2004 characters.

- Generate Alert Manager email messages with JIS2004 Unicode text attachments (for example, a Job Log).
- Execute job scripts using scripts created in previous releases of Arcserve Backup.
Note: This capability is limited to scripts created using Arcserve Backup r12.
- Execute Arcserve Backup command line operations using JIS2004 characters.

Arcserve Backup Applications Supporting JIS2004 Unicode Characters

The Arcserve Backup applications listed below support JIS2004 Unicode characters.

- All releases of the Arcserve Backup base product from Arcserve Backup r12 Service Pack 1 to this release on x86 and x64 systems.
Note: Unicode support is only applicable to local backup, restore, and compare operations.
- All releases of the Arcserve Backup Client Agent for Windows from Arcserve Backup r12 Service Pack 1 to this release on x86 and x64 systems.
- All releases of the Arcserve Backup Agent for Microsoft Exchange Server from Arcserve Backup r12 Service Pack 1 to this release on x86 and x64 systems.
- All releases of the Arcserve Backup Agent for Microsoft SQL Server from Arcserve Backup r12 Service Pack 1 to this release on x86 and x64 systems.
- All releases of the Arcserve Backup Agent for Microsoft SharePoint Server 2007 from Arcserve Backup r12 Service Pack 1 to this release on x86 and x64 systems.
- All releases of the Arcserve Backup Agent for Virtual Machines from Arcserve Backup r12 Service Pack 1 to this release on VCB Proxy systems.
- All releases of the Arcserve Backup Agent for Open Files from Arcserve Backup r12 Service Pack 1 to this release on x86 and x64 systems.
Note: Unicode support is only applicable to VSS-based backup and restore operations.
- Arcserve Replication and High Availability agents on x86 and x64 systems.
- All releases of the Arcserve Backup Disaster Recovery Option.

Note: Some non-English language-based characters may appear as garbled characters when the Arcserve Backup base product is installed on any Windows operating system. This limitation does not affect the outcome of backup and restore jobs.

Limitations of Using JIS2004 Unicode Characters with Arcserve Backup

The following limitations apply when using JIS2004 Unicode characters with Arcserve Backup:

Operating Systems

The operating systems that the following agents protect do not support JIS2004 Unicode characters:

- Client Agent for UNIX
- Client Agent for Linux
- Client Agent for Mac OS X
- Client Agent for AS400
- Client Agent for OpenVMS

Arcserve Backup Agents and Options

The Arcserve Backup agents and options listed below do not support JIS2004 Unicode characters.

You can work around this limitation by backing up and restoring data from the parent-level path instead of the path containing the JIS2004 characters.

Important! To display ANSI character-based agents properly, you must set the locale and the product language options on the system hosting these agents to the same locale and product language as that of the system hosting the Arcserve Backup database before you install Arcserve Backup.

- Agent for Informix: all supported releases and service packs
- Agent for Lotus Domino: all supported releases and service packs
- Agent for Microsoft Exchange Server 2000 and 2003: all supported releases and service packs
- Agent for Microsoft Exchange Server 2007 on Windows Server 2003: all supported releases and service packs
- Agent for Microsoft Exchange Server 2007 Service Pack 1 on Windows Server 2008: r12
- Agent for Microsoft SharePoint Server 2003: all supported releases and service packs
- Agent for Microsoft SharePoint Server 2007: r12
- Agent for Microsoft SQL Server: r12
- Agent for Open Files: r12
- Agent for Oracle: r12

- Agent for Sybase: all supported releases and service packs
- Client Agent for Windows: r12
- Enterprise Option for SAP R3 for Oracle: all supported releases and service packs
- NDMP NAS Option: all supported releases and service packs
Note: EMC/Celera and NetApp NAS filers do not support JIS2004 Unicode characters.
- Image Option-based backups and restores: all supported releases and service packs
Note: The Arcserve Backup Enterprise Module is a prerequisite component for the Image Option.

Arcserve Backup Components

The Arcserve Backup components listed below do not support JIS2004 Unicode characters.

- Alert Manager
- Alert options specified as Global Options in the Arcserve Backup managers and utilities that follow:
 - Backup Manager
 - Restore Manager
 - Media Assure & Scan Utility
 - Merge Utility
 - Count Utility
 - Purge Utility
- BConfig.exe
Note: This component lets you configure the Arcserve Backup server when you are installing or upgrading Arcserve Backup.
- DBAConfig.exe
Note: This component lets Arcserve Backup configure database instances during the installation process.
- Discovery Configuration
- License Management dialog
Note: This component lets you manage Arcserve Backup licenses. You can open the License Management dialog by clicking Manage Licenses on the Help, About Arcserve Backup dialog.
- Report Writer
- Server Configuration Wizard

- Server Migration Component
- SetupSQL.exe

Note: This component lets the installation wizard create the Arcserve Backup database with Microsoft SQL Server when you are installing Arcserve Backup or upgrading Arcserve Backup from a previous release.

