

BrightStor® ARCserve® Backup for UNIX

Administrator Guide **r11.5**



Computer Associates®

D01195-2E

This documentation and related computer software program (hereinafter referred to as the "Documentation") is for the end user's informational purposes only and is subject to change or withdrawal by Computer Associates International, Inc. ("CA") at any time.

This documentation may not be copied, transferred, reproduced, disclosed or duplicated, in whole or in part, without the prior written consent of CA. This documentation is proprietary information of CA and protected by the copyright laws of the United States and international treaties.

Notwithstanding the foregoing, licensed users may print a reasonable number of copies of this documentation for their own internal use, provided that all CA copyright notices and legends are affixed to each reproduced copy. Only authorized employees, consultants, or agents of the user who are bound by the confidentiality provisions of the license for the software are permitted to have access to such copies.

This right to print copies is limited to the period during which the license for the product remains in full force and effect. Should the license terminate for any reason, it shall be the user's responsibility to return to CA the reproduced copies or to certify to CA that same have been destroyed.

To the extent permitted by applicable law, CA 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 CA be liable to the end user or any third party for any loss or damage, direct or indirect, from the use of this documentation, including without limitation, lost profits, business interruption, goodwill, or lost data, even if CA is expressly advised of such loss or damage.

The use of any product referenced in this documentation and this documentation is governed by the end user's applicable license agreement.

The manufacturer of this documentation is Computer Associates International, Inc.

Provided with "Restricted Rights" as set forth in 48 C.F.R. Section 12.212, 48 C.F.R. Sections 52.227-19(c)(1) and (2) or DFARS Section 252.227-7013(c)(1)(ii) or applicable successor provisions.

© 2005 Computer Associates International, Inc.

All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

Contents

Chapter 1: Introducing BrightStor ARCserve Backup 13

BrightStor ARCserve Backup Main Components	13
BrightStor ARCserve Backup Manager Functionality	14
BrightStor ARCserve Backup Server Functionality	14
BrightStor ARCserve Backup Core Services.....	15
BrightStor ARCserve Backup Domain Functionality.....	16
bab Command	18
BrightStor ARCserve Backup Functionality	18
BrightStor ARCserve Backup Managers	18
Back-End Services Access	19
caroot User Profile	21
caroot Equivalences List	21
User Profile Manager Toolbar.....	21
The ca_auth Command.....	22
Disk Staging Option.....	22
Disk Staging Option Overview	23
Disk Staging Option Features	24
BrightStor Portal Integration.....	26
Unicenter NSM Integration	27
Unicenter Console.....	27
Logger Email Alert Messages	28
Workload Management Integration.....	29

Chapter 2: Planning Your Storage Environment 31

Enterprise Storage Requirements	31
Budget Considerations.....	32
Network and Computer Infrastructure Requirements.....	32
Data Transfer Requirements	33
Backup Schedule Requirements	33
Data Backup Window Considerations	34
Hardware Data Transfer Rates	34
Network Bandwidth Considerations	36
Data Transfer Requirements and Resources Calculations	36
Data Path Considerations.....	37
Alternate Data Path Considerations	38
Parallel Storage Operations (Multiple Streaming).....	41
Storage Capacity Requirements.....	41

Online Recovery Data Storage Requirements	41
Backup Data Storage Requirements	42
Storage Capacities and Resources	42
Testing Plans and Assumptions	43
Catastrophic Events.....	44
Risk Assessment.....	44
Off-Site Repository Considerations	45
Disaster Recovery Archive Considerations	46
Disaster Recovery Testing	46
Sample Calculations	46
Transfer Rate for Clients and Servers on a 100Base-T Ethernet LAN With No Subnets	47
Transfer Rate for Clients and Servers on Two 100Base-T Ethernet Subnets	48
Transfer Rate for Clients and Servers on a Gigabit Ethernet Network.....	48
Transfer Rate For a Server With No Clients.....	49
Transfer Rate For Server With SAN Option.....	49
Storage Capacity For Two Sets of Recovery Data, One Full and One Incremental Backup.....	50

Chapter 3: Backing Up Data **53**

Backup Manager	53
Backup Manager Source Tab.....	54
Entire Node Backups	56
Multiple Code Pages Support	57
Node Level Options	58
Volume Backup Options	62
Object Information Tab for Volumes, Directories, and Files	62
Filter Tab at the Volume Level	62
Session Password at the Volume Level.....	63
Access Tab at the Volume Level	64
Volume Option Tab	64
Disk Staging Option.....	65
Staging Operations	66
How the Max Number of Streams Option Affects Backup and Restore Operations	66
Disk Staging Option and Rotations	67
Staging Tab	67
Staging Configuration.....	68
Backing Up Data Using the Disk Staging Option	74
Backup Manager Destination Tab.....	80
Drive and Media Selection.....	80
Multistreaming	81
Backup Session Chunking to a File System Device (FSD)	81
Method/Schedule Tab	82
Backup Scheduling Methods	83

Backup Methods	83
Custom Backup Schedules.....	84
Repeat Intervals for Custom Backups	84
Media Rules for Custom Backups.....	84
Rotation Schedules.....	87
Simple Rotation Schedules	87
GFS Rotation Schedules.....	87
Backup Job Filters	88
Backup Options	88
Miscellaneous Options	89
Session Password Options.....	90
Media Exporting Options.....	92
Verification Options.....	92
Load Options	93
Advanced Backup Options.....	94
Pre/Post Options.....	95
Backup Manager Log Options	96
Backup Manager Database Options.....	97
File Retry/Sharing Options	97
The ca_backup Command	99
Multiplexing	100
Features Supported by Multiplexing.....	101
Multiplexing Job Option	102
Preflight Checks for your Backups.....	104

Chapter 4: Restoring Data 105

Restore Manager.....	105
Restore Manager Source Tab	106
Restore Find Functions	107
Find Button.....	107
Version History Button	109
Duplicate Backup Sessions.....	109
Smart Restore.....	110
Find Files in Current Path.....	111
Restore Methods	112
Restore Data Backed Up Using Staging	113
The Role of the Database in the Restoration Process	114
Multiple Code Pages	114
Specify Code Pages in the Restore Manager Window	114
Restore Manager Destination Tab	115
Agent View Tab	116
Restore Manager Object Information Tab	116

Restore Manager Security Tab	117
Restore Job Filters.....	117
Restore Options	117
Pre/Post Options.....	118
Restore Manager Log Options	120
Restore Manager Media Rules Options	121
Destination Options	121
The ca_restore Command	123

Chapter 5: Customizing Your Jobs 125

Job Filters	125
Filtering Precedence.....	126
Include Filters	128
Exclude Filters.....	128
File Pattern Filters.....	128
Directory Pattern Filters.....	129
File Attributes Filters	131
File Modified Filters	131
File Changed Filters	132
File Accessed Filters	133
File Size Filters.....	133
Rotation Methods	134
Choose a Backup Method	134
Simple Rotation Schedule	135
GFS Rotation Schedule	136
Custom Rotation Schedule	136
Job Scheduling and Submission.....	140
Submit a Job Using the Run Now Option	141
Schedule Option	141
Submit Job on Hold Option.....	142
Job Scripts	142
Script Development	143
Add a Script to the Job Queue	143
Generic Job Manager	144
Submit a Job Using the Generic Job Manager	145

Chapter 6: Using the Job Status Manager 149

Servers and Groups Configuration	149
Job Queue	150
Job Management.....	151
Job Management Using the File Menu	151

Job Management Using the Manager Console	151
Job Status Manager Toolbar	151
Job Summary Report	156
Job Log Tab	156
The ca_qmgr Command	157

Chapter 7: Managing Devices and Media 159

Device Manager	159
Manager Console	160
Device Manager Views	160
Device Manager Summary Tab	161
Device Manager Detail Tab	162
Device Manager Report Tab	163
Device Groups	169
Device Manager Toolbar	169
Multistreaming	171
Automated Media Spanning	171
Slot Sharing Among Groups	171
Format Media Option	172
Expiration Dates	173
Erase Media Option	174
Methods for Copying Media	175
Tapecopy Tool	175
Retention Media Option	177
Compression Option	177
Eject Media Option	178
Enable/Disable Option	178
Mount/Dismount Option	178
Load/Unload Option	179
Import/Export Option	179
Clean Tape Heads Option	180
Online/Offline Option	180
Enable Automatic Performance Tuning	181
Large Library Support	181
The ca_devmgr Command	182
Media Pool Manager	182
Save Set	183
Scratch Set	184
Serial Numbers	185
Simple and GFS Rotation Media Pools	186
Media Pool Manager Toolbar	187
Create a Media Pool	188

Modify Media Pools	189
Move Media	189
Assign Media to a Media Pool.....	190
Delete a Media Pool	190
Maximum Number of Sessions on a Single Tape	190
The ca_dbmgr Command	191
Media Management Administrator.....	191
Media Management and Tape Service	192
How the MMO Works with Domains	192
Media Management Administrator Terms.....	192
MMO Admin Window	193
MMO Admin Toolbar	194
Log in to the MMO Admin.....	194
MMO Admin Features and Functionality.....	195
Maximum Number of Reports	200
Find Media in Vault Object	200
Current Status Object	201
Vault Processing Status Changes	201
How the Media Management Process Works	201
How the Vault Cycle Process Works	202
Vault Cycle Simulation.....	202
Vault Cycle Reports.....	203
Tape Volume Movement Scheduling	205
Tape Volume Retention Rules	208
View Slot Information.....	210
Special Tape Volume Movement.....	211
Find Media in Vaults.....	213
Back Up and Restore the MMO Primary Server.....	213
How the MMO Primary Server Works	214
Add a Server to the BrightStor ARCserve Backup MMO Domain	214
Back Up Primary MMO Member Server Data.....	215
Restore Primary MMO Member Server Data to an Alternate MMO Member Server	216
Demote an MMO Primary Server to an MMO Member Server	219
Reinstitute the MMO Primary Member Server	219
Resynchronize an MMO Member Server.....	219
Remove an MMO Member Server from the BrightStor ARCserve Backup MMO Domain	220
The ca_mmomgr Command.....	222
Configure Your Firewall to Optimize Communication	222

Chapter 8: Managing the Database and Reporting 225

databaseBackup.log	225
Database Manager	226

Database Manager Toolbar.....	226
Configure Database Options	226
Add Clients	227
Record Modification.....	228
Delete Records.....	228
Database Views.....	229
The ca_dbmgr Command	233
Recover and Maintain the Advantage Ingres Database.....	234
Recover the Advantage Ingres Database Offline	234
Recover the Advantage Ingres Database Online	236
Best Practices for Maintaining the Advantage Ingres Database	238
Compress the Advantage Ingres Database using the ca_dbadmin Command	246
Extend Database Utility (extend_db)	247
Merge Manager.....	249
Merge Manager Toolbar	250
Merge Manager File Menu	250
Merge Manager Window.....	251
Merge Options.....	252
Submit a Merge Job	256
The ca_merge Command	256
Scan Manager	257
Scan Manager Toolbar.....	257
Scan Manager File Menu	258
Scan Manager Window	258
Scan Manager Object Information Tab	259
Scan Options	259
Submit a Scan Job	261
The ca_scan Command.....	262
Report Manager	263
Report Manager Toolbar	264
Report Manager File Menu.....	264
Report Manager Logs and Reports.....	265
Report Manager Database Reports	267
The ca_log Command	268
Diagnostic Utility.....	269
Diagnostic Wizard.....	269
cadiag Command Line Utility	269
Run the Diagnostic Utility.....	270
BrightStor Portal Reporting.....	271

Chapter 9: Recovering Data from a Disaster 273

Disaster Recovery Data.....	273
-----------------------------	-----

Manual Disaster Recovery for Sun Solaris	274
Prepare for Disaster Recovery on Sun Solaris Systems	274
Recover From a Disaster on Sun Solaris Systems	274
Manual Disaster Recovery for HP-UX	278
Prepare for Disaster Recovery on HP-UX Systems.....	279
Recover From a Disaster on HP-UX Systems	279

Appendix A: Using Command Line Utilities **285**

Available Command Line Utilities	285
Usage, Syntax, and Argument Information	287
bab Command	288
bab Syntax	288
bab Options	288
ca_auth Command	289
ca_auth Syntax	289
ca_auth Options	290
ca_backup Command	291
ca_backup Syntax	292
ca_backup Oracle RMAN Specific Syntax	292
ca_backup Usage.....	292
ca_dbadmin Command	309
ca_dbadmin Syntax	309
ca_dbadmin Options	310
Compress the Advantage Ingres Database using the ca_dbadmin Command	312
ca_dbadmin Hidden Options.....	313
ca_dbmgr Command	313
ca_dbmgr Syntax	313
ca_dbmgr Options	314
ca_devmgr Command.....	317
ca_devmgr Syntax.....	317
ca_devmgr Usage.....	317
Staging Command Line Query Tool.....	328
Staging Command Line Purge Tool	329
ca_generic Command	330
ca_generic Syntax	331
ca_generic Global Options.....	331
ca_jobstat Command.....	334
ca_jobstat Syntax.....	334
ca_jobstat Options.....	334
ca_log Command	335
ca_log Syntax	336
ca_mediarep Command	338

ca_mediarep Syntax	338
ca_mediarep Options	338
ca_merge Command	339
ca_merge Syntax	339
ca_merge Usage.....	340
ca_mmomgr Command.....	343
ca_mmomgr Syntax.....	343
ca_mmomgr Options.....	343
ca_qmgr Command	346
ca_qmgr Syntax	347
ca_qmgr Commands	347
ca_recoveryrep Command	349
ca_recoveryrep Syntax.....	349
ca_recoveryrep Options.....	349
ca_restore Command	350
ca_restore Syntax	350
ca_restore Oracle RMAN Specific Syntax	351
ca_restore Usage.....	351
ca_scan Command	358
ca_scan Syntax	358
ca_scan Usage	358
ca_stagingrep Command.....	363
ca_stagingrep Syntax.....	364
ca_stagingrep Options.....	364
ca_summaryrep Command.....	365
ca_summaryrep Syntax.....	366
ca_summaryrep Options.....	366
ca_utilizationrep Command	367
ca_utilizationrep Syntax	367
ca_utilizationrep Options	368
ca_vaultrep Command	369
ca_vaultrep Syntax.....	369
ca_vaultrep Options	369
cadiag Command	370
cadiag Syntax	370
cadiag Options	370
pfc Command.....	371
pfc Syntax.....	371
pfc Options	372
tapecopy Command.....	373
tapecopy Syntax.....	373
tapecopy Options.....	374

Appendix B: Acknowledgements	379
RSA Data Security, Inc. Acknowledgement	379
Apache Acknowledgement	380
OpenSSL Acknowledgement	384
 Index	 389

Chapter 1: Introducing BrightStor ARCserve Backup

BrightStor® ARCserve® Backup is a comprehensive, scalable 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, NetWare, and Linux) using optional client agents. BrightStor ARCserve Backup also provides media and device management utilities.

BrightStor ARCserve Backup provides a Java manager to manage BrightStor ARCserve Backup servers through Mozilla or Internet Explorer browsers. It can support small-scale or large-scale enterprise environments comprised of one machine or many, across different platforms and organizations.

BrightStor ARCserve Backup Main Components

BrightStor ARCserve Backup comprises the following two main components that work together to back up, copy, and restore your data:

- BrightStor ARCserve Backup Manager
- BrightStor ARCserve Backup Server

BrightStor ARCserve Backup Manager Functionality

Use the BrightStor ARCserve Backup Manager to perform tasks such as submitting backup and restore jobs, managing your database, and searching reports. You can control all BrightStor ARCserve Backup operations from a single machine.

The manager includes Java classes, icons, images, the httpd daemon (back-end service that attends to various tasks without human intervention), and the Communicator program which functions as an information bridge between the browser and back-end services.

The core components of the BrightStor ARCserve Backup Manager are as follows:

- Web Server (httpd)—Serves the BrightStor ARCserve Backup Java-based interface.
- cacommd—Behaves as a gateway and controls all requests between the BrightStor ARCserve Backup Java-based interface and the BrightStor ARCserve Backup server.

You can stop and start these services using the stopgui and startgui commands.

BrightStor ARCserve Backup Server Functionality

The server runs on a UNIX-based system and consists of back-end services that process all jobs and update information about these jobs in the activity logs and database. The core components of BrightStor ARCserve Backup server are:

- Discovery Service (cadiscovd)—Collects and maintains information on the availability of the BrightStor ARCserve Backup servers on the network, and supplies this information to clients requesting access to BrightStor ARCserve Backup resources. In addition, it maintains the user information obtained from the BrightStor ARCserve Backup servers on the network.
- BrightStor ARCserve Backup Loader (caservd)—Loads and unloads all BrightStor ARCserve Backup services.
- Authentication Service (cauthd)—Provides a centralized point of control for security issues concerning the entire network of resources made available by BrightStor ARCserve Backup. It verifies and controls the actions performed on BrightStor ARCserve Backup resources using a predefined set of privileges (permissions).

- Job Scheduler (caqd)—Manages the jobs scheduled through the manager or command line. Related to this service is the cprocess service. The cprocess service processes jobs relating backup, restore, merge, and scan jobs. For example, during a back up, the cprocess service communicates with the source (agent) and destination (media), and transfers the data from the agent to the media.
- Media Server (camediad)—Handles the transfer of data to and from media.
- Database Server (cadbd)—Keeps track of the jobs run and all the files backed up or restored by BrightStor ARCserve Backup.
- Event Logger (cloggerd)—Handles message logging and message retrieving. You can specify to send messages to numerous destinations, such as printers or email.
- dbclean—Purges and prunes BrightStor ARCserve Backup database information.
- Mergecat—Merges backup data information in the BrightStor ARCserve Backup database.

BrightStor ARCserve Backup Core Services

The following services work in the background to perform the workload for the BrightStor ARCserve Backup Managers:

- The Queue Service processes jobs at a specified date and time, using the Job Engine to scan the job queue for a job that is ready to run and send it to the appropriate handler.

Note: If your jobs start running an hour or so before or after their scheduled time, there may be a discrepancy between your Java GUI manager and your BrightStor ARCserve Backup server system time. This can occur because of daylight saving changes or if there is a change in the time zone, date, or time on the system when the BrightStor ARCserve Backup server is up and running. To fix this, recycle caqd using the following commands:

```
bab -unload caqd
bab -load caqd
```

- The Media Service communicates with, and controls, storage devices, using the Tape Engine to select the device needed for a job.
- The Database Service stores information contained in the database and makes it available for centralized reports, job logs, and management of shared devices in your BrightStor ARCserve Backup environment.

Database information includes:

- Files
- Directories
- Drives
- Machines that BrightStor ARCserve Backup has backed up or copied
- Records of the job types
- Logs
- Final results
- The starting and ending time of jobs processed by BrightStor ARCserve Backup
- Information about media used by BrightStor ARCserve Backup, such as type, name, date of first formatting, expiration date, and the sessions it contains.

You can monitor and control each of these services from one of the BrightStor ARCserve Backup managers, or from the BrightStor ARCserve Backup command line, using the bab command. For more information on BrightStor ARCserve Backup managers, see BrightStor ARCserve Backup Managers in this chapter.

BrightStor ARCserve Backup Domain Functionality

BrightStor ARCserve Backup domains are logical collections, or groupings, of servers that enable easier administration of BrightStor ARCserve Backup servers and users. BrightStor ARCserve Backup domains provide the following:

- A single logon to multiple servers
- The same access to all the servers in the domain for the same user
- Security and efficient management of the entire BrightStor ARCserve Backup network

Each BrightStor ARCserve Backup domain has a name and a list of servers belonging to it. You can manage the domain to select any server from the BrightStor ARCserve Backup domain on which to perform database management, tape and device management, and backup strategy and schedule management, without having to log into each BrightStor ARCserve Backup server separately. You can perform a particular operation on one server and can perform the same operation on any server in the domain.

In addition, each domain has a mandatory designated primary server, and an optional secondary server. The primary server synchronizes information to the secondary server, which provides fault tolerance in the event of a primary server failure. BrightStor ARCserve Backup users need not be administrators (root superuser) at an operating system level. Users can start and stop BrightStor ARCserve Backup services on any server in the BrightStor ARCserve Backup domain.

When you open one of the BrightStor ARCserve Backup managers from the home page, the server that appears in the host server field is the machine that will run all your operations. Because you can select different host servers from the manager window, you can manage multiple BrightStor ARCserve Backup servers from one or more locations. From here you can select any of the BrightStor ARCserve Backup managers.

The host server determines the available devices for backup and restores jobs. The device groups attached to the host server are those that are available to you for the backup or restore destination. If the currently selected host server does not contain the media device groups you want to use, you can change the host server to one that does.

Modify BrightStor ARCserve Backup Domain Settings

If you want to modify domain settings, such as the primary server, secondary server, and domain name, use the following procedure:

1. Make sure that no jobs are running, and then use the following command to stop the BrightStor ARCserve Backup services:

```
cstop
```

2. Run the bab_configure script.
3. When prompted, specify your primary and secondary servers and the BrightStor ARCserve Backup domain name.
4. Use the following command to start the BrightStor ARCserve Backup services:

```
cstart
```

bab Command

Use the bab command to monitor and control BrightStor ARCserve Backup services.

You can use the bab command options and switches to start or stop BrightStor ARCserve Backup services, check the status of each service, view the configuration of a selected service, or instruct BrightStor ARCserve Backup to re-configure the BrightStor ARCserve Backup service or services.

For a complete list of the options and switches for this command, see the appendix "Using Command Line Utilities."

BrightStor ARCserve Backup Functionality

BrightStor ARCserve Backup provides the components and functions required by Network Managers to obtain and manage network operations. When you start BrightStor ARCserve Backup, the BrightStor ARCserve Backup Home Page opens. From the Home Page, you can access any of the BrightStor ARCserve Backup managers.

BrightStor ARCserve Backup Managers

BrightStor ARCserve Backup managers provide the front-end interface used to perform all BrightStor ARCserve Backup functions. You can access these managers from the home page. The following is a list of the managers and the functions they perform:

- **Backup Manager**—Schedules and submits backup jobs on any machine in the BrightStor ARCserve Backup domain.
- **Device Manager**—Displays information about storage devices and media, and lets you compress, format, erase, and eject media. You can also manage media rotations using this manager by assigning unique serial numbers.
- **Job Status Manager**—Monitors, reschedules, and submits all pending, completed, and active jobs. Log information is provided for each completed job.

- Restore Manager—Schedules and submits restore jobs to perform a complete or partial restore of backed up data.
- Database Manager—Maintains information such as the jobs processed by BrightStor ARCserve Backup, the media used, data applications backed up, the devices you are using, error logs, session information, disk usage, and client information.
- Report Manager—Interfaces with the BrightStor ARCserve Backup database to supply the reports you need about BrightStor ARCserve Backup activity. This includes job logs, media error logs, session information, job scheduling logs, and user logs.
- Generic Job Manager—Schedules generic jobs to run on the BrightStor ARCserve Backup server.
- Merge Manager—Merges information from BrightStor ARCserve Backup media into the BrightStor ARCserve Backup database.
- Scan Manager—Scans media for information on your backup sessions. You can scan a single session or the entire media. You can view results of the media scan in the Report Manager under the Activity Log listing, or under the User Log listing if an additional log file is created.
- User Profile Manager—Lets the BrightStor ARCserve Backup root user customize user access to BrightStor ARCserve Backup.
- Media Management Admin—Provides the tools you need to organize tape movement to off-site storage locations and protect, control, and manage media resources.
- Media Pool Manager—Creates logical groupings of media to quickly identify and maintain the retention of its data. You can also choose specific media rotation schemes to suit your archival needs.
- Tapecopy Tool—Enables you to copy from tape to tape. Use this tool to duplicate tapes for off-site storage. You can make copies between different kinds of media and copy specific sessions.

Note: You cannot copy data to VM:Tape media using the Tape Copy Manager.

Back-End Services Access

You can access back-end services from a BrightStor ARCserve Backup manager window. You can use the BrightStor ARCserve Backup File Menu or the status icons found on the BrightStor ARCserve Backup Status Bar.

How You Can Control Back-End Services Using the File Menu

In addition to controlling BrightStor ARCserve Backup services from the command line, using the bab command, you can control the Queue services, Media services, and Database services from any of the managers. You can also start or stop all of the back-end services running on a host server by selecting Start All Services or Stop All Services.

From the BrightStor ARCserve Backup File Menu, select the back-end service you want to control:

- If you select Start All Services or Stop All Services, BrightStor ARCserve Backup applies the elected action to the services (Queue, Media and Database) on the current host server.
- If you select one of the individual services, you can use a drop-down menu to select Start Service or Stop Service, depending on the required action.

Status Bar Service Icons

 Service Icons—The Status Bar at the bottom of each BrightStor ARCserve Backup manager displays an icon for each of the back-end services:

- Queue service (left icon)
- Media service (middle icon)
- Database service (right icon)

Depending on their color, the icons indicate one of three states:

- Green: The service is running.
- Yellow: The cacommd daemon is either not running or that the connection to it is broken
- Red: The service is stopped or not running

You can use the status bar icons to control the selected service. Click an icon to stop the service, stop or start all services, or view detailed information about the status of the service you selected.

caroot User Profile

The default caroot superuser profile has root privileges for all BrightStor ARCserve Backup functions. You can set a password for the caroot profile during the configuration of the software, or after configuration using the User Profile Manager. You can also create additional user profiles using the User Profile Manager.

Note: BrightStor ARCserve Backup User Names control access only to BrightStor ARCserve Backup-related functions, and should not be confused with the operating system-required login name and password.

caroot Equivalences List

By creating an equivalence list, all clients can use BrightStor ARCserve Backup without requiring the user to log into the domain. BrightStor ARCserve Backup can validate whether the current user has equivalent access to the domain. Effectively, the OS access guarantees BrightStor ARCserve Backup access.

For example, define user *xyz@machine1* as equivalent to user caroot on BrightStor ARCserve Backup domain ABC (caroot is the predefined root, administrator, user in the BrightStor ARCserve Backup domain), then when user *xyz* logs into system *machine1*, the user can execute BrightStor ARCserve Backup commands as a root user of the BrightStor ARCserve Backup domain, and is automatically authenticated.

User Profile Manager Toolbar

Administrators can change user passwords and add and delete user profiles, if necessary, using the buttons on the User Profile Manager. The following features are available from the User Profile Manager toolbar:



Add User—Opens the Add User dialog to create a new user profile. Enter the user's name and a password, confirm the password, and click OK.



Change User Password—Opens the Change User Password dialog. Administrators can use this dialog to change another user's password. Enter the old and the new passwords, confirm the new password, and click OK.



Delete User—Opens the Delete User dialog to delete a user profile. Select a user and click the Delete User button. BrightStor ARCserve Backup prompts you to confirm that you want to delete this user. Click OK to confirm.



Refresh—Refreshes the information displayed in the User Profile Manager.

The ca_auth Command

Use the `ca_auth` command to create new users, delete existing users, change user passwords, and establish authentication equivalences associated with a particular user from the command prompt. This command provides an alternate method of using BrightStor ARCserve Backup without accessing the User Profile Manager.

You can perform the following functions using the options and switches available for the `ca_auth` command:

- User Manipulation options—Use these options to add, modify, or delete a user account.
- Equivalence Management options—Use these options to set equivalences for users.


For a complete list of the options and switches available for this command, see the appendix “Using Command Line Utilities.”

Disk Staging Option

The Disk Staging Option allows you to 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, the Disk Staging Option also allows you restore data directly from the staging area.

The Disk Staging Option is basically a two-part data backup process.

- Backup Process-Backs up data from the source to the staging area.
- Copy Process-Copies or migrates the backed-up data from the staging area to the final destination.



Disk Staging Option Overview

First, this option allows you to back up data to a file system device (FSD), which is used as a temporary staging area. A staging job can divide your backup job into several sub-jobs that run simultaneously. The Disk Staging Option allows you to utilize simultaneous streaming to send multiple streams of data to the FSD at the same time. Since the work is split up among several different drives, 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 it. 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 backup-to-FSD process, if the FSD gets full or reaches the specified volume threshold, the Disk Staging Option allows you to 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, the Disk Staging Option also allows you to 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, the Disk Staging Option automatically purges the data from the FSD, and reclaims valuable disk space so that additional backups can continue.

For rotation jobs or GFS rotation jobs, the Disk Staging Option allows you to specify policies to 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.

Disk Staging Option Features

In addition, the Disk Staging Option provides or supports the following features:

- File System Device Capacity Management--The Disk Staging Option allows you to specify capacity thresholds of the file system device. The threshold can be represented as either the absolute value or as a percentage of the total volume capacity.
- Ensures that BrightStor 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 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 direct BrightStor ARCserve Backup to create a makeup job to back up directly to tape if an exceeds threshold condition occurs, and to create a makeup job on hold if a data migration failure occurs.
- Allows you to perform full, incremental, and differential backups.

- **Pause Data Migration**--The Disk Staging Option allows you to 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 sub-jobs that run simultaneously. The Disk Staging Option 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 sessions (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.

Note: The Disk Staging Option provides you with the capability of streaming multiple jobs simultaneously to the FSD. The unlicensed Disk Staging Option allows you to stream two jobs simultaneously. To stream more than two jobs, you must license the Disk Staging Option. After you license the Disk Staging Option, you can stream up to 32 jobs simultaneously to the FSD.

- **SnapLock Support**--SnapLock™ is technology from Network Appliance that provides non-erasable, non-rewritable, Write Once Read Many (WORM) data protection. The Disk Staging Option allows you to enable 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 can not 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, BrightStor ARCserve Backup write-protects the data, however, the data can be deleted from the device.

- **Copy Image Tracking**--BrightStor 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), the Disk Staging Option 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--BrightStor ARCserve Backup provides a transparent Smart Restore feature, which is further enhanced by the Disk Staging Option feature of providing multiple locations for the backed-up data. If during the restore process from either the FSD or from the final destination, a media or drive error occurs, BrightStor ARCserve Backup internally finds the alternate media and starts restoring the data from the alternate media. This increases the success rate of restores in the event of any hardware problems.
- Optimize Restore Option--If, during a restore operation, BrightStor 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 BrightStor ARCserve Backup to restore the data from the session that resides on the file system device.
- Command Line Support--BrightStor ARCserve Backup allows you create backups to an FSD 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 FSD, you can use the Purge tool to remove data and free extra space on the FSD.
- Disk Staging Option Reports--BrightStor ARCserve Backup provides the capability to generate additional reports dedicated to the Disk Staging Option. Using these reports you can find the backup to disk status of every session, 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.

BrightStor Portal Integration

BrightStor ARCserve Backup uses the iGateway and iSponsor components to facilitate communication with BrightStor Portal. These components let you collect information from BrightStor ARCserve Backup and manage it using BrightStor Portal.

BrightStor Portal lets you view this information from any web browser, empowering administrators to make rule-based decisions without installing the BrightStor ARCserve Backup Manager on their local machines.

The iGateway and iSponsor components are installed, by default, in the opt/CA/SharedComponents/iTechnology directory.

For more information about BrightStor Portal, see the BrightStor Portal Documentation.

Note: iGateway and iSponsor are supported on all UNIX-based platforms, with the exception of Tru64 platforms.

Unicenter NSM Integration

BrightStor ARCserve Backup integrates with the Unicenter Console and Workload Management components of Unicenter® Network and Systems Management (NSM) (formerly known as Unicenter® TNG). The following sections include information on integration with each of these components.

Unicenter Console

BrightStor ARCserve Backup offers the ability to route messages from BrightStor ARCserve Backup jobs to the Unicenter Console using the BrightStor ARCserve Backup Unicenter NSM or Simple Network Management Protocol (SNMP) Alert features. This simplifies overall system management. iGateway and iSponsor are available for all supported UNIX platforms so that you can view messages from both BrightStor ARCserve Backup and Unicenter from one central location.

Note: The Unicenter Console component on UNIX does not have a graphical user interface (GUI). If you want to access it using a GUI, you must use a Microsoft-based computer (for example, using Unicenter EM Classic).

The following sections include information on both Unicenter NSM Alert and SNMP Alert.

Unicenter NSM Alert Considerations

Review the following considerations before using Unicenter NSM Alert:

- To use Unicenter NSM Alert, you must install CA Common Services™ or a full Unicenter NSM installation on the same system.
- After you install, update the **UNICENTER_NSM_NODES** parameter in the **caloggerd.cfg** configuration file to enable the Unicenter NSM Alert feature.
- If you want to send BrightStor ARCserve Backup messages to Unicenter NSM machines, include a whitespace-separated list of these machines in the **UNICENTER_NSM_NODES** parameter in the **caloggerd.cfg** configuration file.
- If you make a change to the **caloggerd.cfg** configuration file after BrightStor ARCserve Backup starts, you must reload the BrightStor ARCserve Backup logger daemon for the change to take effect.

Messages from BrightStor ARCserve Backup jobs in the Unicenter NSM Console will have a prefix **CABAB** to distinguish them from other Unicenter component messages.

SNMP Alert

To use SNMP Alert, you must set up the Unicenter Console to receive SNMP trap messages. After you set this up, BrightStor ARCserve Backup parses the system SNMP configuration file to get the community name and the trap server name. When you submit a job using the SNMP Alert option, its related messages are sent to the configured trap server using the User Datagram Protocol (UDP) and to the standard SNMP trap port 162.

Usually the Network Management Stations are configured to receive the trap messages only from a specific "**COMMUNITY**." The default community is public. As a result, you must modify the standard SNMP configuration file on the BrightStor ARCserve Backup server to reflect this community name and the trap server name.

Examples:

Note: The IP address in the following examples can be replaced either with the host name or the IP address of SNMP trap server. The community name can be replaced with the one that is known to trap server.

- **SUN SPARC**—in `/etc/snmp/conf/snmpd.conf`, add lines similar to the following:

```
trap      172.16.0.0  172.16.0.0
trap-community public
```

- **HP-UX 11.0, 11.11**— in `/etc/snmpd.conf`, add lines similar to the following:

```
get-community-name: public
trap-dest: 172.16.0.0
```

- **IBM AIX**—in `/etc/snmpd.conf`, add lines similar to following:

```
trap      public      172.16.0.0  1.2.3  fe # trap server
```

- **TRU64/Digital**—in `/etc/snmpd.conf`, add lines similar to following:

```
trap public 172.16.0.0
```

Logger Email Alert Messages

If a critical event occurs, BrightStor ARCserve Backup can send an email alert notification to a customizable list of email addresses. In addition to the email message itself, this feature can attach a job-specific file to the email alert notification.

You can customize the list of email addresses by modifying the `caloggerd.cfg` file located in the `$BAB_HOME/config` directory.

The following is an example of the message contents of an email alert notification:

```
-----
                        BrightStor ARCserve Backup Event
-----
Time       : 12/04/03 12:56:20
Job ID    : 28
Session   : 12
-----
GRP1: Load media
-----
```

BrightStor ARCserve Backup sends a file attachment to the email alert if there is a specific file that corresponds to a component. For example, if BrightStor ARCserve Backup performs a full database backup, the event triggers an email alert message, attaches the database backup log (databaseBackup.log) to the email alert notification, and then sends the email to all of the addresses specified in the caloggerd.cfg file. In this case, the databaseBackup.log is needed for disaster recovery.

To create and modify the list of email addresses, access the EMAIL_ALERT section of the caloggerd.cfg file and use the syntax displayed below. You must separate each email address by one space character.

```
EMAIL_ALERT = name01@companyA.com name02@companyB.com name03@companyC.com
```

Note: Although BrightStor ARCserve Backup supports specifying an unlimited number of email addresses, you should specify the absolute minimum number of email addresses. If, for example, the alert notification requires a large file attachment, the process of sending the message and its attachment to all email addresses can be a performance bottleneck.

Workload Management Integration

BrightStor ARCserve Backup integrates with Workload Management when you submit a backup job from the command line using the utilities `ca_backup -waitForJobStatus` and `ca_merge -waitForJobStatus`. For example:

```
ca_backup [-cahost <hostname>] [global options] [global filters] {source args}
[destination args] [schedule args] [run job args] -waitForJobStatus
```

When you use the command line utilities, BrightStor ARCserve Backup waits until the operation is completed, and then exits with a return code that indicates the success or fail outcome of the job:

- **0**—successful
- **1**—failed
- **2**—incomplete

- **3**—cancelled
- **4**—unknown
- **5**—active

For more information about `ca_backup` and `ca_merge`, see the appendix “Using Command Line Utilities.”

Chapter 2: Planning Your Storage Environment

Protecting your data and managing your backup storage is fundamentally a policy issue rather than a technical problem. Technology can implement policy, but it cannot tell you what your policy should be.

Before you can use BrightStor ARCserve Backup software effectively, you need to analyze your organization's data storage requirements. You need to do the following:

- Understand how your organization's data resources are used.
- Understand how security and availability at any given time can affect your corporation's bottom line.
- Develop a comprehensive, high-level storage plan before you purchase additional hardware or configure BrightStor ARCserve Backup.

After you have a clear idea of your storage needs, this chapter can help you to develop a comprehensive implementation plan that allows for:

- Fast recovery of user-deleted files and directories, and database-related data.
- Centralized, single-point backup administration for networked systems.
- Backup operations that do not interfere significantly with normal business operations.
- Adequate quantities of media and adequate numbers of devices for your needs.
- Full recovery from catastrophic data loss.

Enterprise Storage Requirements

To determine your need for vault space, storage hardware, and storage media, you have to translate your high-level plan into a set of concrete requirements. You need to decide:

- How much you have to spend on media, hardware, and network improvements?
- How much data you really need to protect?
- When can you run backups without interfering with other work?

- How much traffic your network can handle during backup periods?
- How long you can wait for an average file or file system to be restored following a data loss?

The following sections discuss these issues in more detail.

Budget Considerations

Sometimes it pays to stress the obvious early in the planning of a major project: each of the parameters discussed in this chapter comes with a price tag attached. If you need speed, you need a faster, higher-bandwidth network and more and faster backup devices. Both require premium prices.

To meet your speed or data security requirements, you may need to buy more media. Media elements are surprisingly expensive, particularly for newer and faster backup devices.

You need to decide how much your organization can afford:

- To spend on a backup and recovery solution
- To lose in lost data and staff time

Then, do the following:

- Decide what you are prepared to do in order to keep both kinds of costs in bounds.
- Decide whether performance or economy is your primary concern.
- Evaluate the trade-offs discussed in the next section in light of this initial decision.

Network and Computer Infrastructure Requirements

If you have not already done so, you should familiarize yourself with the hardware, network, and site configuration that your backup and recovery plan supports. You should know:

- The numbers and types of computers and workstations you need to back up.
- The identities of computers that have media libraries or devices attached (these are the BrightStor ARCserve Backup servers).
- The type of SCSI or fiber cabling connecting each library to its server and the transfer rate of the cabling.
- The type of library on each server.
- The type of devices in each library and their transfer rate.

- The degree of data compression that you plan to use, if any.
- The types and capacities of your network, subnets, routers, and so on.

Data Transfer Requirements

The overall data transfer rate for your backup and recovery system sets the amount of time required for storage operations. You have to balance your backup window, backup data, and recovery speed requirements against the capabilities of your existing infrastructure and the budgetary constraints of your organization.

After you have quantified the amount of data that you have and the times when you can back it up, you can roughly estimate the minimum data transfer rate that you must achieve to fully back up the data in the allotted time. Use this requirement as a starting point for the decisions you make later in this chapter.

To calculate a rough, minimum transfer rate, divide the amount of data by the amount of time available to back up the data:

$$\text{databackedup} \div \text{backup_window} = \text{required_rate}$$

For example, if you have 1 Terabyte to back up and 5 hours available each night and you intend to back up everything in one session, you need to achieve a rate of 200 GB per hour.

Backup Schedule Requirements

The more data you have, the more time, hardware, media, and network bandwidth you require.

You need to decide:

- Whether you need to back up user data only.
- Whether you must also include system configurations and installed applications.
- Estimate the total size for the data that you must back up, allowing a reasonable margin for growth based on past experience in your organization.

Data Backup Window Considerations

As well as the amount of data that you have to back up, your infrastructure and management requirements will depend on the time that is available for backup operations in any given period. Ask yourself the following questions:

- Can you run backups during non-working hours, at night or on weekends?
- Do you have to run backups concurrently with normal business operations because your network is in use round the clock?

Identify the blocks of time that are available during the day and the week. If your organization shuts down for any long periods during the month or year, you might consider these times as well.

Hardware Data Transfer Rates

Your backup hardware is unlikely to be a limiting factor in reaching your target data transfer rate. Most devices are very fast. However, you should evaluate hardware speed at the planning stage. At a minimum, you must have enough hardware, or fast enough hardware, to write your data to storage media within the time allowed. Smaller numbers of fast devices or larger numbers of slower devices can often achieve the same total throughput. Use the information that follows to estimate the aggregate data transfer rate for your hardware.

SCSI or Fibre Interface Considerations

No device is faster than its connection to its data source. Current backup devices connect using standard SCSI or fibre interfaces. The following table lists the common varieties.

Version	Bus Width	Approximate Maximum Data-transfer Rate
Wide Ultra SCSI	16 bits	40 MB/seconds=144 GB/hour
Ultra2 SCSI	8 bits	40 MB/seconds=144 GB/hour
Wide Ultra2 SCSI	16 bits	80 MB/seconds=288 GB/hour
Ultra 160 SCSI	16 bits	160 MB/seconds=576 GB/hour
Ultra 320 SCSI	16 bits	320 MB/seconds=1152 GB/hour
Fibre Channel	1 Gb	100 MB/seconds=360 GB/hour
Fibre Channel	2 Gb	200 MB/seconds=720 GB/hour

You can see that many of the SCSI interfaces and fibre interfaces will be able to handle your requirement of 200 GB per hour. For example, if you are using a Wide Ultra2 SCSI you can achieve 200 GB in less than an hour. Even if you are using a slower SCSI controller you can use multiple SCSI controllers to achieve the aggregate data transfer rate of 200 GB per hour.

Obviously, the SCSI bus or fibre interface should seldom limit your ability to achieve your required data transfer rate. Any of these SCSI varieties could easily meet the 40 GB per hour requirement in our example. Indeed, most could handle the whole 200-GB job in under two hours. A Wide Ultra 160 SCSI could do it in about 30 minutes.

Tape Drive Considerations

There are many kinds of devices. A few of the most common are listed in the following table.

Device type	Approximate Transfer rate 2:1 (compressed data)	Maximum Capacity (compressed data)
DDS-4	6.0 MB/seconds=21.0 GB/hour	40 GB
AIT-2	12.0 MB/seconds=43.2 GB/hour	100 GB
AIT-3	31.2 MB/seconds=112.3 GB/hour	260 GB
DLT 7000	10.0 MB/seconds=36.0 GB/hour	70 GB
DLT 8000	12.0 MB/seconds=43.2 GB/hour	80 GB
Super DLT	24.0 MB/seconds=86.4 GB/hour	220 GB
Mammoth-2	24.0 MB/seconds=86.4 GB/hour	160 GB
Ultrium (LTO)	30.0 MB/seconds=108.0 GB/hour	200 GB
IBM 9890	20.0 MB/seconds=72.0 GB/hour	40 GB
IBM 3590E	15.0 MB/seconds=54.0 GB/hour	60 GB

Even though a single device may not be able to give the data transfer rate of 200 GB per hour set by our example, using multiple media devices should be able to achieve this aggregate transfer rate. For example, if you are using Ultrium tape drives, you need 2 tape drives to achieve 200 GB per hour, or 5 DLT 8000 drives to achieve the same throughput.

Network Bandwidth Considerations

Now you need to consider your network. More than any other factor, your available network bandwidth determines the amount of data that you can realistically transfer during a backup period. The following table compares the performance of different types of networks. As you can see, network performance can significantly impede large backup operations.

Network Type	Theoretical Transfer Rate	Realistic Throughput	Realistic Transfer Rate*
10Base-T Ethernet	10 mbps =1.25 MB/seconds	40-50%	500 KB/seconds=1.8 GB/hour
100Base-T Ethernet	100 mbps=12.5 MB/seconds	80%	10 MB/seconds=36 GB/hour
1 Gigabit Ethernet	1000 mbps=125 MB/seconds	70%	87.5 MB/seconds=315 GB/hour

* If you are backing up concurrently with other operations, remember that your backup operations will not achieve the maximum, realistic transfer rate listed.

Data Transfer Requirements and Resources Calculations

If the preliminary calculations outlined in the preceding sections show that your required data transfer rate is feasible given your existing infrastructure, you may be able to stop here. However, preliminary calculations usually uncover conflicts between stated requirements and available time and resources.

If minbandwidth is the amount of data that can be sent in a given time through the narrowest, slowest bottleneck in the path from the backup source to the backup media and if backupwindow is the time available, then the backup process is governed by the following equation:

$$\text{datatransferred} = \text{backupwindow} \times \text{minbandwidth}$$

In our example, we have a 5-hour window, fast storage devices, and 100Base-T Ethernet. So the Ethernet LAN is our weakest link, and the following equation is true:

$$\text{datatransferred} = 5 \text{ hrs} \times 36 \text{ GB/hour} = 180 \text{ GB}$$

Therefore, to back up 1 Terabyte of data, you have to do at least one of the following tasks:

- Increase the amount of time available to back up data.
- Increase the bandwidth available at the narrowest part of the data path.
- Reduce the size of *datatransferred* by backing up our 1 Terabyte in a series of smaller, independent operations.

The following sections suggest several possible alternatives that will achieve one or more of the above tasks.

Data Path Considerations

If you cannot decrease the amount of data that you need to move in the time available, then a possible solution is to increase the available bandwidth. You can do this either on the network that links data hosts to the BrightStor ARCserve Backup server or in the hardware that connects the server and the backup media.