Arcserve Backup Tasks

Arcserve Backup does not support performing the following tasks:

- Installing Arcserve Backup agents on remote systems using Agent Deployment. Agent Deployment does not support using JIS2004 Unicode characters for host names, user names, and passwords.
- Using JIS2004 Unicode characters to specify customer information on the Customer Information dialog when installing Arcserve Backup.
- Browsing JIS2004 Unicode-based agent machine names, user names, and passwords.
- Specifying JIS2004 Unicode-based names on Arcserve Backup host names, primary server names, stand-alone server names, and member server names.
- Specifying JIS2004 Unicode-based directory paths for installing Arcserve Backup or any Arcserve Backup component, specifying file system device paths, catalog database path, and so on.
- Importing a host list from a text file with a file name that contains Unicode characters when you install and upgrade Arcserve Backup, agents, and options on remote systems.
- Importing a host list from a text file with a file name that contains Unicode characters when running Remote Agent Deployment.
- Specifying JIS2004 Unicode-based names on ARCserve-specific objects. For example, job names, Device Group names, Media Pool names, media names, location names, and session passwords for encryption.
- Specifying file names and file paths as criteria for creating a Vault Schedule using the Media Management Administrator using JIS2004 Unicode characters.
- Specifying JIS2004 Unicode characters on network directory and file shares.
- Specifying JIS2004 Unicode characters for the caroot password using the Arcserve Backup Server Configuration Wizard. The Server Configuration Wizard interprets Unicode characters as the question mark symbol "?," which is an acceptable character for the caroot password. However, after you set the caroot password using the Server Configuration Wizard, you will not be able to change the caroot password from the Manager Console.

- Sending Alert email messages with file attachments when the file attachments are stored in directories that are named with Unicode characters.

Note: The file attachment itself can contain Unicode characters.

- Renaming and editing XML content in the report template files using JIS2004 Unicode Characters. The generated reports will not display properly. The report template files are stored in the following directory:

ARCserve_Home\template\reports

Arcserve Backup Reports

Arcserve Backup does not support creating the following reports when using an ARCserve server that is running Arcserve Backup for Windows r15 connect remotely to an ARCserve server that is running Arcserve Backup for Windows r12 Service Pack 1 or Arcserve Backup for Windows r12.

- 7 days Backup Status Report.xml
- 7 days Job Status Report.xml
- 7 days Media Usage History Report.xml
- Backup Client Data Size Report.xml
- Backup Clients And Job Associations Report.xml
- Backup Window And Throughput Comparison Report.xml
- Daily Backup Status Report.xml
- Daily Failed Backups Report.xml
- Daily Job Status Report.xml
- Detailed Media Pool Report.xml
- Detailed Media Usage By Backup Clients Report.xml
- Failed Backups Report.xml
- Media Required For Data Recovery Report.xml
- Staging Migration Report.xml
- Staging Purge Failed Report.xml
- Staging SnapLock Report.xml
- Staging Summary Report.xml
- Custom Report (New Report)

Appendix D: Protecting Hyper-V Systems Using the Hyper-V VSS Writer

This section contains the following topics:

[Overview of Protecting Hyper-V VMs Using the Hyper-V VSS Writer](#) (see page 955)

[Prerequisite Components for Hyper-V VSS Writer Protection](#) (see page 955)

[Configure Arcserve Backup to Detect Hyper-V VMs](#) (see page 957)

[How Back Up Using Saved State Works](#) (see page 958)

[How Back Up Using Child Partition Snapshot Works](#) (see page 958)

[Back Up Hyper-V VMs Using the Hyper-V VSS Writer](#) (see page 958)

[Restore Data to Its Original Location](#) (see page 959)

Overview of Protecting Hyper-V VMs Using the Hyper-V VSS Writer

Arcserve Backup lets you protect Hyper-V VMs using the Arcserve Volume Shadow Copy Service (VSS) agent. You can protect Microsoft Hyper-V data with VSS Writers using Volume Shadow Copy Service technologies.

The sections that follow describe how to configure, back up, and restore Hyper-V VMs using the Hyper-V VSS writer. The processes described are applicable to Arcserve Backup for Windows r12 SP1 installations, and can be used protect Hyper-V systems in Arcserve Backup for Windows r12.5, Arcserve Backup for Windows r15, and this release of Arcserve Backup.

Limitations and Considerations

- You cannot restore data at file level granularity from raw (full VM) backup data.
- You cannot perform mixed mode backups, which consist of raw (full VM) weekly backups and file mode daily backups.
- You can protect Hyper-V VMs that are in a powered off state when you execute Arcserve Hyper-V Configuration Tool.