Network Enhancements

The network is usually the most significant source of delays in the enterprise-backup environment. If a faster technology is available or feasible, an upgrade may be a good investment.

For example, if we have a 100Base-T Ethernet LAN and the same data transfer requirement as in the example we have been using so far (200 GB per hour), we cannot get backups done in the time allowed (5 hours). It would take approximately six times as long as we have to back everything up. A Gigabit Ethernet network would back up everything with time to spare and would benefit other business operations as well.

Storage Area Networks

A Storage Area Network (SAN) can improve backup performance significantly by moving data over the high-speed fibre connections rather than the slower network connections. In addition to the performance benefits derived from the high bandwidth fibre connectivity and low host CPU utilization, a SAN also improves the overall network performance by off loading the backup data transfer from the enterprise network to a dedicated storage network.

Though a SAN is expensive to implement and maintain, benefits go beyond just backup. A careful analysis of your requirements is necessary before a decision is made to implement a SAN. For information on how BrightStor ARCserve Backup can help you take advantage of a SAN, see the *Storage Area Network (SAN) Option Guide*.

SCSI Bus and Device Enhancements

In cases where poor device throughput is the limiting factor or when you have excess capacity on a fast network, you may need higher performance devices or more of your existing devices. If you use an older, slower drive technology, it may pay to upgrade to higher speed devices and faster SCSI buses. But in many cases, it may be better to add devices and, where necessary, libraries. You can then run storage operations in parallel using several devices at once.

Alternate Data Path Considerations

If you cannot upgrade the network or expand the time available for backups, you can almost always reduce the size of the data set that has to be handled during any particular instance of your backup. You achieve this by doing one of the following tasks:

- Segment your network.
- Segment your data so that it is backed up during a series of successive backups.
- Restrict the scope of your backups such that they only store data that has changed since the data set was last stored.

Segment Your Network

In many cases, you can make better use of your existing network bandwidth by placing BrightStor ARCserve Backup servers on different subnets.

- In the absence of subnets, all backup data has to cross a single network to reach the BrightStor ARCserve Backup servers. In effect, every piece of data travels sequentially to every node on the network.
- When you subnet your network, in effect you create two or more networks of equal speed, each of which handles a fraction of the backup data. Data travels in parallel.

In our example, if we backed up 500 GB on two subnets instead of 1 Terabyte on the entire network, we could back up twice as fast. Each subnet could transfer its 500 GB at 36 GB per hour for a total elapsed time of 14 hours (versus 28 hours). In our 5-hour backup window, we could transfer 360 GB, which, though not enough, is still far better than the 180 GB we could attain over a network that is not subnetted.

Segment Data

Nothing forces you to treat all of your organization's data as a single unit. It often makes better sense to *segment* the data into logically related chunks before trying to back it up. This reduces the time required for any single storage operation, makes better use of short backup periods and works better on slow networks. You still back up all of your data. You just do it in a series of shorter operations spread over several days.

We might, for instance, back up 20% of the 1 Terabyte of data in our example each night, Monday through Saturday. In the course of a week, this approach would back up our entire 1 Terabyte across the 100Base-T network, without exceeding the daily 5-hour backup period. As an added benefit, the compact backup elements make locating and restoring our data faster and easier by reducing the scope of searches.

The downside of this approach is that the entire data will not be backed up daily. Most organizations cannot afford to not have daily backups of complete data; therefore, this approach may not be suitable.

You might segment your data for backup purposes in any of the following ways:

- Business function (such as accounting, engineering, personnel management, sales, and shipping)
- Geographical location (such California development lab, St. Louis distribution center, New York business office, Miami business office, Tokyo business office, and Paris distribution center)
- Network location (such as NA005, NA002, NA003, JP001, and EU001)

Your segmentation scheme should, however, group the data into reasonably contiguous backup sources, so that the speed you gain is not lost in lengthy searches and additional network traffic.

Backup Scope

After you have segmented your data, you can further reduce the required data transfer rate by reducing the scope of some backups. Typically, a relatively small percentage of your data changes from day to day. While these changes need to be saved, a full backup is usually unnecessary.

For example, if you try to back up everything daily and only 10% of the data changes in the course of a day, you are spending 95% of your limited backup time storing data that is already backed up. When you include media consumption and wear and tear on your backup devices, this can be an expensive proposition.

You should consider backing up everything weekly, after 50% or more of your data has changed. You could then use the longer, weekend backup period for your longest storage operation. On a daily basis, you could back up the changes only. This would let you stay within the short, nightly back up window and would economize on media.

BrightStor ARCserve Backup provides options for you to address this issue with the following types of backups.

- Full backups—stores everything, regardless of when the data last changed.
- Differential backups—stores files that have changed since the last full backup.
- Incremental backups—stores files that have changed since the last full or incremental backup.

Creating the right mix of full and partial backup operations is something of a balancing act. Ideally, you want each version of each piece of data backed up once. You want to minimize unnecessary duplication that consumes media and time. Therefore, you should keep the following considerations in mind:

- Full backups store all of your data at once. They produce a complete, coherent image of the data as it was at the time of the backup. They also store the backed up data together in a single, easily managed storage object. As a result, backup strategies that rely exclusively on full backups are usually inefficient because the relative percentage of new data in the overall data set is generally small. Full backups save too many files that are already adequately backed up by a previous storage operation.

In exceptional situations, however, where the bulk of an organization's data changes substantially over short periods, a plan that relies on full backups exclusively may be the best choice. Because, in this case, most of the data is fresh at any given time, the full backup may actually be less prone to needless duplication than a mix of full and partial storage operations.

- Incremental and differential backups let you avoid network congestion and excessive media consumption. They better fit your existing hardware and bandwidth constraints and mesh better with your users' working hours. Incremental and differential backups are faster than full backups. If you do several of them between full backups, many files are still backed up more than once, because the differential backup backs up all files that have changed since the last full backup. This redundancy means that you can restore quickly, because all the data you need for a full recovery is stored in, at most, two data sets (the full and the last incremental).

Incremental and differential backups are only economical when the volume of changes is small compared to the volume of the data set as a whole. When this is the case, you can store changes on a small amount of media that is rewritten frequently.

Parallel Storage Operations (Multiple Streaming)

If device transfer rates limit your operations and if the necessary network bandwidth is available, you may want to set up your operations to use all of the available devices at once. By distributing the data across parallel streams, this approach greatly reduces the time required for backup operations. It does, however, consume more network bandwidth. Recovery after a catastrophic loss may be faster, since all available devices collaborate to restore all or most of the backup data at once. BrightStor ARCserve Backup has the capability to automatically create multiple streams based on the availability of tape devices.

Storage Capacity Requirements

So far, we have discussed factors that affect the speed with which backup and restore operations can be performed. But you also need to consider the volume of online data storage that you require.

Online Recovery Data Storage Requirements

You need to figure out how much recovery data you need to store online, in your robotic libraries. Data that is used primarily for archival purposes or for recovery after a catastrophe can be stored offline in a repository or vault. It is unlikely to be needed quickly. But recent backup data generally has to be available in a robotic library so that users can easily locate and swiftly recover the most recent, intact copies of the files they are most likely to lose.

To calculate the amount of recovery data you must store online, perform the following steps:

1. Estimate the size of an average, full backup.
2. Add the estimated size of an average incremental backup.
3. Multiply by the number of backup sets that your organization wants to have immediately available ("1" for the most recent, "2" for the two most recent, and so on). This is the amount of recovery data you need to keep online:

$$\text{recoverydata} = (\text{avgsizefull} + \text{avgsizeincrements}) \times \text{numberbackupskept}$$

Backup Data Storage Requirements

You need to reserve online storage space for scheduled backup operations.

To calculate this space:

1. Estimate the size of an average, full backup.
2. Add the average, percent growth of the data set during a typical, full backup cycle.
3. Add the estimated size of an average incremental backup.
4. Add the average percent growth of the data set during a typical, incremental backup cycle.

Storage Capacities and Resources

Your ability to meet your storage-capacity requirements depends on the following criteria:

- The types of libraries you have
- The number of each type you have
- The types of media each library uses

After you have identified types and numbers of libraries that will be available, you can calculate the capacity of each library using the following formula:

`totalcapacity = numberslotsavailable × mediaelementcapacity`

In this formula, the `numberslotsavailable` is the number of slots available in the robotic library and `mediaelementcapacity` is the capacity of the media elements used by the installed drives.

Media Capacities

The raw capacity of the media varies with the type of drives, the type of media, and the degree of data compression that you are using. You should deduct the following from the raw capacity to arrive at the real data capacity:

Deduct ~10% for overhead.

This allows for the BrightStor ARCserve Backup media header and various engine-specific overhead information. Note that the overhead may be more if you are backing up a large number of very small files.

For example, if you try to back up 1 Terabyte on ten media elements that hold 100 GB each (after deducting overhead), media usage will require 100% efficient every time you back up. Because this is unlikely, you need to use eleven media elements. On the other hand, you can back up 1 Terabyte to six cartridges that hold 200 GB each (after deducting overhead), because you have a healthy 200-GB (20%) cushion.

The allowances specified above are important. If you do not set aside space for overhead and variations in media usage, you may run out of media during a backup operation and may, consequently, not have a timely and complete backup.

Factors Affecting Storage Capacity Calculations

Media elements have lifetimes that are usually specified in usage time or numbers of uses or passes across the media. Make sure you take media aging into account when calculating the number of tapes required. Consult the manufacturer's recommendations.

Restrictive media-selection criteria and extensive off-site storage can increase your need for media well beyond the minimums calculated previously.

Finally, the overall size of the data you need to back up usually increases over time. The amount of data increases faster in some organizations than it does in others, but the total amount almost always increases. The preceding calculations assume a more-or-less constant amount of data. So, when you estimate how much you need to back up (1 terabyte in the examples), always allow for growth. Then check periodically to be sure that you always have enough extra storage to accommodate emerging needs.

Testing Plans and Assumptions

After you have made the required estimates, performed all the necessary calculations, and formulated a plan that should work for your organization, you should test it. Set up a pilot test configuration using a scaled down environment and run tests.

Note: You can simplify the pilot tests by using file system devices. You can set file system devices to `/dev/null`, thereby eliminating the requirement of dedicated disk space for pilot tests.

Using the BrightStor ARCserve Backup logs, you can see how good your estimates were. Use the backup logs to:

- Determine if you estimated the correct amount of backup data correctly by checking the size of a full backup generated by your plan.
- Check your estimate of the average percent change in your data by checking the size of the incremental backups.
- Make sure that all the data that should be backed up is backed up.
- Verify if your data and network segmentation tactics have worked as intended.

Catastrophic Events

So far, we have focused on the major threat to your data—routine losses due to equipment failure or operator error—and on the processes common to all backup and recovery efforts. But there are some additional considerations when you are planning your organization's recovery from a major catastrophe.

A catastrophe is a natural or man-made disaster, such as a fire or flood that results in the loss of multiple hosts, a data center, or an entire network, including locally stored backup media and hardware. To handle an extreme emergency, you must provide secure, off-site storage for some of your backup media, and you must keep the off-site data current.

Risk Assessment

Before going further, decide what sorts of disaster you can realistically prepare for, given the importance of your data, the expense of protecting it, the magnitude of the risk, and the corporate policies that apply to your sites.

Consider the following questions.

- What is the likelihood that your organization will face a large-scale disaster that affects the whole region or metropolitan area? Such catastrophes might include earthquakes, large floods, or acts of war.
- What is the likelihood of smaller disasters, such as building fires, localized flooding, or vandalism?
- How much data would you lose in a large disaster? In a small disaster?
- How severely would the loss affect your organization in each case?
- How much is your organization prepared to spend to defend against each of the risks you identify?

Off-Site Repository Considerations

In storage management, the selection of an off-site repository or *vault* is the result of a series of trade-offs.

Vault Security Considerations

The vault should be isolated enough from your main facility to protect the off-site data from the kind of catastrophes you are prepared to guard against.

For example:

- If earthquakes are the biggest threat you need to deal with, the vault should be in an earthquake-resistant building at some distance from your main site or even in another city or a different seismic zone.
- If fire or local flooding is the danger, a storage room in an upper floor of the building across the street might be enough.

Vault Accessibility Considerations

Measures that isolate your data repository from your primary site also make it harder (and more expensive) to keep the data in the remote repository current. To be of use, off-site data has to be reasonably up-to-date, which means it has to be reasonably accessible. A vault in a distant city might protect the data against even the most extreme disasters, but it might be impractical to ship media there on a daily basis.

Vault Expense Considerations

In general, the more secure a vault is, the more expensive it is to use. You pay more for more secure storage facilities. It often takes longer to get media to and from these facilities. The more media you store off-site, the more you have to buy for your main site.

Disaster Recovery Archive Considerations

Because catastrophes will, by definition, strike your infrastructure as well as your backup media, you should assume that you will have to rebuild systems completely before you can start the actual data recovery. For this reason, you should always maintain the following off site:

- Media elements that contain bootable operating systems for the BrightStor ARCserve Backup servers.
- A current, complete backup of the file systems, databases, and mail servers supported by BrightStor ARCserve Backup.

You may want to include BrightStor ARCserve Backup distribution media and a text file that lists your hardware configuration parameters.

Disaster Recovery Testing

To be sure that your data is available after a disaster, you have to periodically test the data that you are archiving. Routine file-backup routines get tested every time a user cannot restore a deleted file. You soon hear about problems and, in general, the results are not too costly. But disasters are, by definition, rare and expensive. When your data center has just burned down, it is too late to find out that your backup routine does not work. So be sure to test these infrequently used processes on a regular basis.

Whenever you install new software or hardware, or change existing procedures, complete the following tests:

- Backup to media as you would for off-site storage and disaster recovery.
- Verify that the backup operation stored all the specified data successfully.
- Simulate a post-catastrophe recovery operation using the backup media from the test.

You should also run brief, simulated, backup and restore operations whenever the opportunity arises. Routine testing lets you exercise and assess your storage processes on an ongoing basis.


Sample Calculations

The examples below illustrate some representative situations that a backup and recovery plan has to deal with.

Note: It is assumed that the backup server has enough CPU power and memory, and the hard disk speed on the client or server is adequate.


Transfer Rate for Clients and Servers on a 100Base-T Ethernet LAN With No Subnets

In this configuration, data cannot move across the network faster than 36 GB per hour, regardless of the number of servers and libraries available. To back up 1 Terabyte of data, the backup operation must run for 28 hrs.




Transfer Rate for Clients and Servers on Two 100Base-T Ethernet Subnets

In this configuration, you can move twice as much data at the 36 GB per hour 100Base-T data rate. To back up 1 Terabyte of data, each subnet has to handle only 500 GB, so the operation takes 14 hours. Some performance is lost because the network cannot keep the media drives in each library streaming along at their combined 36 GB per hour optimum speed.




Transfer Rate for Clients and Servers on a Gigabit Ethernet Network

In this configuration, you move data at 315 GB per hour data ratio. To back up 1 Terabyte of data, the backup operation must run for 3 hours.




Transfer Rate For a Server With No Clients

In this case, the 216 GB per hour drives are the limiting factor, assuming that disk system or server is not the bottleneck. The system would take 5 hours to back up 1 Terabyte.



Transfer Rate For Server With SAN Option

In this configuration, local backups of each server on the SAN can achieve a data transfer rate of 432 GB per hour.



Storage Capacity For Two Sets of Recovery Data, One Full and One Incremental Backup

Assume the following:

- You have to do a full backup of 1 Terabyte of user data per week.
- You have to do daily incremental backups.
- About 10% of the data changes daily.
- The data from the last two backup cycles are available, online, for fast recovery.
- You are using LTO tape drives with 2:1 compression in a library with 20 slots.
- All media are used as efficiently as possible.

First, calculate the amount of capacity you need to store the output of the current backup operations. LTO media elements have a raw capacity of 200 GB with 2:1 compression. After you deduct 10% for overhead, the real capacity is close to 180 GB. The 1 Terabyte full backup thus requires:

$$1 \text{ Terabyte} \div 180 \text{ GB / media element} = 6 \text{ media elements}$$

Using the above equation, you can also calculate the safety margin as follows:

$$(6 \times 180 - 1000) / 1000 = 8\%$$

Because six tapes (1 Terabyte) provide an 8% safety margin, you do not need to add extra tapes. In this example, you need only 6 LTO tapes to store a full backup. Based on the rate of change you estimated, the incremental backups amount to:

$$1 \text{ Terabyte} \times 10\% \text{ changed / incremental} \times 5 \text{ incrementals} = 500 \text{ GB changed}$$

Therefore, at a minimum, you need the following:

$$500 \text{ GB} \div 180 \text{ GB / media element} = 3 \text{ media elements}$$

Because three tapes (500 GB) provides a 9% safety margin, you do not need to add extra tapes. You need only three tapes to store a single set of incremental backup data.

Next, calculate the amount of storage space you need for your online recovery data. You need to retain the last two backup sets in the library, so you need 9 tapes for the oldest set of recovery data and 9 tapes for the newest set. To store your recovery data you need 18 tapes.

Therefore, your total storage requirement is as follows:

9 tapes for current backup + 18 tapes for recovery = 27 tapes

Next, you calculate the capacity of the library by deducting cleaning slots:

20 slots/library - 1 cleaning slot = 19 available slots

Therefore, you have a deficit of $27 - 19 = 8$ slots and must do one of the following:

- Add a library.
- Compress the stored data.
- Store only one set of recovery data online.

Chapter 3: Backing Up Data

You can use BrightStor ARCserve Backup to back up data from most machines attached to your network. With the optional BrightStor ARCserve Backup client agents, you can communicate with remote workstations in various environments. This provides complete system backups, including system information from non-UNIX systems, such as Linux, Windows and NetWare. BrightStor ARCserve Backup provides great flexibility in specifying options, filters, and scheduling information for your backup jobs.

Backup Manager

Using the Backup Manager you can customize your backup jobs using filters, options, and scheduling. For procedural information on how to submit backup jobs using the Backup Manager, see the online help.

You can use the Backup Manager to perform the following tasks:

- Back up to various media.
- Back up an entire node containing database files. For more information about this feature, see Entire Node Backups in this chapter
- Create a customized backup scheme.
- Use filters to selectively exclude or include directories and files from backup jobs.
- Create an automated backup scheme using a Grandfather-Father-Son (GFS) rotation scheme.
- Apply backup options to local source objects (such as volumes and nodes) or globally to the entire backup job, or to both at the same time.

Each backup job requires a source, a destination (media), and schedule or method. The Backup Manager window provides three tabs to define your backup job:

- Source
- Destination
- Method/Schedule

The buttons on the Backup Manager toolbar let you further define the backup job with additional criteria. The available buttons are:



Submit—Submit your job to run immediately, or schedule it for a later time. For more information, see the chapter “Customizing Your Jobs.”



Option—Select Option to customize your job. The options available are discussed in the Backup Options section in this chapter.



Filter—Select Filter to include or exclude specific files or directories from your job. For information about filters for backup jobs, see Backup Job Filters in the chapter “Customizing Your Jobs.”




File System Device Group Configuration—Configures file system device groups. You must use this toolbar option to run Disk Staging backup jobs.

Backup Manager Source Tab

The source is the path that you want to back up. You can easily find the files you want to back up by browsing through the directory tree to select the user-shared volumes, directories, and files.

To start the Backup Manager, click the Backup Manager icon on the BrightStor ARCserve Backup home page. The Source tab of the Backup Manager opens as shown in the following example:



Note: To perform a backup, you must have user access to the server you want to back up.

On the Source tab, the Host and Path fields display the host and path associated with the object currently selected and support the search for a specific object. With the Host field you can search for a particular BrightStor ARCserve Backup server, while the Path field supports the search of a directory or file. You can use these entry fields together to search for an object on a particular server.

Note: When using these fields to search for an object, the information entered must make the information an exact match. In addition, the object found is only highlighted, not selected.

When viewing workstations, you can use the default view option, View Agents Which Are Running, to browse and select objects associated with machines running the appropriate agent software and entered into the BrightStor ARCserve Backup agent database.

The View All File System Agents option shows all workstations entered into the current host server's Remote Client database, regardless of whether Agent software is running. Select this option to prepare a backup job by selecting an entire node without browsing its contents.

■ Not selected—A partially filled green square indicates that you have not selected the object.

■ Selected—A solid green square indicates the you selected the object.

Note: The client agent software must be running on all selected machines before submitting backup jobs. If the agent service is not available, the square next to the node is inaccessible. If you restarted the client agent, the connection must be refreshed. To refresh the connection, collapse and expand the branch for the workstation.

Entire Node Backups

If you want to back up an entire node, BrightStor 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 BrightStor ARCserve Backup to back up a selected node and all of its contents with a single click in the Backup Manager directory tree. BrightStor ARCserve Backup will back 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, BrightStor ARCserve Backup detects the new drive automatically and backs up the entire node when you run the backup job.

How BrightStor ARCserve Backup Authenticates Entire Node Backups

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, BrightStor 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 when you click the Submit toolbar button, or if you select Save or Save As from the File menu on the Backup Manager window.

Security and Agent Information Dialog

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.

Back Up an Entire Node Containing Database Files

To back up an entire node containing database files, perform the following procedure:

1. From the Source tab of the Backup Manager, select the node that you want to back up.
2. Click the Submit toolbar button.

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.

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.

Note: In the Security dialog, you must specify User Name and Password with backup rights on that machine. For example, Administrator or root.

4. Click OK.

The Submit dialog opens. For procedural information on how to submit backup jobs using the Backup Manager, see the online help.

Multiple Code Pages Support

The following sections describe how BrightStor ARCserve Backup supports the use of multiple code pages.

How BrightStor ARCserve Backup Supports Multiple Code Pages

A code page is a map of characters as they relate to a particular language. If the BrightStor ARCserve Backup server resides in an environment where different languages and their character sets are running on other computers, the Backup Manager and the Restore Manager may not be able to interpret and display recognizable text in the source tree.

When you encounter this situation, you can specify any code page supported in your environment. The code page lets BrightStor ARCserve Backup interpret the information and display the text in a format that is recognizable to you.

When you specify a code page at the node or volume level, BrightStor ARCserve Backup applies the characteristics of the code page to all child volumes, directories, and so on. Code pages do not affect BrightStor ARCserve Backup functionality.

Specify Code Pages in the Backup Manager Window

You can change the code page on all tree items in the source tree. To specify a code page, use the following steps:

1. From Source tab on the Backup Manager window, right-click the node, volume, or directory for which you want to specify a code page.
2. From the Encoding pop-up menu, point to and click the code page that you want to display.

BrightStor ARCserve Backup applies the new code page settings immediately.

Node Level Options

When you select a host (node) object from the directory tree in the Backup Manager, BrightStor ARCserve Backup presents you with node object tabs. You can use node object tabs to set options and filters at the node level, and display node-level information on the Object Information tab.

You can set backup options for any host (node) you include in a backup job, using the Node Option tabs (Object Information, Filter, Security, and Node Option). The following sections discuss the options available using these tabs.

Backup Manager Object Information Tab

The Object Information tab displays the properties of the object selected in the left pane. The information displayed varies according to the type of source you select.

The Object Information includes the following:

- Host Name—The name of the selected host machine
- IP Address—The network address of the selected node
- System Name—The name of the operating system
- Release—The release number of the operating system
- Version—The version number of the operating system
- Machine—The hardware model

The Object Information tab displays read-only information. You cannot change any of the fields displayed on this tab.

Filter Tab

Using filters you can include or exclude specific files and directories from your backup jobs. Use the filters to help focus on the files you want. You can 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. Because you can perform filtering on a per node basis, you can, for example, include a directory from one node and exclude the same directory from another node.

- To apply a filter locally to the source, select the Filter tab on the object's source tab. The filters you set here apply only to the selected object, not to the entire backup job.
- To apply global filters, click the Filter button found on the Backup Manager toolbar.

Important! *Exercise caution when using filters for your backup operation! If you incorrectly apply filters, you could back up the incorrect data, you can lose data, and you can waste time.*

For more information on filters, see the chapter “Customizing Your Jobs.”

Backup Manager Security Tab

To back up a machine, you must log in to that machine. Use the Security tab to enter the User Name and User Password for your backup operation. This information enables BrightStor ARCserve Backup to store the information, and use it each time you select the machine for browsing.

Important! *After you successfully log in to the target machine and submit a backup job, BrightStor ARCserve Backup retains the authentication information for the target machine on the Backup Manager Security tab. If you delete or change the information contained on the Security tab to incorrect information after you submit the backup job, the job will fail.*

Node Option Tab

The Node Option tab allows you to modify the way BrightStor ARCserve Backup handles specific node options for the backup job. You can set node options for any host you include in your backup job using these tabs, which appear when a node object is selected.

The Node Option tab has four secondary tabs that you can use to set various options to be applied to the node specified in the left panel of the Backup Manager. Use the secondary tabs (Priority, Format, Misc, and NDS Option) to set the following options:

- The node priority of the backup.
- The format in which to back up selected files.

- Whether to have BrightStor ARCserve Backup traverse symbolic links and NFS mounted file systems.
- Control over file access times and backup estimation.

Priority Tab

Use the Priority tab to set the priority of your backup operation. When selecting multiple nodes for backup, you can set each node's priority thus establishing which nodes to back up first. The range is from 1 (the highest priority) to the default setting of 255 (the lowest priority).

Format Tab

By default, BrightStor ARCserve Backup backs up data in its original, native format. Alternatively, you can back up your files in tar or cpio format, allowing the data to be accessed using standard UNIX backup formats. Choose from the following formats:

- BrightStor format
- Posix tar
- Posix cpio
- Posix dump

BrightStor ARCserve Backup uses the posix header when backing up files in *cpio* format. On some UNIX systems, the *cpio* utility cannot always read this type of header. When this scenario occurs, you should apply the -H switch to the *cpio* command line.

Note: When backing up data on Windows systems, BrightStor ARCserve Backup does not support the use of the Posix tar, Posix cpio, and Posix dump formats. As such, it uses the BrightStor tape format when backing up data on Windows systems.

Miscellaneous Tab

Using the Misc tab, you can set the following options in relation to the selected node:

- Traverse Symbolic Link File (UNIX only)—BrightStor ARCserve Backup follows symbolic links and backs up the linked files.
- Traverse Across File System (UNIX only)—BrightStor ARCserve Backup automatically includes locally mounted file systems.

- **Preserve File Access Time (UNIX only)**—This option directs BrightStor 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 full backup 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 backed up), you need to preserve the original access time.

- If this option is selected (check in box), BrightStor 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: This option has no effect on the Modified Time.

- **Disable File Estimation**—Disables the default setting to run an estimation for the backup job.
- **Traverse NFS**—Backs up NFS mounted drives. (UNIX only).

NDS Option Tab

You can enter information on the NDS Option tab only if you are backing up a NetWare server. You must enter the necessary information in the fields on this tab for the Novell Directory Services (NDS) to successfully back up a NetWare server. BrightStor ARCserve Backup can back up your base schema and your extended schema.

To back up an NDS tree, highlight an entire Volume from the selected NetWare machine, and select the check box Backup NDS (NetWare only) located on the NDS Option tab. Enter the appropriate NDS information as follows:

- **NDS Tree Name**—The NDS tree name you want to back up.
- **NDS Login Name**—The user's full context name is required.
- **NDS Password**—The password associated with the login name.
- **Server Name**—The name of the server selected to back up the NDS tree. Enter the name of the currently selected NetWare machine, or any NetWare machine in the NDS tree running the Backup Agent for NetWare.
- **Server IP Address**—The IP address of the server entered in the Server Name field.

Note: This option is only available for machines running the BrightStor® ARCserve® Backup Client Agent for NetWare.

Volume Backup Options

When you select a volume object from the directory tree of the Backup Manager, BrightStor ARCserve Backup presents you with Volume Object tabs, so that you can set volume level filters and options, and displays volume-related information on the Object Information tab.

You can set volume backup options for any host you include in a backup job, using the Volume Option tabs. The following sections discuss the options available using the Volume Option tabs.

Object Information Tab for Volumes, Directories, and Files

When you select the Object Information tab, the information displayed varies according to the type of source selected. When you select a volume, directory, or file object from the source browser, the following object information displays:

- Path—Path representing the device file for the raw partition.
- Name—The name of the object.
- Mode—File system information, such as object type and permissions.
- nlink—Number of links to the object.
- User ID—User ID number.
- Group ID—Group ID number.
- rdev—Raw device (Character Special or Block Special).
- dev—The device on which the object resides.
- Last Access Time—The time the volume, directory, or file was last accessed.
- Last Modified Time—The time the volume, directory, or file was last modified.
- Last Change Time—The time the volume, directory, or file was changed.
- File Size—The size of the volume, directory, or file.

Filter Tab at the Volume Level

Using filters you can include or exclude files and directories from your backup or restore jobs. Use the filters to help focus on the files you want. When accessing the source option Filter tab, the filters that you set apply only to the selected object, not the entire backup job.

For more information on filters, see the chapter “Customizing Your Jobs.”

Important! *Exercise caution when using filters for your backup operation! If you incorrectly apply filters, you could back up the incorrect data, you can lose data, and you can waste time.*

Session Password at the Volume Level

The session password, when set using the Session Password tab, is directly associated with the selected volume and all objects in that volume.

When you specify a session password, BrightStor ARCserve Backup uses 168-bit data encryption. For more information, see Encryption Methodology in this chapter.

When backing up data, you can use the following session password options:

- Session/Encryption Password—To enhance network security, specify a session/encryption password to be required when restoring this data from media. If you enable the Encrypt Files Before Backup option, you must set the encryption key in this field.

Note: Make sure that you remember your password so that it is available when you want to restore your data. If you do not know your password, you will not be able to access the data.

- Compression/Encryption
 - Compress Files Before Backup—Enable the compression of files before backup, to increase the effectiveness of your media and increase data speed. If the tape drive does not support compression, or if the hardware compression is disabled, BrightStor ARCserve Backup does not compress the data.
 - Encrypt Files Before Backup—Specify whether to encrypt files before backup. Encryption encodes your data so that it is not intelligible. BrightStor ARCserve Backup uses the Session/Encryption Password, set on this tab, as the encryption key.

Note: Compression/Encryption options are available for UNIX and Windows agents only.

Important! *Remember your session password! If you try to restore your data and you do not know your password, you will not be able to access the data.*

Access Tab at the Volume Level

Some backup scenarios include different types of file systems residing on one host (standard file systems, a database partition, NetWare volumes, and so on). You can use the options on the Access tab to back up different types of file systems without submitting multiple jobs. Methods are available only if the selected object applies to that method.

The following are the available backup options that you can apply to volumes:

- **Regular File System Backup**—Sets BrightStor ARCserve Backup to traverse each file system during a backup. This is the default backup option.
- **Raw Partition Backup**—Displays a list of device files associated with disk partitions. This option reads the hard disk block and sets BrightStor ARCserve Backup to back up an entire hard disk partition. Although this is the fastest backup option, selection is restricted to device files associated with disk partitions. Alternatively, you can enter the full path of a disk device file.
- **Stack File System Backup**—Sets BrightStor ARCserve Backup to back up only StackFS mounted file systems. Use this option for backing up open files. This option is available only if the file system is StackFS mounted.
- **Network File System Backup**—Specify this setting if the object selected is an NFS mounted file system. The backup option used is the Regular File System Backup.

Volume Option Tab

The Volume Option tab has two secondary tabs, Priority and Verification, that you can use to set the priority of the backup job, and how BrightStor ARCserve Backup will verify the backed up data.

Job Priority Levels - Priority Tab, Volume Level

When selecting multiple volumes to back up, you can specify the priority level for each volume. The priority range is 1 (the highest priority) to the default priority, 255 (the lowest priority).

Backup Job Verification at the Volume Level

BrightStor ARCserve Backup provides you with three options for backup verification. Backup verification ensures that the data backed up to media matches what is on your hard disk. The three options are:

- None—No verification is performed. This is the default setting.
- Scan Media Contents—BrightStor ARCserve Backup checks the header of each file on the backup media.
 - If BrightStor ARCserve Backup can read the header, it assumes the data is correct.
 - If BrightStor ARCserve Backup cannot read the header, it updates the Activity Log with this information.
- Compare Media to Disk—BrightStor ARCserve Backup reads blocks of data from the media, and compares the data, byte for byte, against the files on the source. This ensures that all data stored on the backup media is exactly as it was on disk. If BrightStor ARCserve Backup finds a mismatch, it updates the Activity Log with this information.

Note: The Compare Media to Disk option is not supported with agents for databases and applications.

Important! A verification option set on the volume level takes precedence over the same verification option set using the global Options button on the Backup Manager Toolbar. As a rule, if a local source verification option is set, it takes precedence over the same option set globally.

Disk Staging Option

The following sections provide conceptual information about backing up data using the Disk Staging Option.

Note: For information about configuring the Disk Staging Option, see the section Staging Configuration in this chapter. For information about using the Disk Staging Option, see the section Backing Up Data Using the Disk Staging Option in this chapter.

Staging Operations

The operations and tasks associated with the Disk Staging Option include the following:

- Specify and configure file system devices.
- Configure a file system device as a staging group and configure staging group parameters.
- Submit backup jobs to a staging group.
- Define policies for managing backup, data migration, and purge operations. For additional information about staging policies, see the section Staging Policies in this chapter.
- Perform simultaneous backup operations to a file system device 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 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 direct BrightStor ARCserve Backup to restore the data from either location.
- Run command line tools that can analyze and purge data stored on a FSD in a staging group.

How the Max Number of Streams Option Affects Backup and Restore Operations

Simultaneous streaming is a process that divides your backup jobs into several sub-jobs that run simultaneously. The Disk Staging Option allows you to utilize the simultaneous streaming feature to send multiple streams of data to a file system device (FSD) in a staging group. Since the work is split up among several different drives, simultaneous streaming-enabled backup jobs can be completed significantly faster than regular backup jobs.

BrightStor ARCserve Backup provides you with the capability of streaming multiple jobs simultaneously to the FSD. The base product allows you to write a maximum of two streams per job simultaneously, as well as two streams per staging group simultaneously. Licensing the Disk Staging Option enables you to increase the simultaneous streams to 32 (for each job and each staging group).

When you back up data using the Disk Staging Option, 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 depending upon whether the backup job is a node-level or a volume-level backup job. 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.

For example, 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.

Disk Staging Option and Rotations

When you back up data using regular or GFS rotation rules, BrightStor 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.

For example, if you detect 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.

Staging Tab

To access the information and options on the staging tab, start the Backup Manager and select the staging tab.

The Staging tab contains the following options and informational fields:

Enable Staging

Click the Enable Staging check box to enable or disable staging backup operations.

Group Field

Displays the name of the group selected for the job.

Note: A staging group must be selected in a staging job. Specifying a "*" group is not allowed for staging.

Policy

Opens the Staging Policy dialog. Using the Staging Policy dialog, you can specify staging policies for full, incremental, and differential backup operations. Staging policies allow you to specify copy policies, purge policies, enabling SnapLock protection, and other miscellaneous policies.

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.

Staging Groups Directory Tree

Displays the names of the groups which were configured as staging groups.

Staging Configuration

Before you can back up data using the Disk Staging Option, you must perform the following tasks:

- Create file systems devices. First, you must specify the file system devices in your environment that you will use for staging purposes. To create file system devices, you must run the BrightStor ARCserve Backup csetup script. For more information about running the csetup script, see the *Getting Started*.
- Configure staging groups. After specifying the file system devices in your environment, you must configure the file system device group to a staging group. For more information, see Configure Staging Groups in this chapter.
- Configure staging policies. To perform backup operations using staging, you must define the backup, copy, and purge policies that BrightStor ARCserve Backup will use to manage data stored on staging devices. For more information, see Staging Policies in this chapter.

The following sections provide you with information about how to configure the Disk Staging Option.

Configure Staging Groups

To configure staging groups:

1. Open the Device Manager or the Backup Manager and click the Staging Group Configuration toolbar button.

The Staging Group Configuration dialog opens.

2. 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:

- **Threshold**--Specifies file system device capacity thresholds. The threshold can be represented as either the total number of MB or GB used, or as a percentage of the disk's total capacity used.
- **Maximum # Streams per Group**--Specifies the maximum number of simultaneous streams to the selected file system device group.
- **Enable SnapLock**--Enables SnapLock WORM protection on the file system device.

Note: The device must support SnapLock technology. If you enable SnapLock on a device that does not support SnapLock WORM protection, BrightStor ARCserve Backup write-protects the data, however, the data can be deleted from the device.

- **Pause data migration**--Pauses the data migration operation.
- **Staging Chunk Size**--Specify the amount data to be written to the staging device per write operation.

Note: A higher chunk size can decrease the level of fragmentation on your staging device. However, a higher chunk size can adversely affect the throughput of data to the disk.

3. Repeat Step 2 as necessary to configure other file system device groups.
4. Click OK.

Staging Policies

The following sections describe defining staging copy, purge, and SnapLock policies for full backups, incremental and differential backups, and other miscellaneous options.

Staging Copy Policies

After BrightStor ARCserve Backup completes the backup to disk phase, copy policies let you specify when to copy the data to its final destination media.

The following information applies to Full, and Incremental/Differential copy policies.

- **Do not copy data option**—Choose this option if you do not want to copy 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.
- **Copy data After option**—Choose this option to direct BrightStor ARCserve Backup to start the copy from disk to final destination media operation after specified length of time elapses. BrightStor ARCserve Backup starts the copy to media operation based upon the occurrence of one of the following events:
 - **After job starts option**—Choose this option if you want to start the copy to media operation at a fixed point in time after the backup to disk operation starts.
 - **After job ends option**—Choose this option if you want to start the copy to media 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 option**—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 BrightStor 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, BrightStor 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.

- **Copy data At Option**—Choose this option to direct BrightStor ARCserve Backup to start the copy to media operation at a specific time of day. When you use this option you can direct BrightStor 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 BrightStor ARCserve Backup from copying sessions from disk to tape while the backup operation is in progress.

Staging Purge Policies

If you direct BrightStor ARCserve Backup to use copy policies, the data stored on your file system device will remain on the FSD until such time that the data is copied to final destination media.

Note: You can use the `ca_devmgr` command line utility to forcefully purge sessions that were not copied or not purged from a file system device. For more information, see the section "Staging Command Line Purge Tool" in the appendix Using Command Line Utilities.

Use the following information to determine how BrightStor ARCserve Backup will process backed up data stored on a file system device. This information applies to Full and Differential/Incremental backup to staging device operations.

- Purge data **After** option—Choose this option to direct BrightStor ARCserve Backup to start the purge operation after specified length of time elapses. BrightStor ARCserve Backup starts the purge operation based upon the occurrence of one of the following events:
 - **After job starts option**—Choose this option to direct BrightStor ARCserve Backup start the purge data from disk operation at a specified time after the backup to staging device operation starts.
 - **After job ends option**—Choose this option to direct BrightStor ARCserve Backup to start the purge data from disk operation at a specified time after the backup to staging device operation ends.

For example, if 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 At 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..

- Purge data **At** option—Choose this option to direct BrightStor ARCserve Backup to start the purge data from disk operation at a specific time of day. Use the spin box to specify the time of day that you want the operation to start.

Consider the following scenario: You have a backup job 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 you can schedule the purge operation to start before the next backup operation starts. This approach helps to ensure that you have freed enough disk to prevent the backup job from failing.

Important! *If you specify a copy policy, BrightStor ARCserve Backup does not start the purge operation until after the copy to final media operation is finished.*

Staging Miscellaneous Policies

The Disk Staging Option supports the following Miscellaneous options:

Purge canceled sessions from disk immediately

- Use this option to direct BrightStor 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 immediately

- Use this option to direct BrightStor 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 BrightStor 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, BrightStor 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 a makeup job on hold if a data migration job fails

- Use this option to direct BrightStor 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.

Configure Staging Policies

To configure Staging policies, use the following steps.

1. From the Staging tab on the Backup Manager, select the Staging group that you want to configure and then check the Enable Staging check box.
2. Click the Policy button to open the Staging Policy dialog.
3. Select the Full Backup tab or the Differential/Incremental Backup tab.
4. Specify the Copy Policies for this job. For more information, see the section Staging Copy Policies in this chapter.

5. Specify the Purge Policies for this job. For more information, see the section Staging Purge Policies in this chapter.
6. If you want to delete-protect backup data, click the Enable SnapLock check box.

Note: The file system device must support SnapLock functionality to delete-protect data. BrightStor ARCserve Backup cannot detect if a file system device is capable of providing SnapLock security until the backup job is complete. As such, if you specify “Enable SnapLock” on a backup job and the file system device does not support SnapLock security, the backup job will run properly.

7. Select the Miscellaneous tab and choose the options that you want to apply to the job. For more information, see the section Staging Miscellaneous Policies in this chapter.
8. Click OK.

Backing Up Data Using the Disk Staging Option

The following sections provide you with information about how to back up data and use the Disk Staging Option.

Submit a Staging Backup Job Using the Backup Manager

Prior to performing a backup job using the Disk Staging Option, you must have already configured the staging groups. If you did not configure BrightStor ARCserve Backup to use the Disk Staging Option, see Staging Configuration Tasks in this chapter.

BrightStor 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 disk staging backup job using the command line utility, see Submit a Staging Backup Job (Command Line) in this chapter.

To submit a backup job using the Disk Staging Option, use the following steps:

1. From the Backup Manager, access the Source tab and specify the volumes that you want to back up.
2. Access the Staging tab, check the Enable Staging check box, and specify the Staging Group that you want to use for staging.

Note: To run a staging backup job, you cannot specify a asterisk (*) in the Group field.

3. To modify the copy and purge policies, click the Policy button and modify the policies for the job, as necessary.
4. To modify the number of simultaneous streams during the backup operation, use the spin box to change the Max Number of Streams.

Note: BrightStor ARCserve Backup provides you with the capability of streaming multiple jobs simultaneously to the FSD. The base product allows you to write a maximum of two streams per job simultaneously, as well as two streams per staging group simultaneously. Installing the Disk Staging Option enables you to increase the simultaneous streams to 32 (for each job and each staging group).

5. Click the Submit toolbar button to submit the backup job.

Submit a Staging Backup Job Using the Command Line

BrightStor ARCserve Backup provides you with the capability to submit a backup job using either the Backup Manager or the command line. The information contained in this section describes how to perform a staging backup job using the command line utility.

For information about how to submit a staging backup job using the Backup Manager, see the section Submit a Backup Job Using the Disk Staging Option in this chapter.

ca_backup Syntax

To submit a staging backup job using the ca_backup command line utility, use the following syntax:

```
ca_backup [-diskstage <groupname>]
[-purgeFailedSessions]
[-purgeCancelledSessions]
[-makeupJobToTape]
[-createDMJMakeupJobOnHold]
[-chunkSize <size>]
[-maxStreams <Max # Streams>]
[-fullbackup
  [[-DONOTCOPY] |
  [-copyDataToDestination
    [afterjobstarts <weeks> <days> <hours> <minutes>] |
    [afterjobends <weeks> <days> <hours> <minutes>] |
    [aftersessionends <weeks> <days> <hours> <minutes>] |
    [at <hh:mm:ss> [afterjobends]]]]]
  [-purgeData
    [afterjobstarts <weeks> <days> <hours> <minutes>] |
    [afterjobends <weeks> <days> <hours> <minutes>] |
    [at <hh:mm:ss>]]]
  [-ENABLESNAPLOCK] ]

[-incdiffbackup
  [[-DONOTCOPY] |
  [-copyDataToDestination
    [afterjobstarts <weeks> <days> <hours> <minutes>] |
    [afterjobends <weeks> <days> <hours> <minutes>] |
    [aftersessionends <weeks> <days> <hours> <minutes>] |
    [at <hh:mm:ss> [afterjobends]]]]]
  [-purgeData
    [afterjobstarts <weeks> <days> <hours> <minutes>] |
    [afterjobends <weeks> <days> <hours> <minutes>] |
    [at <hh:mm:ss>]]]
  [-ENABLESNAPLOCK] ]
```

Note: For a detailed description of the staging command line options, see the section ca_backup in the appendix "Using Command Line Utilities."

Modify a Staging Rotation Scheme

If you are using rotation or GFS rotation disk staging jobs, BrightStor 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 Method/Schedule tab.
2. Select either the Rotation option or the GFS Rotation option, and then select the Cycle Table tab.
3. Select the scheme from the Scheme Name drop-down list.

The Staging column displays the current status of staging as it applies to your rotation scheme.

4. Select and double-click the schedule that you want to modify.

The Job Unit dialog opens.

5. From the Staging list, select Enable or Disabled.
6. Click OK.

Note: To disable staging for a staging group, see the section [Disable Staging](#) in this chapter.

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. For example, if your tape library is offline or you need to perform maintenance on the library, you can pause the data migration process and then restart it after the library is back online.

1. From the Backup Manager, click the Staging Group Configuration toolbar button.

The Staging 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 blue flag. The groups that are not enabled for staging display with a corresponding red flag.

2. From the Groups field, select the groups that you want to pause.
3. Check the Pause data migration check box and click OK.

Disable Staging

BrightStor 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 Staging Group Configuration.

Backup Manager

To disable backup to staging device operations from the Backup Manager, perform the following steps:

1. From the Backup Manager, open the Method/Schedule tab.
2. Select the Scheme Name from the drop-down list.
3. Click the Rotation Rules tab. Select and double-click the rotation that you want to disable.

The Job Unit dialog opens.

4. From the Staging drop-down list on the Job Unit dialog, select Disabled.
5. Click OK.

Staging Group Configuration

To disable back up to staging device group operations using Staging Group Configuration:

1. From the Backup Manager, click the Staging Group Configuration toolbar button.

The Staging 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 blue flag. The groups that are not enabled for staging display with a corresponding red flag.

2. In the Groups pane, select the group that you want to disable staging.
3. Clear the check mark from the Enable Staging check box.
4. Click OK.

Manage Staged Data When the Database Fails

When you use the Disk Staging Option to back up data, the information about the backup jobs, sessions, staging policies, and so on is stored in the BrightStor ARCserve Backup database. If the database fails, and you need to re-initialize the BrightStor ARCserve Backup database, the staging policies for the data residing on the file system device (FSD) that specify when to copy the data to the final destination media and when to purge the data from the FSD are no longer available.

If this situation occurs:

- BrightStor ARCserve Backup cannot copy (migrate) the data on the FSD to its final destination media.
- BrightStor ARCserve Backup cannot purge data from the FSD to reclaim disk space.
- Future backup jobs will probably fail due to an insufficient amount free disk space on the FSD.

There are two approaches you can use to remedy this situation:


- To retain all of the backup data stored on the FSD, you can use the `tapecopy` command line utility to copy all the backup data from the FSD 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.
- If you do not want to copy all of the data from the FSD to its final destination media, or purge all of the data from the FSD, you can use the Merge Manager to quickly merge only the session details (not the file details) from the FSD to the BrightStor ARCserve Backup database.

The process of merging session details is much faster than merging file details. To merge session details only, open the Merge Manager and select the "Record Job and Session Information Only" option before you run the merge job. You can access this option from the Database tab of the Global Options dialog.

After the Merge job is complete, this approach lets you specify the sessions that you want to copy to final destination media using the `tapecopy` command line utility. Then, to reclaim disk space, you can use the `ca_devmgr -purge` option to specify the sessions that you want to purge from the FSD.

Backup Manager Destination Tab

The destination is the backup media device or disk. Browse through the left pane to select the groups and device. Place an asterisk in the Group or Media Name field to use the first available drive and media in the group. If you want to use any available group, enable the option Any Group. If you want to use a specific media in a specific device group, select the media in the left pane.



Drive and Media Selection

When you use the Destination tab to select a drive and media to use for backup, you can perform the following actions:

- Back up your data to hard disks (local or mounted) rather than the traditional magnetic media.
- Configure a volume as the destination for all backups. Multiple disk media can be created, and this data can be automatically staged to tape media. In addition, you can perform a restore from this volume.
- Create file system devices in locally mounted file systems of the BrightStor ARCserve Backup server, or NFS mounted file systems, if the BrightStor ARCserve Backup server has root NFS access on the remote system.
- Specify that BrightStor ARCserve Backup use multistreaming for backups.

Multistreaming

Multistreaming offers you the advantage of using all the available tape devices on the system by splitting a single backup job into multiple jobs using all the tape devices. As a result, it will increase the overall backup throughput compared with the sequential method. Since the work is split up among several different drives, you can complete multistreaming-enabled backup jobs significantly faster than regular backup jobs.

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.

Multistreaming Levels

You can enable multistreaming on one of two levels—the node level or the volume level. You can also specify the maximum number of streams into which a job can be split. For more information on these options, see Advanced Options in this chapter.

Note: File system backups will be performed using the multistreaming level you have selected (node level or volume level). However, other types of backups (such as those that involve databases) can be split up on a level other than the one specified.

To copy data to another media, disk or tape in a different group when the backup job is completed, you can choose the Copy Tape Session After Backup option. To enable this function, you must enter a destination group in the `$BAB_HOME/config/tapecopy.cfg` file. Check the log file in `$BAB_HOME/logs/tapecopy.log` to find the status of your copy.

Backup Session Chunking to a File System Device (FSD)

When you back up data to an FSD, the overall size of each backup session can become very large. If the size of a single backup session becomes greater than the maximum file size supported by the operating system, BrightStor ARCserve Backup will chunk (or slice) this large session into smaller sessions such that each slice is less than the maximum file size supported by the operating system.

Backup session chunking is an automated process that is managed by the Media Server daemon (camediad). This functionality lets you create backups for any job, regardless of size, without having to consider operating system limitations.

Default Operation

By default, BrightStor ARCserve Backup activates the chunking process when the size of the backup session reaches the lesser of 2048 MB (2 GB) or the maximum file size supported by the operating system.

Examples:

- The maximum file size supported by the operating system is 1024 MB (1 GB). BrightStor ARCserve Backup activates the chunking process when the size of the backup session reaches 1024 MB (1 GB).
- The maximum file size supported by the operating system is 4096 MB (4 GB). BrightStor ARCserve Backup activates the chunking process when the size of the backup session reaches 2048 MB (2 GB).

When you browse a directory of an FSD that contains session slices, the file names of the session slices display using the following file naming convention:

`S#####.CTF.???`

Where ##### represents the session number and ??? represents the extension number of the session.

Examples:

The following are examples of valid session slice file names.

S0009999.CTF

The file name of the first slice.

S0009999.CTF.0001

The file name of the next slice.

S0009999.CTF.0002

The file name of the next slice, and so on.

Method/Schedule Tab

In the Backup Manager window, click the Method/Schedule tab. Use the Method/Schedule tab to determine the schedule and method of your backup operation. You can use one of the BrightStor ARCserve Backup predefined backup strategies, or you can customize a backup strategy to suit your environment's needs. BrightStor ARCserve Backup also lets you determine the type of backup to be performed, and the rules governing overwriting or appending your data to media.

Backup Scheduling Methods

BrightStor ARCserve Backup provides three backup scheduling methods from which to choose. The method you choose determines when your backups are run, the type of backup to be done on particular days, and the rotation of the backup media.

There are three methods from which you can choose as follows:

- Custom—Run a backup job once or on a repeating basis. Customize your backup job using the Repeating Interval and Media Rules tabs for each backup job submitted. You can choose to exclude particular days from the backup job. Use this option if you want to perform an unscheduled full backup without affecting your scheduled incremental or differential backups.
- Rotation—Maintain a weekly rotation of media using a combination of two of the three different backup methods: full, differential, and incremental. Customize the backup job using the Cycle Table, Media Rules, Calendar View, and Exception View tabs for each backup job submitted.
- GFS Rotation—Maintain a monthly rotation of media using a combination of two of the three different backup methods: full, differential, and incremental. Customize the backup job using the Cycle Table, Media Rules, Calendar View, and Exception View tabs for each backup job submitted.

Backup Methods

The backup method determines which files or directories are backed up. The following types of backup methods are available in all rotation methods:

- Full—All source files are backed up. You can run a full backup once or you can run a full backup job daily, weekly, or monthly.
- Incremental—The incremental backup includes only files modified since the last backup of any kind, determined by the date and time of the last saved changes. After each incremental backup, the date and time are recorded for each file. Files with unchanged dates and times are not backed up in the next incremental backup job.
- Differential—The differential backup includes all files modified since the last full backup. When you run a full backup, the dates and times are recorded. A differential backup backs up all files with modification times that are more recent than the full backup, regardless of whether these files were backed up in the last differential backup. Date and time are not recorded after differential backups.

Note: Files are flagged for backup based on the modified date and time.

In general, if you have large amounts of data to be backed up every day, regular differential, or incremental backup jobs after a full backup, using one of the default rotation schedules, can be your best solution. For more information about selecting a rotation schedule, see the chapter “Customizing Your Jobs.”

Custom Backup Schedules

You can use the Custom Backup Schedule Type to schedule a backup job to run once or at repeating intervals, without specifying a rotation. You can use the Custom Schedule Type to create a schedule to suit your needs (for example, to schedule a backup to run more than once a day, if necessary) or you can use this option if you want to perform an unscheduled backup without affecting your regular rotation schedule.

Repeat Intervals for Custom Backups

When you select the custom backup method, you can schedule your backup job at repeating intervals. Specify a Repeat Interval in months, days, hours, and minutes, and identify any days of the week to exclude. You can create a unique backup schedule to suit your environment’s needs. If your backup is to be run once, leave these fields blank. You can also enable the Retry Missed Targets feature to reschedule the backup of any missed workstations or file servers to a specified time.

Media Rules for Custom Backups

Using the Media Rules tab, you can specify various options for the media used in your custom backup job while you are configuring the job. With the options available on the Media Rules tab, you can identify the media pool, if any, with which the destination media is associated, and to define the rules for overwriting or appending to the first media and any additional media.

- **Select Media Pool**—Unlike with rotation and GFS rotation, the use of a specific media pool in a custom backup is not required, but you can submit a backup job to a given media pool, provided you are authorized to access that pool. Select one of the following:
 - **Existing Media Pool**—Select a previously created media pool to use with a current backup job, or to verify that the destination media is associated with the media pool to which it belongs.
 - **New Media Pool**—Specify a new media pool. The new media pool is created after the job is saved or submitted.

Note: If your destination media is associated with a media pool, you must set the Select Media Pool field correctly, or the job will fail. If it is not associated with a media pool, use the default setting.

- **Eject Media**—Specify to eject the media from the drive after the job finishes, to prevent the possibility of another job overwriting information on this media.

First Media Options

By default, data from the backup operation is appended to the specified media. With the options available for First Media, you can define the Append or Overwrite rules for the first media used for the backup job. If your destination media is associated with a media pool, the First Media options reflect the rules governing the media pool and cannot be changed here. For more information about media pool options, see the chapter “Managing Devices and Media.”

The options relating to the rules for the first media are as follows:

- **Append to Media**—The default setting. BrightStor ARCserve Backup adds job sessions to the selected media.
- **Overwrite Same Media Name, or Blank Media**—BrightStor ARCserve Backup overwrites the media in the drive only if it is the media specified for the job, set in the Media Name field, or if it is blank. If media with a different name is in the drive when the job takes off, the job fails.
- **Overwrite Same Media Name or Blank Media First, then Any Media**—BrightStor ARCserve Backup overwrites any media in the drive. Select this option to instruct BrightStor ARCserve Backup to check if the media in the drive is the one specified for the job, or if it is blank. If it is neither, BrightStor ARCserve Backup reformats any media it finds in the device with the name from the Media Name field and starts backing up files at the beginning of the media.

Note: Rotation and GFS rotation media are not overwritten. If BrightStor ARCserve Backup finds that the media in the drive is a rotation/GFS rotation media, the job fails.

- **Overwrite Same Media Name, or Any Media First, then Blank Media—**BrightStor ARCserve Backup overwrites any media in the drive with the same name as that in the Media Name field, or any media. If the media in the drive is a rotation or GFS rotation media, the job fails.
- **Timeout for First Media—**Specify the number of minutes
BrightStor ARCserve Backup waits for a media to be inserted into a drive before it cancels the job.

Span Media Options

Span Media options define the overwrite rules for jobs that require additional media. You must specify which media BrightStor ARCserve Backup can use when your job uses more than one media. If your destination media is associated with a media pool, the Span Media options reflect the rules governing the media pool, and cannot be changed here. For more information about media pool options, see the chapter “Managing Devices and Media.”

The options relating to the rules for additional media used in a backup job are as follows:

- **Overwrite Same Media Name, or Blank Media—**Overwrite any media with the same media name as the first media, or if it is blank. As long as the ID is different, BrightStor ARCserve Backup reformats the media, giving it the same name and ID as the first media. Only the sequence number is different.
- **Overwrite Same Media Name or Blank Media First, then Any Media—**Overwrites any media found in the device with an ID different from the first media’s ID. All subsequent media are reformatted with the same name and ID as the first media. Only the sequence number will vary.

Be careful when using this option when you have grouped drives. Media left in a drive will be reformatted if you specify that BrightStor ARCserve Backup span to any media.
- **Overwrite Same Media Name, or Any Media First, then Blank Media—**Use this option to overwrite any media in the drive with the same name as in the Media Name field, or any media. If the media in the drive is a rotation or GFS rotation media, the job fails.
- **Timeout for Span Media—**The number of minutes
BrightStor ARCserve Backup waits for an additional media to be inserted into a drive before it cancels a job.

Rotation Schedules

In addition to the Custom Schedule option, BrightStor ARCserve Backup has two types of rotation schedules available through the Backup Manager. You can choose a simple rotation or GFS rotation. The following sections briefly discuss these two rotation schedules. For more information, see the chapter “Customizing Your Jobs.”

Simple Rotation Schedules

With the simple rotation policy, you can maintain a regular, weekly rotation of media, and use full, incremental, or differential backup methods. Unlike the custom schedule, you can assign a media pool to the rotation scheme.

GFS Rotation Schedules

The GFS rotation schedule is a set of predefined backup jobs consisting of full backups combined with incremental or differential jobs. Using the GFS rotation schedule, you can maintain a monthly rotation of media, using full backups combined with differential or incremental backup methods. Unlike the custom schedule, GFS rotation media are automatically assigned to a media pool, and, unlike the simple rotation schedule, data on your media is overwritten by default; however, you can choose to append data to media. You cannot change the name associated with the media within the assigned pool. BrightStor ARCserve Backup automatically names the media.

Setting up a GFS rotation schedule offers several advantages:

- Guarantees data security. BrightStor ARCserve Backup automatically names media, and prevents the media from being overwritten until you have the correct number of media in your Save Set, and the minimum retention period has been met.
- Maintains media integrity by ensuring that media is not overused.
- Keeps a current and accurate backup of the selected source.

After you set up your GFS rotation scheme, you need only make sure the right media is in the drive for each day of the week. BrightStor ARCserve Backup provides this information for you.

For more information about selecting or defining a GFS rotation schedule, see the chapter “Customizing Your Jobs.”

Backup Job Filters

For backup jobs, you can filter on a global basis as well as a node-level basis, and you can set node-level and job-level filters for the same job. Node-level filters apply to one specific node, not the entire job. If you want to add a filter for the entire job, click the Filters button on the Backup Manager toolbar to open the global Filters dialog.

Note: If you select the include directory pattern filter and do not specify an absolute path, empty directories for all the directories that do not match the user provided criteria will be backed up. To avoid creating these empty directories during restore, disable the global restore option Create Empty Directories when creating your restore job. For more information on this option, see Create Empty Directories Option in the chapter “Restoring Data.”

For more information on selecting and applying filters using the Filter button on the toolbar, see the chapter “Customizing Your Jobs.”

Important! *Exercise caution when using filters for your backup operation! If you incorrectly apply filters, you could back up the incorrect data, you can lose data, and you can waste time.*

Backup Options

With Backup options, you can further customize your backup job. Using BrightStor ARCserve Backup you can apply backup options to local source objects such as volumes and nodes, or globally to the entire backup job, or to both at the same time.

- If an option is to be applied locally to the source, you can access it from that object's source tab.
- If it is to be applied globally, you can use the Options button found on the Backup Manager toolbar to open the Option dialog.

The available backup option tabs are as follows:

- Misc—Specify various actions to occur after your backup job is complete.
- Session Password—Specify password, compression, and encryption options.
- Virus Scan—Specify a rule for BrightStor ARCserve Backup to follow if it detects a virus.

- Media Exporting—Specify media export options.
- Verification—Specify the rules BrightStor ARCserve Backup uses to verify the accuracy of your backup operation.
- Load Options—Enable this if you want to display and package deleted files and directories from the job scripts.
- Advanced—Specify multiplexing and multistreaming options.
- Pre/Post—Specify commands to run before or after your backup job is complete.
- Log—Specify how BrightStor ARCserve Backup logs activities and identifies the output devices for the activity logs.
- Database—Specify the type of information BrightStor ARCserve Backup records in the database.
- File Retry/Sharing—Specify the rules BrightStor ARCserve Backup follows when encountering an open file.

Note: In order to have the file retry function available, you must have the -I parameter included in the Client Agent Configuration File (uag.cfg). By default, this parameter is not included in this configuration file. For more information, see the topic Client Agent Configuration File in the online help.

The following sections provide information about the various tabs and their options.

Miscellaneous Options

The Miscellaneous tab determines actions that occur during or after the backup. The options are:

- Disable File Estimate—By default, before backing up to media, BrightStor ARCserve Backup estimates how long the job will take. Select this option if you want to skip this function.
- Clear Archive Bit (DOS and Windows Only)—Resets the archive bit after the file has been backed up.

- **Delete Files After Backup**—Removes the files from the hard disk after the file backup is completed. This option deletes only the files from unprotected directories. Files in protected directories will not be deleted. By default, the following system directories are protected for UNIX and Linux platforms:

- /
- /bin
- /etc
- /lib
- /sbin
- /usr/bin
- /usr/sbin

You can protect other directories by modifying the applicable file system agent configuration file, `groom.cntl`, located on the client agent machine. The agent configuration file, `groom.cntl`, is specific to UNIX and Linux client agents.

By default, the following system directories are protected for NetWare platforms:

- SYSTEM
- PUBLIC
- LOGIN
- ETC
- MAIL

On Windows platforms, System State and system protected files are not deleted.

Session Password Options

To ensure security throughout your network, you can set a session/encryption password and compression/encryption options using the Session Password tab.

Choose from the following methods:

- **Session/Encryption Password**—To enhance network security, specify a session or encryption password to be required when restoring this data from media. If you enable the Encrypt Files Before Backup option, you must set the encryption key in this field.

Note: Make sure that you remember your password so that you can restore your data. If you do not know your password, you will not be able to access the data.

- Compression/Encryption
 - Compress Files Before Backup—Enable the compression of files before backup, to increase the effectiveness of your media and increase data speed. If the tape drive does not support compression, or if the hardware compression is disabled, BrightStor ARCserve Backup does not compress the data.
 - Encrypt Files Before Backup—Specify whether to encrypt files before backup. Encryption encodes your data so that it is not intelligible. BrightStor ARCserve Backup uses the Session/Encryption Password, set on this tab, as the encryption key.

Note: Compression/encryption options are available for UNIX and Windows agents only.

Encryption Methodology

When you specify data encryption, BrightStor ARCserve Backup uses 168-bit encryption to back up files. To use data encryption, you must specify a session password.

168-bit encryption is the default and preferred method of data encryption. Although 168-bit encryption is a highly secure method of data protection, it can have a slight adverse affect on backup performance. However, you can disable the default encryption method by modifying the ENHANCED_ENCRYPTION section of the cprocess.cfg file. The cprocess.cfg file is stored in the \$BAB_HOME/config directory. For more information about how to disable enhanced encryption, see Disable 168-bit Encryption in this chapter.

Disable 168-bit Encryption

To disable 168-bit encryption, use the following steps:

1. Open the cprocess.cfg file located in the \$BAB_HOME/config directory.
2. In the ENHANCED_ENCRYPTION section, insert a # immediately before ENHANCED_ENCRYPTION.
3. Save the file.

Note: This is also known as “commenting out” a parameter.

Media Exporting Options

Note: This feature is available only for devices that support importing and exporting media.

At the end of a backup job, you can export or move media out of the library or to an off-site location for safe storage. If the job includes verification, the export is done at the end of the verification. BrightStor ARCserve Backup provides the following media exporting options:

- None—No media exporting takes place at the end of a backup job.
- Export All Media after Job—BrightStor 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. If the library is a single slot library and the slot is full, BrightStor ARCserve Backup notifies you that there are not enough slots available to support the request and then completes the remaining export tasks.
 - If the operator does not move the media, BrightStor ARCserve Backup writes this information in the activity log.

Note: Some libraries have slots that allow media to be loaded and unloaded without opening the library door. Mail slots are used for the Import/Export feature.

Verification Options

Using BrightStor ARCserve Backup you can 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. Any options selected for the drive override the global verification options (applied to the entire job).

- None—BrightStor ARCserve Backup does not verify the backup job.
- Scan Media Contents—Check the proprietary BrightStor ARCserve Backup data area (the header) of each file on the backup media. If it is readable, BrightStor 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.
 - Record Detail Information—Enable this option to include additional information about the file being scanned in the daily log of a GFS or rotation job. This information includes the file's name, size, and full path.

- Compare Media to Disk—BrightStor ARCserve Backup reads and compares data from the backup media byte for byte against the source files. This option takes time, but ensures that all data on the backup media is identical to the data on the disk. If BrightStor ARCserve Backup finds a mismatch, the errors are recorded in the Activity Log.

Note: The Compare Media to Disk option is not supported with agents for databases and applications.

Load Options

If you enable Load Options, directories in your job script that are no longer available for backup will be unavailable in the Source tree. This reduces the chances that your backup job will fail because BrightStor ARCserve Backup cannot find a directory specified in your job script. Conversely, if your backup job fails and you did not have Load Options enabled, you can enable Load Options to quickly see if there are missing directories that could cause your job to fail.

A directory can be unavailable for a number of reasons:

- The directory has been deleted.
- The directory has been renamed.
- The directory resides on a shared drive which is no longer available.

When you see an unavailable directory, you can re-create the directory, re-establish a connection to the directory, or remove the directory from your job script.

If you do not enable Load Options, only the directories that will be backed up appear in the Source tree and missing directories do not appear. You will not be informed of any missing directories unless the backup job fails.

Advanced Backup Options

BrightStor ARCserve Backup offers the following multiplexing and multistreaming global options:

Multiplexing

Multiplexing—Multiplexing is a process in which data from multiple sources is written to the same media simultaneously. For more information on multiplexing, see Multiplexing in this chapter. If you use multiplexing, you can select the following options:

- **Chunk Size**—Sets the performance of restore operations and memory usage. The chunk size value determines the amount of contiguous data written for one session before data from another session is multiplexed. The higher the value, the faster the restore on some drives, but at the cost of memory size during backup. For most drives, the default value of 1 MB is recommended.
- **Maximum Number of Streams Per Drive**—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

Multistreaming—Multistreaming is a process that divides your backup jobs into several sub-jobs that run simultaneously. If you use multistreaming, you can select the following options:

- **Stream Level**—Select the level at which you want the job split:
 - **Node**—When node-level multistreaming is selected, jobs that involve multiple nodes (computers) are split into multiple streams (one stream per node). If node-level multistreaming is selected, and a backup job involving only one node is run, BrightStor ARCserve Backup attempts to complete the job using volume-level multistreaming.
 - **Volume**—When volume-level multistreaming is selected, jobs in which data is being backed up from multiple volumes are split into multiple streams. For example, if you had a backup job that involved three volumes from the same computer, or a volume each from two computers, the job would be split into multiple streams. A job involving one volume on a single computer would not be split into multiple streams.
- **Max # Drives**—Defines the maximum number of drives to which a job can be split.

Pre/Post Options

BrightStor ARCserve Backup allows you to specify commands to perform various operations immediately before or after your data is merged based on the exit codes received. You can specify pre/post commands:

- For the entire merge job. These commands are executed at the beginning or end of the job.
- Run commands, executable programs, and shell scripts.

For example, you may want to specify this option to load a virus scanning application before files are merged and then run the batch file you created that sends a detailed report to the printer.

Note: If you specify pre/post commands for the entire job, as well as for machines in the job, the global pre/post commands are executed before or after the job starts or ends, and local commands are executed when the selected machine is backed up.

The Pre/Post option allows you to run a command on your BrightStor ARCserve Backup server machine before, after, or before and after the job is executed. The outcome of Pre/Post commands can be found in caqd.log. The available Pre/Post options are defined in the following sections.

Backup Manager Run Command Before Job Options

Enter the path and name of the application to be executed on the machine before the job starts.

- On Exit Code— If you want BrightStor ARCserve Backup to detect exit codes of any application, enable this option, select the condition for detecting exit codes (Equal To, Greater Than, Less Than, or Not Equal to), specify the exit codes you want detected, and then select how you want BrightStor ARCserve Backup to respond after exit codes are detected:

Note: If you select the condition Greater Than or Less Than, only one exit code should be specified.

- Skip Delay—If the exit code conditions are met, run the job immediately, ignoring the delay.
- Skip Job—If the exit code conditions are met, skip the job.
- Skip Post Application—If the exit code conditions are met, enabling this skips commands scheduled to run after the job completes.

The skip delay, skip job, and skip post application options will be activated only if BrightStor ARCserve Backup detects that the exit codes meet the specified condition (Equal To, Greater Than, Less Than, or Not Equal To).

- **Delay in Minutes**—Specify the delay in which BrightStor ARCserve Backup waits before running a job when the appropriate exit code is detected. This gives the specified application time to finish processing before your job begins.

Note: Ensure that the path to the command you want to run is correct. For example, if you want to run the script `pre_exec.ksh`, and the command `pre_exec.ksh` is in the `usr` directory, enter this:

```
/usr/pre_exec.ksh
```

Backup Manager Run Command After Job Options

BrightStor ARCserve Backup runs the command after the backup job finishes. Enter the path to, and name of, the application to be executed on the machine after the job completes. As with the Run Command Before Job option, you must ensure that the path to the command is correct.

Backup Manager Run Before/After Command As Options

Enter the User Name and Password to run Pre/Post commands. The system on the selected host server requires User Name and Password to check the system privileges on that server. The user is authenticated using the File System agent running on the server machine. The File System Agent must also be installed and enabled on the server machine to run Pre/Post commands.

Do not confuse the User Name and Password entered into these fields with the BrightStor ARCserve Backup User Name and Password.

Backup Manager Log Options

Log options determine the level of detail included in the log report for the operation and the devices to which the log will be sent. You can view the log report from the Job Status Manager or Database Manager (Job Records). BrightStor ARCserve Backup provides the following log options:

- **Log File Name**—Enter a name for the log file.
Note: If you are attempting to back up BrightStor ARCserve Backup database agent files, the Backup Manager does not support specifying log file names.
- **Log All Activities**—Record all of the activity that occurs while the job runs in the Job Log.
- **Log Summary Only**—Record summary information for the job (including source, destination, session number, and totals) and errors.

You can also select the output devices for the activity log. Select any or all of the following:

- Unicenter NSM Alert—Use this option to send a message to the Unicenter Console when an alert is generated.
- SNMP Alert—Use this option to send messages to your SNMP messaging console.
- Printer Name—Use this option to print messages to a designated local printer.
- Internet Email—Use this option to send email messages to specific email addresses.

Note: You can enable the activity log destination options (NSM alert messages, printers, and email addresses) by modifying the configuration file named `caloggerd.cfg`. This configuration file is located at: `$BAB_HOME/config`.

Backup Manager Database Options

The Database options determine the level of detail included in the database for the backup job. The options are:

- Record Detail Information—Records detailed information about your backup job, including the session, and each file name in that session.
- Record Job and Session Information Only—Records only the backup job and the sessions in that job.

File Retry/Sharing Options

The options on the File Retry/Sharing tab determine the rules BrightStor ARCserve Backup follows when encountering an open file, or one in use by an application.

Note: In order to have the file retry function available, you must have the `-I` parameter included in the Client Agent Configuration File (`uag.cfg`). By default, this parameter is not included in this configuration file. For more information, see the topic Client Agent Configuration File in the online help.

File Retry Options

When BrightStor ARCserve Backup encounters an open file, you can select one of the following options:

Note: These options are only available for UNIX and NetWare agents.

- **Retry Immediately**—Back up or copy the file again, immediately after the first attempt failed. If the file is still unavailable, BrightStor ARCserve Backup writes information to the Activity Log, and labels the job "Incomplete."
- **Retry After Session**—Back up or copy the file again after all the other source files have been backed up. If the file is still unavailable, BrightStor ARCserve Backup writes information to the Activity Log, and labels the job "Incomplete."

For each of these options you must specify the following:

- **Number of Retries**—Number of times you want to try to back up or copy the file.
- **Retry Interval (sec)**—Period of time you want to wait between attempts.

File Sharing Options

File Sharing determines how BrightStor ARCserve Backup shares files with other applications when backing up or copying a file.

Note: The following file sharing options only apply to Windows and DOS operating systems.

- **Use Deny None if Deny Write Fails**—[default] Attempt to place the file in "Deny Write" mode. If this is not possible (because the file is already open), then place the file into "Deny None" mode.
- **Use Lock Mode if Deny Write Fails**—Attempt to place the file in "Deny Write" mode. If this is not possible (because the file is already open), then lock the file completely (prohibiting any user from opening or writing to the file). This option ensures the most recent version of the file is backed up or copied.

- **Deny Write**—Prevents another process from writing to the file while BrightStor ARCserve Backup has it open. If another process opens the file before BrightStor ARCserve Backup can open it, BrightStor ARCserve Backup does not back up the file unless you specified an Open File Retry option.
- **Deny None**—Other processes can read or write to the file, regardless of whether BrightStor ARCserve Backup opens the file first or opens it after another process already has it open. This option ensures that your files are up-to-date, although the backed up or copied file may not be the most recent version.

Note: With the Deny None method, as long as no other process is writing to these files during the job, the backup or copy will be consistent. If you want to ensure that BrightStor ARCserve Backup copies or backs up only the most current version of every file, select a Deny Write or Lock Mode option.

The `ca_backup` Command

In addition to the Backup Manager, BrightStor ARCserve Backup provides a command line interface that lets you perform a variety of backup functions without using the Backup Manager, giving you an alternate method of accessing almost all of the operations available through the Backup Manager from the command prompt.

With the options and switches for the `ca_backup` command, you can:

- Set global options and filters.
- Specify the source, destination, rotation schedule, and run schedule for your backup operation.
- Submit the backup job to run immediately or at a scheduled time.

To build a backup operation, you must set one category of options at a time, in the order shown.

The command normally submits the backup job to the job queue, then returns immediately with the status of whether the job was submitted successfully to the queue or not. To integrate with Unicenter Workload, the command must return only after the job itself has completed execution, not simply when the job has been placed into the queue.

For a complete list of the options and switches available for this command, see the appendix “Using Command Line Utilities.”

Multiplexing

Multiplexing is a process where data from multiple sources is written to the same media simultaneously. When a job that has multiple sources is submitted with the multiplexing option enabled, it is broken into child jobs—one for each source. Child jobs write data to the same media simultaneously.

Note: When using multiplexing, you can select the maximum number of streams that can write a tape at the same time. For more information about this setting, see *Set Multiplexing Global Options* in this chapter.

- 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 significantly 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 decreases the amount of time it takes to perform a backup while increasing the life of the hardware.

Features Supported by Multiplexing


The following table includes the features that are supported and not supported with multiplexing:

Supported	Not Supported
■ 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.	■ Configuration of two drives within the same device group. ■ 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. ■ Staging operations 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 on Optical Changers and DVD drives.

Multiplexing Job Option

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



Set Multiplexing Global Options


BrightStor ARCserve Backup offers the following multiplexing global options:

- **Chunk Size**—Sets the performance of restore operations and memory usage. The chunk size value determines the amount of contiguous data written for one session before data from another session is multiplexed. The higher the value, the faster the restore on some drives, but at the cost of memory size during backup. For most drives, the default value of 1 MB is recommended.
- **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.

To set the multiplexing global options, perform the following steps:

1. From the Backup Manager, click the Options toolbar button.
2. Click the Advanced tab.


The Advanced tab opens as shown in the following example:



3. Move the Multiplexing Chunk Size slider to the desired value.
4. Select the maximum number of streams value.
5. Click OK.

Multiplexing Job Status

After submitting a multiplexing job, you can monitor its status from the Job Summary tab in the Job Status Manager. In the job summary, multiplexing jobs appear in levels so that you can view the parent and child job relationship—children jobs appear as secondary jobs below the parent job, as shown in the following example:



In addition, the status of the parent job is the highest severity status of a child. For example, if Child 1 is successful, Child 2 is incomplete, and Child 3 has failed, BrightStor ARCserve Backup labels the parent job "Failed."

Preflight Checks for your Backups

The Preflight Check (PFC) utility is a command line utility that you can use to run vital checks on the BrightStor ARCserve Backup server and agents to detect conditions that can cause backup jobs to fail. The checks performed by PFC fall into four categories: system checks, BrightStor checks, agent checks, and media checks:

- **System Checks**—These include checking system requirements for the server, available disk space for the database, and RPC service registration.
- **BrightStor Checks**—These include checking the BrightStor ARCserve Backup system account and its privileges, the status of the BrightStor engines, SAN server connectivity (if the BrightStor 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.
- **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 several hours before your jobs are scheduled to run so that you can have ample time to correct any problems that appear in the PFC report. For more information on the PFC utility and its associated options, see the appendix "Using Command Line Utilities."

You can access the PFC report in the \$BAB_HOME/logs/pfclogs directory. BrightStor ARCserve Backup labels the PFC log files using the following labeling convention:

`pfc_ + hostname + string + .log`

Where the string is a six-digit number that increments by 1 after each execution of the PFC utility.

Chapter 4: Restoring Data

BrightStor ARCserve Backup provides you various tools and options to restore your data. This chapter includes information about how you can safely and efficiently restore your data.

Restore Manager

The objective of running a successful restore job is to quickly identify the data you need and to retrieve it from the appropriate backup media.

With BrightStor ARCserve Backup you can restore data to most machines attached to your network.

- Each restore job requires a source and destination.
- The files selected as your source must originate from backup media created by BrightStor ARCserve Backup, and the destination can be a destination of your choice (for example, a hard drive).

The Restore Manager screen contains two tabs to customize your restore job:

- Source
- Destination

Using the Source tab you can specify the files to be restored and using the Destination tab you can specify the device or file system to which your files will be restored.

You can use the optional BrightStor ARCserve Backup client agents to communicate with remote workstations in various environments to restore data to non-UNIX systems, such as Linux, Windows, and NetWare.

Similarly, you can use the optional database agents to restore online databases and applications such as Lotus Domino, Oracle, and IBM Informix.

For procedural information on how to submit a basic restore job, see the online help.

Use the buttons on the Restore Manager toolbar to customize your restore job. The available buttons are:



Submit—Submit the restore job to the Job Queue to run immediately or at a scheduled time. For information about the Submit button and scheduling options, see the chapter “Customizing Your Jobs.”



Option—Select Option to customize your job. The options available from the Options button are discussed in the section Restore Options in this chapter.




Filter—Set Filters to include or exclude specific files or directories from your job. For information about Filters, see the chapter “Customizing Your Jobs.”

Restore Manager Source Tab

BrightStor ARCserve Backup makes it easy to find the files you want to restore. You can restore entire hosts, file systems, drives, and volumes, depending on your needs.

When you select the Restore By File System view, the Source tab lists all available types of agents, as shown below. You can expand an agent object to view its associated nodes or machines on which you can run a restore. You can click a node to traverse through the volumes, directories, and files you can restore.



Note: When you expand a source tree and BrightStor ARCserve Backup detects that there is an excessive amount of files within the selected directory, instead of attempting to display all of the associated files, a pop-up message will first appear notifying you of this condition. After you click the OK button, the Find Files in <Current Path> dialog will then appear. The Find Files in <Current Path> dialog will allow you to select the target files that will be displayed in source tree.

When restoring data from media, you need not specify the device group. By default, BrightStor ARCserve Backup looks for the specified media in each host's associated device group. However, BrightStor ARCserve Backup can restore from any device group on your host. You can change device groups if another group is already being used for a job.

Restore Find Functions

BrightStor ARCserve Backup can help you find the volume, drive, directory, and files you need to restore your data. If you are unsure of what you want to restore, use the Find button to have BrightStor ARCserve Backup perform a search of the database to match the pattern you specify.

Note: You can only use Find and Version History when you are using Restore by File System as your source view.

Find Button

By clicking the Find button on the Restore Manager source tab, the Find File dialog will appear.

There are four tabs associated with the Find File dialog:

- Host
- File
- Date
- Options

Host Tab

You can select the host you want to search.

File Tab

You can specify the name of the file to start the search. You can include wildcard characters * and ? in the file name.

Date Tab

You can select date ranges to find files that were modified between two specified dates or during a specified previous number of days, months, or years.

Options Tab

You can specify the Maximum Returned Records during the find process. The default value is 1024.

Expressions Supported When Searching For Files

BrightStor ARCserve Backup supports all regular expressions. For example, you could search for anything beginning with report:

`report*`

The following search patterns are acceptable as well:

- `report?` - files with names of report1, or reports match this pattern
- `[123]report` - files with names of 1report, 2report, or 3report match this pattern

Specify your search criteria in the Find File dialog and click Find Now. BrightStor ARCserve Backup searches the database and returns the names of all files matching the specified pattern.

Note: If the find process detects more files that match the search pattern than the number specified in the Maximum Returned Records field (on the Options tab), a pop-up message will appear indicating that if displayed list does not contain the files you are looking for, you can either refine the search pattern criteria or increase the Maximum Returned Records value.

Version History Button

If you have several backups of a file, directory, drive, or volume, use the Version History button to display a list all the versions you have backed up to help you can find the one that you need.

Note: Restoring data from a disk is generally faster than restoring from a tape, because there are no delays due to the tape load and seek latency. If 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.

- You can restore directly from the final destination by selecting a session and clicking OK to start the restore process.
- You can restore from a different location by clicking the Duplicates button.

The Duplicate Sessions dialog opens displaying any sessions which are duplicates or clones of each other (including the original session). After selecting the session and clicking OK, the restore process will start.

Duplicate Backup Sessions

When you use the Disk Staging Option to back up data or copy media using Tapecopy or the tapecopy 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 quickly 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, you can direct BrightStor ARCserve Backup to use the media that is on site to facilitate the restore operation.

To search for duplicate sessions, click the Duplicates button on the Version History dialog. The Duplicates Sessions dialog displays the original backup session and all of its copies. If duplicates for a session exist, you can direct BrightStor ARCserve Backup to use the session that allows you to restore the session as quickly as possible.

Smart Restore

BrightStor 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, BrightStor ARCserve Backup searches for an alternate media to use to complete the restore job. Consider the following scenario:

During a restore job, the restore source media jams and disables the library. BrightStor 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.

Find Files in Current Path

The Find Files in <Current Path> dialog is used to specify multiple files to be displayed in the Restore Manager directory source tree. When you expand a source tree and BrightStor ARCserve Backup detects that there is an excessive amount of files within the selected directory, instead of attempting to display all of the associated files, a pop-up message will first appear notifying you of this condition. After you click the OK button, the Find Files in <Current Path> dialog will then appear. The Find Files in <Current Path> dialog will allow you to select the target files that will be displayed in source tree.

You can use the Check All, Clear, and Select buttons to specify which of the found files will be displayed in the source tree.

There are three tabs associated with the Find Files in <Current Path> dialog:

- File
- Date
- Options

File Tab

You can specify the name of the file to start the search. You can include wildcard characters * and ? in the file name.

Date Tab

You can select date ranges to find files that were modified between two specified dates or during a specified previous number of days, months, or years.

Options Tab

You can specify the Maximum Returned Records during the find process. The default value is 1024.

Note: If the find process detects more files that match the search pattern than the number specified in the Maximum Returned Records field (on the Options tab), a pop-up message will appear indicating that if displayed list does not contain the files you are looking for, you can either refine the search pattern criteria or increase the Maximum Returned Records value.

Restore Methods

BrightStor ARCserve Backup has three convenient methods to select the files to restore. When selecting the data (the source) to restore, you can use any of the following methods:

- **Restore by File System**—Select the files to restore and determine the media containing these files.
 - This method restores a specific directory or drive from a display of files and directories backed up with BrightStor ARCserve Backup. Select a client to display the drives or volumes, directories, and files that were backed up.
 - Use this method when you do not know which media contains the data you need, but you have a general idea of what you need to restore and which machine it came from.
 - Restore by File System is the default method. This view provides you with the best overall picture of your network and the files you wish to restore.
- **Restore by Session**—Using Restore by Session you can select the session, and 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 method uses the BrightStor ARCserve Backup database to locate sessions; if the database is stopped, this method of restore will not work.
- **Restore by Backup Media**—Place your source media in the drive and select an entire media session to restore.
 - This method displays all the device groups and media currently attached to the host. You select the device group and media, and specify the sessions on the media to restore.
 - This method restores a complete backup session from a specified media in a storage device. Because the Restore by Backup Media method does not use the BrightStor ARCserve Backup database, you cannot pick individual files and directories to restore from a directory tree.
 - All files in the session are restored to the destination, unless filters are added to the restore job.
 - Use this method if a tape was created by a different version of BrightStor ARCserve Backup, if it was created in Alexandria 4.5, or if the database does not recognize it.

Restore Data Backed Up Using Staging

The process for restoring data that was backed up using the Disk Staging Option 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 the Disk Staging Option, and the backed up data has been copied to its final destination media, the data can reside in two locations (the file system 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 faster than tape-based restores.

To restore data that was backed up using staging, perform the following steps:

1. Open the Restore Manager and select the Restore by File System method.
2. In the left pane of the Restore Manager, select the volume, drive, directory, or file you want to restore.
3. Click the Version History button.

BrightStor ARCserve Backup searches the databases and the Version History dialog opens displaying a list of all backed up versions of this file, directory, drive, or volume.

4. From this list, select the version you want to restore.

Note: Restoring data from a disk is generally faster than restoring from a tape, because there are no delays due to the tape load and seek latency. If 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.

- If you want to restore directly from the final destination, click OK to start the restore process.
- If the you want to restore from a different location rather than from the final destination, click the Duplicates button.

The Duplicate Sessions dialog opens displaying any sessions which are duplicates or clones of each other (including the original session). The clones could have been created by the staging process or by the tapecopy process. If the selected session has no duplicates, the Duplicates field will be blank.

For each copy of the selected session, the Duplicates Session dialog displays the Modified Date, Size, Media Name, Backup Time, Session #, Type, and Media Type to help you decide the location from where you want to restore from.

After you select the session and click OK, the restore process will start.

The Role of the Database in the Restoration Process

The three restore methods display source information either by using the BrightStor ARCserve Backup database, or by reading what devices are currently attached to the BrightStor ARCserve Backup host. Restore by File System and Restore by Session use the database; Restore by Backup Media does not.

When using Restore by File System or Restore by Session, the information for the source tree is taken from the database and includes information based on all the backups made since you started using BrightStor ARCserve Backup. BrightStor ARCserve Backup can help you find the volume, drive, directory, and files you need to restore. Unless you have purged or pruned your Job records or media records from the database, you should be able to restore all files from backups you have performed with BrightStor ARCserve Backup.

The method you choose depends on what you know about the files you want to restore and the media you will need to use.

You must use the Restore by Backup Media View to save restore jobs to a script. You cannot use scripts when using either the Restore by File System, or Restore by Session.

Multiple Code Pages

For more information about code pages, see How BrightStor ARCserve Backup Supports Multiple Code Pages in the chapter "Backing Up Data."

Specify Code Pages in the Restore Manager Window

You can change the code page on all tree items in the source tree. To specify a code page, use the following steps:

1. From Source tab on the Restore Manager window, right-click the node, volume, or directory for which you want to specify a code page.
2. From the Display Encoding right-click menu, select the desired code page.


BrightStor ARCserve Backup applies the new code page settings immediately.

Restore Manager Destination Tab

Using the Destination tab, you can either enable the Restore files to their original location(s) option or select directories and drives as the destination to which you want to restore.

Important! *The Restore files to their original locations option is not supported for local Full restores. In addition, do not restore volumes on servers that have BrightStor ARCserve Backup running. Doing so will overwrite existing files and may cause BrightStor ARCserve Backup to crash.*

Before you can submit a restore job, the client agent software must be running on all selected machines. If the agent service is not available, the square next to the node is inaccessible. If you restarted the client agent, the connection must be refreshed. To refresh the connection, collapse and expand the branch for the workstation.



If you select directories and drives as the destination, the Host and Path fields display the host and path associated with the object you select. You can also use these fields to search for a specific object; The Host field allows you to search for a particular BrightStor ARCserve Backup node, while the Path field supports the search of a directory or file. These entry fields can be used together to search for an object on a particular server.

Notes:

- When using these fields to search for an object, the information you enter must be an exact match. In addition, the object found is only highlighted, not selected.
- If you enable the Restore files to their original location(s) option, BrightStor ARCserve Backup automatically restores your files to their original machine and path.
- The Restore files to their original location(s) option is supported only for UNIX file system, Linux file system, NT file system, and NT System State restores. It is not supported when you are attempting to restore backed up BrightStor ARCserve Backup database agent files, or a database from the command line.

Agent View Tab

Before you select an object in the left pane of the Restore Manager, BrightStor ARCserve Backup displays the Agent View tab in the right pane. The following view options are available:

- View all file system agents—(Default) Shows all workstations entered into the current host server's Remote Client database, regardless of whether Agent software is running. Select this option to select an entire node without browsing its contents.
- View Agents which are running—Allows you to browse and select objects associated with machines running the appropriate Agent software and entered into the BrightStor ARCserve Backup Agent database.

Note: The agent software must be running on all selected machines prior to submitting any restore jobs.

Restore Manager Object Information Tab

When you select an object in the left pane of the Restore Manager, the right pane displays two tabs, Object Information and Security. The Object Information tab gives you information about the object you select in the left pane of the Restore Manager. The information on the Object Information tab varies depending on the destination you choose in the left pane of the Restore Manager. The type of information provided is identical to the information provided on the Object Information tab on the Source Tab. This tab is purely informational in nature and requires no input.

Restore Manager Security Tab

The Security tab allows you to enter the User Name and User Password for your restore operation. The Security tab is associated with the selection of a node or machine, and is used to enter your system password for the server currently highlighted.

Restore Job Filters

For restore jobs, filtering is performed on a global basis. To select filters for the entire job, click the Filters button on the Restore Manager toolbar to invoke the Filters dialog.

If you select the include directory pattern filter and do not specify an absolute path, empty directories for all the directories that do not match the user provided criteria will be restored. To avoid creating these empty directories during restore, disable the global restore option Create Empty Directories when creating your restore job. For more information on this option, see Create Empty Directories Option in this chapter.

Note: The Restore Manager does not support the use of File Changed filters.

For more information on selecting and applying filters, see the chapter "Customizing Your Jobs."