Prerequisite Components for Hyper-V VSS Writer Protection

The prerequisite components for Hyper-V VSS Writer protection are identical to that of standard VSS Writer requirements. The applications that follow are required to deploy Hyper-V VSS technology in your Arcserve Backup environment:

- This release of the Arcserve Backup for Windows base product

- This release of Arcserve Backup for Windows Client Agent for Windows

The Arcserve Backup Client Agent for Windows must be installed in partition zero (0) on the Hyper-V server machine. Partition zero (0) is reserved for the host operating system and its applications. All other partitions, for example, partition 1, 2, and so on, are reserved for child partitions or virtual machines (VMs).

- This release of Arcserve Backup for Windows Agent for Open Files

You must register the license for the Agent for Open Files for the Hyper-V host system with the Arcserve Backup server.

Note: Optionally, you can register the license for the Agent for Open Files using the license key for the Agent for Open Files for Virtual Machines on Windows.

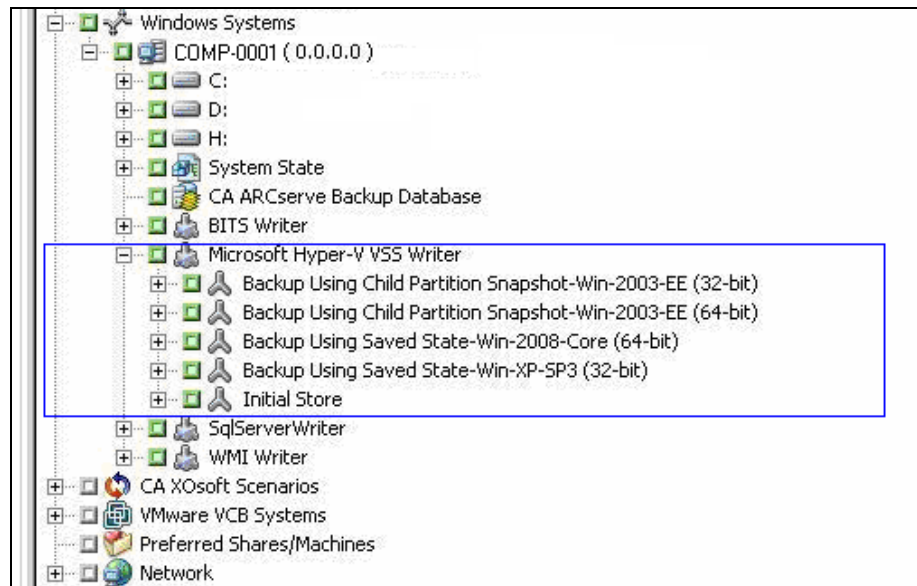
Configure Arcserve Backup to Detect Hyper-V VMs

To perform backup and restore operations on machines using the Hyper-V VSS writer, you must configure Arcserve Backup to detect the Hyper-V server.

To configure Arcserve Backup to detect Hyper-V VMs

1. Based on the configuration in your backup environment, complete one of the following actions and then go to the next step.
 - If the Arcserve Backup server components are installed on the Hyper-V server system, add the local Hyper-V server into the Backup Manager.
 - If the Arcserve Backup server components are not installed in the Hyper-V server, add the remote Hyper-V server into the Backup Manager by completing the steps that follow:
 - a. From the Source tree in the Backup Manager, right click the Windows Systems object and select Add Machine/Object from the pop-up menu.
The Add Agent dialog opens.
 - b. From the Add Agent dialog, specify the name of the Hyper-V server in the Host Name field or the IP address in the IP address field, and then click Add.

After you add the Hyper-V server system into the Backup Manager, expand the Hyper-V server to display the Microsoft Hyper-V VSS Writer as illustrated by the following screen.



How Back Up Using Saved State Works

Back up Using Saved State is a backup operation that places VMs into a saved state before the backup is performed. This state lets you perform point-in-time backups of guest operating systems. It is a stateful, data inconsistent backup. Back up Using Saved State presents the following limitations on VM backups:

- The virtual hard disk in the backup cannot be offline mounted to retrieve specific files.
- The applications in the VM will not be aware that a backup, a restore, or both occurred when you restore the backed up data.

Note: For more information about these limitations, see the Microsoft website.

How Back Up Using Child Partition Snapshot Works

Back up Using Child Partition Snapshot is a backup operation that lets the VSS Writer take a snapshot of the data from the guest operating system in the VM. Backups of this type let you back up VMs that support VSS and have the Integration components installed and enabled. It is a stateless, data consistent backup.