Important! *Be very careful when specifying filters for your restore operation! Incorrectly applied filters may not restore the data you need, and can result in lost data and wasted time!*

Restore Options

When creating a restore job, you can use the Option tabs to customize the restore job.

To open the Option dialog, click the Option toolbar button. The available Option tabs are as follows:

- Pre/Post—Specify commands to run before or after your restore job is complete.
- Log—Specify how BrightStor ARCserve Backup logs activities and identifies the output devices for the activity logs.

- **Media Rules**—Set timeout periods that control how long BrightStor ARCserve Backup waits for appropriate media to be inserted into devices.
- **Destination**—Specify how the directory structure is to be created on the destination during a restore job and to set the rules that control how BrightStor ARCserve Backup responds when encountering a file name conflict between the files on the source and the files on the destination.

Pre/Post Options

BrightStor ARCserve Backup allows you to specify commands to perform various operations immediately before or after your data is merged based on the exit codes received. You can specify pre/post commands:

- For the entire merge job. These commands are executed at the beginning or end of the job.
- Run commands, executable programs, and shell scripts.

For example, you may want to specify this option to load a virus scanning application before files are merged and then run the batch file you created that sends a detailed report to the printer.

Note: If you specify pre/post commands for the entire job, as well as for machines in the job, the global pre/post commands are executed before or after the job starts or ends, and local commands are executed when the selected machine is backed up.

The Pre/Post option allows you to run a command on your BrightStor ARCserve Backup server machine before, after, or before and after the job is executed. The outcome of Pre/Post commands can be found in caqd.log. The available Pre/Post options are defined in the following sections.

Restore Manager Run Command Before Job Options

Enter the path and name of the application to be executed on the machine before the job starts.

- On Exit Code— If you want BrightStor ARCserve Backup to detect exit codes of any application, enable this option, select the condition for detecting exit codes (Equal To, Greater Than, Less Than, or Not Equal to), specify the exit codes you want detected, and then select how you want BrightStor ARCserve Backup to respond after exit codes are detected:

Note: If you select the condition Greater Than or Less Than, only one exit code should be specified.

- Skip Delay—If the exit code conditions are met, run the job immediately, ignoring the delay.
- Skip Job—If the exit code conditions are met, skip the job.
- Skip Post Application—If the exit code conditions are met, enabling this skips commands scheduled to run after the job completes.

- Delay in Minutes—Specify the delay in which BrightStor ARCserve Backup waits before running a job when the appropriate exit code is detected. This gives the specified application time to finish processing before your job begins.

Note: Ensure that the path to the command is correct. For example, to run the script `post_exec.ksh`, and the command `post_exec.ksh` is in the `usr` directory, enter this:

```
/usr/post_exec.ksh
```

Restore Manager Run Command After Job Options

BrightStor ARCserve Backup runs the command after the restore job finishes. Enter the path to, and name of, the application to be executed on the machine after the job completes. As with the Run Command Before Job option, you must ensure that the path to the command is correct.

Restore Manager Run Before/After Command As Options

Enter the User Name and Password to run Pre/Post commands. The system on the selected host server requires the User Name and Password to check the system privileges on that server. The user is authenticated using the File System agent running on the server machine. The File System agent must also be installed and enabled on the server machine to run Pre/Post commands.

Do not confuse the User Name and User Password entered into these fields with the BrightStor ARCserve Backup User Name and Password.

Restore Manager Log Options

Log options determine the level of detail included in the log report for the operation and the devices to which the log will be sent. You can view the log report from the Job Status Manager or Database Manager (Job Records). BrightStor ARCserve Backup provides the following log options:

- Log File Name—Enter a name for the log file.
- Log All Activities—Record all of the activity that occurs while the job runs in the Job Log.
- Log Summary Only—Record summary information for the job (including source, destination, session number, and totals) and errors.

You can also select the output devices for the activity log. Select any or all of the following:

- Unicenter NSM Alert—Use this option to send a message to the Unicenter Console when an alert is generated.
- SNMP Alert—Use this option to send messages to your SNMP messaging console.
- Printer Name—Use this option to send messages to a specified printer local to the BrightStor ARCserve Backup server. The BrightStor ARCserve Backup server selected to send print, must be configured to do so, for this option to work.
- Internet Email—Use this option to send messages to the specified email address. The BrightStor ARCserve Backup server selected to send messages via email, must be configured to do so, for this option to work.

Note: You can enable the activity log destination options (NSM alert messages, printers, and email addresses) by modifying the configuration file named `caloggerd.cfg`. This configuration file is located at: `$BAB_HOME/config`.

Restore Manager Media Rules Options

The Restore Manager supports the following media rules options:

Timeout Options

Specify a timeout period, during which BrightStor ARCserve Backup waits for you to provide the media you need to restore your data. Available Media options are:

- **Timeout for the First Media**—Period of time
BrightStor ARCserve Backup waits for you to insert the first media required for your restore job. If the time expires and do not provide the media, the job fails.

By default, BrightStor ARCserve Backup waits five minutes.

- **Timeout for Additional Media**—Period of time
BrightStor ARCserve Backup waits for you to provide any additional media required for your restore job. By default, there is no timeout.

Optimize Restore

If, during a restore operation, BrightStor 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 BrightStor 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.

Destination Options

Using Destination options, you can determine how the directory structure is created on the destination when BrightStor ARCserve Backup restores files, and which files (if any) BrightStor ARCserve Backup can overwrite.

Directory Structure Options

Select the method to be used to create directories on your destination. The Directory Structure options are:

- **Create Directories from the Base**—Default setting. Creates the destination path beginning from the base directories on the destination. A base directory is considered the directory in which the selected file or directory resides in the source path.
- **Create Entire Path from the Root**—Create the entire source path, except the root drive or volume name, on the destination.
BrightStor ARCserve Backup does not restore files from a parent directory. BrightStor ARCserve Backup creates only the directory path to the base directory.
- **Do Not Create the Base Directories** —Do not create the base directory on the destination path, but create all subdirectories below the source base directory.

File Conflict Resolution Options

Select the method that BrightStor ARCserve Backup should use when there are files on the destination disk that have the same name as files being copied from the source:

- **Overwrite All Files**—Default setting. Restore all source files to the destination regardless of conflicting file names. The files from the source overwrites existing files on the destination.
- **Rename Files**—Copy the source file to the destination with the same file name but a different extension. The extension maintains the first two characters of the original, but the last character will be 1, 2, 3..., depending on how many files BrightStor ARCserve Backup encounters with the same name. BrightStor ARCserve Backup renames files without extensions with the following extensions: .AS1, .AS2, and so on.
- **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.

Create Empty Directories Option

An empty directory is a directory that does not contain subordinate directories or files. Using the Create Empty Directories option you can direct BrightStor ARCserve Backup to create an empty directory during the restore job when the directory is empty. When you disable this option, BrightStor ARCserve Backup does not create an empty directory during the restore job.

The ca_restore Command

BrightStor ARCserve Backup provides a command line interface that lets you perform a variety of restore functions without using the Restore Manager, giving you an alternate method of accessing almost all of the operations available through the Restore Manager from the command prompt. You can use this command to create and submit restore jobs to the BrightStor ARCserve Backup queue, and to set all associated options.

You can use the options and switches for the ca_restore command to set global options and filters, select your source and destination for the restore job, and submit the restore job to run immediately or at a scheduled time. You must specify them in the stated order.

For a complete list of the options and switches available for this command, see the appendix "Using Command Line Utilities."

Chapter 5: Customizing Your Jobs

BrightStor ARCserve Backup provides a number of methods to customize your jobs to suit your needs.

- Filters allow you to select the files and directories to be included in, or excluded from, your backup and restore jobs, based on a wide variety of criteria.
- Rotation schedules allow you to define standard and consistent intervals at which to rotate and retire backup media.
- Scheduling options provide you with the ability to schedule your jobs to run immediately, later, or on a regular basis.
- 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.
- The Generic Job Manager is a powerful tool that allows you to schedule generic jobs to run on the BrightStor ARCserve Backup server.

This chapter discusses these customizations methods in further detail.

Job Filters

Filters allow you to include or exclude files and directories from your backup and restore jobs. The filters applied to backup and restore jobs are identical, giving you the ability to use a variety of criteria to filter the data you back up and restore.

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.

In general, click the Filter button on the Manager toolbar to access global filters for your job. Access node-level filters for backup jobs from the Filter tab of the object's Source tab.

You can include or exclude files based on the following criteria:

- Specific file names, patterns or attributes.
- Specific directory names or patterns.
- Files accessed before, after, between, or within a specific date range.
- Files modified before, after, or between a specific date range.
- Files changed before, after, or between a specific date range.

Note: On Windows-based file systems, this filter will behave as a file **created** before, after, or between a specific date range filter.

BrightStor ARCserve Backup uses wild card or replacement characters, except when it detects that an absolute path is specified. If a valid absolute path is specified, BrightStor ARCserve Backup will only exclude (or include) the absolute path specified, rather than excluding (or including) more directories, as it would for regular expression.

Important! *Be very careful 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!*

Filtering Precedence

When you use a combination of filters, the following filtering precedence applies:

- Directory Pattern filters take precedence over File Pattern filters.
- Exclude filters take precedence over Include filters.

For more information about Include, Exclude, File Pattern, and Directory Pattern filters, see Include Filters, Exclude Filters, File Pattern Filters, and Directory Pattern Filters in this chapter.

Examples

The following table describes examples of filtering conditions and the expected results:

Condition	Result
Directory Pattern specified	Both node level and volume level filters are favorable. Otherwise, the directory and its sub-directories are missed.
File Pattern specified	Files specified will be backed up and restored when node level and volume level filters are favorable.
Directory Pattern and File Pattern specified	Node level directory filters and the volume level directory filters are favorable. Node level file filters and volume level file filters are favorable. Exception: If you specify an Include Directory Pattern for the directory, explicitly or implicitly, all files in the directory are backed up or restored irrespective of the File Pattern filter.
Time Pattern specified	Node level time filters and volume level time filters are favorable for file backups and restores. Exception: If you specify an Include Directory Pattern for the directory, explicitly or implicitly, all files in the directory are backed up or restored irrespective of the Time Pattern filter.
File Pattern and Time Pattern specified	All node level and volume level filters are favorable. Exception: If you specify an Include Directory Pattern for the directory, explicitly or implicitly, all files in the directory are backed up or restored irrespective of the Time Pattern filter.
Directory Pattern, File Pattern, and Time Pattern specified	All the node level and volume level filters are favorable. Exception: If you specify an Include Directory Pattern for the directory, explicitly or implicitly, all files in the directory are backed up or restored irrespective of the Time Pattern filter.

Include Filters

Use an Include filter to ensure that you only back up or restore needed files. This reduces the time it takes to perform a backup or restore operation because the results contain only those files that satisfy the filter specifications.

For example, if, in the source area, you specify a backup of your entire local host, and you then set a Directory filter to include files in the /opt directory, BrightStor ARCserve Backup only backs up files from your /opt directory. No other files are backed up.

Note: If you select the include directory pattern filter and do not specify an absolute path, empty directories for all the directories that do not match the user provided criteria will be backed up. To avoid creating these empty directories during restore, disable the global restore option Create Empty Directories when creating your restore job. For more information on this option, see Create Empty Directories Option in the chapter "Restoring Data."

Exclude Filters

You can set Exclude filters to remove unnecessary files from your backup and restore operations, to ensure that only needed files are backed up or restored. Excluding unnecessary files saves time, CPU usage, and media.

Exclude filters always take precedence over Include filters. If you set a filter to include all files of a particular type, and another filter to exclude all files from a particular directory, the specified type files in the excluded directory are excluded from the operation.

File Pattern Filters

Use File Pattern filters to include or exclude certain files from a job. Specify a particular file name, or use wildcards to specify a file pattern.

For example, you can exclude all files with a .bak extension. If you do not know the specific file name, you can provide as much of the file name as you know, and use wildcards to fill in the blanks.

Note: The File Pattern filter can be applied locally (per host) or globally (per job).

Syntax for Entering Path and File Patterns

The following sections show you how to specify path and file names for UNIX and non-UNIX servers when using a file pattern filter.

UNIX Host Server File Filtering

To specify a file in any directory, use FILE. For example:

```
memo
```

where the PATH is defined as a series of directories following the long file name format (max of 255 chars), separated by forward slashes (/), and the FILE is defined as the name of the file to filter, following the long file name format (max of 255 chars).

When applying filters to a host server, to specify an absolute path and file name, use the pattern /PATH/FILE. For example:

```
/tmp/memo
```

Non-UNIX Server File Filtering

When applying filters to a non-UNIX server, to specify an absolute path on the server volume, use the pattern VOLUME:\PATH\FILE. For example, for NetWare:

```
SYS:\SYSTEM\TSA\*.NLM  
SYS:\PUBLIC\LOGIN.EXE
```

To specify a file pattern when the file could be in any directory on the specified server volume, use the pattern VOLUME:FILE. For example, for NetWare:

```
SYS:*.NLM  
LOGIN.EXE
```

where VOLUME is defined as a non-UNIX server volume, PATH is defined as a series of directories separated by backslashes (\), and FILE is defined as the name of the file to filter, following the 8.3 format.

To specify an absolute path on a server, use the pattern \PATH\FILE. For example, for NetWare:

```
\SYSTEM\TSA\*.NLM  
\PUBLIC\LOGIN.EXE
```

Directory Pattern Filters

Use a Directory Pattern filter to include or exclude specific directories from a job. This filter can be applied locally (per host) or globally (per job). All files under the directory you specify are included or excluded, depending on your choice. For example, if you set a filter to include the /user directory, the filter includes all of the files in the /user directory with your job.

Directory Pattern filters always take precedence over File Pattern filters. For example, if you select the Include Directory Pattern filter and the Exclude File Attribute filter, all files in the directory will be included in the job.

As with the File pattern filter, you can enter an entire directory name or provide a pattern the directory name follows. If you enter `use*`, directories starting with the letters *use* are included with your job (user, users, userdoc, and so on).

Path and Directory Pattern Syntax

The following sections show you how to specify path and directory names for UNIX hosts and non-UNIX servers when using a directory pattern filter.

UNIX Host Directory Filtering

To specify that the directory to filter may be in any other directory, use the pattern `DIR_NAME`. For example:

```
pete
```

where `PATH` is defined as a series of directories following the long file name format for (maximum of 255 chars), separated by forward slashes (/), and `DIR_NAME` is defined as the name of the directory to filter, following the file name format (maximum of 255 chars).

When applying filters to a UNIX host directory, use the pattern `/PATH/DIR_NAME` to specify an absolute path to the directory to filter. For example:

```
/usr/pete
```

Non-UNIX Server Directory Filtering

When applying directory filters to a non-UNIX server, to specify an absolute path on the server volume, use the pattern `VOL:\PATH`. For example:

```
SYS:\BrightStor ARCserve Backup\DATABASE  
SYS:\BrightStor ARCserve Backup\*BASE
```

To specify an absolute path on a server, use `\PATH`. For example:

```
\BrightStor ARCserve Backup\DATABASE  
\CAS*\D*
```

To specify that the directory could be in any other directory on the specified server volume, use the pattern `VOL:DIR_NAME`. For example:

```
SYS:DATABASE  
SYS:DATA*
```

To specify that the directory may be in any other directory on any file server, use DIR_NAME. For example:

```
DATABASE  
D*BASE
```

where VOL represents a non-UNIX server volume, PATH is defined as a series of directories separated by backslashes (\), and DIR_NAME is defined as the name of the directory to filter, following the DOS file name format (8.3).

File Attributes Filters

Use the File Attributes Filter to include or exclude files that match the attributes you specify. You can specify more than one attribute for this filter, but only those files that match all the attributes you specify will be included or excluded. You can specify the following file attributes:

- Hidden—Files that do not appear in a directory listing.
- 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.

Important! *The File Attributes Filter can be used only for non-UNIX agents.*

File Modified Filters

Use the files last modified attribute to include or exclude files based on the time a file's contents 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.
- On or 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 Changed Filters

Use the files last changed attribute to include or exclude files based on the time a file's contents or attributes (for example, permissions, owner information, and so on) changed. 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.
- **On or 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.

Notes:

- On Windows-based file systems, file **changed** filters behave as file **created** before, after, or between a specific date range filters.
- The Restore Manager does not support the use of File Changed filters.
- Incremental backups and differential Backups are based upon the File Modified date rather than the File Changed date to filter the files to be backed up.

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.
- On or 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.

When setting filters, remember that not all file systems record access dates. Therefore, the File Accessed Filter may not be available for your job.

Important! *Incremental backups and differential backups are based upon the File Modified date rather than the File Accessed date to filter the files to be backed up.*

File Size Filters

File size filters let you to filter files based on the file size.

Filter Types

There are four types of filters from which to choose:

- Equal to--Include or exclude files equal to the specified size.
- Greater than--Include or exclude files greater than the specified size.
- Less than--Include or exclude files less than the specified size.
- Between--Include or exclude files whose size falls between the two file sizes specified.

Filter Sizes

There are four types of filter sizes to specify:

- Bytes
- Kilobytes
- Megabytes
- Gigabytes

Rotation Methods

Rotation methods provide you with predefined patterns for your backup operations, including the schedule for performing backups, the method to use for each backup, and the rotation of the backup media. The backup method determines which files in a selected drive or directories are backed up and the rotation schedule defines the rules for use and retention of media.

BrightStor ARCserve Backup provides three backup schedule strategies from which to choose. You can use one of the BrightStor ARCserve Backup predefined backup methods, simple rotation, or Grandfather-Father-Son (GFS) rotation, or you can customize a backup scheme to suit your environment's needs. You determine the type of backup to be performed, and the rules governing overwriting, or appending your data to media. For more information about the Custom backup method, see the chapter "Backing Up Data."

Choose a Backup Method

The backup method you use depends on your needs. Examine each method and each rotation, and decide which is the best method for your environment. For example, a full backup is easiest to restore, since all the most recent data is on one media. However, full backups take longer to perform than incremental or differential backups, and if you have a large amount of data to back up, it may be unreasonable to perform full backups every day.

Some of the questions to consider when selecting a backup rotation method are:

- How much data is backed up?
- What percentage of data in my set changes each day?
- How long do I want the backup to take?
- If I have to recover files, how long do I want to spend restoring?
- How much media do I need?

Incremental Backups

With the incremental method, your backups take less time because of the way the incremental method flags data for backup. If you have large amounts of data to back up every day, the best solution may be to choose daily incremental backups.

An incremental backup includes only the files that have changed since the last backup (full, incremental, or differential). In general, an incremental backup is the quickest method for backing up. However, to recover a client, you need the last full backup, and all subsequent incremental backups.

For example, if you perform full backups on Monday, and incremental backups Tuesday through Friday, and on Friday you have to restore a host, you need Monday's full backup and all of the incremental backups from Tuesday, Wednesday, and Thursday to recover all the files for that host.

Differential Backups

If you want to recover files quickly and are not concerned about how long the backup takes, choose the differential method. Differential backups restore files quickly, because only files that have changed since the last full backup are backed up.

Since differential backup jobs are based on the file's modification date since the last full backup, files that were backed up in the last differential job are backed up again.

For example, using the same backup schedule as in the previous example, but with differential backups, on Friday you would need Monday's full backup, and Thursday's differential backup media to restore all the files for that host. This is because Tuesday's differential backup contains all the files that have changed since Monday's backup, Wednesday's differential backup contains all the files that have changed since Monday's backup (thus, all the files from Tuesday's backup will also be on Wednesday's backup), and so on.

Simple Rotation Schedule

A simple rotation schedule allows you to define a regular pattern of backup jobs, maintaining a weekly rotation of media using a combination of two of the three available backup methods—full, differential, and incremental. You can choose to use a five or seven-day rotation schedule, set exceptions to the schedule as well as days when no backup will be performed, specify a media pool for your media, and choose whether to append backup data to your media or overwrite the media.

Note: For GFS rotations, you cannot change the media name.

GFS Rotation Schedule

BrightStor ARCserve Backup allows you to select from pre-defined GFS rotation schemes consisting of full weekly backup jobs combined with daily incremental and differential jobs. The GFS schemes are methods of maintaining backups on a daily, weekly, and monthly basis.

The primary purpose of the GFS scheme is to suggest a minimum standard and consistent interval at which to rotate and retire your 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.

A five-day GFS rotation policy requires approximately 21 media per year, while a seven-day policy requires approximately 23 media per year. For both of these schedules, the amount of media needed can vary depending upon your 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.

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).

Note: For GFS rotations, you cannot change the media name.

Custom Rotation Schedule

Although rotation schedules are predefined, you can modify their rules to suit your needs. Click either the Rotation or GFS Rotation option on the Method/Schedule tab to access the tabs related to rotation rules to customize your simple or GFS rotation jobs. The following section provides information about the available customization options.

Cycle Table for Rotation and GFS Rotation

BrightStor ARCserve Backup lets you select from the six predefined rotation schedules, available on the drop-down menu at the top of the Cycle Table tab:

- 5-day full backup
- 5-day weekly incremental backup, full backup on Friday
- 5-day weekly differential backup, full backup on Friday
- 7-day full backup
- 7-day weekly incremental backup, full backup on Sunday
- 7-day weekly differential backup, full backup on Sunday

When you select one of these rotation schedules, the Cycle Table allows you to view the schedule you have chosen. Columns display the Day of the Week, Execution Time, Method, Media Name, and Mode for each day in the rotation.

To modify the backup method or media name (simple rotation only) of any day of the week in the rotation, double-click the day you want to change to invoke the Job Unit dialog. Modify the available fields in the Job Unit dialog, and click OK. Using the Job Unit dialog to modify your rotation schedule permanently changes the method or media name used on that day of the week for the life of the rotation schedule. You cannot change the day of the week or the execution time using this dialog.

If you are defining a rotation schedule, you can instruct BrightStor ARCserve Backup to append your data to the media or overwrite the data existing on your destination media. However, full backups and the first incremental or differential backup of the week overwrites the destination media.

Missed Targets

When submitting either a simple rotation or a GFS rotation job, the Cycle Table tab also allows you to reschedule the backup of any missed workstations or file servers to a specified time. Select the Retry Missed Targets field and set a specific time for BrightStor ARCserve Backup to try again.

There are several reasons why a GFS rotation job would miss a target workstation:

- The agent is not running.
- The workstation has been turned off.
- The proper media was not mounted before the job timed out.
- The device drive is off.
- The GFS rotation job was stopped while its status was ACTIVE (while the job was running).

Any of these reasons can cause a target to be missed. Always make sure you check your GFS rotation logs and the Activity Log for important information about any GFS rotation job. Missed targets are always recorded in these logs.

Media Rules for Rotation and GFS Rotation

By default, BrightStor ARCserve Backup sets up a media pool for all GFS rotation schedules. The Media Rules tab allows you to define the prefix used for the media pool in this rotation schedule and the media pool information.

- **Pool Name**—The name of the media pool with which the destination media is associated.
- **Owner Name**—The name of the user who created the media pool.
- **Base Serial Number**—The number BrightStor ARCserve Backup uses to begin automatically assigning serial numbers to the media in the media pool. The first media BrightStor ARCserve Backup formats has the same serial number as the Base number. Thereafter, each media's serial number increases by one.
- **Next Serial Number**—The serial number to be assigned to the next media to be added to the media pool.
- **Maximum Serial Number**—The highest possible serial number in the media pool.
- **Serial Number Increments**—The number by which the serial number increases.
- **Minimum Number of Media in Save Set**—The fewest number of media to be retained in the Save Set before the oldest media can be recycled to the Scratch Set (simple rotation only).
- **Retention Period (days)**—The minimum number of days media is kept in the media pool Save set before it can be released to the Scratch set and overwritten (simple rotation only).
- **Timeout for First Media**—Specify the number of minutes BrightStor ARCserve Backup will wait for a media to be inserted into a drive before it cancels the job.
- **Timeout for Span Media**—The number of minutes BrightStor ARCserve Backup waits for an additional media to be inserted into a drive before it cancels a job.
- **Eject Media**—Lets you specify to eject the media from the drive after the job finishes. This prevents another job from overwriting information on the media.

In addition to this information, when you select a GFS rotation schedule, the Media Rules tab allows you to define the Preserve Media options. These options let you modify the default number of media needed for your daily, weekly, and monthly backups or the total number of media needed for your GFS rotation schedule.

Calendar View for Rotation and GFS Rotation

Use the Calendar View tab to view or to customize your rotation schedule according to the types of backups you want on specific dates. Each date in the calendar view displays an icon indicating the type of backup operation scheduled for that day. The method icons are:

-  Off—No backup job is scheduled for this day.
-  Full—A complete backup of all files is scheduled for this day.
-  Incremental—BrightStor ARCserve Backup compares the source file modification dates with the date of the last backup. Only files that have been modified or changed since the last backup are included in the backup job scheduled for this day.
-  Differential—BrightStor ARCserve Backup compares the source file modification dates with the date of the last full backup. Only files that have been modified since the last full backup will be included in the backup job scheduled for this day.

To change the backup method scheduled for a particular date, right-click the method icon to select another type of backup operation. You can only change to a backup method available under the rotation schedule you selected on the Cycle Table tab. For example, if you selected a rotation schedule that includes only differential and full backup methods, you can change a particular date's backup method to an Off, full, or differential backup method. You cannot change the scheduled backup method to incremental, because there are no incremental backups in your rotation schedule.

This feature only allows you to specify exceptions to the backup method scheduled for specific dates (for instance, if a holiday falls on Wednesday, usually a backup day). Modified backup methods appear in red on the Calendar View tab.

Exception View for Rotation and GFS Rotation

The Exception View tab lists any exceptions you may have added to specific dates under the predefined rotation schedule. It does not display any changes you made using the Job Unit dialog on the Cycle Table tab. You can use this tab to modify or delete any previously defined exceptions, or click Add to define particular days on which the backup method or the date should differ from your pre-existing rules.

The Exception dialog lets you select specific days off, or define different backup criteria for days in your rotation on which you want to deviate from the standard schedule, such as on holidays. Any modification you make to the standard backup pattern appears in red on the Calendar View tab.

Job Scheduling and Submission

Any kind of job can be submitted and scheduled in essentially the same way using the Submit button available on the toolbar in each Manager.



Submit—Click the Submit button to open the Submit dialog.

The Submit dialog allows you to instruct BrightStor ARCserve Backup to run your job immediately, schedule your job for a later date or time, or submit the job with an initial status of Hold.

Important! *All scheduled times for BrightStor ARCserve Backup jobs are based upon the time zone where the BrightStor ARCserve Backup server is located. If your agent machine is located in a different time zone than the BrightStor ARCserve Backup server, you will need to calculate the equivalent local time that you want the job to be run.*

You can also enter a job description to better identify your job, and you can save your job as a job script.

Important! *When using the Backup Manager, Restore Manager, or Generic Job Manager to submit a job, BrightStor ARCserve Backup considers the underscore character (_) a special character and cannot be used for a job description. (The Merge Manager and Scan Manager do not have these job description restrictions). If you are attempting to submit a job using Backup Manager, Restore Manager, or Generic Job Manager, do not use an underscore character in the job description.*

Submit a Job Using the Run Now Option

Specify Run Now when the job you are submitting is a one-time job to be executed immediately or you want to monitor the job as it is running.

You should not specify Run Now if your storage device is currently busy. If the device is busy, BrightStor ARCserve Backup reports that the storage device is busy, and the backup job is not submitted to the job queue, and the job will not run successfully when a device group becomes available. Instead, you are asked to wait until the device is available before resubmitting the job. If you want to prevent this situation, schedule your backup job, keeping the current date and time, rather than submitting a Run Now job. This way, when BrightStor ARCserve Backup discovers that the storage device is busy, it automatically retries the backup job until the drive becomes free.

Important! *All scheduled times for BrightStor ARCserve Backup jobs are based upon the time zone where the BrightStor ARCserve Backup server is located. If your agent machine is located in a different time zone than the BrightStor ARCserve Backup server, you will need to calculate the equivalent local time that you want the job to be run.*

Schedule Option

You should schedule your job when:

- You submit a job, but want it to run at a specific time, rather than immediately.
- You submit a job that should run regularly. This is especially useful when setting up a media rotation schedule for your network.
- Your storage device is currently 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.

For details on how to specify a scheduling option, see the online help.

Important! *All scheduled times for BrightStor ARCserve Backup jobs are based upon the time zone where the BrightStor ARCserve Backup server is located. If your agent machine is located in a different time zone than the BrightStor ARCserve Backup server, you will need to calculate the equivalent local time that you want the job to be run.*

Submit Job on Hold Option

If you are not sure of the exact time you want a job to run, BrightStor ARCserve Backup allows you to submit your job on hold. You can select the Submit Jobs on Hold option only if you select the Schedule At option. Hold status indicates the job does not execute until the owner of the job (or a user with rights to operate on that job) changes the status to Ready, using the Job Status Manager. For more information, see the chapter "Using the Job Status Manager."

Job Scripts

A script is a job you saved to a file. It contains the original source, destination, options, and schedule information, if any, for the job. It also contains any filters you created to include and exclude files and directories. You can save any type of job as a script.

Creating a script has the following advantages:

- You can re-use the same settings at a later time.
- You can copy your settings to a different node running BrightStor ARCserve Backup.
- You can quickly resubmit regularly executed jobs after a job has been accidentally deleted.
- You can avoid recreating the same job over again.

All BrightStor ARCserve Backup scripts are stored and saved in the \$BAB_HOME/jobscripsts subdirectory with an extension of the BrightStor ARCserve Backup user who created the script.

Example: User-defined job script

A user logs in to the BrightStor ARCserve Backup domain as "supportuser." This user creates a job script and labels the script backup01. The complete name of the new job script is backup01.supportuser and it is stored in the \$BAB_HOME/jobscripsts directory.

Script Development

You can save any kind of job as a script after you have selected a source and destination for the job, and scheduled and customized the job, if necessary. Select Save, or Save As, from the Manager's File menu. Alternatively, from the Submit dialog, enter a description of the job, if desired, and click the Save button. Either method invokes the Save As dialog. The Save As dialog displays a list of previously saved scripts, identifying the user name, the script name, the job type, and the description, if any. This dialog also allows you to enter a name for your job script. Click OK.

Add a Script to the Job Queue

After you have saved your script, you can run the job by adding the script to the job queue using the Job Status Manager.

Note: In addition to the using the Job Status Manager to open a script, you can add a script from the Backup Manager window. To do this, select Open from the File menu on the Backup Manager window, then select the script that you want to open. For more information, see the online help.

To add a script to the job queue

1. From the Job Status Manager window, click the Add Job button, or select Add Job from the Status menu.

The Add Job dialog appears. The Add Job dialog lists the script previously saved on the selected host server.

2. Select the script you want to add to the Job Queue, and click the Add Job button.

The script you selected is added to the job queue.

Generic Job Manager

The Generic Job Manager lets you define, schedule, and submit generic jobs.

The benefits of using this manager include:

- Jobs can be scheduled and repeated.
- Jobs appear in the Job Status Manager.
- Jobs can be stopped in the Job Status Manager.
- It automates manual tasks that impact your system backups and resources.
- It automates schedules of new tasks based on the status of other job system routines.
- It provides an easy way to quickly package and submit jobs.

Some examples of how you can use the Generic Job Manager include:


- Ensuring media is available before performing backup procedures
- Invoking a tape copy after a backup completes
- Pruning a database in a timely manner

You can create, modify, and save generic job scripts using the File menu options Open and Save As.

Submit a Job Using the Generic Job Manager

To submit a job using the Generic Job Manager:


1. Open the Generic Job Manager.



2. On the Job tab, in the Program/Command to Execute field, either enter the command, enter the path and name of the script to be executed, or click Choose and select from the Program Choice list. For more information on the commands you can choose from the Program Choice list, see the online help.
3. Enter your user name and password. The user name and password correspond to that of the system of the host server selected, and it is required to check the system privileges on that server. Do not confuse the user name and password that you enter into these fields with the BrightStor ARCserve Backup user name and password.

Note: If you enter an incomplete or incorrect user name or password, the generic job does not execute because of an authentication failure. This security feature prevents unauthorized users from executing any commands, while posing as a super user.

4. Click the Schedule tab.



The Schedule tab lets you set up repeating intervals for when the job will run, without specifying a rotation schedule, or perform an unscheduled backup without affecting your regular rotation schedule.

5. If you want to run the job at repeating intervals, specify a repeat interval in months, days hours, and minutes, and identify any days of the week to exclude. If you want to run the job once, leave these fields blank.
6. If you want to customize your generic job by setting up pre and post job execution options that allow you to run a command or script on the host BrightStor ARCserve Backup server before and after the job is executed, or customize the job by setting up events to send messages to people in your organization using various methods of communication, click the Options button. For more information on these options, see the online help, or the section Pre/Post Options in the chapter "Backing Up Data," or the output device information in Log Options in the chapter "Backing Up Data."
7. Click the Submit button and select to run the job immediately or schedule it to run at a specific date and time. If you schedule the job and are not sure of the exact time you want the job to run, enable the feature Submit Job on HOLD.

Important! All scheduled times for BrightStor ARCserve Backup jobs are based upon the time zone where the BrightStor ARCserve Backup server is located. If your agent machine is located in a different time zone than the BrightStor ARCserve Backup server, you will need to calculate the equivalent local time that you want the job to be run.

8. Enter a description for information purposes that explains the type or functionality of the generic job that you are scheduling to run. This description appears in the description column in the Job Queue and can be used to identify the job.
9. Click OK to submit the job.

After you submit a job using the Generic Job Manager, it is labeled as a generic job in the Job Status Manager.

As an alternative to using the Generic Job Manager, you can submit generic jobs from the command line using the `ca_generic` command. For more information about this command line utility, see the appendix "Using Command Line Utilities."

Chapter 6: Using the Job Status Manager

The Job Status Manager is a graphical tool that helps you centrally manage BrightStor ARCserve Backup servers enterprise-wide. You can use the Job Status Manager to:

- Activate or deactivate the job queue.
- View all available BrightStor ARCserve Backup servers, job queues, and activity logs or job logs.
- Manage jobs—add, modify, reschedule, stop, run, or delete jobs. Only active jobs can be stopped and only inactive jobs can be deleted.
- Monitor progress of active jobs—view real-time status of active jobs in the queue.
- View job detail about job log information for all jobs that have been executed.

The left-panel browser navigates through BrightStor ARCserve Backup servers and their objects. All servers are arranged into groups, which you can configure. The right panel displays information on the job queue, job summary, and job log.


Servers and Groups Configuration

The Job Status Manager gives you the ability to set up user-defined groups containing customized views of jobs submitted to the job queue of any BrightStor ARCserve Backup host within a given BrightStor ARCserve Backup domain. Groups can be set up to view jobs operating only on certain servers. This makes it easy to manage your servers in groups. When you select a group, one or more of the following options are activated in the bottom left panel of the Job Status Manager:

- Add Group—Appends a new server group at the end of the list.
- Add Group Member—Adds a server to the selected group.
- Rename Group—Allows you to edit the group name.
- Delete—Deletes a group or a group member.

Job Queue

Every time you run a job with the BrightStor ARCserve Backup Manager, you submit it to the job queue. Information about all jobs (such as execution time, status, and owner) is stored here. BrightStor ARCserve Backup continuously scans the job queue for jobs that are waiting to execute.



The job queue is displayed on the right side in the top pane. It lists the Server, Job ID, Status, Execution Time, Job Type, Last Result, Job Description, and Owner of each job in the queue.

You can create a custom view of the job queue for each group using the Filters dialog to specify that the Job Status Manager show only the information you need when displaying jobs. For more information about job queue filters, see the section Filter Option in this chapter.

In addition, you can organize the jobs displayed in the queue according to your needs. By default, jobs are listed in order of execution time, but BrightStor ARCserve Backup allows you to sort job information using any of the Job Status Manager's column headings. Click a column heading once to reverse the order in which the information displays. For example, if the information currently displays in ascending order, click the column heading for BrightStor ARCserve Backup to sort the information in descending order.

Job Management

The Job Status Manager allows you to manage jobs that have been submitted to a server's job queue. A number of functions are available allowing you to add, delete, stop, run, or monitor a job, or to view information about jobs.

Job Management Using the File Menu

You can use the File menu at the top of the Job Status Manager to control the job queue and the BrightStor ARCserve Backup Queue, Media, and Database back-end services.

In addition to managing back-end services from the File menu, you can use the Status Bar Service icons at the bottom of the right pane to start an individual service, start and stop all services, and view the status of a selected service. For information about the File menu and the Status Bar Service icons, see Job Management Using the File Menu and Status Bar Service Icons in the chapter "Introducing BrightStor ARCserve Backup."

Job Management Using the Manager Console

The BrightStor ARCserve Backup Job Status Manager provides a Manager Console to allow you to view all activity, including Stop, Start, or any service communications failures related to BrightStor ARCserve Backup services such as the Database, Queue, or Media service. To view the Manager Console, select Show Manager Console from the Job Status Manager File menu. If necessary, click the Clear button on the Manager Console to clear the messages.

Job Status Manager Toolbar

A variety of functions are accessible using the buttons on the BrightStor ARCserve Backup toolbar and the Job Status menu. The following functions are available:



Add Job—Quickly submit a job to the queue using a previously saved script. Scripts contain original source, destination, option, and schedule information for the job.



Modify Job—Modify or add options or additional sources to an existing job without creating a new job.



Reschedule Job—Change a job's execution date, time, or status.



Stop Job—Cancel an active job from the queue and reschedule it for its next regular interval, if necessary. If you stop a job, the Last Result field displays Cancelled.



Run now—Available only for jobs that have the Ready status. If a device group is available, this option runs the job immediately.



Delete Job—Deleting an inactive job removes it from the queue completely. If you delete a makeup job, a new backup job is scheduled. Active jobs must be stopped before they can be deleted.



Job Monitor—Display detailed information about an active job being executed, including its progress in percentages.



Job History—Display information about an executed job, such as repeating jobs, rotation or GFS rotation jobs, and so on.



Filter—Open the Filter dialog to change the way jobs display in the queue.

Add Job Option

You can add previously saved job scripts to the job queue using the Add Job button. From the Job Status Manager toolbar, click the Add Job button, or select Add Job from the Job Status menu to invoke the Add Job dialog. Each script you have saved on the selected host server is listed in the Job Script field, identified by the name of the user who created it, the name of the script, the Job Type, and a Job Description. Select the name of the script you want to add and click Add Job. The script you selected is added to the job queue.

You can also use the Add Job dialog to delete a Job Script from the database. Select the script and click the Delete Script button to remove the script. For more information about using scripts, see Job Scripts in the chapter “Customizing Your Jobs.”

Modify Job Option

You can use the Modify Job button to change a job in Ready or Hold status. You cannot modify an active or a completed job. When you click the Modify button, the job you originally created appears in the appropriate Manager window, with the original sources, options, and destinations selected. For example, if the job is a backup or GFS rotation job, the Backup Manager appears with the selected job displayed. If the job is a scan job, the Scan Manager appears. Modify the job in the Manager and click the Submit button. Change or reconfirm the original execution information and confirm that you want to resubmit the job.

Reschedule Job Option

You can reschedule an existing job using the Reschedule Job button. You can only change a job's execution date, time, or status if the job is not active. If the job is active, the Reschedule Job button is inaccessible. However, if the job is a repeating job (specified when you created the job), you can use the Reschedule Job button to reschedule the job for the next automatic interval. The Reschedule Job dialog allows you to change the execution date or time. You can also change the job status to Hold or Ready.

- Ready—Indicates the job will be executed at the specified date and time.
- Hold—Indicates the job will not execute until the owner of the job changes the status to Ready.

If the job you want to reschedule is in active status, see Stop a Job in this chapter.

Stop Job Option

You can only stop jobs in active status in the queue. All stopped jobs remain in the queue and the Last Result column displays a result of Cancelled.

When you stop a job, only the selected instance of the job is stopped. Custom, one-time jobs are removed and do not run, but, for repeating jobs, only the selected instance is stopped. The next scheduled instance of the repeating job appears in the queue, set to run based on the job's execution cycle.

Job Status Manager Run Now Option

Available only for jobs that have the Ready status. 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. This option is useful if you want to run a job earlier than the time it was scheduled to run. It is also useful if a scheduled job does not run because of a hardware problem and you want to run it immediately after the problem is fixed.

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.

Delete Option

To delete a job from the queue, the job must display a status of hold or ready. You cannot delete a job displaying any other status (for example, done). Deleting jobs from the Job Status Manager window completely removes the job from the job queue.

If you delete a job and recreate it or submit it again from a job script, the new job is not associated with the deleted job and does not include its job history. If you delete a makeup job, a new backup job is scheduled.

Job Monitor Option

The Job Monitor displays the current status of a running job and is not available unless the job status is active. The Job Monitor displays information about the job and about what BrightStor ARCserve Backup is currently processing, including the percentage of the job completed and current status information. Messages written to the Job Status Activity Log can be monitored from this window.

Job History Option

The Job History window displays the entire, detailed, history of particular jobs in the job queue. To view the job history for a particular job, highlight the job in the job queue and click the Job History button on the Job Status Manager toolbar. If a history exists, it displays in the Job History window. BrightStor ARCserve Backup alerts you if a job history does not exist.

The Server, Job ID, Execution Time, Job Type, Last Result, and Job Description display in the top portion of the window, and the bottom portion of the Job History window displays any messages from the Job Log regarding the selected job. Use the Job History window to gather together all of the job records for a simple rotation or GFS rotation job in one screen.

Filter Option

Using the Filter dialog, you can choose the status and type of jobs displayed, select the columns to appear in the Job Status Manager, and define the maximum number of records to keep for the job history. Sorting the job queue is for informational purposes and does not affect the processing order of the jobs. You can remove unnecessary information from the job queue, depending on your specific needs.

On the Job Filter tab of the Filter dialog you can specify the following:

- **Job Status**—Select the type of current jobs BrightStor ARCserve Backup displays in the queue. You can choose to display current jobs in the queue with the following status: Active, Ready, Hold, or Owned by Other Users.
- **Done Jobs**—Specify that BrightStor ARCserve Backup displays done jobs in the database with the following results: Finished, Incompleted, Cancelled, Failed, and Unknown. You can also specify the number of days finished jobs remain in the queue using the Executed Within the Following Time Range fields.
- **Job Type**—You can have BrightStor ARCserve Backup display jobs selectively, based on their type or characteristic. Choose to display jobs with the following job types: Backup, Restore, Merge, Scan, Generic, and Data Migration.

On the Display tab of the Filter dialog, specify the columns BrightStor ARCserve Backup displays in the job queue. You can choose to view any or all of the following information about a job: Server, Job Status, Execution Time, Job Type, Last Result, Job Description, or Owner. You can also set the Maximum Number of records to be retained in the Job History window.

Job Status Indicators

When a job is in the queue, the status appears beside it. BrightStor ARCserve Backup has four job status categories:

-  **Ready**--A new, one-time or repeating job waiting to be executed.
-  **Hold**--The job is not scheduled for execution.
-  **Active**--The job is currently being executed.
-  **Done**--The job was executed. The result of a "Done" job can be either **Finished** successfully, **Failed** for some reason, **Cancelled** by the user, or **Incomplete** for some reason.

Completed jobs remain listed in the job queue for a specified number of days, as set on the Job Filter tab of the Filters dialog. The default setting is 3 days. To view jobs after this period expires, access the Filters dialog and change the time period to include the execution time of the jobs you want to see.

Last Result Indicators

The Last Result field displays icons indicating the status of the last executed job. If the last executed job was not successful, the information in the Last Result field helps you determine why the job may have failed. This field contains one of the following icons unless the job is still running:

-  **Finished**--All of the nodes and the job-specific source files were processed.
-  **Incomplete**--The job was partially successful. Some or all of the nodes and the job-specific source files were not processed. Review the Activity Log to check the exact nature of what occurred to prevent job completion.
-  **Canceled**--The job was intentionally canceled.
-  **Unknown**--This status indicates there was something wrong with the job and it could not be completed. Review the job log to identify the problem.
-  **Failed**--The job failed to perform its designated task. This status appears if the job was started, but the manager could not complete it. Review the Activity Log to determine what prevented job completion.
- **Null (blank)**--A null (blank) status indicates that a last result for the specified job does not exist. For example, the last result field will display a null value if the job is a custom job that has not yet been executed, or if it is the initial full backup job in a rotation or GFS rotation job.

Job Summary Report

The Job Summary report displays details about a job in Ready, Active, or Hold status, including the job type, source and destination targets, scheduling information, and options selected. To view the Job Summary, select a job in Ready, Active, or Hold status from the queue and click the Job Summary tab at the bottom of the screen. After a job is completed, the Job Status Manager no longer displays Job Summary information. Instead, it displays Job Log information accessible via the Job Log tab on the bottom of the screen.

Job Log Tab

The Job Log tab, in the bottom panel of the Job Status Manager, allows you to view Activity Log information about each job executed by BrightStor ARCserve Backup. The level of detail that appears depends on the Log Options you selected when you scheduled the job.

When you select a job, a detailed description of the job displays at the bottom of the screen. The information lists the number of complete/incomplete nodes and volumes, the media and media pool used, statistics such as sequence and session number, the serial number of the media, the device group, the total number of files backed up, the total number of files missed, and the total MB processed.

Note: You can also view a detailed session report for each job in `$BAB_HOME/logs/sessionSummary` that lists each session, start time, end time, and throughput for the session.

The ca_qmgr Command

BrightStor ARCserve Backup provides a command line interface that lets you perform a variety of functions from the command prompt, without using the Job Status Manager, giving you an alternate method of accessing almost all of the operations available through the manager. You can use this command to monitor and control jobs and scripts submitted to the BrightStor ARCserve Backup job queue.

For a complete list of the options and switches available for this command, see the appendix "Using Command Line Utilities."

Chapter 7: Managing Devices and Media

BrightStor 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.

This chapter discusses these device and media management tools in greater detail.

Device Manager

When you want information about storage devices, including the tape drives connected to your system, the media in these drives, or the status of a particular storage device, use the Device Manager. When you highlight a media, a storage device, or the adapter card it is configured to, BrightStor ARCserve Backup displays summary information about the media, the adapter card, or the storage device.

In addition to the device and media information available from the Device Manager, BrightStor ARCserve Backup provides features so that you can perform the following maintenance tasks:

- Configure device groups and file system device groups
- Format and copy media
- Erase or compress data
- Retension tapes
- Eject media from the storage device
- Enable or disable devices
- Take a tape library online or offline
- Mount or dismount magazines

- Load or unload media
- Import or export media
- Clean tape heads

Important! Before you use these options, especially the destructive ones (such as formatting and erasing), make absolutely certain you have the right media selected.

Manager Console

Using the Manager Console (a component of the Device Manager) you can view all activity, such as stop, start, or any service communications failures related to BrightStor ARCserve Backup services such as the Database, Queue, or Media service. To use the Manager Console, select Show Manager Console from the Device Manager File menu. If necessary, click the Clear button on the Manager Console to clear the messages.

Device Manager Views

The Device Manager provides three views to help you select the storage devices and media you want to view. The three views display essentially the same information, organized in three different ways. Select a view from the drop-down menu on the top left-hand side of the screen. The devices displayed are associated with the host server.

- Adapter View—Displays information about your storage devices, including the drives connected to your system and the media in these drives, which are organized by the configured adapter cards. The left pane of the Device Manager shows the adapter, the devices configured to it, and the media in those devices.

Note: The default number of devices per adapter and the maximum number of file system devices that appear in the Adapter View is 16. For information on how to modify this and other defaults, see Media Configuration File, `camediad.cfg` in this chapter.

- Device View—Displays information about storage devices, including the drives connected to your system and the media in the drives.
- Group View—Displays information about your devices organized by the device groups to which they are assigned. The left pane of the Device Manager shows the group, the devices assigned to the group, and the media in the devices.

Device Manager Summary Tab

The Summary tab provides basic information about the devices and media you select in the left pane of the Device Manager. The information displayed in the right pane of the Device Manager depends upon the object selected in the expandable tree in the left pane of the Device Manager:

- In Adapter View, select an Adapter in the left pane of the Device Manager, and the Summary Tab displays the name and board number of the adapter. The Device Manager only displays Adapter information in Adapter View.
- When you select a device, the Summary tab displays the Device Information, identifying the adapter number, its SCSI ID, Logical Unit Number (LUN), and the device's vendor, product, and firmware.
- When you select a library, the Summary tab displays information relating to the library, while the left pane of the Device Manager provides information about the slots in the library. Each slot in the library is identified by a number and the media in each of the slots is identified, as are any empty slots, in the left pane of the Manager.
- When you select a media in the specified device, the Summary tab displays the name of the media, the sequence number and ID, the serial number, and if the media is write protected. You can obtain the vendor-specific characteristics of the device using the Summary tab. The Media Characteristics fields display the Media Type, the Density Code and the Format Code, and the Block Size.
- When you select a group in the left pane of the Device Manager, the Summary tab displays information about the group, including the name. If the group is part of a library, the Device Manager displays the name of the library and the status of the group.

Device Manager Detail Tab

The Detail tab displays when you select a device from the Adapter View, Device View, and Group View. This tab provides more information about the adapter, device, or media you selected in the left pane of the Device Manager:

- If you select a device, the Detail tab shows the device properties as shown in the following example:

The screenshot shows the 'Detail' tab selected in the Device Manager. The interface is divided into two main sections: 'Device Information' and 'Device Status'. The 'Device Information' section includes a device icon, a text field for the device name ('/hpFS1 CAFSDDevice Revision 11.5'), and several labeled input fields: 'Group Name' (GROUP0), 'Cartridge Type' (N/A), 'Format Code' (REMOVABLE), 'Volume Capacity' (0), 'Volume Free Space' (830), 'Compression' (Off), 'Block Size' (512), and 'Volume Used Space' (3264(79%)). The 'Device Status' section includes fields for 'Status' (Tape Loaded), 'Operation' (Idle), and 'SCSI Command' (N/A).

Device Information	
/hpFS1 CAFSDDevice Revision 11.5	
Group Name:	GROUP0
Cartridge Type:	N/A
Format Code:	REMOVABLE
Volume Capacity:	0
Volume Free Space:	830
Compression:	Off
Block Size:	512
Volume Used Space:	3264(79%)

Device Status	
Status:	Tape Loaded
Operation:	Idle
SCSI Command:	N/A


- If you select a library, the Detail tab shows the library properties as shown in the following example:

The screenshot shows the 'Detail' tab selected in the Device Manager, displaying library properties. The interface is divided into two main sections: 'Library Information' and 'Library Status'. The 'Library Information' section includes a library icon, a text field for the library name ('OVERLAND NEO SERIES Revision 0421'), and several labeled input fields: 'Drives' (4), 'Slots' (58), 'Magazines' (N/A), 'Import/Export' (Yes), 'Cleaning' (Yes), and 'Bar Code Reader' (Yes). The 'Library Status' section includes fields for 'Status' (Library initialized), 'Operation' (N/A), and 'SCSI Command' (N/A).

Library Information	
OVERLAND NEO SERIES Revision 0421	
Drives:	4
Slots:	58
Magazines:	N/A
Import/Export:	Yes
Cleaning:	Yes
Bar Code Reader:	Yes

Library Status	
Status:	Library initialized
Operation:	N/A
SCSI Command:	N/A

- If you select a media, the Detail tab shows the media properties as shown in the following example:



The screenshot shows the 'Detail' tab of a media management interface. It contains three sections: 'Format Information', 'Usage', and 'Error Count'. Each section has a title bar with a small icon and a list of properties with their values in text boxes.

Format Information	
Expiration Date:	03/19/2007
First Format Date:	03/19/2004
Last Format Date:	03/19/2004
Times Formatted:	1

Usage	
MB Written:	0.0
Write Passes:	0
Usage Time:	0h0m0s
Read Passes:	3


Error Count	
Soft Read:	0
Soft Write:	0
Media Error:	0
Unrecoverable Error:	N/A

Note: Use the Error Count to see the soft read and soft write counts, the media error count, and the number of unrecoverable errors, if any. This information helps you to gauge the age and reliability of the media.

Device Manager Report Tab

The Report tab appears only if you select a library in the Adapter View or the Device View. This tab provides more detailed device information, including the Slot Number, Tape Name, Group Name, Serial Number, Media Pool Name, and the state of the Media Pool Status (Blank, Save Set, or Scratch Set).

For example, in the window below, the first row shows a media in Slot 3, with a Tape Name of CR MON 03-22-04, which belongs to groups CH_GRP0, CH_GRP1, CH_GRP2, and CH_GRP3. This media has a serial number of 41L1. From the Report Tab, you can also select media to perform media operations, such as formatting and erasing.



Media Configuration File, camediad.cfg

If you cannot see some or all of your devices in a certain view, you can edit the MAX_VALUES section in the camediad configuration file. To do this, go to \$BAB_HOME/config/camediad.cfg and check the default MAX_VALUES settings with the rest of the file.

Note: See MAX_VALUES in this chapter for a list and explanation of the default values.

Modify the Media Configuration File

If you find that you have more devices, groups, adapters, changers, number of devices on a single adapter, number of devices configured under one changer, or more cleaning slots configured for a changer than what is listed in the MAX_VALUES section, perform the following procedure:

1. Run `bab -unload camediad`.
2. After all of the camediad processes have stopped (you can check by using `cstatus` command), open the `$BAB_HOME/config/camediad.cfg` file using a text editor.
3. Remove the `;` from the `[MAX_VALUES]` line.
4. Remove the `;` from the corresponding `MAX_ASMS` lines that you need to increase and set the value accordingly.
5. Remove the `;` from the `[MAX_VALUES_END]` line.
6. Save `$BAB_HOME/config/camediad.cfg` and exit the file.
7. Run `bab -load camediad`.

MAX_VALUES

This section describes the default values contained within the MAX_VALUE section of the camediad.cfg file.

- `[MAX_VALUES]` is the starting label for the MAX_VALUES section. This label is required and must have the `;` removed if one or more of the following values are used.
- `MAX_ASMS_DEVICES=32`—If you have more than 32 devices, you must remove the `;` and increase this value.
Note: The maximum number of devices allowed on the system is 128.
- `MAX_ASMS_GROUPS=32`—If you have more than 32 groups, you must remove the `;` and increase this value. Any groups listed in the `[GROUP]` section of the camediad.cfg file will be counted whether you have devices attached at those points or not.
- `MAX_ASMS_ADAPTERS=16`—If you have more than 16 adapters in your system, you must remove the `;` and increase this value.
- `MAX_ASMS_CHANGERS=16`—If you have more than 16 changers attached to the system, you must remove the `;` and increase this value.
- `MAX_ASMS_ADAPTER_DEVICES=16`—If you have more than 16 devices on an adapter, you must remove the `;` and increase this value.
- `MAX_ASMS_CHANGER_DEVICES=16`—If you have more than 16 devices in a changer, you must remove the `;` and increase this value.

- `MAX_ASMS_CLEAN_SLOTS=8`—If you have configured more than 8 cleaning slots in a changer, you must remove the ; and increase this value.
- `[MAX_VALUES_END]`—is the ending label of the `MAX_VALUES` section. You must include this label and have the ; removed if one or more of the values above are used.

Circular Logging

Circular Logging lets you control the size and behavior of the Media Server debug file. Using this feature, you can set a size limit that directs BrightStor ARCserve Backup to chunk the debug file into smaller log files when a user-specified size limit is exceeded. Additionally, you can specify a retention period for debug files. After the retention period elapses, BrightStor ARCserve Backup deletes the chunked debug files.

The Media Server debug file is labeled `camediad.dbg`. It can be found in the `$BAB_HOME\logs` directory.

Note: To configure and use Circular Logging, you must modify the `camediad.cfg` file. For more information, see the section `Specify Circular Logging Settings`. For information about setting the `DEBUG` level, see the topic `ca_devmgr Miscellaneous Commands` in the appendix "Using Command Line Utilities."

Log File Names

If you do not specify Circular Logging settings, BrightStor ARCserve Backup uses the default file name, `camediad.dbg`. If you do specify settings, `camediad.dbg` is still generated, but it is chunked into smaller files and the smaller files are named using the following format:

```
camediad.dbg.####.YYYYMMDD
```

where `####` represents the sequential log number for the day, and `YYYYMMDD` represents the date BrightStor ARCserve Backup captured the information.

Example

For example, on February 1, 2005, the Media Server generates three debug files based upon a file size limit of 100 MB. The log file names are as follows:

```
camediad.dbg
camediad.dbg.0001.20050201
camediad.dbg.0002.20040201
```

How BrightStor ARCserve Backup Labels Debug Files

BrightStor ARCserve Backup labels the log files using the following guidelines:

1. If `camediad.dbg` reaches 100 MB (100 MB is the default value), BrightStor ARCserve Backup renames `camediad.dbg` to `camediad.dbg.0001.20041105`, and creates a new `camediad.dbg` file.
2. If `camediad.dbg` reaches 100 MB for the second time, BrightStor ARCserve Backup renames `camediad.dbg.0001.20041105` to `camediad.dbg.0002.20041105`, renames `camediad.dbg` to `camediad.dbg.0001.20041105`, and creates a new `camediad.dbg` file.
3. If `camediad.dbg` reaches 100 MB for the third time, BrightStor ARCserve Backup renames `camediad.dbg.0002.20041105` to `camediad.dbg.0003.20041105`, renames `camediad.dbg.0001.20041105` to `camediad.dbg.0002.20041105`, renames `camediad.dbg` to `camediad.dbg.0001.20041105`, and creates a new `camediad.dbg` file.

This process continues in a cyclical manner. BrightStor ARCserve Backup always retains the latest three log files.

`camediad.cfg` Circular Logging Parameters

Circular Logging lets you customize the behavior of debug files generated by the Media Server.

You can specify how BrightStor ARCserve Backup generates and maintains debug files by modifying the value of the following parameters in the `camediad.cfg` file:

MAX_LOG_KEEP_TIME

The `MAX_LOG_KEEP_TIME` represents the total number of days that BrightStor ARCserve Backup retains the individual debug files before they are deleted. To enable this option, you must specify a value in the range of 1 to 365.

If this option is not set, the time based debug file purging criteria will not take effect and the debug files will not be purged because of timeout.

MAX_LOG_SIZE

The `MAX_LOG_SIZE` represents the maximum size in KB of a single log file. To direct BrightStor ARCserve Backup to generate more than one debug file, you must specify a value in KB in the range of 10 to 1000000. If you specify zero KB, BrightStor ARCserve Backup creates a single debug file.

This option must be set to enable the circular logging feature. If it is not set, the circular logging feature will be disabled, however BrightStor ARCserve Backup will create a single debug file.

MAX_LOG_NUM

The MAX_LOG_NUM represents the maximum number of debug files that BrightStor ARCserve Backup will generate. To enable this option, you must specify a value in the range of 3 to 32, based on the usage specified in the camediad.cfg file.

Note: For more information about the configuring the camediad.cfg file, see the online help.

If this option is not set, BrightStor ARCserve Backup will use the maximum value of 32 as the default.

Specify Circular Logging Settings

To configure Circular Logging settings:

1. Run `bab -unload camediad`.
2. From the `$BAB_HOME/config` directory, open `camediad.cfg`.
3. In the section labeled CONFIG, enter a value for the `DEBUG_LEVEL` parameter. With `DEBUG_LEVEL` enabled, the `camediad.dbg` file will be created.

DEBUG_LEVEL—Controls the level of debugging information that appears in the `$BAB_HOME/logs/camediad.dbg` file. Confirm that this line is not commented out; otherwise BrightStor ARCserve Backup will not put any debugging information into the file.

To enable this option (and generate the `camediad.dbg` file), you can use the Device Manager command line utility (`ca_devmgr`) to set the debug level using the `SETDEBUG` command. If you modify the `camediad.cfg` file, you must specify a parameter that is zero or a positive number. If you specify a negative value, `camediad` will not start as there is invalid information in the configuration file.

Note: This file can become quite large in a short period of time.

4. In the section labeled CIRCULAR_LOG, enter values for the **MAX_LOG_NUM**, **MAX_LOG_KEEP_TIME**, and **MAX_LOG_SIZE** parameters.

Note: For more information about these parameters, see the section "camediad.cfg Circular Logging Parameters."

5. Save `$BAB_HOME/config/camediad.cfg` and exit the file.
6. Run `bab -load camediad`.

Device Groups

Use device groups to separate your storage devices. If you have more than one storage device connected to your machine, you can separate them into two or more groups. Establishing device groups is a key feature of BrightStor ARCserve Backup—flexibility and efficiency.

With device grouping, you can take advantage of parallel streaming. With parallel streaming you can have several operations occurring simultaneously, one at each device group configured for your system.

By default, BrightStor ARCserve Backup is installed with each storage device assigned to its own group. If identical storage devices (same make, model, and firmware) are found, they are automatically placed in the same group. Later, you can use the Configure function of the Device Manager to regroup your devices.

By grouping devices, you can also share slots between device groups in the library. For more information about parallel streaming, automated media spanning, and sharing slots among groups, see Device Group Configuration in this chapter.

Device Manager Toolbar

The buttons on the Device Manager toolbar provide you with a number of different options to help you manage and maintain groups, devices, and media.



Device Group Configuration—Creates a new device group, assign or remove a device from a device group, or rename or delete a device group.



Staging Group Configuration—Configures file system device groups. You must use this toolbar option to configure groups that you will use for staging backup jobs.



Options—Displays the slots in libraries based on a range of criteria.



Format—Formats blank media, but you can use this option to manually format your media. Formatting destroys all data on your media.



Erase—Erases all data from media. Unlike reformatting, this option also erases all references to the contents of the media from the BrightStor ARCserve Backup database.



Media Copy—Copies the contents of one media to another blank media. However, it applies only to media of the same device type, model, and firmware.



Retension—Ensures that your media is evenly and properly tensioned. This option is especially important if you are having trouble writing to or reading from a media.



Compression—Enables or disables compression, if your tape drive supports compression.



Eject—Ejects media from a storage device, if your device supports this feature.



Disable Device—Enables or disables a device for security purposes. No one can use a disabled device, so you can use this feature to protect the media used for backup.

If you have a tape library, the following buttons appear on the toolbar:



Mount/Dismount—Loads or removes a magazine from the library.



Load/Unload—Loads or unloads specific media from the slots of a library.



Import/Export—Adds new media to a library by specifying an empty slot to which you can import the media (add) to a library, or export (remove) the media from a library.



Clean—Cleans the heads of any drive in your library.



Offline Library—Takes a tape library offline to secure backup media from being overwritten.

Multistreaming

If you have more than one device group configured on your machine, you can take advantage of multistreaming, so that you can have more than one job running at a time. For example, if you have two device groups, GROUP0 and GROUP1, you could have a backup job running on GROUP0, while another backup or restore job runs on GROUP1.

Automated Media Spanning

Media spanning means that when the media is full, the session transfers to another media. If you have a device group consisting of two or more standalone devices, media spanning occurs across devices, so that, when one media in a group drive becomes full, the job transfers to another media in another device within the same device group. (This process is unnecessary if you have a changer.)

For example, suppose you have two device groups, GROUP0 and GROUP1. The GROUP0 device group has one storage device and GROUP1 has two standalone storage devices. If you submit a backup job that takes up two media, you can insert blank media in each GROUP1 drive and BrightStor ARCserve Backup spans the media for you.

If you used GROUP0 for the same backup job, you would have to eject the first media manually and then insert a second media. Automated spanning means that BrightStor ARCserve Backup does most of the work for you.

Slot Sharing Among Groups

Rather than associating the groups in the tape library with exclusive sets of slots, you can share slots among groups so that all groups can access all the slots in the library.

You can share slots among groups, however, you cannot share tape drives among groups. Each group has only one drive associated with it. If a job is submitted to a specific group, the associated drive must be available; the job is not automatically redirected to another group if the drive is busy.

The benefits of sharing slots among groups are:

- A backup job is not limited to a certain range of slots. Regular backup can use as many tapes as are available in the library.

- You can perform device management operations on slots from one group while another group is busy. You can format, erase, mount, or import to any slot from any group. For example, if you want to format a tape in SLOT1 but GROUP1 is busy, you can use GROUP2.
- If you want to export a tape that is not being used, you can do it from any group that is not busy, without waiting for the group that is tied to its slot to become available. You only have to wait if all the groups are busy.
- Relating media pools and slot groups is easier: as long as media pools do not contain tapes across libraries. Media pools can contain any sets of tapes because all groups can access all the tapes in the library. It does not matter if the tapes reside in a particular group or if a tape belongs to a media pool that is outside the slot range of the group on which the job is submitted.
- You can avoid a problem if you submit a restore job from a * Group, and the group that contains the tape associated with the restore is busy. (The problem occurs when the restore job tried the next group and prompted for the tape in that group. But this required exporting the tape from one group and importing it in another group, which could not be done because the first group was busy.)
- In the case of multistreaming jobs, there is a better chance that the same tape will be used for the next incremental backup, even if the group is busy. This improves the use of the media.

Note: Restrictions apply if the multistreaming job spans multiple libraries.

Format Media Option

Although BrightStor ARCserve Backup automatically formats blank media, you can use the Format 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. For more information, see Media Pool Manager in this chapter.

File System Devices (FSD) that are part of a staging group cannot be formatted using the Format Media Option. To prevent accidental formatting of an FSD prior to the data being migrated to a final destination media, the Format toolbar button on the Device Manager window is disabled. 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.

Important! *Use Format with care! The Format option permanently destroys the data contained within a media and any job sessions associated with it.*

Using the Device Manager to format media, you can specify media pooling information associated with the media. Use the Format dialog to define the following media pool information:

- **New Media Name**—You can change the name of the media inserted in the storage if it already has a name. If the media is blank, you can assign the media a name. If the media is blank but has an assigned serial number, you can assign it to any existing media pool.
- **Expiration Date**— Sets the expiration date for the media. When this date passes, you will be reminded that the media has exceeded its expiration date. For more information about setting expiration dates, see Expiration Dates in this chapter.
- **Media Pool**—Select an existing media pool, previously created using the Media Pool Manager.

Note: To use this option, you must select the Overwrite Serial Number check box.

- **Serial Number**—The serial number is the current serial number, the bar code label, or the next available serial number (based on predefined rules). You can manually enter a different serial number only if the tape supports bar codes, bar code usage is enabled, and the serial number is not already in use.

Expiration Dates

Media life is based on passes. A pass is defined as the storage device 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. Media manufacturers rate useful media life to be 500 to 1500 passes. This does not mean that the media is unusable after it reaches the maximum number, only that it is more susceptible to errors at this point.

Specify an expiration date based on how you will use the media.

- If you plan to use the media often, for example, a few times a week, you should set the expiration date to a year from now, or sooner.
- If you plan to use the media only once or twice a month, you can set the expiration date to two or three years from the current date.

When a media reaches its expiration date, you can still use it, but when you make a backup, for example, a note is made in the Activity Log that this media is expired. The expiration date is a way of tracking how long media has been in service so you can stop using it before it reaches the end of its useful life.

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.

Erase Media Option

Use the Erase option to erase all data from media. If you use Erase, rather than Format, BrightStor ARCserve Backup deletes all references to the contents of this media from the database. Reformatting a media does not destroy its physical history (read and write passes, and so on). To remove all data and references from your media, choose Erase.

File System Devices (FSD) that are part of a staging group cannot be erased using the Erase Media option. To prevent accidental erasing of staged data from an FSD prior to it being migrated to a final destination media, the Erase toolbar button on the Device Manager window is disabled. If you want to erase the data from the FSD, you can either use the command line (`ca_devmgr`) or disable the staging option for the selected FSD.

BrightStor ARCserve Backup has the following erase options from which to choose:

- Quick Erase—Quick Erase requires less time than a Long Erase because it overwrites only the current media label. Although technically there is still data on the media, the data is effectively inaccessible without the media label. Quick Erase is useful if you want to re-use BrightStor ARCserve Backup media, but you do not have the time to wait for a Long Erase.
- Quick Erase Plus—Similar to Quick Erase, Quick Erase Plus erases the current media label. In addition, Quick Erase Plus erases the serial number, so that you can assign a new serial number to the media.
Note: If the tape uses bar codes, you cannot erase the serial number.
- Long Erase—Long Erase completely removes all data from the media. It requires more time than a Quick Erase, however, the media is considered blank, as if it were formatted. For security reasons, if you want to make sure that the data on the media is permanently erased, use Long Erase.

Important! *Use Erase with care! The Erase option permanently deletes data on media.*

Methods for Copying Media

BrightStor ARCserve Backup provides three methods to copy the contents of your media to another media. You can use the Tapecopy Tool on the home page, the tapecopy command at the command prompt, and the Media Copy option in the Device Manager.

Note: You cannot copy data to VM:Tape media using the Tapecopy Tool.

To make a media to media copy using different types media, run the tapecopy command or open the Tape Copy Manager by clicking the Tapecopy Tool icon on the BrightStor ARCserve Backup home page. For more information about the Tapecopy Tool, see Tapecopy Tool in this chapter.

Copy Media

Open the Device Manager and click the Media Copy toolbar button to copy the contents of one media to another blank media. It applies only to media of the same device type, model, and firmware.

Copy Media Using the tapecopy Command

To make a media to media copy with two different type media, or to make a session level copy, use the tapecopy command at the command prompt. For a complete list of the options and switches available for this command, see the appendix "Using Command Line Utilities."

Tapecopy Tool

To make a media to media copy with two different media types or to make a session level copy, use the Tapecopy Tool.

Note: You cannot use the Tapecopy Tool to copy to VM:Tape media.

Tape Copy Manager Source Tab

Use this tab to select the source tapes you want to copy. Browse the tree structure in the left pane to locate the tape. The right pane provides information for group and tape objects such as write protection, block size and expiration date. You can copy all sessions or enter a range of sessions. If you want to use tape session information from a database that has only merged information, merged the source tape information into the database prior to executing the tape copy. This increases efficiency for the merged source tape.

Tape Copy Manager Destination Tab

Use the Destination tab to select the tapes you want to copy to. Locate the tape by navigating through the tree in the left pane. The right pane contains information for group and tape objects such as tape ID, density code, and sequence. Select a group or tape to display its specific object information. If you want to use any available group, tape, or both, do not select a group or tape and leave the default wildcard character '*' in the Group and Media Name fields.

Note: You cannot copy data to VM:Tape media using the Tape Copy Manager.

Tape Copy Manager Media Rules Tab

This tab shows the media rules for the destination tape and the media options.

Tape Copy Manager Copy Settings

The Copy Settings are as follows:

- Append to Media—Writes to the end of the media.
- Overwrite Same Media Name, or Blank Media—Overwrite the media in the drive only if it is the one specified for the job, set in the Media Name field, or if the media is blank. If media with a different name is in the drive when the job takes off, the job fails.
- Overwrite Same Media Name, or Blank Media First, then Any Media—(Overwrite any media in the drive.) Select this option to check whether the media in the drive is the one specified for the job. If not, BrightStor ARCserve Backup checks whether the media is blank. If it is not blank, the media found is reformatted in the device with the name in the Media Name field, and starts backing up files at the beginning of the media.
- GFS rotation media is not overwritten. If the media found in the drive is a GFS rotation media, the job fails.
- Overwrite Same Media Name, or Any Media First, then Blank Media—Use this option to overwrite any media in the drive with the same name as in the Media Name field, or any media. If media in the drive is a rotation or GFS rotation media, the job fails.
- If you choose overwrite a media name, you can use a new media name for formatting the destination tape.

Tape Copy Manager Media Settings

The Media Settings are as follows:

- Erase Source Media—Use this to erase all data from your source media. BrightStor ARCserve Backup also erases all references to the contents of this media from the database. When you re-format the media, its physical history (read and write passes) is carried over.
- Merge Destination Media—Merges information from media containing one or more backup sessions into your BrightStor ARCserve Backup database. The database information from the media is appended to your existing database files.
- Export Destination Media—Allows you to export your media after your copy job finishes. You can move your media out of the library or to an off-site location.
- Off-line Destination Media—Prevents overwriting by taking media off-line.

Retension Media Option

Use the Retension option to make sure media is evenly wound and properly tensioned. Retension a media, especially, if you are having trouble writing to it or reading from it. When a media becomes unevenly wound, it is prone to errors, may jam, or worse yet, break.

Note: The Retension option applies primarily to Quarter Inch Cartridge tapes.

Compression Option

You can use the Compression option only if your storage device supports tape compression. If it does not, the Compression toolbar button will be disabled.

Under most circumstances, you should leave compression turned on. You should only turn it off 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.

Important! *You can only change compression when a blank tape is in the drive. This prevents mixing of uncompressed and compressed data between sessions on a tape.*

Eject Media Option

Use the Eject option to eject media from a storage device. This option protects media from being accidentally overwritten. You can use the Eject option if the storage device supports this option.

If you try to eject media from a drive that does not support the eject feature, BrightStor ARCserve Backup registers that the media has been ejected even though it is still physically in the drive. You will need to eject the media from the drive manually and put it back in the drive for BrightStor ARCserve Backup to see the media again.

Note: When accessed from the Device Manager, the Eject toolbar button is disabled for device drives that are part of a changer (library) and will only be enabled for individual device drives that are manually loaded and unloaded.

Enable/Disable Option

Use the Enable/Disable option to enable or disable a device for security purposes. If you disable a device, no one can use that device. This option effectively protects the media used for backup. To use the device again, you must enable it.

Mount/Dismount Option

Note: For multiple-drive changers, the Mount/Dismount option is available only with the BrightStor ARCserve Backup Tape Library Option.

Use the Mount/Dismount to load or 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 taken by this process varies according to the number of media in the magazine you mount or dismount.

Note: You must mount magazines for library operations to commence. Magazines should be dismounted before physically removing them.

For more information about mounting and dismounting magazines, see the *Tape Library Option Guide*.

Load/Unload Option

Note: For multiple-drive changers, the Load/Unload option is available only with the BrightStor ARCserve Backup Tape Library Option.

Use the Load/Unload option load a specific media into a tape drive, and to unload media from a tape drive.

For more information about loading and unloading media, see the *Tape Library Option Guide*.

Import/Export Option

Note: For multiple-drive changers, the Import/Export option is available only with the BrightStor ARCserve Backup Tape Library Option.

Use the Import/Export option to add new media to a library by specifying an empty slot to which the media can be imported, or by directing BrightStor ARCserve Backup to locate available slots. You can import one tape at a time or many tapes simultaneously. When you import media, the library reads the media and adds it to its inventory. Use the Export function to remove media for off-site storage, or if you suspect it is defective.

The Import/Export dialog contains the following fields and buttons:

- Import—Click this button to import the media.
- Export—Click this button to export the media.
- Group Name—The name of the group that you want to import tapes to or export tapes from.
- Import any slots—Choose this option to direct BrightStor ARCserve Backup to scan the group for available slots and import each tape to the next available slot. This function eliminates the need for you to scroll the Import/Export dialog to find available slots and import many tapes simultaneously. Use the spin box to specify the number of tapes that you want to import.

- Select all slots—Choose this option to direct BrightStor ARCserve Backup to import all of the slots in your library to the specified group.
- Slot—Select an empty slot to which to import a media, or choose the slot containing the media you want to export. If you know the slot that you want to import the tape to, check the check box corresponding the slot or media name and then click Import.
- OK—Click this button when you are done.

Note: If your device does not support importing and exporting, the Import/Export toolbar button is inaccessible.

For more information about importing and exporting media, see the *Tape Library Option Guide*.

Clean Tape Heads Option

Note: For multiple-drive changers, the Clean Tape Heads option is available only with the BrightStor ARCserve Backup Tape Library Option.

Use the Clean Tape Heads option to can clean the heads of any drive in your library. You must have a cleaning tape installed in the tape cleaning slot you specified during setup to use this option.

For more information about cleaning tape heads, see the *Tape Library Option Guide*.

Online/Offline Option

Note: For multiple-drive changers, the Online/Offline option is available only with the BrightStor ARCserve Backup Tape Library Option.

By taking a tape library offline, you secure your backup media from being overwritten by others. After you use the offline option, all the tape library slots are unavailable. To use the library again, you must bring it online.

For more information about bringing libraries online and taking libraries offline, see the *Tape Library Option Guide*.

Enable Automatic Performance Tuning

To optimize performance when writing to a tape drive, BrightStor ARCserve Backup can dynamically adjust the amount of data in a single read or write request to tape drive based on the drive's capability. However, due to the wide variety of tape drives used in today's markets, this feature is disabled (by default) when you install BrightStor ARCserve Backup.

If you are using a tape drive with a cartridge type of LTO, LTO2, SDLT, or Sony SAIT, you can enable this feature by modifying the `camediad` configuration file (`camediad.cfg`).

To enable automatic performance tuning

1. Stop the Media Server.

To stop the Media Server, access the command line utility and execute the following command:

```
bab -unload camediad
```

2. Access `camediad.cfg` located in the `$BAB_HOME/config` directory.
3. In the `CONFIG` section, add the following parameter:

```
ENABLE_DYNAMIC_RWSHOTS = 1
```

4. Save the file and restart Media Server.

To restart the Media Server, access the command line utility and execute the following command:

```
bab -load camediad
```

Note: To disable automatic performance tuning, delete this parameter entirely, or, change the `ENABLE_DYNAMIC_RWSHOTS` value from 1 to 0.

Large Library Support

The BrightStor ARCserve Backup definition of Large Library is any library with an unlimited number of slots. There are no limitations on the number of drives or slots BrightStor ARCserve Backup can use. BrightStor ARCserve Backup can also correlate media names to specific physical slot numbers in media views, so that all users can find media based on a variety of criteria.

If your library contains fewer than fifty (50) slots, each slot is shown. However, if your library contains fifty slots or more, BrightStor ARCserve Backup groups the slots in groups of 50. To view slots within these groups, expand the group.

The ca_devmgr Command

BrightStor ARCserve Backup provides a command line interface that you can use to perform a variety of backup functions without using the Device Manager. The command promptly gives you an alternate method of accessing almost all of the operations available through the Device Manager.

You can use the ca_devmgr command to:

- Obtain information.
- Manipulate the tape or library device.
- Control storage devices.
- Format and erase media in drives or changers.

All of the features available from the Device Manager are available from the command line.

Although you can perform these operations using the Device Manager, the commands are useful if you do not have access to a browser. To use these commands, BrightStor ARCserve Backup must be running and your user account must be authorized in the BrightStor ARCserve Backup authentication database. If your user account is not authorized, run ca_auth to authorize it.

For a complete list of the options and switches available for this command, see the appendix "Using Command Line Utilities."

Media Pool Manager

A media pool is a collection of media managed as a unit. Each media pool is assigned a name, and is organized by the serial number range of the media the pool contains. The assigned serial numbers are permanent. If you are using a device with a bar code reader, bar code labels are used for the serial number of the media. Media pools are divided into two sets, the Save Set and the Scratch Set.

The Media Pool Manager provides the following information when you select a media pool in the left pane of the Media Pool Manager:

- **Media Pool Name**—The name of the selected media pool. If the media pool has the extension _DLY, _WLY, or _MLY, the pool contains the media for a GFS rotation schedule. For more information, see Simple and GFS Rotation Media Pools in this chapter.
- **Owner Name**—The name of the user who created the media pool.

- **Create Date**—The date the media pool was created.
- **Base Serial Number**—The number BrightStor ARCserve Backup uses to begin automatically assigning serial numbers to the media in the media pool. The first media BrightStor ARCserve Backup formats has the same serial number as the Base number. Thereafter, each media serial number increases by one. The default of the first serial number of the first media pool is 1000000, the second media pool is 1100000, the third media pool is 1200000, and so on.
- **Next Serial Number**—The next serial number to be assigned. This serial number will be given to the next media assigned to the media pool. When you create the media pool, before any media is assigned to it, this number is the same as the Base Serial Number.
- **Maximum Serial Number**—The highest possible serial number in the media pool. The default first media pool is 1099999, the second media pool is 1199999, the third media pool is 1299999, and so on.
- **Serial Number Increments**—The number by which the serial numbers assigned by BrightStor ARCserve Backup increase. For example, if you set the Serial Number Increments to 5, the first three serial numbers BrightStor ARCserve Backup assigned in the first media pool are 1000000, 1000005, and 1000010. The default is 1.
- **Minimum Save Set Copies**—The minimum number of media that must be held in the Save Set before any media can be released to the Scratch Set. The default is 4.
- **Retention Period (Days)**—The minimum time a media must be held in the Save Set before it can be released to the Scratch Set. The default is 7 days.
- **Prune Retention Period (Days)**—The minimum number of days a media must be retained before it can be pruned. Pruning removes detail records but retains job and session records. The default is 0.

Save Set

The media pool Save Set is a set of media containing important data that cannot be overwritten. Media in the Save Set will not be overwritten, reformatted, or erased until the applicable retention criteria are met and the media is moved into the Scratch Set.

The Media Pool Manager displays the media pool name, set name, owner name, creation date, minimum number of media in the Save Set, retention period, and identifies this as the Save Set in the right pane of the Manager when you select the Save Set in the left pane of the Manager.

BrightStor ARCserve Backup performs media pool maintenance at the beginning of a job, and does not allow media in the Save Set to be moved to the Scratch Set until the oldest tape in the Save Set exceeds the Retention Period, and the Save Set contains the minimum number of Save Set media required.

Minimum Save Set Copies

You can specify the minimum number of media to be retained in the media pool Save Set before the oldest media is recycled to the Scratch Set. This is a safeguard for preventing data loss in case backups are not done for extended periods of time. It prevents media from being mixed up and ensures that the media with the oldest data is overwritten first.

You can move media that can be re-used and overwritten from the Save Set to the Scratch Set, and you can move media from one media pool Save Set to another media pool Save Set. Each time a media in the Scratch Set is written to, the media moves from the Scratch Set to the Save Set. That media moves back to the Scratch Set once the specified retention criteria have been met.

Retention Period

The retention period is the number of days, weeks, or months during which a media is not used. This retention period must expire before the media is moved into the Scratch Set. For example, if you specify a retention period of 14 days, a media remains in the Save Set for 14 days after the last time it was used. If the media has not been used for 14 days, it is moved to the Scratch Set.

Scratch Set

The media pool Scratch Set is a set of media that has been recycled from the Save Set after its retention period has passed. All media in the Scratch Set will be overwritten the next time they are used. The oldest media, those that have not been used for the longest period of time, are used first. When you select a media pool's Scratch Set in the left pane of the Media Pool Manager the right pane displays the media pool name, the set name, the owner name, and the date the Scratch Set was created.

The Scratch Set can be an individual scratch set for each media pool or a *global Scratch Set*. The global Scratch Set treats all the scratch tapes in all media pools as one big Scratch Set. This ensures that the backup job never fails because a scratch tape is not available in its own media pool. As long as there is a scratch tape existing in the global Scratch Set, any backup job can move tapes from any media pool to the Scratch Set before starting its backup operation.

To configure how Scratch Set tapes are managed by the media pool, access the following configuration setting in \$BAB_HOME/config/cadbd.cfg:

POOL_SCRATCH_SET_MODE = GLOBAL

- GLOBAL, the default entry, specifies that if no tapes are available in the scratch set of the current pool, BrightStor ARCserve Backup should search for an available tape in any other media pool scratch set.
- BY_POOL specifies that only the scratch set in the current pool should be searched. If you modify this setting, you must restart cadbd.

Serial Numbers

The media serial number is one of the factors used in categorizing media pools. You can use one of the following two methods to create a serial number for media:

- Bar Code—A number is read from a bar code label, assigning the serial number. A changer with a bar code reader is required for this method. Bar codes override any previously defined media pool settings.
- Automatic—Automatically assigns a serial number for the media, based on the Base and Range of serial numbers set when the pool was created.

You cannot change the serial number of media. However, you can change the following serial number information on the Add Media Pool dialog:

- Base Serial Number—The number BrightStor ARCserve Backup uses to begin automatically assigning serial numbers to the media in the media pool. The first media BrightStor ARCserve Backup formats has the same serial number as the Base number. Thereafter, each media's serial number increases by one. The default is 1000000.
- Next Serial Number—The next serial number to be assigned. This serial number will be given to the next media to be assigned to the media pool. When you create the media pool, before any media is assigned to it, this number is the same as the Base Serial Number.
- Maximum Serial Number—The highest possible serial number in the media pool. The default is 1099999.

When you select a media in the left pane of the Media Pool Manager, the right pane displays information about that media, including the media name, ID, sequence number, type, media status, the location status, and the serial number.

Note: When you select a media in a Save Set, the Sequence Number field displays one entry for each tape sequence instead of one entry for the entire tape sequence.

Simple and GFS Rotation Media Pools

All rotation backup jobs create their own media pools, based on the name entered in the Media Pool Name field (Rotation schedule) or the Media Pool Name Prefix field (GFS Rotation schedule) in the Method/Schedule tab in the Backup Manager. Each media pool has a default retention period and number of media to save. You can choose to assign a Custom schedule to a media pool, as well. For further information about Simple and GFS rotation schedule media pools, see the chapter “Customizing Your Jobs.”

You can assign GFS (Grandfather-Father-Son) rotation schedules to a particular group of media, which are recycled and reused throughout the cycle of the rotation schedule. The GFS backup uses three media pools—Daily, Weekly, and Monthly. BrightStor ARCserve Backup names these pools according to the Media Pool Name Prefix entered using the Media Rules tab on the Method/Schedule tab. Each media pool has its own Scratch Set and Save Set with a default retention period and minimum number of media to save.

For example, for the media pool BK, the following media pools are created in a GFS rotation schedule:

Pool name	Description
BK_DLY	Daily incremental, differential, or full backup media for the BK media pool.
BK_WLY	Weekly full backup media for the BK media pool.
BK_MLY	Monthly full backup media for the BK media pool.

Important! *You must provide a Media Pool Name Prefix for your GFS schedules.*

GFS Save Sets and Scratch Sets

Media is assigned to a Save Set as soon as it contains GFS backup data, and stays there until it is eligible to be assigned to a Scratch Set after the expiration of the retention period. Media in the Scratch Set is eligible for recycling, and BrightStor ARCserve Backup re-uses the oldest media in a Scratch Set first. In a GFS rotation schedule, a set number of media in a media pool is used for your backup jobs. If there is no media in a Scratch Set, you are prompted for blank media.

GFS Retention Periods

The retention period for media, whether for a five-day or seven-day GFS rotation schedule, depends on the type of GFS rotation schedule you are using, and the media pool to which the media is assigned, whether daily (_DLY), weekly (_WLY), or monthly (_MLY).

By default, you can re-use daily media after six days. Weekly media can be overwritten after five weeks retention. Monthly media is saved for one year before re-use. Monthly media should be taken off-site for storage. You can change any of these media rotation defaults to suit your particular environment using the Media Rules tab of the Backup Manager Method/Schedule tab.

Minimum Save Set and Retention Periods

The following are the formulas used to calculate the number of media in the Save Sets and the retention times for the GFS media pools:

- Daily pool—Media for daily backup jobs. The default retention period is six days for a five-day rotation schedule. 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—Media for weekly backup jobs. The retention period equals the number of weekly media times seven, minus one (# of weeklies * 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—Media for monthly backup jobs. The retention period equals the number of monthly media times 29 minus five (# of monthlies * 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).

Note: If you preserve one media each month (12 monthly media), the retention period is 343 days, using the $(12 * 29) - 5$ formula, which takes into account months with differing number of days. You can change retention times in the Media Pool Manager.

Media Pool Manager Toolbar

The Media Pool Manager toolbar provides the following features:



New Pool—Adds new media pools to use with custom backup jobs.



Modify—Modifies the number of minimum Save Set copies and the retention period of an existing media pool.



Delete—Deletes an existing media pool from the BrightStor ARCserve Backup database.



Move—Moves media from a Scratch Set to a Save Set, or vice versa, in an existing media pool.



Assign Media—Assigns media to an existing media pool.

Create a Media Pool

To create a media pool, using the Media Pool Manager, supply or modify the following information:

- **Media Pool Name**—The name of the pool with which to associate the media. Make the media pool name from one to 15 uppercase characters.
- **Base Serial Number**—The number BrightStor ARCserve Backup uses to begin automatically assigning serial numbers to the media in the media pool. The first media BrightStor ARCserve Backup formats has the same serial number as the Base number. Thereafter, each media serial number increases by one. The default of the first serial number of the first media pool is 1000000, the second media pool is 1100000, the third media pool is 1200000, and so on.
- **Next Serial Number**—The serial number to be given to the next media to be assigned to this media pool. When you create a media pool, this number is the same as the Base Serial Number.
- **Maximum Serial Number**—The highest possible serial number in the media pool. The default first media pool is 1099999, the second media pool is 1199999, the third media pool is 1299999, and so on.
- **Serial Number Increments**—The amount by which the serial numbers in this media pool will increase. The default is 1.
- **Minimum Save Set Copies**—The minimum number of media that must be in the Save Set before the oldest media can be recycled to the Scratch Set. The default is 4.

- Retention Period (Days)—The minimum number of days media is kept in the media pool Save set before it can be released to the Scratch set and overwritten. The default is 7.
- Prune Retention Period (Days)—The minimum number of days a media must be retained before it can be pruned. Pruning removes detail records but retains job and session records. The default is 0.

BrightStor ARCserve Backup automatically creates a media pool when you submit a backup job either with rotation or GFS rotation. For more information about media pools and GFS rotation schedules, see Simple and GFS Rotation Media Pools in this chapter and the chapter “Customizing Your Jobs.”

Modify Media Pools

You can change certain information for an existing media pool. Although you cannot change the Media Pool and Owner Names, the Create Date, and Serial Number information, you can modify the number of minimum Save Set copies, the retention period, and the prune retention period of an existing media pool.

If you want to increase or decrease the minimum number of media to be retained in the media pool’s Save Set, or lengthen or shorten the retention period, click the Modify button to open the Modify Media Pool dialog. The dialog displays information about the media pool. Modify the appropriate available fields and click OK.

Move Media

Use the Move option to relocate media from one media pool to another, or from a media pool’s Scratch Set to the Save Set, or vice versa. You can use this feature to remove media from a media pool before you delete the media pool.

To move media, use the following steps:

1. Specify the media you want to move.
2. Click the Move toolbar button.

The Move Media dialog opens and presents a list of available media pools on the current host server.

3. Select the media pool to which you want to move the media, and specify whether the media is to be moved into the media pool’s Save Set or Scratch Set.
4. Click OK.

Assign Media to a Media Pool

Use the Assign Media option to assign media to an existing media pool.

Note: If the media you want to assign to a particular media pool is already assigned to another media pool you must use the Move option, rather than the Assign Media option, to change the media pool to which a media is assigned.

To assign media to an existing media pool, use the following steps:

1. Click the Assign Media toolbar button.

BrightStor ARCserve Backup presents you with the names of available media that have not yet been assigned to a media pool, and a list of existing media pools.

2. Specify the desired media and media pool.
3. Click OK.

Delete a Media Pool

Use the Delete option to remove existing media pools from the BrightStor ARCserve Backup database. This feature is useful when, for example, a GFS rotation backup job is permanently removed from the job queue. Before you use the Delete option, you should reassign the media in the media pool to another media pool. If you do not, when you delete the media pool you will unassign the media as well.