Back up Using Child Partition Snapshot presents the following advantages on VM backups:

- You can offline mount the virtual hard disk from this backup to retrieve specific files.
- The VSS capable applications residing in the VM will detect that the backup or restore of the VM is taking place, and they will participate in that backup or restore process to ensure that the application data is consistent.

Note: For more information, see the Microsoft website.

Back Up Hyper-V VMs Using the Hyper-V VSS Writer

The Hyper-V VSS Writer lets you back up VMs that are in an online and offline state. These operations are transparent to Arcserve Backup.

Note: The Hyper-V VSS Writer supports only full backups.

The following steps describe how to back up Hyper-V VMs using the Hyper-V VSS Writer. For information about backing up data using the VSS Writer, see the *Arcserve Backup for Windows Microsoft Volume Shadow Copy Service Guide*.

To back up Hyper-V VMs using the Hyper-V VSS writer

1. Open the Backup Manager, select the Source tab, and select the Microsoft Hyper-V VSS Writer object.

All Hyper-V settings and virtual machines are specified for backup. If you do not want to back up all of the VMs, expand the Microsoft Hyper-V VSS Writer object (to display all servers) and clear the check mark next to the server that you do not want to back up.

2. (Optional) Right-click the Microsoft Hyper-V VSS Writer object and select Writer Options from the pop-up menu.
3. Click the destination tab to specify the destination for backup.
4. Click the Submit on the toolbar submit the job.

The Submit Job dialog opens.

5. Complete the required fields on the Submit Job dialog and click OK.

The job is submitted.

Restore Data to Its Original Location

This method lets you restore the Hyper-V configuration, the VM configurations, and the backup data to its original location. The current Hyper-V configuration and VMs configuration and data will be restored to the state they were in when they were backed up.

Limitations and Considerations

- The Hyper-V servers can be in an online state or offline state during the restore operation.
- The Hyper-V VSS Writer ensures that the backup data is properly restored to its original location.
- You do not need to perform additional steps during the restore or after the restore is complete.
- The VM can be used as soon as the restore is complete.


Using Arcserve Backup to restore Hyper-V server data, you can recover data in the following scenarios:

- You can restore Hyper-V server backup data to its original location.
- You can restore VM backup data to its original location.
- You can recover a guest operating system in a VM to its original location.


Note: For information about using the VSS Writer, see the *Arcserve Backup for Windows Microsoft Volume Shadow Copy Service Guide*.

To restore data to its original location

1. Open the Restore Manager and do one of the following:
 - From the drop-down list, select the Restore by tree method, expand the Windows Systems object, browse to Microsoft Hyper-V VSS Writer, and specify one or more VMs that you want to restore.



- From the drop-down list, select Restore by session, browse to and specify a session to restore.



2. Click the Destination tab.
Click the Restore files to their original location option.
3. Click the Submit on the toolbar to submit the job.
The Submit Job dialog opens.
4. Complete the required fields on the Submit Job dialog and click OK.
The job is submitted.

Note: After the restore is complete, the restored VMs will be in a Saved state. In other words, the online restore places the VMs in an offline state when the restore is complete. You must then start the VMs manually to bring them back online.

Glossary

archive job schedules

An archive job schedule lets you configure your archive job by using the Repeat Method on the Schedule tab in the Archive Manager. The schedule lets you run an archive job on a repeating basis.

archive manager

The Archive Manager lets you customize your archive jobs using filters, specifying options, and setting schedules.

cloud connection configuration

A cloud connection is a user-defined configuration that contains the information that Arcserve Backup needs to communicate with the specified cloud vendor. The cloud connection helps ensure that your backup data is stored under the account that you created with the cloud vendor.

cloud storage

A cloud storage is the final destination where your cloud devices are stored.

cloud-based device

A cloud-based device is a virtual, Arcserve Backup device that you create in association with an Arcserve Backup cloud connection. The cloud-based device lets you direct Arcserve Backup to store data with the cloud vendor specified by the Arcserve Backup cloud connection.

data deduplication

Data deduplication is technology that allows you to fit more backups on the same physical media, retain backups for longer periods of time, and speed up data recovery.

data mover server

Arcserve Backup data mover servers facilitate the transfer of data to local storage devices. The storage devices include shared libraries and file system devices. Data mover servers are supported on UNIX or Linux operating systems. Arcserve Backup manages data mover servers from a single, centralized, primary server. Arcserve Backup data mover servers function in a manner that is similar to member servers.