To delete a media pool from the database, use the following steps:

1. Specify the media pool that you want to delete.
2. Click the Delete toolbar button to remove the tape from the media pool whether it is in the save set or the scratch set.

BrightStor ARCserve Backup asks you to confirm that you want to delete this media pool and unassign all the media it contains.

3. Click OK.

Maximum Number of Sessions on a Single Tape

BrightStor ARCserve Backup allows up to 20000 sessions on a single tape and up to 25 sequences of a series of spanned tapes. For a file system device (FSD), the limitation is 65535 sessions on a single FSD. Keep this in mind when planning your backups, because, if your sessions are small, you can reach 20000 sessions very quickly. If you have a large amount of data to back up, you can quickly exceed 25 sequences, depending upon how much data each tape can hold.

The ca_dbmgr Command

The ca_dbmgr command, in addition to being the command line interface to the Database Manager, provides the functions of the Media Pool Manager from the command line. Use the ca_dbmgr command to configure media pools, add, modify, delete, move, and assign media to media pools. All of the options available from the Media Pool Manager are available from the command line.

The media pool options available through the ca_dbmgr command are:

- Show—Display information about the specified pools or pool media.
- Add—Create a new media pool. Switches to specify information about the serial number and the retention period.
- Modify—Modify previously entered information about the media pool.
- Delete—Delete a specified media from a media pool.
- Move—Move and assign media from media pool to media pool. Use this command to move tapes from the Scratch Set to the Save Set, or back.

For a complete list of the options and switches available for this command, see the appendix “Using Command Line Utilities.”

Media Management Administrator

Using the Media Management Administrator (MMO), you can perform the following tasks:

- Control, manage, and protect your media resources.
- Organize tape movement to off-site storage locations.
- Define retention rules to ensure that tapes are protected from premature overwrite.
- Secure access to tape-resident files and maintain a comprehensive inventory of the tape library resource.

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. The MMO works with the 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.

Note: After media leaves a vault for use with BrightStor ARCserve Backup, you cannot send the media back to the vault unless BrightStor ARCserve Backup writes a new session to the media.

You can define vaulting criteria using the MMO. 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.

How the MMO Works with Domains

The MMO is based on the client/server model. The following criteria determine where the media management primary server and client (member) server are applied:

- You can have only one MMO primary server in the BrightStor ARCserve Backup domain, and it must function as the BrightStor ARCserve Backup domain server.
- You can have zero or more MMO member servers. Each must function in the capacity of member server in the BrightStor ARCserve Backup domain.

The MMO primary server and MMO member servers communicate through the Ingres Net service, which is installed by default. The member server synchronizes with the primary server after updating the member server's database.

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, BrightStor 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.

MMO Admin Window

The MMO Admin window makes it easier for you to create vaults, VCDs, and reports, and to schedule rotation. Using the tools provided by the MMO Admin window, you can establish vaulting rules needed for complete media management. To access the MMO Admin window, click the Media Management Admin icon on the BrightStor ARCserve Backup home page.

The MMO Admin window consists of a menu bar, a toolbar, and the Media Management pane. You can access all of the options on the toolbar from the File menu. The left pane of the MMO Admin window displays information corresponding to the primary MMO server in a tree-like structure for easy navigation. The right pane displays information related to the object selected in the left pane, and any output messages and reports generated during your MMO Admin session.

MMO Admin Toolbar

The toolbar on the MMO Admin window contains the following buttons:



Refresh—Refreshes and updates the information displayed in the MMO Admin window.



Create Vault—Opens the Create Vault dialog.



Create Schedule—Opens the Create Schedule dialog.



Start Vault Cycle—Starts the vault cycle process, which updates location information for tape volume sets, indicates movement into or out of a vault, automatically vaults tape volumes, and produces reports.



Simulate Vault Cycle—Generates a Vault Selection Report. You can use the Simulate Vault Cycle option to predict how many tape volumes will be moved, without actually updating location information.



Find Media in Vault—Opens the Find Media in Vault dialog, to search for media by Tape Name or Serial Number.



Print—Prints the information displayed in the right pane of the MMO Admin window.



Print Preview—Previews information before printing.

Log in to the MMO Admin

Before you can use the MMO Admin window, you must log in using BrightStor ARCserve Backup domain login information. Without a valid user name and password, you cannot open the MMO Admin window.


To open the MMO Admin window, use the following steps:

1. From the BrightStor ARCserve Backup home page, click the Media Management icon.

The BrightStor ARCserve Backup Domain Login dialog opens.

2. Enter your BrightStor ARCserve Backup domain login user name and password.
3. Click OK.

After you log in to the BrightStor ARCserve Backup domain, the MMO Admin window opens in your browser as shown in the following example:



MMO Admin Features and Functionality

The objects in the left pane of the MMO Admin window are arranged in an expandable tree. From the expandable tree, you can right-click an object to create a branch (For example, right-click Vault to create a new vault). Click + next to an object to view its branches. For each object, there is unique set of right-click menu options that you can use to configure the object.

Initially, the MMO Admin window displays the currently defined MMO primary management server. Double-click the Media Management Server branch to expand it and access the available objects. The following objects are available:

- Current Server—Displays information about the server you are currently using.
- Vault—Provides information about the vaults.
- Schedule—Lists the names of the previously created schedules. You can access the Vault Criteria Descriptor and Rotation objects.
- Reports—Provides access to the five available reports.
- Find Media in Vault—Opens the Find Media dialog so that you can locate specific media.
- Current Status—Shows the status of the most recent operation.

Vault Object

When you select the Vault object in the left pane of the MMO Admin window, the right pane displays the following information about existing vaults:

- **Vault Name**—The name of a defined storage area or location.
- **Max Slots**—The maximum number of slots this vault can contain. This maximum is set when the vault is created.
- **Active Slots**—The number of slots in use in the vault. Each slot in the vault stores one tape volume.
- **Local**—A Yes in this column indicates that the Use in Local option was selected when the vault was created or updated, and this vault will not be moved to other locations. A No indicates that this option was not selected. This option can be enabled at any time, if you decide to move a particular vault off-site.
- **Create Date**—The date the vault was created.
- **Description**—A description or comment about the vault, entered when the vault was created or updated.

Expand the Vault object to access individual vaults, listed under the Vault object in the left pane of the MMO Admin window. When you select a particular vault, the right pane of the MMO Admin window displays information about that vault and its slots, created when tape volumes are moved into the vault. Each media in the vault is listed, identified by slot number and media name, and the status of the slot and the date the slot was created are also displayed. Right-click an existing vault to update it.

To create a new vault, right-click the Vault object. For further information about creating, updating, or deleting a vault, see *Create a Vault* in this chapter.

Schedule Object

Use the Schedule object to create new schedules and view information about previously defined schedules. You must create a schedule before you define the Vault Criteria Descriptor and Rotation that determine the selection and retention rules for your vault.

When you select the Schedule object, the right pane of the MMO Admin window displays the names of previously defined schedules. These schedules are also listed under the Schedule object in the left pane.

Create a New Schedule

To create a new schedule, right-click the Schedule object and choose Create from the pop-up menu. To delete a schedule, right-click a specific schedule and choose Delete from the pop-up menu. For more information about creating or deleting a schedule, see Schedule Tape Volume Movement in this chapter.

After you have named and created a schedule, the Vault Criteria Descriptor (VCD) and Rotation objects appear in the left pane of the MMO Admin window.

Vault Criteria Descriptor Object

Use the Vault Criteria Descriptor to set source information, which governs the tape volumes assigned to a vault. You can select either a Media Pool Name or a File Name, including a full path, as the controlling data set. When this data set is vaulted, the tape volumes on which it resides are assigned to slots in the vault.

When you select the Vault Criteria Descriptor object, the right pane of the MMO Admin window displays columns listing the following information for existing VCDs:

- VCD Name—The name of the Vault Criteria Descriptor.
- VCD Type—Whether the controlling data set is a Media Pool or a File Name.
- Media Pool—If the controlling data set is a Media Pool, the name of the Media Pool appears in this column.
- Host Name—If the controlling data set is a File Name, the host on which 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.

Previously defined VCDs also appear under the Vault Criteria Descriptor object in the left pane of the window. For previously defined VCDs, you can perform the following tasks:

- Update or delete a previously defined VCD. To update or delete a previously defined VCD, right-click the existing VCD and choose Update or Delete from the pop-up menu.
- Create a new VCD. To create a new VCD, right-click the Vault Criteria Descriptor object in the left pane of the MMO Admin window and choose Create from the pop-up menu.

For more information about creating, updating and deleting a VCD, see [Create a Vault Criteria Descriptor](#), [Modify a Vault Criteria Descriptor](#), and [Delete a Vault Criteria Descriptor](#) in this chapter.

Rotation Object

The MMO relies upon user-defined rotation schedules determine when and where tape volumes are to be moved. Use the Rotation object to set or update the retention rules 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 MMO Admin window lists all previously defined rotations, and displays columns listing the following information for existing rotations:

- **Rotation Name**—The name of the rotation.
- **Vault Name**—The name of the vault with which this rotation is associated.
- **Retention Hold Days**—Indicates the specific number of days tape volumes are held in this rotation.
- **Retention Keep for Cycles**—Indicates the specific number of vault cycles tape volumes 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 left pane of the MMO Admin window under the Rotation object. For existing rotations, you can perform the following tasks:

- **Update an existing rotation.** To update an existing rotation, right-click the existing rotation and choose **Update** or **Delete** from the pop-up menu.
- **Create a new rotation.** To create a new rotation, right-click the Rotation object and choose **Create** from the pop-up menu.

For more information about creating, updating, or deleting rotations, see [Tape Volume Retention Rules](#) in this chapter.

Reports Object

Although updating tape volume location information in the database is automated, the physical movement of tape volumes is performed manually. Media Management generates reports indicating the current location and the destination to which the tape volumes must be moved, so that you can route these tape volumes to and from 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. Expand the Reports object in the left pane of the MMO Admin window to view the following report types:

- Vault Selection Report—Contains a list of tape volumes to be selected for moving into or out of the vaults.
- Shipping Report—Contains a list of tape volumes to be pulled from each of the vaults.
- Shipping Content Report—Contains all of the information included on the Shipping Report and additional information such as the session number, source path, start date, size, and number of files.
- Receiving Report—Contains a list of tape volumes to be distributed to the vaults.
- Receiving Content Report—Contains all of the information included on the Receiving Report and additional information such as the session number, source path, start date, size, and number of files.
- Inventory By Vault Report—Lists tape volumes grouped by the vault in which they currently reside.
- Inventory By Media Report—Lists tape volumes grouped by vault, and shows media name in front.

When you select a report type in the left pane of the MMO Admin window the right pane displays the Contents, listing the available reports of that type identified by date. Double-click a report type to access the specific reports for viewing. You can print any of these reports using the Print button on the toolbar. On the day they are generated, you can also select to send reports by email to a customizable list of recipients that you can define in the `cadbd.cfg` file. To define the list of recipients, go to `$BAB_HOME/config/cadbd.cfg` and enter the email addresses next to the `EMAIL_MMOREPORTS` setting.

- The inventory reports are built from information contained in the Slot table, and can be generated at any time. The Shipping Report and the Receiving Report are built from the movement records generated during a vaulting cycle, and are updated after each vault cycle process ends.
- The Vault Selection Report 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.

Maximum Number of Reports

When you start a vault cycle, Media Management generates a group of reports, the Vault Selection Report, Shipping Report, Receiving Report, Inventory Report By Vault, and Inventory Report By Media. If you start more than one vault cycle on the same day, information for each cycle is appended to the same group of reports. By default, the MMO can retain the last 30 groups of reports. You can customize this setting in `$BAB_HOME/config/cadbd.cfg` by increasing or decreasing the value for `MAX_MMO_REPORT_NUM`.

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.

Current Status Object

The MMO can run one vault cycle at a time. To monitor the progress of the vault cycle, or to obtain current online status, you can click the Current Status object in the left pane of the MMO Admin window to view the following information:

- **Current Status**—The status of the current operation displays 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.

Vault Processing Status Changes

You can use the MMO to manually reset the current status of Vault Processing if something went wrong during the vault cycle, such as the MMO database becoming corrupted while the vault cycle process is running. You can use the `ca_mmomgr` command line utility to reset the current status from the command prompt. For more information about `ca_mmomgr`, see "The `ca_mmomgr` Command" in this chapter. After the status is reset, you should be able to restart another vault cycle.

How the Media Management Process Works

The Media Management process involves setting vaulting rules, scheduling tape volume movement, selecting tape volumes, defining retention rules, executing the vault cycle, and moving the media to the proper location.

After you set your vaulting rules, select the tape volume sets to be vaulted according to the Vault Criteria Descriptor information, and set the retention rules, the vaulting rotation process begins. We recommend that you 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.

How the Vault Cycle Process Works

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 the Start Vault Cycle button from the MMO Admin window toolbar, or by selecting Start Vault Cycle from the File menu. You can also initiate the vault cycle using the `ca_mmomgr` command from the command prompt.

Execute the Start Vault Cycle process to generate reports detailing the movement of the tape volumes and updating locality information. The slots that already contain tape volumes and the new slots to 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 created and tape volumes 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. Any tape volume sets that remain unvaulted are returned to Tape Service.

Note: For information about how run Vault Cycle commands from the command line, see the appendix "Using Command Line Utilities."

Vault Cycle Simulation

Use the Simulate Vault Cycle command to produce simulation versions of the Receiving Report, Receiving Content Report, Shipping Report, Shipping Content Report, and the Vault Selection Report. You can print any of these reports using the Print button on the toolbar. On the day they are generated, you can also select to send simulated reports by email to a customizable list of recipients that you can define in the `cadbd.cfg` file.

To define the list of recipients, go to `$BAB_HOME/config/cadbd.cfg` and enter the email addresses next to the `EMAIL_MMOREPORTS` setting. You can enable the option to send reports by email either from the command line or by selecting the **Send the report by email** option from the File menu. For more information about emailing reports from the command line, see Vault Cycle Commands in the appendix "Command Line Utilities."

These reports are identical in format to their non-simulation counterparts. Their contents, however, will report on the actions that the MMO will perform at the next vault cycle if it were to be run on the same day that the simulate vault cycle was run. It is also possible to run a simulate vault cycle on one day to report on a vault cycle that will be run x number of days into the future. This command enables you to predict how tape volumes will be moved without actually updating their location information. To run this command, start the Media Management Admin and click the Simulate Vault Cycle toolbar button, or, from the File menu, choose Simulate Vault Cycle.

You can also use the Simulate Vault Cycle command to predict media movement, based upon the current media status and rotation schedule, for a user-specified number of days into the future. The Simulate Vault Cycle command generates custom reports that you can use to inform your media vaulting services provider about tapes required in-house and on what date.

The Simulate Vault Cycle command supports simulating media movement from one to an unlimited number of days into the future. However, best practice is to run the Simulate Vault Cycle command for one to several days into the future. For example, if you run the Simulate Vault Cycle command five days into the future, and on day two your media status changes, the original information for day five is probably incorrect.

Note: For information about how run Vault Cycle commands from the command line, see the appendix, "Using Command Line Utilities."

Vault Cycle Reports

The vault cycle generates the Shipping Report, the Receiving Report, and the Vault Selection Report. The reports list the old and new locations of the tape volume set, to provide you with the information you need to manage your media.

- The Shipping Report details the media to pull manually, and where to send it on the current date.
- The Shipping Simulate Report, generated by the Simulate Vault Cycle command, details the media you will need to pull manually, where to send it on the user-specified date.
- The Shipping Content Report contains all of the information included on the Shipping Report and additional information such as the session number, source path, start date, size, and number of files.
- The Receiving Report details the media that is coming in to each particular vault on the current date.
- The Receiving Content Report contains all of the information included on the Receiving Report and additional information such as the session number, source path, start date, size, and number of files.
- The Receiving Simulate Report, generated by the Simulate Vault Cycle command, details the media that will be coming in to each particular vault on a user-specified date.
- The Vault Selection Simulation Report details a list of tape volumes to be selected for moving into or out of the vaults on a user-specified date.
- The Vault Selection Report contains a list of tape volumes to be selected for moving into or out of the vaults.

You can print any of these reports using the Print button on the toolbar. On the day they are generated, you can also select to send reports by email to a customizable list of recipients that you can define in the `cadbd.cfg` file. To define the list of recipients, go to `$BAB_HOME/config/cadbd.cfg` and enter the email addresses next to the `EMAIL_MMOREPORTS` setting.


You can enable the option to send reports by email either from the command line or by selecting the **Send the report by email** option from the File menu. For more information about emailing reports from the command line, see Vault Cycle Commands in the appendix "Command Line Utilities."

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. Since all vaulted tape volumes are placed in checked out status, if you need to retrieve tape volumes, you must check the tape volumes back in to the Tape Service before they can be used. For more information about checking in and checking out media, see Special Tape Volume Movement in this chapter. For information about how to generate these reports using the command line, see the appendix, "Using Command Line Utilities."

Create a Vault

The first step in establishing vaulting rules is to create a vault. Use the Vault object to create a new vault and obtain information about existing vaults on the server.

To create a new vault, right-click the Vault object and choose Create from the pop-up menu to open the Create Vault dialog. Enter a name for the new vault, a description, and designate a maximum number of slots for the vault. An example of the Create Vault dialog is shown next:



Each slot in the vault stores one tape volume. Slots are created when tape volumes are moved into the vault (called *vaulting a tape volume*). By default, the maximum is 32,000 slots, but you can assign a number between 1 and 2,000,000,000 when you define the vault.

The Use in Local option helps you to keep track of the physical location of the tape volumes in your vault. The following guidelines apply to the Use in Local option:

- Select the Use in Local option if this vault will not be moved to another location.
- Do not select Use in Local if the tape volumes in this vault are to be maintained off-site.


Click OK when you have defined the vault. The new vault is saved and added to the Vault branch in the MMO management window.

Modify a Vault

To update or modify information for a vault, use the following steps:

1. Double-click the Vault object in the left pane of the MMO Admin window to expand the list of vaults.
2. Select and right-click the vault you want to update from the list, and then choose Update from the pop-up menu.

The Edit Vault dialog opens as shown in the following example.



3. Update or modify the desired settings.
4. Click OK.

Tape Volume Movement Scheduling

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 MM Admin window or you can define new rotation policies and vaulting criteria.

Create a Schedule

To create a tape volume movement schedule, use the following steps:

1. Click the Schedule object.
2. Right-click the Schedule object and choose Create from the pop-up menu.

The Create Schedule dialog opens.

3. Enter a Schedule Name and click OK.

The new schedule is saved and added to the Schedule branch in the MMO management window.

After you have named and created a schedule, the Vault Criteria Descriptor (VCD) and Rotation objects, which allow you to set media selection and retention rules, appear in the left pane of the MMO Admin window.

Delete a Schedule

To delete a tape volume movement schedule, use the following steps:

Note: Before you can delete a schedule, you must first ensure that any VCD and rotation for the schedule have been deleted.

1. Expand the list of schedules below the Schedule object.
2. Click the schedule you want to delete.
3. Delete the VCD and rotation for this schedule.
4. Right-click the schedule that you want to delete and choose Delete from the pop-up menu.
5. Click OK.

Create a Vault Criteria Descriptor

To assign media to vaults you must specify a Vault Criteria Descriptor (VCD) and rotation. Either a Media Pool or file name can be selected to become the controlling data set. When this data set is vaulted, the tape volume set on which it resides is placed in slots in that vault. Slot number assignment is based on the rotation records belonging to the rotation schedule selected, defined using the MMO Admin interface.

After you have created a schedule, you can describe the rules for media selection by creating a Vault Criteria Descriptor (VCD).

To create a VCD, use the following steps:

1. Expand the list of schedules under the Schedule object.
2. Select and right-click the Vault Criteria Descriptor object, and then choose Create from the pop-up menu.

The Create VCD dialog opens.

3. Choose one of the following options:
 - 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.
 - 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 /opt/doc/readme.txt, in the appropriate fields. Browse through the Database or Restore Manager to obtain Path or File information. MMO will find all tapes used for the backup of this directory or file.
4. Click OK to save the VCD and add it to the Vault Criteria Descriptor branch in the MMO Admin window.

Modify a Vault Criteria Descriptor

To update or modify an existing VCD, use the following steps

1. Expand the list of schedules under the Schedule object.
2. Selected a schedule from the list.
3. Expand the schedule to display the Vault Criteria Descriptor and Rotation objects.
4. Right-click on the Vault Criteria Descriptor object, and choose Update from the pop-up menu.
5. Update or modify the Media Pool Name or File Name fields, and click OK to save the new setting for this VCD.

Delete a Vault Criteria Descriptor

To delete a schedule, you must first delete the associated rotation and VCD.

To delete a VCD

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 Rules

After you have created a schedule, you must set the rules governing tape volume retention for your vault. The Rotation object, along with the Vault Criteria Descriptor, appears in the left pane of the MMO Admin window. After you define the schedule, you can access the Create Rotation dialog.

Create a Rotation

To create a rotation, use the following steps:

1. Expand the Schedule object to display the list of schedules.
2. Select a schedule from the list.
3. Double-click the schedule to access the Rotation object.

4. Right-click the Rotation object and choose Create from the pop-up menu.

The Create Rotation dialog opens. Use the Create Rotation dialog to set the following values:

- Rotation Sequence Number—MMO can generate a sequence number for your rotation automatically, or you can select the Sequence Number option, and specify a particular sequence number for your rotation. Vault cycles start with the lowest sequence number. The default for a new rotation is 10. The next new rotation is assigned sequence number 20 by default, unless you specify a different number.

- Vault Name—You must specify the vault name for each rotation. You can select the name of a vault from the drop-down vault list.

- In the Retention fields, you can set any of the following conditions:

Hold Days—The number of days tape volumes will be retained.

Keep for Cycles—The number of vault cycles tape volumes will be retained in this rotation.

Days Elapsed from First Format Date—The length of time a tape volume will be retained is calculated from the date the tape volume was first formatted.

By Date—Tape volumes will be retained in this rotation until the specified date is reached.

By Tape Expiration Date—Tape volumes will be retained in this rotation until their expiration date passes.

Permanent—All tape volumes will be retained in this rotation permanently.

Note: If a tape volume meets one of these conditions, it will remain in the same rotation. When the Retention period for a tape volume expires, you must manually start a vault cycle to unvault the tape and return it to the Tape Service for reuse.

5. Click OK to save and add the new rotation to the Rotation branch.

Modify a Rotation

To update or modify a rotation, use the following steps:

1. Expand the Schedule object to display the list of schedules.
2. Select a schedule from the list.
3. Double-click the schedule to access the Rotation object.

4. Double-click the Rotation object, and select a specific rotation.
5. Right-click the rotation, and choose Update from the pop-up menu.
The Edit Rotation dialog opens.
6. Edit the desired fields.
7. Click OK to save the new setting for this rotation.

Delete a Rotation

To delete a rotation, use the following steps:

1. Expand the Schedule object to display the list of schedules.
2. Select a schedule from the list.
3. Double-click the schedule to access the Rotation object.
4. Select the rotation that you want to delete.
5. Right-click the rotation and choose Delete from the pop-up menu.
6. Click Yes to delete the rotation.

View Slot Information

After tape volumes have been assigned to slots in a vault, Media Management window displays slot information for the vault. Select the Vault object in the left pane of the MMO Admin window and double-click to expand it. When you select a particular media from the list, the right pane of the MMO Admin window displays a view of the vault and its slots. This view provides the following information:

- Media Name, ID, Sequence Number, and Serial Number
- Media Type
- Media Class
- Last Write Date
- Last Read Date
- Create Date
- Slot status

Because slots are automatically created when a tape volume is vaulted, you typically have no reason to update slot information.

Special Tape Volume Movement

You can move tapes between the MMO and the Tape Service to inform the MMO that you are about to use a vaulted tape (for example, in a restore operation), and that the tape will be moved out of the vault, either temporarily or permanently.

- You can use temporary movement if you want to restore a file from one of the vaulted tape volumes, and return it to the vault when done.
- You can use permanent movement if you want to permanently remove a tape volume from the vault.

Use the Temporary Check In or Permanent Check In to check tape volumes back into Tape Service on a temporary or permanent basis. To check tape volumes out of Tape Service permanently, select the Permanent option on the Create Rotation dialog.

Temporary Tape Volume Movement

If you need to temporarily move a tape volume from a vault for use in a restore job, and want to return it to the vault when the job is finished, use Temporary Check In.

All tape volumes that are vaulted are in *checked out* status. Use the Temporary Check In feature to change this status to *checked in* so to track the tape volume while it is temporarily in use during 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*.

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.

Note: The Temporary Check in is only for tracking tapes that are temporarily check in from the vault, and is not a requirement for the actual tape movement; if you do not use this feature, you can still manually move a tape volume from a vault and return it when a job is finished. We strongly recommend that you do use this feature, however, 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 MMO and the actual location of the tape.

Permanent Tape Volume Movement

For permanent retention, the slots, and the tape volumes they contain, are vaulted permanently, unless you change the vault status back to the default. After a tape volume is vaulted, it is never returned to Tape Service.

The Permanent Check In feature returns tape volumes to Tape Service permanently. You need to reformat or erase those tapes before they can be checked out again. Use the Permanent Check In feature to remove tape volumes from your vaults when you no longer need to protect the data on them.

Check In a Tape Volume

To specify a tape volume for Temporary Check In or Permanent Check In, use the MMO Admin window.

To check in a tape volume, use the following steps:

1. Click the Vault object from MMO Admin window and expand the list of vaults.
2. Expand the desired vault and select a volume.
3. Right-click the volume and choose Temporary Check In or Permanent Check In from the pop-up menu.

BrightStor ARCserve Backup asks you to confirm that you are removing the tape volume for Temporary Check In or Permanent Check In.

4. Click Yes to confirm, or click No to cancel the task.

When you select the associated vault object, the Slot Status column displays the status Temporary Check In for the tape volumes you have checked in, and a notation appears on the Shipping, Receiving, and Inventory Reports, identifying the media as having been checked in to Tape Service temporarily.

To return a temporarily checked in tape volume to the vault, click Start Vault Cycle. The tape volume is selected according to the VCD, as it was originally, and returned to its slot in the vault.

The fastest way to find media for Temporary or Permanent Check In, particularly when you do not know what vault a needed tape volume is in, is to use the Find Media in Vault feature.


Find Media in Vaults

Use the Find Media in Vault object to locate media in your vaults. This feature can be very helpful when you want to locate a tape volume quickly, if you know the tape name or serial number of the tape volume. If this information is not available, you can obtain it using the Database Manager.

To find media in a vault, use the following steps:

1. From the MMO Admin window, right-click the Find Media In Vault object and choose Find from the pop-up menu.

The Find Media In Vault dialog opens as shown in the following example:



2. Choose one of the following options:
 - Find by Tape Name—Enter the Tape Name, the Random ID, and the Sequence Number to identify the tape you want BrightStor ARCserve Backup to find.
 - Find by Serial Number—Enter the serial number of the desired media. Note that this is case sensitive. For example, serial number ABC123 is different from serial number abc123.
3. Click the Find button to begin the search. When the search is finished, the specified vault and slot information appears in the right pane of the MMO Admin window.

Back Up and Restore the MMO Primary Server

Since the MMO primary server has the critical responsibility of centralizing all MMO information, BrightStor ARCserve Backup takes an extra measure to ensure that your data is safe in the event your primary server fails. If the MMO primary server fails, you can either restore it or promote one of its MMO member servers to act as a new MMO primary server.

Promoting a member server is the process of assigning the primary server's functional roles and responsibilities and moving data to a different server on the BrightStor ARCserve Backup MMO domain. Promoting an MMO member server to an MMO primary server is beneficial if you:

- Cannot afford to wait the time necessary to restart the failed MMO primary server.
- Find that your MMO primary server is not powerful enough and you want to upgrade servers.

Before you can harness the security BrightStor ARCserve Backup offers with MMO primary server failover protection, you must first complete the following tasks:

- Back up the MMO primary server related data.
- Restore data from the last full backup on the MMO primary server.
- Assign your old primary server as an MMO member server (optional).

The following sections describe the concepts involved and how to perform these tasks.

How the MMO Primary Server Works

In environments with multiple BrightStor ARCserve Backup installations, a BrightStor ARCserve Backup server functions as the MMO primary server while all others function as member servers. The primary server is the server used to:

- Centralize and collect all information required by MMO, including all other BrightStor ARCserve Backup server's tape and media pool information.
- Collect and synchronize media management information with the other MMO member servers in the environment.
- Perform all MMO related operations.

All MMO member servers, and the MMO primary server, must be running the same version of BrightStor ARCserve Backup. If they are not, MMO members will have problems accessing MMO primary database because the database schema is not compatible.

Add a Server to the BrightStor ARCserve Backup MMO Domain

To add a server to the BrightStor ARCserve Backup MMO domain, run the `cammo_setup` script and follow the on-screen directions.

Back Up Primary MMO Member Server Data

The first task that you must perform to prepare for MMO primary server failover is back up MMO related data on the primary server. You should do this on a regular basis. Automated MMO backups will occur as part of the BrightStor ARCserve Backup database backup, but you can also schedule backups separately if you want to have more frequent MMO backups than database backups.

- To schedule more frequent MMO backups, include the backup as part of a cron (script) or Unicenter Workload job to extract the MMO related data from the database in a format that can be easily extracted and applied to your new primary database.

- You can back up the MMO information by running the following command:

```
backupMMOPrimary[-MMODataDir <path> (default: $BAB_HOME/dbase/MMOBackup)][-keepSet <nbSetToKeep> | Default is 3][-log <fullLogFilePath>] [-h (for this display)]
```

The following table describes the function of each variable in this command:

Note: Since the command backs up the MMO related information to disk, make sure you save the data to tape after the backup completes.

Variable	Description
-MMODataDir	Use this variable to specify the directory location where you want your MMO backups to be stored. By default, all MMO backups are stored in separate directories under the \$BAB_HOME/dbase/MMOBackup directory.
-keepSet <nbSetToKeep>	Use this variable to specify the number of MMO backups BrightStor ARCserve Backup should preserve. By default, BrightStor ARCserve Backup automatically stores the last three backups. Note: The keepSet value must be between 1 and 9999.
-log	Use this variable to specify the location where you want BrightStor ARCserve Backup to write MMO messages. The default location is STDOUT.
-h	Use this option to display help.

BrightStor ARCserve Backup includes automated MMO backups when you back up the Ingres database. During Ingres backups, cprocess detects if MMO data must also be backed up by extracting MMO configuration information from the discovery.cfg file. See the BrightStor ARCserve Backup discovery.cfg MAN page for information about MMO settings, and see the cadbd.cfg file for information about configuration options for automated MMO backups (MMO_BACKUP_DIR, KEEP_MMO_BACKUP).

Restore Primary MMO Member Server Data to an Alternate MMO Member Server

To restore your MMO primary data to an alternate MMO member, you must first restore your MMO backup from tape. After you restore from tape, the data will be available to the MMO member server you want to promote.

To load your data into the MMO member database, use the following command:

```
moveMMOPrimary[-MMODataDir <path> | default: $BAB_HOME/dbase/MMOBackup][-  
useSetName <dirname> | default: most recent from -MMODataDir][-  
oldPrimaryObjectOwner <hostname> | default is old primary host][-  
loadOldPrimaryObjectOnly][ -rollback ] [commit atEnd | atTable ] [-h]
```

The following table describes the function of each variable in this command:

Variable	Description
-MMODataDir	Use this variable to specify the directory location from which BrightStor ARCserve Backup should use to restore your data. By default, BrightStor ARCserve Backup looks to restore your data from separate directories under the \$BAB_HOME/dbase/MMOBackup directory (the default backup location when you use the backupMMOPrimary command).

Variable	Description
-useSetName <dirname>	Use this variable to specify the backup you want BrightStor ARCserve Backup to restore from. If you do not specify a location, the most recent backup located under -MMODataDir is automatically selected. You will be asked to confirm before this operation proceeds. If the most recent backup is not valid and you do not use the -useSetName parameter, a warning appears and BrightStor ARCserve Backup automatically selects a previous and valid backup. If you specify -useSetName and identify a non-valid backup set, an error message appears and the operation aborted.
-oldPrimaryObjectOwner <hostname>	Use this variable to specify the name of the machine that you want to assign data to. Use this variable whenever you: Want to assign tapes and media pools that used to belong to the old MMO primary to another machine other than your old MMO primary Know that your original MMO primary server will not act as an MMO member later.
-loadOldPrimaryObjectOnly	Use this variable if you want to load tapes and media pool information into an MMO member (not the old MMO primary) that will now own the tapes and media pools that used to belong to the old MMO primary. This is useful if you know that your old MMO primary will not run BrightStor ARCserve Backup anymore, but you want to keep its tapes and media pools. The moveMMOPrimary command, along with the -loadOldPrimaryObjectOnly command, must be run on the machine you specified as -oldPrimaryObjectOwner when you ran the moveMMOPrimary command on the new MMO primary.
-rollback	If a failed execution of moveMMOPrimary occurs, this option is used to revert to the state of the data prior to the failure.

Variable	Description
-commit atEnd atTable	Commit preserves the changes that have been made to your data. Before a commit occurs, changes are preserved in the Ingres transaction log. The default setting, atEnd, specifies that a commit is automatically performed after you run the moveMMOPrimary command and your data has been successfully loaded in the database. If you change the commit setting to atTable, a commit is performed for each single table that would have been loaded. This setting should be used in environments where the Ingres transaction log is not large enough to update everything before committing the work. If you use atTable and a failure occurs and some tables are not properly updated, a rollback might be necessary to retrieve the state of how the tables were before the initial moveMMOPrimary took place. You can do this by running the moveMMOPrimary -rollback command.
-h	Use this option to display help information.

After running the moveMMOPrimary command, you must perform the following tasks:

- Modify the etc/hosts file on all machines in your system. On the member server, you must specify the IP address of the primary server, and on the primary server you must specify the IP address of the member server. You can access the etc/hosts file by entering "more etc/hosts" at the command line on all machines in your system.
- Run the cammo_setup script to configure a new MMO primary and to configure all MMO members. The script prompts you for the new MMO primary hostname. You must run this script on each BrightStor ARCserve Backup server, including the new primary, and then restart the cadbd process.

Demote an MMO Primary Server to an MMO Member Server

After promoting an MMO member server to primary, you can assign your old primary to function in the capacity of an MMO member server. This lets you continue to use the old primary server inside the domain, but as an MMO member rather than MMO primary member. To do this, run the following command on the old MMO primary that you want to act as an MMO member:

```
cleanMMOPrimary
```

Next, you must configure this machine to connect to the new MMO primary by running the `cammo_setup` script. After running this script, you must restart the `cadbd` process.

Reinstitute the MMO Primary Member Server

If you want to reinstitute the former MMO primary server as the MMO primary server after the machine is available, you must first perform a backup of the MMO data from the current MMO primary member. After you perform the backup, you must do the following:

1. On the former MMO primary that you want to reinstitute as the primary member, run the following commands:
 - `cleanMMOPrimary`
 - `moveMMOPrimary` using the backup of the current MMO primary
 - `cammo_setup` to configure this machine as the MMO primary
2. On the temporarily assigned MMO primary that you want to reassign as an MMO member, run the following commands:
 - `cleanMMOPrimary`
 - `cammo_setup` to configure this as an MMO member
3. On all the other MMO members, run the `cammo_setup` command and enter the proper MMO primary hostname when prompted.

Note: After running the `cammo_setup` script on all servers, you must restart the `cadbd` process.

Resynchronize an MMO Member Server

In a multiple BrightStor ARCserve Backup server environment, BrightStor ARCserve Backup synchronizes media management data between the MMO primary server and all MMO member servers. If any BrightStor ARCserve Backup MMO member server needs to be recovered from an earlier backup, the media management data becomes asynchronous.

There is a one-to-one relationship between tapes and BrightStor ARCserve Backup servers in a multiple MMO server environment.

- The first member server to write to a specific tape is the only member server that can read and write to the specific tape in the future.
- In a Storage Area Network (SAN), many member servers can detect and access all tapes mounted in any tape drive in the SAN. However, BrightStor ARCserve Backup enforces the exclusive member server to tape relationship by tracking and synchronizing the member to tape relationships in the aspool and astape tables.
- The aspool and astape tables consist of media management data that is updated by member servers, and is synchronized with and maintained in the MMO primary server.

If a situation arises that requires you to restore an MMO member server, the recovered server will be in an earlier state than the other BrightStor ARCserve Backup servers in the environment. For example, media management records maintained in the aspool and astape tables may have been modified. Therefore, to insure data integrity, the restored member server must be resynchronized with the MMO primary server.

Records of backup tasks performed by the restored MMO member server after its last full backup will not be available to the restored member server immediately after it is restored. You can recover the unavailable records using the Merge Manager after restoring the failed MMO member server. For more information about using the Merge Manager, see Merge Manager in the chapter "Managing the Database and Reporting."

To resynchronize the MMO member server with the MMO primary server, run the following command:

```
resynchMMOMember
```

If the dbserver main process (cabdb) was running when you executed the resynchronization script, you must restart the main process because the resynchMMOMember only command identifies the member server as the server that needs to be resynchronized. When you run the cstart command immediately after running the resynchMMOMember command, BrightStor ARCserve Backup executes the resynchronization process on the member server.

Remove an MMO Member Server from the BrightStor ARCserve Backup MMO Domain

BrightStor ARCserve Backup provides you with the capability to remove servers from your MMO domain such that they can be used in other capacities to other BrightStor ARCserve Backup MMO domains.

To remove an MMO member server from the BrightStor ARCserve Backup MMO domain, you must first break the associations of the member server's pools, tapes, and vault management data from the primary MMO server and the other member MMO servers in the MMO domain. You can use the Media Pool Manager to break the tape and media pool associations and the Media Management Administrator to break the vault relationships with the tapes and media pools.

After breaking the associations between the member server and the BrightStor ARCserve Backup MMO domain, perform the following tasks:

1. Stop all BrightStor ARCserve Backup services on the MMO member server that you want to remove from the MMO domain by running the `cstop` command.
2. Synchronize the member server that you are removing from the MMO domain with primary MMO server using the following command:

```
resynchMMOMember
```

3. On the member server, execute the `cstart` command. Allow the command two to three minutes to complete the process.

Important! *The `resynchMMOMember` command identifies the member server as the server that needs to be resynchronized. When you run the `cstart` command immediately after running the `resynchMMOMember` command, BrightStor ARCserve Backup executes the resynchronization process on the member server.*

4. Run the `cstop` command on the member server to stop all BrightStor ARCserve Backup services.
5. Reconfigure the member server that you are removing from the MMO domain using the `cammo_setup` command. Failure to do so will cause BrightStor ARCserve Backup to re-register the member server with the primary MMO server when you restart the member server.
6. Remove the pool, tape and vault data from the server that you are removing from the MMO primary server's database using the following command:

```
cleanMMOMember
```

The ca_mmomgr Command

This command is the command line interface to the MMO Admin window from the UNIX prompt. Use the ca_mmomgr command to start or simulate vault cycles, display reports by date, and reset the status. Many of the features available from the MMO Admin window are available from the command line. Using the ca_mmomgr command you can:

- Start, Simulate, or Reset Vault Cycle.
- Monitor current Media Management reports.
- View Media Management reports for specific dates.

ca_mmomgr has three kinds of commands:

- Vault Cycle commands—Provides for control of the vault cycle.
- Report commands—Displays the current Media Management reports.
- Report by Date commands—Displays the Media Management reports for a particular date.

You can automate Media Management operations by saving any of these commands as scripts.

For a complete list of the options and switches available for this command, see the appendix “Using Command Line Utilities.”

Configure Your Firewall to Optimize Communication

Note: BrightStor ARCserve Backup for Mainframe Linux does not support firewall configuration.

In an environment where you are using multiple BrightStor ARCserve Backup servers that reside across a firewall, or there is a firewall within a Storage Area Network (SAN) fibre loop, you must configure your servers to ensure the use of fixed ports and interfaces. The configuration on your BrightStor ARCserve Backup servers must match your firewall configuration so that BrightStor ARCserve Backup servers can communicate with each other.

A BrightStor ARCserve Backup server communicates with other BrightStor ARCserve Backup servers using a set of Remote Procedure Call (RPC) services. Each service can be identified by an interface (IP address) and a port. When you share data and tape libraries between BrightStor ARCserve Backup servers, the services communicate with each other using the interface and port information provided by the RPC infrastructure. RPC infrastructure, however, does not ensure specific port assignment. Therefore, you must know your RPC infrastructure and port number assignments to configure your firewall properly. To achieve static binding, additional configuration is required.

BrightStor ARCserve Backup uses RPC communication for the following types of servers:

- Primary and distributed SAN servers
- Primary and member Media Management Option (MMO) servers
- Primary and member domain authentication servers

Note: Although primary and distributed or member servers perform similar tasks, primary servers store shared data while distributed or member servers rely on primary servers to provide the same data.

You can specify IP addresses that your BrightStor ARCserve Backup servers should use to communicate with each other. Using custom IP addresses, computers residing in a private or public domain can communicate with SAN, MMO, and domain authentication servers that use static RPC port configurations.

You can create custom settings by modifying specific parameters in the RPC configuration file, `rpc.cfg`, found in the `$BAB_HOME/config` directory on all of your BrightStor ARCserve Backup servers. You must modify the port and IP address settings on all of your BrightStor ARCserve Backup servers that use RPC communication to match your firewall communication rules.

The following table lists the services that BrightStor ARCserve Backup components use for RPC communication:

Component	Service
SAN	Primary server: sanpeer1, sanpeer2, sanghost Distributed server: sanpeer1, sanpeer2
MMO	None. By default, MMO servers use port 28336.
Domain authentication	Primary server: cadiscovd

To modify the `rpc.cfg` file, use the following steps:

Note: Lines preceded with `#` are ignored. For more examples of RPC settings, see the `rpc.cfg` file.

1. Open the RPC configuration file located in the `$BAB_HOME/config` directory using a text editor.
2. To allow a service to communicate using a specific port or IP address, add a line, or as many as necessary, to the `rpc.cfg` file using the following format:

```
ServiceName(%s)      PortNumber(%d)
```

or

```
ServiceName(%s)      Address(%d.%d.%d.%d)
```

Note: You can specify a valid hostname rather than an IP address.

3. To specify a port range, use the following format:

```
PortNumberBegin(%d) - [ PortNumberEnd(%d)
```

4. Close the `rpc.cfg` file and save your changes.

Note: If the primary host is behind a firewall, you must perform the following steps, before running `csetup`, to ensure that the `sanpeer` service starts.

- a. Open the RPC configuration file, `rpc.cfg`, found in the `$BAB_HOME/config` directory.
- b. Check the language setting in `$BAB_HOME/lib/nls/nls.cfg`. Get the value of `CA_NLS_LANG`. It is the first word after `CA_NLS_LANG` as described in the following example:

If in `$BAB_HOME/lib/nls/nls.cfg`, you can find: `CA_NLS_LANG tc zh_TW.BIG5`, then the `CA_NLS_LANG` is `tc`.

- c. Copy `rpc.cfg` using the following command:
- d.

```
cp $BAB_HOME/lib/nls/${CA_NLS_LANG}/newconfig/rpc.cfg  
$BAB_HOME/config/rpc.cfg
```
- e. Replace `${CA_NLS_LANG}` with the actual value of `CA_NLS_LANG` in the command.
- f. Modify the port and IP address settings to correspond with your firewall communication rules.

Chapter 8: Managing the Database and Reporting

The BrightStor ARCserve Backup Database records job information for each job it runs, including media information. This information allows you to perform intelligent restores by keeping track of each file and directory backed up to media. When you want to restore a specific file, the database determines which media a file is stored on. The database information is also used to generate the various reports and logs available with BrightStor ARCserve Backup. You can monitor and manage the information in the BrightStor ARCserve Backup database using the Database, Merge, Scan, and Report Managers.

- The Database Manager lets you manage the database from a central console. Using the Database Manager, you can monitor database usage; view job, media, and client information; and configure database options.
- The Merge Manager allows you to add database information to your existing database files. Use the Merge Manager when you want to restore data from a different server than the one where your BrightStor ARCserve Backup database resides.
- The Scan Manager provides you with information about your media sessions. You can scan a single session or the entire media. Results of the media scan can be viewed in the Report Manager under the Activity Log listing or under the User Log listing if an additional log file is created.
- The Report Manager allows you to view the information stored in the BrightStor ARCserve Backup database using the available logs and reports. The Report Manager lets you monitor and manage these logs and reports.

databaseBackup.log

Each time a backup job successfully executes, BrightStor ARCserve Backup records information in its databases about the machines, directories, and files that have been backed up, and the media that were used. This allows you to locate files easily and quickly when you need to restore them. This database information is backed up by default whenever you back up your Advantage Ingres home directory. A log file, `databaseBackup.log`, located under the `$BAB_HOME/logs` directory, records information about every BrightStor ARCserve Backup database backup, such as time of backup, tape used, session number, sequence, and tape ID. Each time the BrightStor ARCserve Backup database is backed up, this information is appended to the log file.

Database Manager

BrightStor ARCserve Backup stores information in its database about the backup jobs you have run, the media you have used for backups, and the clients available for you to back up. You manage and control this information using the Database Manager. The Database Manager allows you to:

- Keep track of the location of your media.
- Determine the session number of your backup.
- Determine if media should be retired.
- View detailed, logged information about jobs you have run.
- Delete old records from the database.
- Graphically compare the size of your database to the total available disk space.

Database Manager Toolbar

The features available from the Database Manager Toolbar let you control the database, set options for your database, and manage database records. The following functions are available from the toolbar:



Configure—Allows you to set database options.