Grandfather-Father-Son rotation

A Grandfather-Father-Son (GFS) rotation is a backup schedule policy that uses daily (Son), weekly (Father), and monthly (Grandfather) backup media sets (tapes).

media pool

A media pool is a collection of backup media (tapes) that is set aside for a specific job and managed as a unit.

member server

Member servers function as worker servers to a primary server. Member servers process jobs dispatched by the primary server. Using primary and member servers you can have a single point of management for multiple Arcserve Backup servers in your environment. You can then use the Manager Console on the primary server to manage its member servers.

multiplexing

Multiplexing is a process in which data from multiple sources is written to the same media (tapes) simultaneously.

multistreaming

Multistreaming is a process that lets you divide backup jobs into multiple sub-jobs (streams) that run simultaneously and sends data to the destination media (tape device or file system device).

primary server

Primary servers function as a master server that controls itself and one or more member servers and data mover servers. With primary servers you can manage and monitor backup, restore, and other jobs that run on primary servers, member servers, and data mover servers. Using primary, member, and data mover servers, you can have a single point of management for multiple Arcserve Backup servers in your environment. You can then use the Manager Console to manage the primary server.

staging

Staging is a method of protecting data that lets you back up data to a temporary data storage location, and then based on selected policy options, migrate (copy) the data to the final destination media.

The disk staging method utilizes a disk as the staging area and is commonly referred to as Backup to Disk to Tape (D2D2T).

The tape staging method utilizes a tape library or a Arcserve Backup virtual tape library as the staging area and is commonly referred to as Backup to Tape to Tape (D2T2T).

synthetic full backup

Synthetic Full Backup, or SFB, is a backup for Windows Client Agents that is executed as part of a Normal backup with staging option, Deduplication backup with staging option, or Deduplication backup that synthesizes, or combines, an initial full session and its subsequent incremental backups into one session. Each synthetic full session becomes the parent full session for the next SFB job (unless a real full backup session is executed after the SFB session).

virtual tape library

A virtual tape library (VTL) is a hard disk that is configured to behave like a tape library. You must configure the hard disk using Arcserve Backup.

writer

A writer is application-specific software that lets Arcserve Backup create shadow copies of application-based data. A *volume* shadow copy is copy of the data that resides in a volume at a specific point in time.

Index

6

64-bit Windows platform support • 38

A

Activity Log

- data • 329, 330, 668, 715
- organizing • 331
- pruning • 331

add Arcserve Backup users • 94

add Windows users • 93

administrative shared drives • 133

Agent for Open Files, managing • 189

Alert Manager • 707

- configuration options • 709

Append X days option • 317, 318

Archive Manager

- How to Specify Source Data for Archiving Using the Classic View and the Group View • 355
- Options on the Archive Manager Destination Tab • 358
- Options on the Archive Manager Policies Tab • 356
- overview • 353

archiving data

- Advanced Options for Archiving • 369
- How Arcserve Backup Archives Data • 351
- overview • 351
- Submit Archive Jobs • 351

Arcserve D2D

- starting • 128

Arcserve database

- database protection job, about • 613
- database protection job, deleting • 624
- database protection job, modifying • 614
- database protection job, recreating • 625
- database protection job, starting • 623
- device records • 605
- error reporting • 605
- MS SQL index rebuilding • 605
- protecting • 607
- pruning • 604
- restoring • 628

Arcserve Replication • 130

Audit Log Reports • 672

B

backing up data

- backing up data, deduplication • 738
- differential backup • 307
- entire nodes • 193
- incremental backup • 307
- multiplexing • 106, 109, 189
- multistreaming, defined • 103
- schedules and rotations • 152
- setting up jobs • 140
- submitting backup jobs • 136

Backup Manager • 137

- global backup options • 154

bar codes • 448

broadcast option • 710

C

ca_auth command • 36, 37, 538, 539

ca_devmgr command • 204, 228

ca_jobsecmgr command • 40

ca_log command • 330

ca_mmo command • 464, 465, 469, 474

cabatch command • 35

CAReports command • 679, 681

catalog database

- about • 650
- enabling media pool maintenance • 607
- enabling the catalog database • 653

Central Agent Admin • 717

central management

- administering Arcserve servers • 52
- managing devices • 52
- managing jobs • 47
- managing licenses • 53
- managing the ARCserve database • 49
- monitoring jobs • 48
- using alerts • 51
- using job history • 55
- using logs • 50
- using reports • 50

change password, system account • 520

change the Arcserve system account password • 573

changing an ARCserve computer name
about • 542

-
- changing a Manager Console system • 553
 - changing a member server computer name • 550
 - changing a primary server computer name • 543
 - changing a stand-alone server computer name • 553
 - classic view, Backup Manager • 140, 145, 306
 - clean media • 412
 - cleaning tape heads • 404, 412
 - cloud storage
 - How Arcserve Backup Purges Data from Cloud Storage • 350
 - How to Manage Cloud Storage • 477, 478, 480, 483, 485, 486, 487
 - Merge Jobs to Cloud Storage • 349
 - Migrate Data to Cloud Storage in a Staging Backup Job • 487
 - Scan Jobs for Cloud Storage • 349
 - Specify Encryption and Compression Options by Migrating Jobs to Cloud Storage • 162
 - clusters
 - failover • 908
 - mirrored disk • 911
 - overview • 905
 - protection • 913
 - MSCS application protection • 917
 - MSCS self-protection • 916
 - NEC CLUSTERPRO application protection • 929
 - NEC CLUSTERPRO self-protection • 929
 - quorum disk • 911
 - resource group • 909
 - shared disk • 910
 - virtual name • 909
 - clusters, MSCS clusters
 - change cluster domain • 927
 - change database • 924
 - cdelete cluster resources • 923
 - clerebuild cluster resources • 922
 - demote primary server • 924
 - failover support • 912
 - promote member server • 924
 - stop HA service monitoring • 917
 - clusters, NEC clusters
 - change cluster domain • 931
 - change database • 933
 - demote primary server • 933
 - disable cluster scripts • 935
 - enable cluster scripts • 938
 - failover support • 912
 - promote member server • 933
 - stop cluster groups • 934
 - stop HA service monitoring • 930
 - clusters, Windows Server 2008 • 923
 - compressing media • 400
 - configuration tools
 - Device Configuration • 210
 - device group configuration • 212
 - configure node tiers • 521
 - configure SQL Server as the ARCserve database • 662
 - configure SQL Server Express as the Arcserve database • 667
 - Copy Utility • 33
 - cstart command • 493
 - cstop command • 493
- ## D
- database
 - agents • 133, 257
 - Database Engine • 491, 514
 - Database Manager • 604
 - MS SQL configuration • 658, 661
 - MS SQL index rebuilding • 605
 - ODBC data source configuration • 660
 - pruning • 328, 331, 604
 - Database Manager • 604
 - database protection job
 - about • 613
 - deleting • 624
 - modifying • 614
 - recreating • 625
 - starting • 623
 - decryption, data • 110
 - deduplication
 - configure groups option • 393
 - create device • 735
 - defined • 729
 - modify device • 387
 - remove device • 387
 - demote a primary server to a member server • 568
 - demote a primary server to a member server • 568
 - demote a primary server to a member server (MSCS clusters) • 924
 - demote a primary server to a member server (NEC cluster) • 933
 - Destination Tab for Archiving • 358
 - Device Configuration • 210
-

- Device Configuration, deduplication device configuration • 393, 735
- Device Wizard • 116
- RAID device configuration • 379
- Removable Storage Management, controlling devices • 380
- replacing devices • 417
- tape library configuration • 373
- virtual library configuration • 380
- device group configuration • 212
- Device Manager • 373, 395
- Device Wizard • 116
- Diagnostic Wizard • 682
 - viewing reports, Diagnostic Wizard • 686
- Discovery Configuration • 592
- DLTSage error handling
 - curing errors • 433
 - overview • 429
- Duplicate Sessions dialog • 280

E

- ejecting media • 396, 400
- email option • 710
- encryption algorithm • 111
- encryption, data,
 - about • 110
 - at the agent server • 113
 - during backups • 114
 - during migration • 114
 - encryption, data, with deduplication • 742
- engines
 - about • 491
 - Database Engine • 491, 514
 - Job Engine • 491, 500
 - service state icons • 493
 - status • 492
 - Tape Engine • 491, 503
- Enterprise Module • 40
- equivalence • 36, 539
- erasing media • 396, 398
 - options described • 398
- error log, devices • 605
- exclude objects from jobs • 302
- exclude the Arcserve database from jobs • 303

F

- failover • 908

- Federal Information Processing Standards (FIPS) • 110, 111
- filters
 - filtering options • 314
 - filters, wildcards • 312
 - job filters • 312
 - types • 315
- firewall configuration • 602
- formatting media • 396
- full backup • 307

G

- GFS media pools • 449
- GFS rotations • 122, 307, 449, 451, 452
- global backup options • 154
 - advanced backup options • 156
 - alert options • 154
 - backup media options • 164
 - global backup options, encryption options • 159
 - job log options • 184
 - media exporting options • 155
 - operation options • 167
 - verification options • 166
 - Volume Shadow Copy Service options • 162
- Global Options for Archiving • 359
 - Advanced Options • 369
 - Job Log Options • 366
 - Media Options • 360
 - Operation Options • 362
 - Pre/Post Options • 364
- global password changes • 40
- global restore options • 271
 - backup media options • 271
 - destination options • 272
 - job log options • 277
 - pre/post options • 276
- group view, Backup Manager • 140, 145, 306

I

- importing and exporting media • 404
- installation considerations
 - Microsoft SQL Server • 658
- interact with the desktop • 512
- inventory slots • 404
- inventorying slots • 407