Add—Add a client to the database.



Modify—Modify a record in the database.



Delete—Delete a record from the database.



Filter—Select job filter options.

Configure Database Options

Click the Configure button to access the Configure dialog and set options for the database. Using this dialog, you can specify database grooming options and set maximum size limits for the database.

BrightStor ARCserve Backup has two methods of database grooming—pruning and purging. By default, BrightStor ARCserve Backup does not prune or purge the database. You must choose these options to enable database grooming.

Pruning preserves all job and session records but removes all details about the files included in the backup. If you have old job records in the database, but you think you may someday need to restore data from those old jobs, select Auto Prune Jobs. Although all of the file details will be removed from the database, you will still be able to restore entire sessions. Select the Auto Prune Jobs option to enable pruning, specify the maximum age of records before they are automatically pruned, and schedule a daily pruning operation.

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

Purging allows you to delete all traces of a database record, including session and file information from the database, with the exception of GFS rotation jobs. If you need to restore files from a job that has been purged, you can use the Merge Manager to merge the media information back into the database. Select the Auto Purge Jobs option to enable purging of your database, specify the maximum age of records before they are automatically purged, and schedule your daily purging operation.

Note: BrightStor ARCserve Backup uses the Advantage Ingres embedded database Version 2.6, Service Pack 2 with UNIX and Mainframe Linux platforms, and Version 3.0.2 with Linux platforms. For more information about how to maintain the Advantage Ingres embedded database, see the Advantage Ingres Embedded Edition Administrator's Guide included with your documentation.

Add Clients


For BrightStor ARCserve Backup to display a machine in its browser, the machine must have been entered into the BrightStor ARCserve Backup database. The server local to BrightStor ARCserve Backup is automatically entered into the database during installation, but you must add information for additional machines. Use the Add Client dialog to enter information about remote clients running agent software.

Click the Add button to invoke the Add Client dialog. The fields on the Add Client dialog allow you to enter information about the machine, including the name, the IP address, and the type of operating system.

If you do not know the IP address of the machine you are adding, enter only the host name. You can further describe the client by entering the following information:

- Networking Type
- Hardware Type
- OS Description
- OS Version.

You can view this information on the Client tab of the Database Manager.



Note: If you are adding a BrightStor Client Agent for NetWare, you must use the Novell server name as the Machine Name.

Record Modification

To modify information about a job in the BrightStor ARCserve Backup database, select the job in the left pane of the Database Manager and click the Modify button to invoke the Modify Job Comment dialog. This dialog displays information about the job from the database in each of the fields, allowing you to modify incorrect information, if available, or to enter a comment about the job in the Comments field.

Delete Records

You can use the Delete button on the Database Manager Toolbar to delete a record from the database. When you click the Delete button, BrightStor ARCserve Backup will confirm that you want to delete the selected record. You should not use this button to perform database grooming.

Note: When you delete a media, the entry remains, but the media is labeled destroyed. For more information about managing and maintaining the Ingres database, see the section Ingres Database Maintenance in this chapter.

Database Views

BrightStor ARCserve Backup has four database views from which to choose. You can change the view by selecting from the available tabs:

- **Summary**—Displays information about the database. This tab provides a summary of the percentage of the file system being used by the database and by all other files, and how much space is free, in addition to a records summary.
- **Job**—Displays information about all jobs processed. This tab gives you a view of all jobs you have run through BrightStor ARCserve Backup, including details about sessions and individual files.
- **Media**—Displays a view of all media you have used with BrightStor ARCserve Backup, including media statistics and session information for each media. This view provides information about the media used, formatting, how much the media has been used since it was put into service, and the number of errors (if any) that have occurred.
- **Clients**—Displays information about clients or machines added to the database. This tab provides a view of all the clients BrightStor ARCserve Backup can access for backup and restore operations.

Note: You must install the appropriate agent software on each host you will be accessing with BrightStor ARCserve Backup. You must also add the Node Name to the BrightStor ARCserve Backup Client database for each host on which you install the agent software.


Database Manager Summary Tab

When you open the Database Manager, the Summary view appears. This tab is informational only. You cannot modify the information that appears on this tab.

The window displays information about your host file system and the total number of records in each BrightStor ARCserve Backup database. The left pane graphically displays the amount of space the database uses on your hard disk. The same information is displayed numerically in the right pane.

The Record Summary, in the right panel, provides you with basic information about your database— the number of jobs, number of media, number of workstations, number of media pools, and number of detail records.

Note: If your database has grown very large, you may want to set options to groom old records from the database. For more information, see Configure Database Options in this chapter.



Job Tab

The Job tab lets you view information about every job processed by BrightStor ARCserve Backup. Depending on what you select in the left-pane, this includes either summary, job, session, or file information:

- Select the Job Records object in the left pane to display a basic summary of all jobs run by BrightStor ARCserve Backup. This includes the job ID, Type, Status, Owner name, and Set Name of every job listed in the left pane of the Database Manager.
- Select a job in the left pane to view the job ID, Start Time, End Time, Comment, Set Name, Owner Name, Status, and Type.

- Select a session in the left pane to view the Source Directory, Host Name, Session Number, Start Time, End Time, Session Type, Session Flags, Session Method, and Session Status.
- Select a file in the left pane to view its Path, Name, Modification Time, and Size.

On the Job tab, you can click the Modify or Delete button to modify the job comment or delete it from the database.

Media Tab

Select the Media tab to see the Media Records view of the BrightStor ARCserve Backup database. The Media Records view tracks information about the media used with BrightStor ARCserve Backup. This includes information about formatting, how often each particular media has been used since it was put into service, and the number of errors, if any, that have occurred while BrightStor ARCserve Backup was using that media.

Each time a media is formatted or merged into the database, a new record is added to the Media Records. In addition, media associated with other BrightStor ARCserve Backup servers is entered into the database as soon as the Media Server (camediad) reads it.

When you select the Media Records object in the left pane of the Database Manager, media summary information displays, including the media's name, ID, serial and sequence number, and whether or not the media has been destroyed. If you want to remove a media record from the database, click the Delete button.

When you select a media in the left pane of the Database Manager, four tabs appear, providing you with information about the media you selected.

Property Tab

The Property tab displays Media Information about the media you selected, including Name, Serial Number, ID, Sequence Number, Media Type, Format Code, Location Status, number of times the media has been formatted, and whether it has been destroyed.

Date Tab

The Date tab displays Life Information about the media, including:

- First Format Date—The date the media was first formatted.
- Last Format Date—The date of the last time the media was formatted.

- **Expiration Date**—The date on which the media is expected to have reached the end of its usable life. This does not mean that the media is unusable after it reaches this date, only that it is more susceptible to errors at this point.
- **Last Accessed Date**—The date of the last time the media was accessed.
- **Last Overwritten Date**—The date of the last time the media was overwritten.
- **Destroyed Date**—The date the media was destroyed, if applicable.

Current Stats Tab

The Current Stats tab displays Current Media Statistics for the media. These statistics include:

- **Medium Error**—Indicates that a read or write operation command terminated with a non-recoverable error condition. This error can be caused by a number of problems, including a flaw on the media, a dirty magnetic tape head, broken or bad media, or if the tape in the drive is a cleaning tape.
- **Soft Read Error**—The tape drive detected a problem when it was trying to read from the media but was able to correct the problem by trying the operation again.
- **Soft Write Error**—The tape drive detected a problem while trying to write to the media but was able to correct the problem by retrying the operation.

The Soft Read Error and Soft Write Error fields help you to determine the quality of your media. A certain number of soft-read and soft-write errors are normal. However, when the number of errors is very high in relation to the amount of data being written to the media, this may indicate that the media should be replaced.

- **KBytes Written**—The amount of data written to the media during the job.
- **Usage Time**—The amount of time the media was used in the job.
- **Read Passes**—The number of times the drive head has passed over a given point on the media when reading the media.
- **Write Passes**—The number of times the drive head passes over a given point on the media. For example, a backup without verification constitutes one pass, whereas a backup with verification constitutes two passes.

Total Stats Tab

The Total Stats tab displays the Total Media Statistics for the media. The type of information displayed on this tab is identical to the information given on the Current Stats tab, but the statistics on this tab are cumulative.

Clients Tab

The Clients view presents information about the clients you can back up and restore. Highlight the Client Records object to display a listing of clients currently in the BrightStor ARCserve Backup database and their IP addresses. Click the Add button to add a client. You can also use the Modify and Delete buttons to modify client information and delete a client from the database.

When you select a client in the left pane of the Database Manager, the right pane displays Machine Information for that client. This information, the Machine Name, IP Address, OS Type, Machine Type, Networking Type, Hardware Type, OS Description, and OS Version, is the same information entered into the database using the Add Client dialog.

Note: In order for BrightStor ARCserve Backup to communicate with any host, you must install and run the appropriate agent software. In addition, for each host on which you install the agent, you must add a record to the database associated with the BrightStor ARCserve Backup host server.

The `ca_dbmgr` Command

This command is the command line interface to the Database Manager. Use the `ca_dbmgr` command to maintain the BrightStor ARCserve Backup database, including configuring media pools. Using this command, you can query database information and set database options. This powerful utility allows other programs to interact easily with backup events. All of the features available from the Database Manager are available from the command line.

The `ca_dbmgr` command has the following types of options:

- Display options—Use these options to display information about the specified database, job, session, media, media pool, or client.
- Media Pool Management options—Use these options to manage and control media pools, including creating, modifying, and deleting media pools, and moving media from one media pool to another.
- Database Management options—Use these options to manage and control the database, including adding and removing clients, setting maximum database size, and setting database grooming options.

For a complete list of the options and switches available for this command, see the appendix “Using Command Line Utilities.”

Recover and Maintain the Advantage Ingres Database

The following sections describe how you can recover the Advantage Ingres database, and best practices for maintaining the database.

Recover the Advantage Ingres Database Offline

To recover the BrightStor ARCserve Backup database when Advantage Ingres has stopped working entirely, perform the following steps to reinstall Advantage Ingres and restore the latest data available. If only the BrightStor ARCserve Backup database, cadbase, is affected, see Recover the Database Online in this chapter.

1. Stop all processes by running `cstop` and `stopingres`.
2. Remove everything from the Advantage Ingres home directory and all extended locations, for example:

```
# rm -rf /disc2/ingresii/ingres
```

Note: You can obtain the home directory of Advantage Ingres from the `INGRES_HOME_PATH` variable in `$BAB_HOME/config/cadbd.cfg` or through the symbolic link `$BAB_HOME/dbase/ingres`. You can obtain all extended locations from `$BAB_HOME/dbase/extend_db.log`.

3. Run `$BAB_HOME/bin/cadbase_setup` to reinstall Advantage Ingres to the same location (for example, `/disc2/ingresii/ingres`), extend the database to the same locations, and run `access_db`. Advantage Ingres should be installed exactly as it was before the disaster.
4. Check the files in the `$BAB_HOME/logs/sessionSummary/` directory to find the tape and session information you need to perform data recovery.
5. Create a directory for the restoration (`/disc1/cadb-restore`) and subdirectories for each session (`/disc1/cadb-restore/s1`).

Ensure that you have enough disk space. Create these directories on the file system that originally contained the database to reduce the chances of running out of space.

6. Run `cstart`.
7. Use the Restore Manager to restore all of the data sessions and the dump session to your restoration subdirectories (`/disc1/cadb-restore/s1`, `/disc1/cadb-restore/s2`, ...), not to their original locations. Use the Restore By Media option and be sure to use the correct session number.

If you have multiple extended locations, check `BrightStor.log` and the files in the `$BAB_HOME/logs/sessionSummary/` directory to obtain the correct session numbers and restore to the correct directory. This might not be in the order you expect but the correct order is important for successful recovery.

8. Run `cstop`, and stop Advantage Ingres by executing `$BAB_HOME/sbin/stoppingres` and change to the user, `ingres`.
9. Copy the dump files from your restoration subdirectory to the Advantage Ingres installation directory. For example, if the dump session is restored to `/disc1/cadb-restore/s6`, enter the following commands:

```
# cd /disc1/cadb-restore/s6/ingres/dmp
# cp -r * /disc1/ingresii/ingres/dmp/
```
10. Remove all files from the data directories in which the `iidbdb` database data is stored (`II_SYSTEM/ingres/data/default/iidbdb`) to allow the files to be restored, for example:

```
# cd /disc1/ingresii/ingres/data/default/iidbdb
# rm -r *
```
11. Copy the `iidbdb` database configuration file from the `dmp` directory to the data directory, for example:

```
# cd /disc1/ingresii/ingres/dmp/default/iidbdb
# cp aaaaaaa.cnf/disc1/ingresii/ingres/data/default/iidbdb
```
12. Copy the `cadbase` database configuration file from the `dmp` directory to the data directory. This should be done on the first location only (for example, `$II_SYSTEM/ingres/data/default/cadbase`):

```
# cd /disc1/ingresii/ingres/dmp/default/cadbase
# cp aaaaaaa.cnf/disc1/ingresii/ingres/data/default/cadbase
```
13. Change to the root user.
14. Start Advantage Ingres by executing `$BAB_HOME/sbin/startingres`.

Important! *Starting Ingres at this point results in some processes not being started. These processes generate error messages. These messages are normal at this point and can be safely ignored. If necessary, press the Enter key or Ctrl+C to exit the error messages.*

15. To recover data, proceed as follows:

- a. Run `cadbutil` to recover `iidbdb` (`cadbutil -r iidbdb -c`). You are prompted for the directory that contains the restored data. Provide the subdirectory that contains the restored data for the `iidbdb` session (for example, `/disc1/cadb-restore/s3/iidbdb/`).
- b. Make sure that the MergeCat and the `dbclean` processes are not running.
- c. Run `cadbutil` to recover `cadbase` (`cadbutil -r cadbase -c`). You are prompted for the subdirectory that contains the restored `cadbase` session (for example, `/disc1/cadb-restore/s4/cadbase/`).

Note: For extended locations for databases, you are prompted for the sessions one by one. Enter the correct directory.

16. The database is successfully recovered. Run `stopingres`, and run `cstart` to start BrightStor ARCserve Backup. Alternately, you can run `startingres` to start only Advantage Ingres.

Recover the Advantage Ingres Database Online

If only `cadbase` (the BrightStor ARCserve Backup database) was lost or corrupted and Advantage Ingres itself is still running, you do not need to reinstall Advantage Ingres. The following procedure restores the `cadbase` database only:

1. Stop all processes by running `ctest`.
2. As the user, `ingres`, run the following command to destroy the `cadbase` database:

```
# destroydb cadbase -ucadbase
```

If the destruction of `cadbase` is successful, go to Step 3.

If the destruction of `cadbase` is unsuccessful, perform the following procedure:

- a. Remove all `cadbase` directories from the Advantage Ingres home location and all extended locations.

Note: You can obtain the home directory of Advantage Ingres from the `INGRES_HOME_PATH` variable in `$BAB_HOME/config/cadbd.cfg` or through the symbolic link `$BAB_HOME/dbase/ingres`. You can obtain all extended locations from `$BAB_HOME/dbase/extend_db.log`.

- b. Run `destroydb` again. It should be successful.
- c. As the root user, run `access_db` to delete all of the previous extended locations, if necessary. To do so, in the `access_db` window, enter the location, select the extended location, and enter Delete to delete it.

3. As the root user, run `$BAB_HOME/bin/cadbase_setup`. You are prompted to specify if you want to extend the database. If, in the previous step, the destruction of cadbase was successful, answer No.

If, in the previous step, the initial destruction of cadbase was unsuccessful, answer Yes, and use the same directory and locations as the extended locations.

4. Run `access_db` to assign all extended locations.
5. Run `extend_dbf`.
6. Check the files in the `$BAB_HOME/logs/sessionSummary/` directory to find the tape and session information for the backed up database sessions. This information is needed to perform data recovery.
7. Create a directory for the restoration (`/disc1/cadb-restore`) and subdirectories for each session (`/disc1/cadb-restore/s1`).

Ensure that you have enough disk space. Create these directories on the file system that originally contained the database to reduce the chances of running out of space.

8. Run `cstart`.
9. Use the Restore Manager to restore all the data sessions and the dump session to your restoration subdirectories (for example, `/disc1/cadb-restore/s1`, `/disc1/cadb-restore/s2`, ...), not to their original locations. Use the Restore By Media option and be sure to use the correct session number.

If you have multiple extended locations, check `BrightStor.log` and the files in the `$BAB_HOME/logs/sessionSummary/` directory to obtain the correct session numbers and restore to the correct directory. This may not be in the order you expect but the correct order is important for successful recovery.

10. Run `cstop`.
11. As the user, `ingres`, copy the cadbase dmp files from the restoration subdirectory to the Advantage Ingres installation directory. For example, when the dump session is restored to `/disc1/cadb-restore/s6`, the commands are:

```
# cd /disc1/cadb-restore/s6/ingres/dmp/default
# cp -r cadbase /disc1/ingresii/ingres/dmp/default
```

12. Copy the cadbase database configuration file from the dmp directory to the data directory. This should be done on the first location only (for example, `$II_SYSTEM/ingres/data/default/cadbase`):

```
# cd /disc1/ingresii/ingres/dmp/default/cadbase
# cp aaaaaaa.cnf/disc1/ingresii/ingres/data/default/cadbase
```

13. To recover data, you must change to the root user.

14. Verify that the MergeCat and the dbclean processes are not running.
15. Run cadbutil to recover cadbase (cadbutil -r cadbase -c). You are prompted for the subdirectory that contains the restored cadbase session (for example, /disc1/cadb-restore/s4/cadbase/).

Note: For extended locations for databases, you are prompted for the sessions one by one. Enter the correct directory.

16. The BrightStor ARCserve Backup database is successfully restored. Run cstart to start BrightStor ARCserve Backup.

Best Practices for Maintaining the Advantage Ingres Database

The following sections describe best practices for using the BrightStor ARCserve Backup database maintenance utility (ca_dbadmin). In these sections you will find information about how to:

- Maintain the BrightStor ARCserve Backup (Advantage Ingres) underlying database.
- Tune the database to improve its overall performance.
- Reclaim free space in the database to prevent performance deterioration.

You can find the ca_dbadmin command line utility in the following directory:

\$BAB_HOME/bin/ca_dbadmin

General Recommendations for Maintaining the Advantage Ingres Database

To improve the overall performance of the BrightStor ARCserve Backup (Advantage Ingres) database:

- Run the maintenance utility to generate reports on a regular basis to identify areas needing maintenance.
- Perform corrective action tasks regularly based upon the recommendations from maintenance reports.
- Provide troubleshooting data upon request when contacting Customer Support.

General Considerations for using the `ca_dbadmin` Utility

The following information must be considered before performing maintenance tasks:

- If you are using Advantage Ingres 2.6, you cannot perform maintenance tasks when the database is active on the following platforms:
 - IBM S/390 Linux
 - Solaris, versions 8, 9 , and 10
 - HP-UX
 - AIX
 - Tru64

When you run the `ca_dbadmin` utility from the command line, and BrightStor ARCserve Backup detects activity in the database, BrightStor ARCserve Backup will exit the maintenance process and notify you as demonstrated by the following example:

At least one of `cadbd` or `MergeCat` or `dbclean` services is running. Please re-run `ca_dbadmin` later.

- If you are using Advantage Ingres r3, you can perform maintenance tasks while the database is active on the following platforms:
 - All Linux platforms, except IBM S/390 for Linux
 - Linux IA-64
 - Solaris 10 AMD 64
- If necessary, you can stop maintenance tasks for the Advantage Ingres database. If maintenance tasks are stopped, all changes to the database are rolled back; however, you must recreate the indexes.
- When you restart a maintenance task, the `ca_dbadmin` utility starts from the beginning of the maintenance task.
- You can automate the process of maintaining the Advantage Ingres database details entity and update database statistics by scheduling a maintenance job using the `ca_db_maintenance` script.

For more information, see the section "ca_db_maintenance Script" in this chapter.

Database Maintenance Report

When to run the database maintenance report:

To insure that the BrightStor ARCserve Backup database is running with optimal performance:

- Run the maintenance utility once every seven days after a full backup is complete, or
- Immediately after purging or pruning the database to generate the maintenance report.

How to run the database maintenance report:

To run the maintenance report, open the command prompt and run the following command:

```
ca_dbadmin -report -type -maintenance -mark -severity critical/warning
```

- The -severity option can be either critical or warning. You can locate the criteria for both conditions in the following configuration file:

```
$BAB_HOME/dbase/maintenance/maintenance.cfg
```

- After you execute this command, the ca_dbadmin utility creates a maintenance report labeled *maintenance_report_mmddyy.txt* and stores it in the following directory:

```
$BAB_HOME/dbase.maintenance/reports
```

How to interpret the results provided by the maintenance report:

The resulting maintenance report presents you with statistical data about categories that affect database performance. The categories include:

- Maintenance on individual BrightStor ARCserve tables--this category identifies specific database tables that require maintenance.
- Diskspace Usage--this category identifies the need for maintenance if the ratio of disk space used compared to the maximum disk space allocated for the database exceeds optimal performance levels.
- Database/Diskspace Usage--this category identifies the need for maintenance if the ratio of free disk space in the database compared to the database's total disk space exceeds optimal performance levels.
- Diskspace Threshold Warning/Critical--this category identifies the need for maintenance if the database needs to be extended to function at optimal performance levels.

Perform Maintenance Tasks

Based upon the results and recommendations presented by the maintenance report, you should perform the maintenance tasks described in this section.

Important! *The target objects should be inactive when you are performing maintenance tasks.* For more information, see the section "General Considerations for using the ca_dbadmin Utility" in this chapter.

Tables Maintenance

After you run the maintenance report, tables marked YES for Marked For Maintenance require maintenance, as shown in the following example:

```
=====
Maintenance Report for Session Related Tables:
Table Name:                               astpses
Entity Name:                               session
Last Maintenance Date:                     01/01/2005
No. Of Rows From Last Maintenance Date:    200,000
Previous Maintenance Date:                 12/24/2004
No. Of Rows From Previous Maintenance Date: 150,000
Last Statistic Date:                       01/01/2005
Marked for Maintenance:                     YES
=====
```

In this example, the Entity Name "session" requires maintenance. To perform maintenance on all marked tables, run the following command:

```
ca_dbadmin -maintain -entity marked
```

Diskspace and Database/Diskspace Maintenance

Many catalog entries can be added and removed from the Advantage Ingres database every day. These events can result in excessive unused and wasted disk space, which can adversely affect the performance of the Advantage Ingres database. The Diskspace and Database/Diskspace categories identify if there is a need to perform maintenance, as shown in the following examples:

```
Diskspace Usage Report:
=====
The Preset Maximum Database Size: 200 MB/Unlimited
Diskspace Usage is at 76%
Is it recommended to perform maintenance as usage reaches 75%.
=====
```

If the Diskspace Usage exceeds the recommended performance threshold, run the following command:

```
ca_dbadmin -maintain -entity details
```

Database/Diskspace Usage Report

```
=====
The ratio of Database/Diskspace is at 65%
It is recommended to perform maintenance as ratio reaches 60%
=====
```

If the Database/Diskspace Usage exceeds the recommended performance threshold, run the following command:

```
ca_dbadmin -maintain -entity details
```

Note: For a detailed description about how to perform Diskspace and Database/Diskspace Maintenance, see the section "Compress the Ingres Database Using the ca_dbadmin Command" in this chapter.

Critical Diskspace Threshold Exceeded: Extending the Database

If the amount of free disk space available for the Advantage Ingres database falls below critical levels, you must extend the disk space for the database to another location. The following is an example of the Diskspace Threshold Warning/Critical Report:

Diskspace Threshold Warning/Critical Report

```
=====
                        Percentage UsagePath
Warning:
Critical:      80                      /export/home0/BAB/ingresii
=====
```

In this example, you must extend the Advantage Ingres database because it has reached the 80% critical level. For information about how to extend the Advantage Ingres database, see the section "Extend the Ingres Database" in this chapter.

Update Advantage Ingres Statistics

When to update Advantage Ingres statistics:

To ensure that the Advantage Ingres database optimizer is running in its most efficient state, you should update the Advantage Ingres statistics:

- Once every two weeks.

How to update Advantage Ingres statistics:

To update Advantage Ingres statistics, run the following command:

```
ca_dbadmin -maintain -type statistics
```

Note: The Advantage Ingres database provides you with the `optimizeingres` utility that you can use to update statistics after a large number of changes to the database have occurred. For more information, see the section "Optimize the Ingres Database" in this chapter.

Export the Advantage Ingres Database

When to export the Advantage Ingres database:

- If you need technical assistance with the Advantage Ingres database, the `ca_dbadmin` utility lets you export the database to a format that you can send to Computer Associates Customer Support for analysis.

How to export the Advantage Ingres database:

To export the Advantage Ingres database, run the following command:

```
ca_dbadmin -exportdb
```

Optimize the Advantage Ingres Database

To increase the performance of the Find operation in the Restore Manager, use the `optimizeingres` utility. Use the utility when a large number of changes in the database have occurred. Changes occur in the database when new hosts are added as clients, tapes have expired, and so on.

The `optimizeingres` utility is located in the following directory:

```
$BAB_HOME/bin/optimizeingres.
```

Note: There are no arguments to the `optimizeingres` utility

The `optimizeingres` utility calls the Advantage Ingres `optimizedb` utility on the more critical tables of the database.

Important! *If the tables contain a lot of data, the execution of the `optimizeingres` utility can take a considerable amount of time.*

Before executing the `optimizeingres` utility for the first time, you should run at least two backups of all client machines for which the BrightStor ARCserve Backup server is responsible. This provides Advantage Ingres a basis for the type of data that will be stored in the database from your normal workload of backups and allows the utility to update its own statistics about those tables at the same time.

ca_db_maintenance Script

The ca_db_maintenance script runs automated maintenance jobs that perform the following tasks:

- Maintain the details entity of the ca_dbadmin utility.
- Update the Advantage Ingres database statistics.

Note: You can also perform these tasks manually by running the ca_dbadmin command line utility. For more information about performing these tasks manually, see the sections "Perform Maintenance Tasks" and "Update Advantage Ingres Statistics" in this chapter.

Script location

You can find the ca_db_maintenance script in the following directory:

\$BAB_HOME/bin

Script syntax

```
ca_db_maintenance -retry -retry_interval
```

-retry

Specifies the number of attempts the script will attempt to run and complete the maintenance job, if a maintenance job failed due to activity in the database. The default value for this option is to retry 12 times.

-retry_interval

Specifies the time interval between retry attempts. The default value for this option is five minutes.

Example

```
ca_db_maintenance -retry 6 -retry_interval 15
```

In this example, if the database is busy, the maintenance task will retry a total of six times with an interval of 15 minutes between each attempt to run the maintenance tasks.

Run the ca_db_maintenance Script

You can run the ca_db_maintenance script manually or using a scheduler, as described in the following sections.

Schedule an automated maintenance job:

You can schedule an automated maintenance job using the following methods:

- Schedule the maintenance job using the Generic Job Manager.
Note: For more information about using the Generic Job Manager, see the section Generic Job Manager in the chapter "Customizing Your Jobs" in this guide.
- Use an external scheduler, such as Autosys, to run the maintenance job.
- Create a cron job.

Important! *You should schedule the maintenance job to run once every 30 days, at a time of day when there is no activity in the Advantage Ingres database.*

Modify or Remove a scheduled, automated maintenance job:

To modify or remove a scheduled, automated maintenance job, do the following:

- Use the Generic Job Manager to modify the options for the maintenance job. For example, the Repeating Interval.
- Use the Job Status Manager to delete a scheduled, automated maintenance job.
- Use the external scheduler to delete a maintenance job created with an external scheduler.
- Remove the maintenance job from the crontab, if you created a cron job.

Failed Maintenance Jobs

A maintenance job will fail if there is activity in the database. If you initiated the maintenance job using the Generic Job Manager, you can specify an option that lets BrightStor ARCserve Backup send an email message to a specified email address to alert you of the failed job. For more information, see Submit a Job Using the Generic Job Manager in the chapter "Customizing Your Jobs."

If a maintenance job fails for any reason, you can view the maintenance job details in the ca_dbadmin log file (labeled ca_dbadmin.log) located in the following directory:

`$BAB_HOME/dbase/maintenance`

Compress the Advantage Ingres Database using the ca_dbadmin Command

When you prune or purge data from the Ingres database, the overall size of the database itself may not decrease.

You should check the size of the Ingres database periodically (once every seven days is recommended) to verify that the overall size does not exceed the recommended levels. If the recommended levels are exceeded, you can use the ca_dbadmin command line utility to compress and reduce the size of the Ingres database.

To compress the size of the Ingres database, perform the following steps:

1. Run the ca_dbadmin command to generate a Maintenance Report as follows:

```
ca_dbadmin -report -type maintenance
```

Note: The default location for the generated Maintenance Report is \$BAB_HOME/dbase/maintenance/reports. The report is labeled maintenance_report.txt.MMDDYY. for example, maintenance_report.txt.010505.

2. Review the Maintenance Report and check the Diskspace Usage Report category and the Database/Diskspace Usage Report category.

Verify that the actual Diskspace Usage and Database/Diskspace Ratio values do not exceed the recommended values in the report.

3. If either of the actual values exceed the corresponding recommended values, you should run the ca_dbadmin command to perform maintenance and compress the databases as follows:

```
ca_dbadmin -maintenance -entity details
```

4. After the database has been compressed, generate another Maintenance Report and verify that the actual Diskspace Usage and Database/Diskspace Ratio values have been reduced to within the recommended values.

Extend Database Utility (extend_db)

The extend_db utility is used to extend the space available in the Advantage Ingres database by creating new locations for the database.

In the database log (\$BAB_HOME/logs/cadbd.log), a warning message is printed when you reach 80% of database usage. This warning message indicates that you should extend the database or verify the available space on the file systems used by the database. If you do not extend the database, it could become full and cause Ingres to send error messages every time you attempt to insert data. To resolve this problem you need to extend the database to additional locations. Extending a database consists of allowing Advantage Ingres to spread the data over multiple file systems. The procedure to extend a database is partially automatic but requires some manual intervention. As with any database modification, it is recommended that you first back up the database before extending it.

The extend_db utility performs the following functions:

- Allows you to extend the database to up to 32 locations—This includes the default location II_DATABASE. The extend_db utility lists all of the locations being used for the database and also specifies how many additional locations you can create. You will be prompted for the number of new locations you want to create and the directory path for each of the new locations. The directory path should be an absolute path that is not being used for other database locations that were created using the extend_db utility. You can create new locations under the same file system or you can use different file systems. The extend_db utility automatically generates the location name.
- Creates new locations—The extend_db utility creates the new locations you specified. After the new locations have been created, you need to manually assign those locations to the database and run a script to extend those locations to the database.
- Generates a log file extend_db.log under \$BAB_HOME/dbase directory—The log file contains the list of locations already being used for the database and the list of locations already created and to be extended.
- Generates a SQL script file \$BAB_HOME/bin/extend_table.sql to modify the tables with the locations just created. This SQL script file is used by the script extend_dbf to extend the database.

Extend the Advantage Ingres Database

If necessary, you can extend the Advantage Ingres database to multiple directory locations to allow for growth, if the underlying file system does not support large files (larger than 2 GB) or does not have enough free space. Additionally, you can spread the directory locations across multiple disks to improve database performance.

Note: If you use Advantage Ingres Version 2.0, you may need to extend the database to circumvent the 2 GB table size limitation. Advantage Ingres Version 2.6 and 3.0 do not have this table size limitation.

The `extend_db` utility can be used to extend the space available in the database by creating new locations for the database. You can extend the database during installation or at a later time.

The `extend_db` utility, located in the `$BAB_HOME/bin` directory, is called automatically by the setup script or manually by the root user if performed after installation. (It is recommended that you call the `extend_db` utility manually if Ingres is already installed). The amount of information in the database will be a factor in the amount of time required by the `extend_db` utility to run.

Important! *It is recommended that you back up the database prior to extending it.*

Advantage Ingres Version 2.6 or 3.0

If you are using Advantage Ingres Version 2.6 or 3.0, the process to extend the database is performed automatically by running the `extend_db` utility.

Advantage Ingres Version 2.0

If you are using Advantage Ingres Version 2.0, you must manually perform some additional steps to complete the database extension. If you extend the database during setup, perform this procedure after installation:

1. Run `access_db`, located in the `$BAB_HOME/bin` directory, to invoke an ASCII graphical interface where you can assign the extended locations to the database. Use terminal type VT100.

To enter a command in the interface, press the ESC key. To move up, press Ctrl+K; to move down, press Ctrl+J.

2. Choose Select Databases from the menu from each terminal window to complete the process.

3. Select the Extend option and perform the following steps:

- a. Select ListChoices.
- b. Select an extended directory using Ctrl+J or Ctrl+K from the list and choose Select from the menu.

Repeat this step as necessary until you have assigned all of the extended directories.

- c. Select Save from the menu.
 - d. Select End.
 - e. Select Quit to exit.
4. Run `extend_dbf`, located under the `$BAB_HOME/bin` directory, to create the data files in the extended directory.

Important! *You must complete these steps before you can run `extend_db` again.*

Merge Manager

The Merge Manager allows you to merge media containing one or more backup sessions into your BrightStor ARCserve Backup database. The Merge Manager appends database information to your existing database files. Use the Merge Manager when you want to restore data from a different server than the one on which your BrightStor ARCserve Backup database resides. For example, if a backup was created using BrightStor ARCserve Backup on a different machine, you can use the Merge Manager to get the media information into the database in the BrightStor ARCserve Backup home directory.

Merge Manager Toolbar

Two buttons available from the Merge Manager toolbar allow you to select options and submit your job:



Submit—Submit your merge job to run immediately or schedule it for a later time. For more information, see the chapter “Customizing Your Jobs.”



Option—Select options to customize your merge job. The options available are discussed in the Merge Options section in this chapter.


Merge Manager File Menu

The options available on the Merge Manager File menu allow you to perform several important actions. Select **Open** to open an existing job script. To create a job script from your merge job, select **Save** or **Save As** to invoke the **Save As** dialog to create a merge job script. For more information about Job Scripts, see Job Scripts in the chapter “Customizing Your Jobs.”

You can also start and stop all back-end services or individual back-end services using the options available from the File menu. For more information about controlling back-end services, see Back-End Services Access in the chapter “Introducing BrightStor ARCserve Backup.”

Merge Manager Window

The left pane of the Merge Manager displays your media organized by device groups. Expand the device groups to view the media they contain to allow you to locate the media you want to merge into the database. If the media you want to merge is not currently in a device, enter the name of the media in the Media Name field. You will be prompted to insert the media before the job begins.



In the top portion of the Merge Manager, indicate whether you are merging all information from the media or only from a specific session. If you want to merge all information from the media, select the All Session option. By default, BrightStor ARCserve Backup merges only the first sequence on any given media or set of media. To merge all data on the currently loaded media, select the "Selected Media" option.

If you know the sessions you want to merge, select the Range option and specify the session or range of sessions. If you do not know the session number, you can use the Scan Manager to list the contents of the media, including all session numbers. Alternatively, you can use the Media Name field to enter the name of media not currently loaded. BrightStor ARCserve Backup will prompt you for the media before the merge job runs.

Note: BrightStor ARCserve Backup must have formatted the media selected to merge.

Merge Manager Object Information Tab

When you select an object in the left pane of the Merge Manager, the right pane displays the Object Information tab. The information displayed on the Object Information tab varies according to the object you select in the left pane of the Manager. Select a group, and the Object Information tab displays the Group Name and Group Type. When you select a media in the left pane, the Object Information tab displays the Media Name, Media ID, Sequence Number, Device ID, and Serial Number of the media selected in the left pane of the Merge Manager.

Merge Options

The options available for the Merge Manager are, for the most part, identical to those available from the Scan Manager. Click the Options button on the Merge Manager toolbar to open the Option dialog, which allows you to set the following types of options for your job:

- Pre/Post—Run commands or batch files before the job runs, after it finishes, or both.
- Log—Determine the level of detail you want recorded into the Job Log and Specify Output devices for Activity Log.
- Database—Specify the level of detail to be recorded for your merge job. (For Merge jobs only.)
- Media Rules—Specify media options for the job, such as the media time-out period.

Pre/Post Options

BrightStor ARCserve Backup allows you to specify commands to perform various operations immediately before or after your data is merged based on the exit codes received. You can specify pre/post commands:

- For the entire merge job. These commands are executed at the beginning or end of the job.
- Run commands, executable programs, and shell scripts.

For example, you may want to specify this option to load a virus scanning application before files are merged and then run the batch file you created that sends a detailed report to the printer.

Note: If you specify pre/post commands for the entire job, as well as for machines in the job, the global pre/post commands are executed before or after the job starts or ends, and local commands are executed when the selected machine is backed up.

The Pre/Post option allows you to run a command on your BrightStor ARCserve Backup server machine before, after, or before and after the job is executed. The outcome of Pre/Post commands can be found in caqd.log. The available Pre/Post options are defined in the following sections.

Merge Manager Run Command Before Job Options

Enter the path and name of the application to be executed on the machine before the job starts.

- On Exit Code— If you want BrightStor ARCserve Backup to detect exit codes of any application, enable this option, select the condition for detecting exit codes (Equal To, Greater Than, Less Than, or Not Equal to), specify the exit codes you want detected, and then select how you want BrightStor ARCserve Backup to respond after exit codes are detected:
 - Note:** If you select the condition Greater Than or Less Than, only one exit code should be specified.
 - Skip Delay—If the exit code conditions are met, run the job immediately, ignoring the delay.
 - Skip Job—If the exit code conditions are met, skip the job.
 - Skip Post Application—If the exit code conditions are met, enabling this skips commands scheduled to run after the job completes.
- Delay in Minutes—Specify the delay in which BrightStor ARCserve Backup waits before running a job when the appropriate exit code is detected. This gives the specified application time to finish processing before your job begins.

Note: Ensure that the path to the command you want to run is correct. For example, if you want to run the script post_exec.ksh, and the command post_exec.ksh is in the usr directory, enter this:

```
/usr/post_exec.ksh
```

Merge Manager Run Command After Job Options

BrightStor ARCserve Backup runs the command after the backup job finishes. Enter the path to, and name of, the application to be executed on the machine after the job is completed. As with the Run Command Before Job option, you must ensure that the path to the command is correct.

Merge Manager Run Before/After Command As Options

Enter the User Name and Password to run Pre/Post commands. The User Name and Password correspond to the system of the host server selected, and are required to check the system privileges on that server. The user is authenticated using the File System Agent running on the server machine. The File System Agent must also be installed and enabled on the server machine to run Pre/Post commands.

The User Name and Password entered into these fields should not be confused with the BrightStor ARCserve Backup User Name and Password.

Merge Manager Log Options

Log options determine the level of detail included in the log report for the operation and the devices to which the log is sent. Unlike the Activity log, which displays the results of all jobs on a given server, Log options allow you to create a log file containing only the information associated with the specific job. The log file is accessed via the Report Manager and is shown under the User Log listing with the name it has been assigned.

When creating a log file for a merge job, you must enter a name for the file into the Log File Name field. You cannot use forward slashes (/) or backslashes (\) when entering the Log File Name. If they are used, the log file does not display in the Report Manager window.

After a name is entered for the log file, the following options can be specified:

- Log All Activities—Record all activity that occurs while the job is running in the Job Log.
- Log Summary Only—Record only summary information for the job (including source, destination, session number, and totals) and errors.

You can also select the output devices for the activity log. Select any or all of the following:

- **Unicenter NSM Alert**—This option allows you to send messages to the Unicenter Console when an alert is generated.
- **SNMP Alert**—This option allows you to send messages to your SNMP messaging console. The server running BrightStor ARCserve Backup must have SNMP configured for this option to work.
- **Printer Name**—This option sends messages to the specified printer local to the BrightStor ARCserve Backup server. The BrightStor ARCserve Backup server selected to print must be configured to do so, for this option to work.
- **Internet Email**—This option sends messages to the specified email address. The BrightStor ARCserve Backup server selected to send messages via email must be configured to do so, for this option to work.

Note: You can enable the activity log destination options (NSM alert messages, printers, and email addresses) by modifying the configuration file named `caloggerd.cfg`. This configuration file is located at: `$BAB_HOME/config`.

Media Rules Options

Using Media Rules options you can specify the timeout period BrightStor ARCserve Backup waits for you to provide the media you want to merge:

- **Timeout for the First Media**—Specify the number of minutes BrightStor ARCserve Backup waits for the first media required for your merge job to be inserted into a drive before it cancels the job. By default, BrightStor ARCserve Backup waits five minutes.
- **Timeout for Additional Media**—The number of minutes BrightStor ARCserve Backup waits for an additional media to be inserted into a drive before it cancels a job. By default, there is no timeout.

Merge Manager Database Options

Using Database options you can determine the level of detail to be included in the information recorded in the database for your merge operation. You can choose one of the following:

- **Record Detail Information**—Records detailed information about each job, including the session, and each file name in that session.
- **Record Job and Session Information Only**—Records only the job and the sessions in that job.

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

Submit a Merge Job

Select the Submit button to run the job. The result of the merge job displays in the Job Status Job Log and the Activity Log file. Both of these log files are viewed with the Report Manager.

The Submit button allows you to submit your job to run immediately or schedule it for a later time. For more information about scheduling jobs and the Submit button, see the chapter “Customizing Your Jobs.”

The ca_merge Command

The ca_merge command is the command line interface to the BrightStor ARCserve Backup Merge Manager. Use this command to create and submit merge jobs to the Job Queue. You can merge database information from backup media into your BrightStor ARCserve Backup database from the command prompt. All of the features available from the Merge Manager are available from the command line.

The commands and switches available for the ca_merge command allow you to control the following:

- Source arguments—These options specify the data to be merged. Use these options to identify the group, tape, and sessions to be used in your merge operation.
- Run Job arguments—These options allow you to submit the merge job to be run immediately, or to submit the job on Hold, or to schedule the job for a later date and time.
- Options—These switches allow you to specify Pre/Post command options, including passwords, log options and output devices, first and span media options, and to save the merge job as a script.

For a complete list of the options and switches available for this command, see the appendix “Using Command Line Utilities.”

Scan Manager

The Scan Manager provides you with information about your media sessions. You can scan a single session or the entire media. Results of the media scan can be viewed in the Report Manager under the Activity Log listing or under the User Log listing if an additional log file is created.

Note: BrightStor ARCserve Backup must have formatted the media selected to scan.

You can use the Scan Manager to find out what information is on a media. You might want to do this if you are trying to recover a BrightStor ARCserve Backup host and you need to find the most recent backup of the BrightStor ARCserve Backup database so that you can restore it.

You can also use the information on the Scan Manager to help you with other operations. For example, when merging information from a single session into the BrightStor ARCserve Backup database, you must specify the session number. If you do not know the session number, you can check the Activity Log. If that is not available, you can use the Scan Manager to list the contents of the media, including all session numbers. You can also use the Scan Manager to view a list of the files that were backed up.

Note: The UNIX client agent or the Linux client agent is required to perform any operations on any target host (including local hosts).

Scan Manager Toolbar

The Scan Manager toolbar contains two buttons, which allow you to select job options and submit the job for execution:



Submit—Submit your scan job to run immediately or schedule it for a later time. For more information, see the chapter “Customizing Your Jobs.”



Option—Select options to customize your scan job. The options available are discussed in the Scan Options section in this chapter.


Scan Manager File Menu

The options available on the Scan Manager File menu allow you to perform several important actions. Select Open to open an existing job script. To create a job script from your scan job, select Save or Save As to invoke the Save As dialog. For more information about Job Scripts, see Job Scripts in the chapter "Customizing Your Jobs."

You can also start and stop all back-end services or individual back-end services using the options available from the File menu. For more information about controlling back-end services, see Back-End Services Access in the chapter "Introducing BrightStor ARCserve Backup."

Scan Manager Window

The left pane of the Scan Manager displays media organized by device groups. If you want to submit a scan job using any device group, enable the Any Group option. Otherwise, expand the device groups to view the media they contain, to locate the media to be scanned. If the media you want to scan is not in a device, enter the name of the media in the Media Name field. BrightStor ARCserve Backup will prompt you for the media before the job begins.



In the top portion of the Scan Manager, indicate whether you want to scan all the information on the media or the information from a specific session only:

- To scan all information from the media, select the All Sessions option. The Scan Media Starts From options allow you to specify where the scan will begin. By default, BrightStor ARCserve Backup scans only the first sequence on any given media or set of media. To scan all data on the currently loaded media, select the "Selected Media" option.
- If you know the sessions you want to scan, select the Range option and specify the session or range of sessions to be scanned.

Scan Manager Object Information Tab

When you select an object in the left pane of the Scan Manager, the right pane displays the Object Information tab. The Object Information tab lets you view certain information about the object you select in the left pane. When you select a group, the Group Name and Group Type appear in the right pane of the Scan Manager. When you select a media in the left pane, the Object Information tab displays the Media Name, Media ID, Sequence Number, Device ID, and Serial Number of the media selected in the left pane of the Scan Manager.

Scan Options

The options available for the Scan Manager are, for the most part, identical to the options available for the Merge Manager. Click the Option button on the Scan Manager toolbar to open the Option dialog, which allows you to set the options for your job.

Pre/Post Options

BrightStor ARCserve Backup allows you to specify commands to perform various operations immediately before or after your data is merged based on the exit codes received. You can specify pre/post commands:

- For the entire merge job. These commands are executed at the beginning or end of the job.
- Run commands, executable programs, and shell scripts.

For example, you may want to specify this option to load a virus scanning application before files are merged and then run the batch file you created that sends a detailed report to the printer.