J

- Job Log Options for Archiving • 366

- job queue, restoring • 557
- Job Scheduler Wizard • 341
 - schedule a custom report, Job Scheduler wizard • 679
- job scripts • 342
- Job Status Manager • 319
 - Job Detail tab • 333, 668
 - Job Log tab • 333
 - job queue • 323
 - job status, types • 325
 - multiple jobs, managing • 322
 - Tape Log tab • 332, 668
- job status, types • 325

L

- labeling • 404
- library functions • 404
 - cleaning tape heads • 404, 412
 - importing and exporting media • 404
 - inventory slots • 404
 - mounting and dismounting media • 404
 - online and offline libraries • 404, 415
- library groups
 - assigning slots • 420
 - creating • 419
 - deleting • 423
 - removing slots • 423
 - renaming • 424
- licensing
 - managing • 523
 - release licenses • 525
- log in to Arcserve Backup • 92
- logs • 711
 - Job Log data • 668
 - Tape Log data • 668

M

- magazines, mounting and dismounting • 409
- Manager Console
 - opening • 90
- managing domain servers • 538
- managing jobs
 - Job Detail tab • 333, 668
 - Job Log data • 668
 - Job Log tab • 333
 - job queue • 323
 - job templates • 343
- media

- cleaning • 412
- consolidation during migration • 441
- expiration Date • 404
- media maximization • 439
- online and offline removable drives • 415
- Media Assure & Scan Utility • 32
- media assure, about • 436
- Media Options for Archiving • 360
- Media Pool Manager • 373, 454
- media pools • 445
 - GFS media pools • 449
 - media pools, with deduplication devices • 754
- member server • 548
- Microsoft SQL Agent Configuration utility • 609
- Microsoft SQL Server
 - Arcserve database, backup options • 619
 - database consistency check • 661
 - installation considerations • 658
 - ODBC configuration • 660
 - SQL connections • 661
- Microsoft SQL Server 2008 Express Edition
 - Arcserve database, backup options • 617
- migration consolidation • 441
- mirrored disk • 911
- MM Admin
 - find media in vault object • 464, 472
 - overview • 373, 457, 458
 - reports object • 462
 - rotation object • 461
 - schedule object • 460
 - special tape volume movement, permanent retention • 471, 472
 - special tape volume movement, temporary check in • 471
 - status object • 464
 - vault criteria descriptor object • 460
- modify Arcserve Backup user properties • 96
- modify Windows user properties • 96
- mounting and dismounting media • 404
- move a member server to a different domain • 572
- moving media • 419
- MSCS clusters
 - change cluster domain • 927
 - change database • 924
 - delete cluster resources • 923
 - demote primary server • 924
 - failover support • 912
 - promote member server • 924
 - rebuild cluster resources • 922

- stop HA service monitoring • 917
- multiple jobs, managing • 322

N

NEC clusters

- change cluster domain • 931
- change database • 933
- demote member server • 933
- disable cluster scripts • 935
- enable cluster scripts • 938
- failover support • 912
- promote member server • 933
- stop cluster groups • 934
- stop HA service monitoring • 930

- network interface cards, configuring multiple • 526
- node tier • 521, 727

O

- online and offline libraries • 404, 415
- Operation Options for Archiving • 362
- Optimize Restore feature • 204

options

- Compare Utility options • 33
- Count Utility options • 33
- Discovery Configuration options • 591, 592
- filtering options • 314
- local backup options • 134, 152
- Media Assure & Scan Utility options • 32
- options, deduplication device group • 393
- Purge Utility options • 34
- staging options • 210, 212, 225, 226, 227

P

- password changes, global • 40
- password management • 79
- pfc command • 126
- Policies Tab for Archiving • 356
- Pre/Post Options for Archiving • 364
- preferred shares, add computers • 119
- Preflight Check Utility • 126
- promote a member server to a primary server • 565
 - promote a member server to a primary server • 565
 - promote a member server to a primary server (MSCS clusters) • 924
 - promote a member server to a primary server (NEC clusters) • 933

Q

- Quick Search • 76
- quorum disk • 911

R

raw backup

- raw backup, definition • 759
- raw backup, enabling • 763
- raw backup, licensing • 759
- raw backup, naming convention • 762
- raw backup, performing backup • 763
- raw backup, restoring, alternate location • 765
- raw backup, restoring, another device • 766
- raw backup, restoring, original location • 765

- rebuild media • 396, 401

- register virtual clusters as domain member servers • 561, 577

- removable drive support • 427

- removable media support • 425, 426

- renaming backup servers

- member server • 548

- primary server • 543

- stand-alone server • 553

- replacing devices • 417

- Report Manager • 669, 671

- custom reports • 671

- report categories • 672

- schedule a custom report, Report Manager • 679

- standard reports • 671

- Report Writer Utility • 34

- creating custom reports • 680

- Restore Manager • 257

- global restore options • 271

- restoring data

- destination options • 269

- duplicate backup sessions • 265

- Smart Restore • 204, 265

- staged backups • 280

- retensioning media • 396, 399

- roles

- about • 83

- add • 97

- remove • 98

- rotation • 471

- creating a rotation • 456

- deleting a rotation • 474

- modifying a rotation • 474

- rotations • 301, 307, 318

run as Administrator • 270

S

SAN

- backup plans • 844
- benefits • 844
- control job run time • 852
- create shared device groups • 849
- data backup and restoration • 850
- device management • 851
- environment • 842
- how the option works • 842
- installation • 847
- installation prerequisites • 845
- licensing • 841
- media management • 851
- operating system compatibility • 845
- reports and logs • 852
- server management • 844
- terminology • 844
- troubleshooting • 890
- using the option • 848
- virtual libraries • 853

save sets • 447

scanning devices • 396, 402

schedule a custom report, Job Scheduler wizard • 679

scheduling jobs • 301, 307, 317, 318

- backup jobs • 152
- restore jobs • 269

scratch sets • 447

Secure Key Manager integration • 433

serial numbers and bar codes • 448

Server Configuration Wizard

- about • 559
- change the Arcserve database application • 662, 667
- change the domain administrator password • 572
- demote a primary server • 568
- how to start • 564
- move a member server • 572
- move the ARCserve database to a different system or instance • 664
- promote a member server • 565
- tasks • 561

service state icons • 493

shared disk • 910

SnapLock • 204

SNMP alerts • 713

Source Tab for Archiving • 355

SRM PKI • 725

Staging

- architecture • 230
- configuration tasks • 233
- disabling staging • 226, 227
- Disk Staging, deduplication • 740
- features • 204
- FSD configuration • 210
- modify a schedule • 225
- multistreaming, Disk Staging Option • 204, 207
- pause migration • 204, 226
- postscripts • 220
- reports • 204
- restoring data, staging backup • 280
- rotation scheme, modifying • 225
- SnapLock • 204
- staging group configuration • 212
- submit a tape staging backup job • 236
- tape staging license requirements • 236

static backups • 198

stopping and starting services • 493, 496

Storage Area Network (SAN) • 842

Submit Archive Jobs • 351

synthetic full backup

- considerations • 836
- how makeup jobs work • 832
- how the purge policy and DB pruning jobs work • 832
- method • 837
- overview • 807
- point in time restore • 833
- specify the synthetic policy • 813
- submitting a D2D2T synthetic full backup job
 - using a custom schedule • 825
- submitting a D2D2T synthetic full backup job
 - using a GFS rot • 819
- submitting a D2D2T synthetic full backup job
 - using a rotation scheme • 120, 813
- submitting a synthetic full backup job • 811
- submitting an on-demand synthetic full backup • 831
- troubleshooting • 877, 878
 - using • 808

system account • 37, 520, 538, 539

system state restore options • 278

T

Tape Engine

- changing the system account • 520
- configuration • 503, 504, 507
- location, log file • 509
- manage the size of the log file, Circular Logging • 510, 511
- message log options • 504
- pruning, tape log file • 511

Tape Log data • 668

Tape Log tab • 332, 668

- tape movement, scheduling MM Admin
 - deleting a schedule • 468

tape staging

- architecture • 230
- migration policies • 233
- miscellaneous policies • 235
- overview • 232, 233
- postscripts • 220
- submit a backup job • 236

tape usage optimization • 439

tapecopy command • 228, 265, 309

tier • 521

U

unicode characters • 947

USB storage devices • 396, 402, 424

User Profile Utility

- adding or deleting a user • 554, 555
- assigning a user to a group • 556
- changing a user password • 555

user profiles

- administrator profile • 36
- equivalence • 36, 37

utilities • 30

- Compare Utility • 33
- Count Utility • 33
- Diagnostic Utility • 682
- Media Assure & Scan Utility • 32
- Preflight Check Utility • 126
- Purge Utility • 34
- Report Writer Utility • 34
- User Profile Utility • 35, 554

V

vault criteria descriptor (VCD) • 469

- creating a VCD • 469

deleting a VCD • 469

modifying a VCD • 469

vault criteria descriptor object • 460

vaults • 466

creating a vault • 466

deleting a vault • 468

modifying a vault • 467

virtual libraries, configuration • 380, 418

VMware • 39

W

wildcards • 312

Windows user authentication • 89

Windows-powered NAS • 346, 347

wizards

Diagnostic Wizard • 682

Job Scheduler Wizard • 341

Server Configuration Wizard • 559

WORM media • 428