Note: If you specify pre/post commands for the entire job, as well as for machines in the job, the global pre/post commands are executed before or after the job starts or ends, and local commands are executed when the selected machine is backed up.

The Pre/Post option allows you to run a command on your BrightStor ARCserve Backup server machine before, after, or before and after the job is executed. The outcome of Pre/Post commands can be found in caqd.log. The available Pre/Post options are defined in the following sections.

Scan Manager Run Command Before Job Options

Enter the path and name of the application to be executed on the machine before the job starts.

- On Exit Code— If you want BrightStor ARCserve Backup to detect exit codes of other programs, enable this option, select the condition for detecting exit codes (Equal To, Greater Than, Less Than, or Not Equal to), specify the exit codes you want detected, and then select how you want BrightStor ARCserve Backup to respond after exit codes are detected:
Note: If you select the condition Greater Than or Less Than, only one exit code should be specified.
 - Skip Delay—If the exit code conditions are met, run the job immediately, ignoring the delay.
 - Skip Job—If the exit code conditions are met, skip the job.
 - Skip Post Application—If the exit code conditions are met, enabling this skips commands scheduled to run after the job completes.
- Delay in Minutes—Specify the delay in which BrightStor ARCserve Backup waits before running a job when the appropriate exit code is detected. This gives the specified application time to finish processing before your job begins.

Note: Ensure that the path to the command you want to run is correct. For example, if you want to run the script pre_exec.ksh, and the command pre_exec.ksh is in the usr directory, enter this:

```
/usr/pre_exec.ksh
```

Scan Manager Run Command After Job Options

BrightStor ARCserve Backup runs the command after the backup job finishes. Enter the path to, and name of, the application to be executed on the machine after the job is completed. As with the Run Command Before Job option, you must ensure that the path to the command is correct.

Scan Manager Run Before/After Command As Options

Enter the User Name and Password to run Pre/Post commands. The User Name and Password correspond to that of the system of the host server selected, and are required to check the system privileges on that server. The user is authenticated using the File System Agent running on the server machine. The File System Agent must also be installed and enabled on the server machine to run Pre/Post commands.

The User Name and Password entered into these fields should not be confused with the BrightStor ARCserve Backup User Name and Password.

Log Tab

These options determine the level of detail recorded in the Job Log, which contains information for your specific job, and is accessed using the Report Manager. Enter a Log File Name to create a log file and specify the level of detail to be recorded in this file. You can choose to record only Summary information or All Activities.

Note: Log all Activities applies to UNIX clients only.

You can also specify output devices for the Activity Log. You can send the activity log to any or all of the following— your SNMP messaging console, a specified UNIX printer local to the BrightStor ARCserve Backup server, the Unicenter Console, or specific email addresses. For these options to work, you must configure the BrightStor ARCserve Backup server for these options.

Media Rules Tab

These options allow you to specify the timeout period BrightStor ARCserve Backup waits for you to provide the media you want to scan. Use this tab to set the timeout period for the first media needed for your scan job and any additional media required.

Submit a Scan Job

Click the Submit button to execute your scan job. The job result displays in the Job Status Job Log, and the Activity Log file. View these log files using the Report Manager.

The Submit button lets you to submit your job to run immediately, or to schedule your job for a later time. For more information about scheduling jobs and the Submit button, see the chapter “Customizing Your Jobs.”

The `ca_scan` Command

This command is the command line interface to the Scan Manager. It allows you to create and submit scan jobs to the Job Queue. All of the features available from the Scan Manager are available from the command line. The `ca_scan` command reports information about one or more backup sessions on media.


The commands and switches available for the `ca_scan` command allow you to control the following:

- **Source arguments**—These options specify the data to be scanned. Use these options to identify the group, tape, and sessions to be used in your scan operation.
- **Run Job arguments**—These options allow you to submit the scan job to be run immediately, or to submit the job on Hold, or to schedule the job for a later date and time.
- **Options**—These switches allow you to specify Pre/Post command options, including passwords, log options and output devices, to save the scan job as a script, and select first and span media options.

For a complete list of the options and switches available for this command, see the appendix “Using Command Line Utilities.”

Report Manager

The information stored in the BrightStor ARCserve Backup database is used to generate many types of logs and reports. BrightStor ARCserve Backup provides access to these logs and reports with the Report Manager.



Note: In addition to the wealth of reports available using the Report Manager, BrightStor ARCserve Backup offers several flexible reporting tools that you use to generate reports using the command line. All of these reporting tools capture data from various information sources, such as the BrightStor ARCserve Backup database, other command line utilities, and logs stored in the \$BAB_HOME/logs directory. For more information about these reporting tools, see the appendix "Using Command Line Utilities."

Report Manager Toolbar

The buttons on the Report Manager toolbar provide functions to help manage both reports and logs. These functions include:



Cancel—Cancels the report currently being generated.



Delete—Deletes the entire user, rotation, or GFS rotation log.



Filter—Allows the filtering of the contents of a log based on its date or job ID. Job ID filtering is only supported when filtering the Activity Log.



Clear—Deletes all of the information from a log file. Additionally, you can use this option clear segments of information from the Activity Log based upon its time period.



Print—Prints the log or report.

The following sections describe the logs and reports available from the Report Manager.

Report Manager File Menu

You can control select services from the Report Manager File menu. You can start or stop all of the back-end services running on a host server, Queue services, Media services, and Database services, by selecting Start All Services or Stop All Services. You can also control individual services from the File menu. Select the back-end service you want to control. A drop-down menu allows you to select Start Service or Stop Service, depending on the required action.

Report Manager Logs and Reports

The left pane of the Report Manager lists the logs and reports you can view using the manager. Select one of the logs in the left pane to display the contents in the right pane of the Report Manager. The following logs are provided:

- **Activity Log (BrightStor.log)**—Logs all BrightStor ARCserve Backup activity. The Activity Log contains comprehensive information about the operations performed by BrightStor ARCserve Backup. It provides an audit trail of all BrightStor ARCserve Backup activity, including every job that has run. Scan this log every day to check if any errors have occurred. You can also use this log to find a session number if you need to restore a specific session. The Activity Log has an organize feature, which allows you to sort the log using filters, message grouping, or message post date.
- **Media Server Log (camediad.log)**—Provides an audit trail of all BrightStor ARCserve Backup Media Server activity, including when the server is loaded and unloaded.
- **Staging Log (staging.log)**—The Staging service monitors purge and migration jobs on staging devices. This log provides details about purge jobs (job creation, execution) and migration jobs (job creation) on staging devices. Furthermore, the Staging log captures information such as the date and time the job started (or ended), details about the staging device, and details about the final destination media.
- **Database Log (cadbd.log)**—Contains information pertaining to all possible database activities including Backups, Media Pool operations, Media formatting, and GFS operations. (For debugging purposes only.)
- **Queue Scheduler Log (caqd.log)**—Provides basic information about all of the jobs in the job queue, including sources, destinations, execution dates and times, and pre/post execution commands. You can view this log to get an inventory of the jobs in the job queue. You can also use this log to see what jobs other users have submitted.
- **Discovery Log (cadiscovd.log)**—Provides all basic domain-related information, including all available servers within a domain.
- **Caservd Log (caservd.log)**—Contains basic starting and stopping information about all BrightStor ARCserve Backup daemon activity, indicating the tasks each daemon is performing.
- **Authentication Log (cauthd.log)**—Contains comprehensive information about the security operations performed by BrightStor ARCserve Backup. It provides an audit trail of all the BrightStor ARCserve Backup user profiles that are added and deleted.

- **Tape Copy Log (tapecopy.log)**—Records all of the tapecopy messages displayed on the console that issued the tapecopy command. The following information about each session is included in the log:
 - Source and Destination Tape Name, ID and sequence number
 - Source and Destination session numbers
 - Amount of data transferred
 - Rate of transfer

In addition, depending on the tapecopy operation involved, the following information may also be available:

- Type of files included in the session (UNIX, NT NTFS, and so on)
- Type of backup the files belong to (full, incremental, and so on)

The Tape Copy Log is a record of all the activities that took place during the execution of the tapecopy command. For example, with a typical query-based tapecopy, the log entries begin with the list of sessions that were copied by tapecopy, followed by the tapecopy messages related to individual sessions.

Each line in the Tape Copy Log includes the process ID of the tapecopy process. When multiple instances of tapecopy are run simultaneously, you can use process ID to identify individual tapecopy operations.

Typically, the Tape Copy Log file needs to be used only when it is necessary to determine what happened during a tapecopy operation.

- **GFS Rotation/Rotation Daily Log**—Helps you manage your GFS rotation backups. It contains all of the Daily Logs for a GFS set, giving you information about your previous backup that can help you to prepare for your next backup. All GFS Daily Log files have a .dly extension. The file name is the media pool set name. The log resides in the \$BAB_HOME/logs/cas_user_logs directory.

If you run GFS rotation jobs, check this log each day so that you can physically label your media and retrieve the appropriate media.

Path:

cas_user_logs/BrightStor ARCserve_Backup_user/Media_Pool_Name

Example:

cas_user_logs/caroot/rotation_job.dly

- **GFS Rotation/Rotation Full Log**—Provides the history of a GFS rotation set, including the dates of each GFS rotation backup and the media used. Each GFS rotation set has its own Full Log. The Full Log does not contain the current Daily Log. All GFS Full Log files have a .ful extension. The file name is the media pool set name. The log resides in the \$BAB_HOME/logs/cas_user_logs directory.

View this log to get a general idea of what has been backed up. You may want to do this if you need to restore files from a GFS backup.

Path:

cas_user_logs/BrightStor ARCserve_Backup_user/Media_Pool_Name

Example:

cas_user_logs/caroot/rotation_job.ful

- **User Log**—Contains a list of any messages generated during any jobs (backup, restore, and so on) submitted by a particular user. This log is user-specific and resides in the \$BAB_HOME/logs/cas_user_logs directory. The user can specify any name.

Path:

cas_user_logs/BrightStor ARCserve Backup_user/Log_name

Example:

cas_user_logs/caroot/scanning_tapeA

Note: User Logs are not supported for NT Agents.

Report Manager Database Reports

The Report Manager allows you to view and print a variety of reports based on information in the database. The left pane of the Report Manager lists the types of reports available. Double-click the desired report to view the contents. The available reports are:

- **Job Records Report**—Provides a brief list of all jobs run by BrightStor ARCserve Backup. It contains information found in the Database Manager on the Job View window. For each job, this report includes the job number, job submitter, start date, status, job type, and media pool prefix.
- **Media Records Report**—Provides information about your BrightStor ARCserve Backup media. It contains information found in the Database Manager on the Media View window. It allows you to generate reports about all media or any individual media. This report includes the media name, ID, and sequence number, usage time, creation, format, and expiration date, amount of data written to the media, and the location status of the device.

- **Media Sessions Report**—Provides information about all of the backup sessions on media. Each source you choose to back up is saved on media as an individual session. The report contains session information found on the Database Manager in the Job View or Media View. It allows you to generate reports about any individual media or about all media. For each media session, this report includes the session number, status, host source selected, start time, end time, session method, flags, and total files/KB in the session.
- **Media Session Detail Report**—Includes all information found in the Media Session Report and lists every file backed up in each session. It contains session information found on the Database Manager in the Job View or Media View. You must select each session separately to view the report.
- **GFS Rotation Job Report**—Provides a brief historical record of all GFS backups performed on a target, listing the set the target belongs to and each of the backups that were performed. For each backup performed, the report includes the target name, target path, media name, sequence, session number, status, number of files backed up, and the amount of data backed up.
- **Rotation Job Report**—Provides a brief historical record of all rotation backups performed on a target, listing the set the target belongs to and each of the backups that were performed. For each backup performed, the report includes the target name, target path, media name, sequence, session number, status, number of files backed up, and the amount of data backed up.

The ca_log Command

The `ca_log` command is the command line interface to the Report Manager. This command allows you to view and maintain BrightStor ARCserve Backup logs. All of the features available from the Report Manager are available from the command line.

The options available for `ca_log` allow you to manage, display, and monitor log files:

- **Clear**—Use this option to reset specific log files.
- **Delete**—Use this option to delete specific log files.
- **Browse**—Use this option to view all available log files.
- **View**—Use this option, and its switches, to view a specific log file and tailor the information displayed.

For a complete list of the options and switches available for this command, see the appendix “Using Command Line Utilities.”

Diagnostic Utility

The BrightStor ARCserve Backup Diagnostic Utility is a convenient tool for gathering and packaging various BrightStor ARCserve Backup and system logs, which may be valuable tools for troubleshooting problems. This utility is installed, by default, on the BrightStor ARCserve Backup server.

The BrightStor ARCserve Backup Diagnostic Utility is comprised of two main components:

- Diagnostic wizard
- cadiag command line utility

Diagnostic Wizard

To run a complete diagnosis and collect local and remote data, you must run the diagnostic wizard. The diagnostic wizard is compressed (.jar) file installed in the BAB_HOME/httpd/diagnostic directory.

You can run the diagnostic wizard in one of the two following modes:

- Express—Collects information about the local machine. This mode does not collect advanced debugging information.
- Advanced—Collects data and generates reports with greater debugging information. When you choose this mode, BrightStor ARCserve Backup prompts you to rerun the relevant job so that the newly specified debug flags can be processed during the job and entered into a report.

Using the advanced mode, you can collect diagnostic information from any machine, provided the diagnostic daemons are running on that particular machine.

Note: Regardless of whether you choose the Express mode or the Advanced mode, you can specify where you want to save the diagnostic information file.

cadiag Command Line Utility

Using the cadiag command line utility, you can start, stop, and check the status of the diagnostic daemons running on the target machine. Additionally, you can use the command line to unpack the results of running the diagnostic wizard to directory on a remote machine, or a local machine where the packed or compressed diagnostic information file was saved.

For more information about syntax and options, see cadiag in the appendix "Using Command Line Utilities."

Run the Diagnostic Utility

To run the Diagnostic Utility:

1. Verify that the diagnostic daemons are running on all the machines that you want to include in the diagnosis. You can use the diagnostic utility command line utility to perform this task.

For more information the `cadiag` command line utility, see `cadiag` in the appendix, "Command Line Utilities."

2. In your browser's address bar, enter the following URL:

Note: `BABwebsite` represents the IP address or host name of the target server.

`http://BABwebsite/diagnostic/index.html`

The Login host dialog opens.

3. Enter the User name and Password, and click OK.

The BrightStor ARCserve Backup Diagnostic Tools Welcome screen opens.

The Welcome screen displays information about BrightStor ARCserve Backup applications and the operating system that is running on the server. The information about BrightStor ARCserve Backup includes version information about BrightStor ARCserve Backup installed components, the host name of the server, the operating system name, file system status, and so on.

4. Click Next.

The Select Diagnostic Type dialog opens.

5. Choose Express or Advanced and then click Next.

Note: For more information about these modes, see Diagnostic Wizard in this chapter.

The What do you want to generate? dialog opens.

6. If you specified Express mode, the local machine name displays on this dialog. If you specified Advanced mode, all machines in your domain with diagnostic daemons running display on this dialog.

Optionally, you can add machines that are not in your domain. To do this, enter the machine name in the Remote Machine text box and click Connect. If the BrightStor ARCserve Backup Diagnostic wizard detects the remote machine, the remote machine name displays on this dialog. Repeat this step as necessary to add other remote machines.

Note: For the diagnostic utility to detect a remote machine, BrightStor ARCserve Backup must be installed on the remote machine and the diagnostic daemons must be running when you start the diagnostic wizard.

7. Select the machines that you want to diagnose and click Next.

The Select debug levels dialog opens.

8. Choose a debug level and click Next.

Note: A Normal Debug Flags setting generates more detailed log information than No Debug Flags, and the Advanced Debug Flags setting generates more detailed log information than the Normal Debug Flags.

The Save the Diagnostic Information File dialog opens.

9. Enter the path to where you want to save the Diagnostic Information File and click Next.

The diagnostic utility performs an analysis of your system, collects the data, and packs the information into the specified Diagnostic Information File with a file extension of .asd.cz.

You can use the command line utility (cadiag) to unpack and view the collected diagnostic information. For more information about the command line syntax and options, see cadiag in the appendix "Command Line Utilities."

BrightStor Portal Reporting

BrightStor Portal is a separately sold management tool that provides a common interface for you to view, implement, report on, analyze, and monitor storage management procedures for many different technologies across a wide variety of platforms. Among these technologies is BrightStor ARCserve Backup.

If you use BrightStor Portal, you can use it to view and customize real-time reports on the backup and recovery tasks performed by BrightStor ARCserve Backup. To do this, you must first configure BrightStor Portal to perform a Host Discovery to discover the network objects that are likely to have storage management applications, assets or both, and an Application Discovery to determine if BrightStor ARCserve Backup is available on these hosts. Then, using the results of the Host and Application Discovery, you must specify the BrightStor ARCserve Backup Methods (or specific BrightStor ARCserve Backup components) you want to run reports on via the Portal. For more information on configuring BrightStor Portal, see the *BrightStor Portal Getting Started*.

After you have configured BrightStor Portal to work with BrightStor ARCserve Backup, click the Knowledge tab and navigate to your server name (the default path for BrightStor ARCserve Backup r11.5 is BrightStor Portal Knowledge\Monitor\Backup\BrightStor ARCserve Backup r11.5\servername) to select one of the BrightStor ARCserve Backup methods in the right-hand pane and view a report.

When the report appears, you can customize it by clicking a column name to sort a column, a scissor icon to delete a row, and the buttons at the bottom of the screen to view your data in different formats (pie or bar charts, line graph, and trend graph). Click the save button to save your customizations or click the default query button to go back to the default view (prior to your customizations). You can also click the publish button to create a customized view of a report with particular users in mind. If you do this, you can specify the location where you want the report to appear in the menu, create a title for the report, and assign the level of access and permissions to the report. For more information on report customizations, see the BrightStor Portal documentation.

Chapter 9: Recovering Data from a Disaster

The most important thing you can do to insure against data loss is to maintain current backups of all of your servers and workstations. If you do not maintain regular backups, BrightStor ARCserve Backup is limited in its ability to recover your data after a disaster. Be sure to create a media rotation schedule to maintain current full backups.

If disaster does strike, BrightStor ARCserve Backup provides you with the ability to recover your system quickly and efficiently. The method varies with your operating system. Disaster recovery is available only for Sun Solaris and HP-UX. For Sun Solaris, BrightStor ARCserve Backup supports a separately-sold, automated Disaster Recovery Option that allows you to recover your system completely, restoring it to its state at the time of the last disaster recovery backup. This lets you avoid performing a full reinstallation and reconfiguration of your operating system and other installed software packages.

If you do not have the Disaster Recovery Option, the following chapter outlines the procedures necessary to recover your system in the event of a disaster on a Sun Solaris or HP-UX platform.

For information about the Disaster Recovery Option, see the *Disaster Recovery Option Guide*.

Disaster Recovery Data

To help you to quickly identify the tapes you need for disaster recovery, BrightStor ARCserve Backup records information in its databases about the machines, directories, and files that have been backed up in every successful backup operation and the media that were used. This helps you to locate files when you need to restore them. Database information is backed up by default whenever you back up your Advantage Ingres home directory.

A log file, `databaseBackup.log`, records information about every BrightStor ARCserve Backup database backup, such as the time of backup, tape used, session number, sequence, and tape ID. Each time the BrightStor ARCserve Backup database is backed up, this information is appended to the log file. The `databaseBackup.log` file is located under the `$BAB_HOME/logs` directory.

Manual Disaster Recovery for Sun Solaris

The following procedure is a manual process for performing disaster recovery for Sun Solaris. The procedure helps you to protect Sun Solaris 2.7, 2.8, 2.9, and Solaris 10 operating systems against a disaster. You will need the Sun Solaris installation CD for the version of the operating system running on the system before the disaster, the BrightStor ARCserve Backup installation CD, and an extra hard disk for the temporary operating system.

Prepare for Disaster Recovery on Sun Solaris Systems

The hard disk layout information for the system to be protected should be saved to an alternate location on the network, or in hard copy, so that the hard disk layout can be reconstructed in the event of a disaster. The following preparatory guidelines will help you to collect critical system-specific information about your machine before a disaster strikes.

- Save a copy of your `/etc/vfstab` file, which provides file system information about each of the mounted volumes and their mount points.
- Using the `format` command, collect slice information for each hard disk on your system.
- You should perform regular full backups of the machine and keep the information about the media containing the last known full backup in a convenient location, for easy access in the event of a disaster.

Recover From a Disaster on Sun Solaris Systems

If disaster strikes, the following procedure will help you perform a recovery operation:

1. Install the temporary operating system on an alternate disk. Use the same version of the operating system CD that you had on the original system.

In most cases, follow the default installation procedure but ensure that you do not include the hard disks that belonged to the original system as a part of the temporary operating system installation.

2. Reboot to the temporary installed operating system on the alternate disk.

3. Use the format command to create slices on the original hard disks. Use the information saved earlier to create these slices.

For example:

- a. Run the format command.
- b. Select the disk device on which slices need to be created.
- c. Enter p or partition to select the partition menu.
- d. Enter the number corresponding to the slice that needs to be created. For example, enter 3 to create slice 3.
- e. Enter l or label to commit the changes that were made to the partition table.

4. Create file systems on these slices. Use the newfs command to create a UFS file system on a disk slice.

For example, use the following command to create a file system on slice 3:

```
#newfs /dev/rdisk/c0t0d0s3
```

5. Make the original root disk bootable. Use the installboot command, as follows:

```
#installboot /usr/platform/<architecture>/lib/fs/ufs/bootblk  
/dev/rdisk/c0t6d0s0
```

To determine the architecture of the system, use the following command:

```
#uname -m
```

6. Create a mount point directory for each of the file systems you created, to which you intend to restore your data.
7. Mount the slices at the mount points you created. For example:

```
#mount /dev/dsk/c0t6d0s0 /root1
```

where root1 is the mount point directory for the root volume of the original system.

8. To Recover a BrightStor ARCserve Backup server and a BrightStor ARCserve Backup client, install BrightStor ARCserve Backup and the BrightStor ARCserve Backup Agent for UNIX.

9. To Recover a BrightStor ARCserve Backup server: Run a Scan Job on the tape containing the last known good full backup of the machine you are recovering. Specify to save summary information from the scan to a file. For more information about scanning, see Scan Manager in the chapter "Managing the Database and Reporting."

Using the Restore by Media option on the BrightStor ARCserve Backup Restore Manager, and the scan job information previously saved to a file, restore the original machine data to the respective file systems you mounted earlier. For more information about restore operations, see the chapter "Restoring Data."

To Restore a BrightStor ARCserve Backup client: Use the BrightStor ARCserve Backup server to restore the last full backup of this machine to the respective file systems you have already mounted.

To Restore the BrightStor ARCserve Backup database: You must restore the BrightStor ARCserve Backup server before you restore the BrightStor ARCserve Backup database. The following procedure details how to recover the BrightStor ARCserve Backup server to restore the Advantage Ingres database:

- a. Create a directory representing your Advantage Ingres installation on your original system.

For example, if INGRES_HOME_PATH was /opt/ingresii/ingres, and you have mounted the disk slice containing the opt file system as /opt1, use the following command:

```
#mkdir /opt1/ingresii/ingres
```

- b. Restore the database sessions to their respective file systems but to an alternate temporary disk place. There will be at least four sessions, which will contain all of the database backups. The first session contains the Ingres binaries (executables, libraries, and so forth). The second session contains the iidbdb database (Ingres master database). The third session contains the cadbase database and could be spread into several sessions if more than one location is being used to receive the data. The last sessions contain the dump directory required to perform a recovery of the database.

10. Dismount the file systems you had mounted for the restore operation.
11. Reboot the machine.

Note: Before the machine reboots, disconnect the alternate hard disk.

12. Connect as "ingres".

Note: In the following steps, we assume that Ingres has been originally installed in the /opt/ingresii directory. We also assume that the first data location (INGRES_DATABASE_PATH) is in this directory. If you specified a different directory during Ingres installation, use that directory information rather than the one in this example.

13. Copy the binaries from the restore directory to the Ingres installation directory. Binaries are located in the first database session restored.

For example, if the binary session is restored to /opt/cadb-restore/s1:

```
# cd /opt/cadb-restore/s1/  
# cp -r * /opt/ingresii
```

14. Copy the dump files from your restore directory to the Ingres installation directory.

For example, if the dump session is restored to /opt/cadb-restore/s6.

```
# cd /opt/cadb-restore/s6/dmp  
# cp -r * /opt/ingresii/ingres/dmp/
```

15. Copy the iidbdb database configuration file from the dmp directory to the data directory.

For example:

If the data used to be stored at the same place as the binary:

```
# cd /opt/ingresii/ingres/dmp/default/iidbdb  
# cp aaaaaaaa.cnf /opt/ingresii/ingres/data/default/iidbdb
```

If the data used to be stored at a different location than the binaries (for example: /Ingres-data):

```
# cd /opt/ingresii/ingres/dmp/default/iidbdb  
# cp aaaaaaaa.cnf /Ingres-data/ingres/data/default/iidbdb
```

16. Copy the cadbase database configuration file from the dmp directory to the data directory. This should be done on the first location only (for example, \$II_SYSTEM/ingres/data/default/cadbase).

For example:

If the data used to be stored at the same place as the binary:

```
# cd /opt/ingresii/ingres/dmp/default/cadbase  
# cp aaaaaaaa.cnf /opt/ingresii/ingres/data/default/cadbase
```

If the data used to be stored at a different location than the binaries (ex: /Ingres-data):

```
# cd /opt/ingresii/ingres/dmp/default/cadbase  
# cp aaaaaaaa.cnf /Ingres-data/ingres/data/default/cadbase
```

17. Exit user ingres.
18. From a root user, start Ingres by executing `$BAB_HOME/sbin/startingres`.

Important! *Starting Ingres at this point results in some processes not being started. These processes generate error messages. These messages are normal at this point and can be safely ignored. If necessary, press the Enter key or Ctrl+C to exit error messages.*

19. To recover data, change the user to root, if you have not already done so.

Run `cadbutil` to recover `iidbdb` (`cadbutil -r iidbdb -c`). You are prompted for the directory containing restored data. Provide the directory containing the restored data for the `iidbdb` session (for example, `/opt/cadb-restore/s3`).

Note: The `-c` option of `cadbutil` will copy the file from the temporary directory you have used to restore the data sessions to the proper place. If the file system does not contain enough disk space to support the copy operation, you can use the `-m` option to move the data instead.

Run `cadbutil` to recover `cadbase` (`cadbutil -r cadbase -c`). You are prompted for the directory containing the restored `cadbase` session (for example, `/disc1/cadb-restore/s4`).

Note: For extended locations for databases, you are prompted for the sessions one by one. Enter the correct directory.

20. The Ingres database is successfully recovered. Run `stopingres`, and run `cstart` to start BrightStor ARCserve Backup. Alternately, you can run `startingres` to start only Ingres.

Manual Disaster Recovery for HP-UX

The following procedure is a manual process for performing a disaster recovery on HP-UX. To perform a disaster recovery operation, you need the HP-UX installation CD for the version of the operating system running on the system before the disaster and the BrightStor ARCserve Backup installation CD. You also need an extra hard disk for the temporary operating system. The following procedures assume that the machine has Logical Volume Manager (LVM) installed.

Note: The following procedure is a manual process for performing a disaster recovery on HP-UX 11.0 only. The syntax of the system commands you must enter may differ depending on your platform and hardware configuration. See your HP-UX documentation to verify the syntax of all system commands before performing this procedure.

Prepare for Disaster Recovery on HP-UX Systems

The LVM information for the system to be protected from disaster should be saved to an alternate location on the network, or in hard copy, so that a machine can be reconstructed in the event of a disaster. The following guidelines help you to collect critical system specific information about your machine before disaster strikes.

1. Save a copy of your `/etc/fstab` file. This file provides you with file system information on each of the mounted volumes and their mount points.
2. Use the `pvdisplay` command for each of your physical volumes, and save the output to a file. This information is required when physical volumes are recreated using `pvcreate`.
3. Use the `vgdisplay` command for each of your volume groups, and save the output to a file. This information is required when volume groups are recreated using `vgcreate`.
4. Use the `lvdisplay` command for each logical volume in each of the volume groups, and save the output to a file. This information is required when logical volumes are recreated using `lvcreate`.
5. Save the output of `lvlnboot -v` to a file. This information is required when recreating root, boot, swap and dump volumes using `lvlnboot`.
6. Note the major number and minor number of the node group in each of the volume groups. For example:

```
#ls -l /dev/vg00/group
```

Tip: You should perform regular full backups of the machine in question and keep information about the tape containing the last known good full backup in a convenient location for easy access in the event of a disaster.

Recover From a Disaster on HP-UX Systems

If disaster strikes your system, the following procedures let you perform a disaster recovery operation:

1. Install the operating system on an alternate disk:
 - a. Install the version of HP-UX you are using on an alternate hard disk.
 - b. In most cases, follow the default installation procedure, but ensure that you do not include the hard disks, which belong to the original system, as part of the temporary operating system installation.
 - c. Manually modify the Volume Group name on the alternate disk so that it does not match any of the volume groups on the machine before it failed. For example, if you had `vg00` on the original system, use `vg01`.
2. Reboot to the temporary installed operating system on the alternate disk.

3. Run the following command for each of the hard disks that existed on your system before it failed:

```
#pvcreate -f -B /dev/rdisk/c0t6d0
```

4. Recreate all the volume groups that existed on the original disks, using the following steps:

- a. Create a directory for each of the volume groups. For example:

```
mkdir /dev/vg00
```

where vg00 is the original volume group.

- b. Create a group node file in each directory you created. Use the major number and minor number of the nodes that existed earlier. For example:

```
mknod /dev/vg00/group c 64 0x000000
```

where 64 and 0x000000 are the major and minor numbers for the volume group vg00.

- c. Create volume groups. For example:

```
vgcreate vg00 /dev/dsk/c0t0d0
```

Note: The vgcreate command might fail with a "File too large" error on hard disks with certain geometry. If this error occurs, see the documentation available on the HP-UX web site to determine what additional parameters to use.

5. Recreate all the logical volumes that existed on your machine before it crashed. For example:

```
lvcreate -L 64 -n lvol1 -C y -r n vg00
```

where the size of the volume lvol1, belonging to the volume group vg00, is 64MB. See the man page for more details.

6. Make the disk bootable. For example:

```
#mkboot -l /dev/dsk/c0t6d0
```

where /dev/dsk/c0t6d0 is the original boot disk.

7. Set the root, boot, swap, and dump partitions using lvinboot. For example:

```
lvlnboot -b /dev/vg00/lvol1
```

```
lvlnboot -r /dev/vg00/lvol1
```

```
lvlnboot -d /dev/vg00/lvol2
```

```
lvlnboot -s /dev/vg00/lvol2
```

where /dev/vg00/lvol1 is the original boot volume, /dev/vg00/lvol1 is the original root volume, and /dev/vg00/lvol2 is the original swap and dump volumes.

8. Create the file system on the volumes you just created. For example:

```
#newfs -F hfs -L /dev/vg00/r1vo11  
#newfs -F vxfs -L /dev/vg00/r1vo13
```

9. Create a mount point directory for each of the file systems you created and to which you intend to restore the data:

```
#mkdir /stand1
```

10. Mount the logical volumes at the mount points that you created. For example:

```
#mount /dev/vg00/lvo11 /stand1
```

11. Install the following to be able to recover data:

To Recover a BrightStor ARCserve Backup server: Install BrightStor ARCserve Backup and the BrightStor ARCserve Backup Agent for UNIX.

To Recover a BrightStor ARCserve Backup client: Install the BrightStor ARCserve Backup Agent for UNIX.

12. To Recover a BrightStor ARCserve Backup server: Run a Scan Job on the tape containing the last known good full backup of the machine you are recovering. Specify to save summary information from the scan to a file. For more information about scanning, see Scan Manager in the chapter "Managing the Database and Reporting."

Using the Restore by Media option on the BrightStor ARCserve Backup Restore Manager and the scan job information previously saved to a file, restore the original machine data to the respective file systems you mounted earlier. For more information about restore operations, see the chapter "Restoring Data."

To Restore a BrightStor ARCserve Backup client: Use the BrightStor ARCserve Backup server to restore the last full backup of this machine to the respective file systems you have already mounted.

To Restore the BrightStor ARCserve Backup database: You must restore the BrightStor ARCserve Backup server before you restore the BrightStor ARCserve Backup database. The following procedure details how to recover the BrightStor ARCserve Backup server to restore the Advantage Ingres database:

- a. Create a directory representing your Advantage Ingres installation on your original system.

For example, if INGRES_HOME_PATH was /opt/ingresii/ingres, and you have mounted the disk slice containing the opt file system as /opt1, use the following command:

```
#mkdir /opt1/ingresii/ingres
```

- b. Restore the database sessions to their respective file systems, but to an alternate temporary disk place. There will be at least four sessions, which will contain all of the database backups. The first session contains the Ingres binaries (executables, libraries, and so forth). The second session contains the iidbdb database (Ingres master database). The third session contains the cadbase database and could be spread into several sessions if more than 1 location is being used to receive the data. The last sessions contain the dump directory required to perform a recovery of the database.

13. Dismount the file systems you had mounted for the restore operation.

14. Reboot the machine.

Note: Before the machine reboots, disconnect the alternate hard disk.

15. Connect as "ingres".

Note: In the following steps, we assume that Ingres has been originally installed in the /opt/ingresii directory. We also assume that the first data location (INGRES_DATABASE_PATH) is in this directory. If you specified a different directory during Ingres installation, use that directory information rather than the one in this example.

16. Copy the binaries from the restore directory to the Ingres installation directory. Binaries are located in the first database session restored.

For example, if the binary session is restored to /opt/cadb-restore/s1:

```
# cd /opt/cadb-restore/s1/
# cp -r * /opt/ingresii
```

17. Copy the dump files from your restore directory to the Ingres installation directory.

For example, if the dump session is restored to /opt/cadb-restore/s6.

```
# cd /opt/cadb-restore/s6/dmp
# cp -r * /opt/ingresii/ingres/dmp/
```

18. Copy the iidbdb database configuration file from the dmp directory to the data directory.

For example:

If the data used to be stored at the same place as the binary:

```
# cd /opt/ingresii/ingres/dmp/default/iidbdb
# cp aaaaaaaa.cnf /opt/ingresii/ingres/data/default/iidbdb
```

If the data used to be stored at a different location than the binaries (ex: /Ingres-data):

```
# cd /opt/ingresii/ingres/dmp/default/iidbdb
# cp aaaaaaaa.cnf /Ingres-data/ingres/data/default/iidbdb
```

19. Copy the cadbase database configuration file from the dmp directory to the data directory. This should be done on the first location only (for example, \$II_SYSTEM/ingres/data/default/cadbase).

For example:

If the data used to be stored at the same place as the binary:

```
# cd /opt/ingresii/ingres/dmp/default/cadbase
# cp aaaaaaaa.cnf /opt/ingresii/ingres/data/default/cadbase
```

If the data used to be stored at a different location than the binaries (ex: /Ingres-data):

```
# cd /opt/ingresii/ingres/dmp/default/cadbase
# cp aaaaaaaa.cnf /Ingres-data/ingres/data/default/cadbase
```

20. Exit user ingres.

21. From a root user, start Ingres by executing \$BAB_HOME/sbin/startingres.

Important! Starting Ingres at this point results in some processes not being started. These processes generate error messages. These messages are normal at this point and can be safely ignored. If necessary, press the Enter key or Ctrl+C to exit error messages.

22. To recover data, change the user to *root*.

Run *cadbutil* to recover *iidbdb* (*cadbutil -r iidbdb -c*). You are prompted for the directory containing restored data. Provide the directory containing the restored data for the *iidbdb* session (for example, */opt/cadb-restore/s3*).

Note: The *-c* option of *cadbutil* will copy the file from the temporary directory you have used to restore the data sessions to the proper place. If the file system does not contain enough disk space to support the copy operation, you can use the *-m* option to move the data instead.

Run *cadbutil* to recover *cadbase* (*cadbutil -r cadbase -c*). You are prompted for the directory containing the restored *cadbase* session (for example, */disc1/cadb-restore/s4*).

Note: For extended locations for databases, you are prompted for the sessions one by one. Enter the correct directory.

23. The Ingres database is successfully recovered. Run *stopingres*, and run *cstart* to start BrightStor ARCserve Backup. Alternately, you can run *startingres* to start only Ingres.

Appendix A: Using Command Line Utilities

BrightStor ARCserve Backup command line utilities allow direct control, using the command prompt, over all operations that can be performed by a BrightStor ARCserve Backup server. The BrightStor ARCserve Backup command line utilities provide an alternative method of accessing almost all of the operations available from the BrightStor ARCserve Backup Managers. The command line also offers the added benefit of creating batch files that can be automatically executed by other programs.

To use the command line feature, the complete BrightStor ARCserve Backup system must be installed on the server and the \$BAB_HOME variable must be set.

Note: As an alternative to using the command line, any command you can enter can also be submitted using the Generic Job Manager. Using the Generic Job Manager provides these benefits:

- The job appears in the Activity Log.
- You can submit repeating jobs.

For more information on submitting jobs using the Generic Job Manager, see Generic Job Manager in the chapter "Customizing Your Jobs."

Available Command Line Utilities

The following table lists the available BrightStor ARCserve Backup command line utilities, and a description of their basic features:

Command Line Utility	Equivalent Manager	Usage
bab	None	Control BrightStor ARCserve Backup services.
ca_auth	User Profile Manager	Create new users, delete users, change user passwords, and establish authentication equivalencies for a particular user.

Command Line Utility	Equivalent Manager	Usage
ca_backup	Backup Manager	Submit backup jobs to the BrightStor ARCserve Backup queue, and set all associated options, filtering, GFS rotation, and rotation jobs.
ca_dbadmin	Database Manager	Maintain the BrightStor ARCserve Backup database and perform administrative and maintenance tasks.
ca_dbmgr	Database and Media Pool Managers	Maintain the BrightStor ARCserve Backup database, including configuring media pools.
ca_devmgr	Device Manager	Control storage devices, including formatting or erasing media in drives or changers.
ca_generic	Generic Job Manager	Define, schedule, and submit generic jobs.
ca_jobstat	None	Generate reports that provide real time job status information about backup jobs.
ca_log	Report Manager	View and maintain BrightStor ARCserve Backup logs.
ca_mediarep	None	Generate reports that provide information about media used.
ca_merge	Merge Manager	Merge database information from backup media into the database.
ca_mmmomgr	Media Management Admin	Control vault cycle, view and monitor reports.
ca_qmgr	Job Status Manager	Monitor and control jobs submitted to the BrightStor ARCserve Backup job queue.
ca_recoveryrep	None	Generates reports about the media required to recover files for a particular node.
ca_restore	Restore Manager	Submit restore jobs to the BrightStor ARCserve Backup queue and set all associated options.

Command Line Utility	Equivalent Manager	Usage
ca_scan	Scan Manager	Report information about one or more backup sessions on media.
ca_stagingrep	None	Generate reports about data that was backed up using the Disk Staging Option.
ca_summaryrep	None	Generate reports that provide backup job summaries.
ca_utilizationrep	None	Generate reports that provide information about media utilization.
ca_vaultrep	None	Generate reports that provide information about MMO shipping and receiving.
cadiag	None	Start, stop, and check the status of the BrightStor ARCserve Backup diagnostic utility daemons. You can also use this command unpack and run the utility on a machine where diagnostic Utility is not installed.
pfc	None	Run vital checks on the BrightStor ARCserve Backup server and Agents to detect conditions that may cause backup jobs to fail.
tapecopy	None	Copy two different type media, or make session-level tape copies.

Usage, Syntax, and Argument Information

Most features available from the various BrightStor ARCserve Backup GUIs are also available from the command line at the command prompt. The following sections provide details about the arguments, switches, and options available for each of the commands.

You can access the syntax and usage associated with each utility by typing the name of the utility at the command prompt. For more detailed information, including arguments and switches, type the utility name followed by an argument.

bab Command

This command provides control over BrightStor ARCserve Backup services.

bab Syntax

```
bab [-cahost <hostname>] -load [options] -unload [options] -show [options] -showcfg [options] -reconfig [options] -status [options] -removehost hostname
```

The [-cahost <hostname>] switch is an optional switch that identifies the name of the system hosting the operation. If you want to execute the operation on a remote system, this switch must be included in the command. If you want to execute this operation on your local system, this switch is not required and should not be included in the command.

Note: If you include -cahost in the command, you must also specify the *hostname* of the system (local or remote) hosting the operation.

bab Options

The bab command line utility supports the following options.

Option	Description
-load [-useUnicenterTape Management] [<i>procId</i> "all"]	Starts BrightStor ARCserve Backup daemons. Use the switch to load services in Unicenter Tape Management mode. All BrightStor ARCserve Backup servers share a common Unicenter tape database and all tapes are either in Save Set, where they are protected, or Scratch Set, where they can be formatted or overwritten.
-unload [-force] [-quiet] [<i>procId</i> "all"]	Stops BrightStor ARCserve Backup daemons.
-show [-v] [<i>procId</i> "all"]	Show status of BrightStor ARCserve Backup daemons. Use the -v switch to show the status of the daemons in verbose output.
-showcfg [-v] [<i>procId</i> "all"]	View the current status of a specific BrightStor ARCserve Backup service or all BrightStor ARCserve Backup services, as well as how the service is configured.

Option	Description
-reconfig [<i>procId</i> "all"]	Bring down a specific BrightStor ARCserve Backup service or services, reread the configuration file, and bring the service up again with new configuration file settings.
-status [<i>procId</i> "all"]	This option checks the status of all BrightStor ARCserve Backup services.
-removehost [hostname]	This option removes the host from your BrightStor ARCserve Backup environment.

ca_auth Command

Use this command to create new users, delete existing users, change user passwords, and establish authentication equivalencies associated with a particular user, without using the User Profile Manager.

ca_auth allows a complete range of operations on the BrightStor ARCserve Backup Authentication database, such as setting ACLs (Access Control Lists). Authentication services are provided by the cauthd process, which must run on every BrightStor ARCserve Backup server. Authentication service can be managed with the ca_auth command or the User Profile Manager.

ca_auth Syntax

```
ca_auth [-cahost <hostname>] -user [user manipulation options] -equiv  
[equivalence management options]
```

The [-cahost <hostname>] switch is an optional switch that identifies the name of the system hosting the operation. If you want to execute the operation on a remote system, this switch must be included in the command. If you want to execute this operation on your local system, this switch is not required and should not be included in the command.

Note: If you include -cahost in the command, you must also specify the *hostname* of the system (local or remote) hosting the operation.

ca_auth Options

The authentication host option allows the user to select the machine to use for authentication, rather than assume that Authentication Service runs on the same machine.

- `cahost hostname`—execute this command using the authentication service running on machine *hostname*.

Miscellaneous Options

Option	Description
-allusage	Displays a list of all ca_auth commands and their switches.
-usage	Displays a list of basic ca_auth commands.

ca_auth User Manipulation Options

Before you can use BrightStor ARCserve Backup, you must have a BrightStor ARCserve Backup account. The first account, `caroot`, has Administrator privileges, and is created by the program at installation. For security reasons, you should set a password on this account, either during `csetup` or by running the `cauth_setup` script, which is executed, typically, as part of initial setup from `csetup`. The user manipulation options or front-end Profiles manager can be used to add, modify, reassign, or delete a user account.

The `ca_auth` command line utility supports the following user manipulation options.

Command	Description
-user add <i>username passwd</i>	Add a BrightStor ARCserve Backup user with the specified password to the database.
-user delete <i>username</i>	Delete the user from authentication database.
-user chgpaswd <i>username [passwd]</i>	Change the user's password to the new password specified.
-user validate <i>username [passwd]</i>	Check if the <i>username</i> and <i>password</i> combination exists, is valid, and can be used to log into the BrightStor ARCserve Backup domain.
-user getall	List all users known to BrightStor ARCserve Backup Authentication Service.

ca_auth Equivalence Management Options

Equivalence allows you to create an equivalent user to caroot, if you know the password for caroot. After you have designated a user on a given host as equivalent to caroot, you can access the entire authentication database as this user. Remember, for users other than caroot to modify the authentication database, they must be granted access to do this through an ACL entry for the CaAdmin resource.

The ca_auth command line utility supports the following equivalence management options.

Command	Description
-equiv add unixUSER hostname BrightStorUser [Brightstor_username password]	Create an equivalence of <i>user</i> on <i>hostname</i> to causer.
-equiv getequiv	View your equivalence level. See -equiv whoami.
-equiv getequiv <i>user</i> <i>hostname</i>	Display equivalence for <i>user</i> on <i>hostname</i> .
-equiv delete <i>user hostname</i> [causer password]	Delete equivalence for <i>user</i> on <i>hostname</i> . Unless current user is equivalent to caroot, credentials (causer and password) for the Administrator's account are required.
-equiv whoami	View how authentication interprets you and acts when you do not provide credentials on the command line.

Note: Only superuser and object owners can modify ACLs. A user with read rights cannot grant somebody else read rights to an object they do not own.

ca_backup Command

This command is the command line interface to the Backup Manager. All of the features available from the Backup Manager are available from the command line. Use this command to submit backup jobs to the BrightStor ARCserve Backup queue, including setting all associated options, filtering, GFS rotation and rotation jobs.

ca_backup Syntax

```
ca_backup [-cahost <hostname>] [global options] [global filters] - source  
[source arguments] [destination arguments] [schedule arguments] [run job  
arguments]
```

The [-cahost <hostname>] switch is an optional switch that identifies the name of the system hosting the operation. If you want to execute the operation on a remote system, this switch must be included in the command. If you want to execute this operation on your local system, this switch is not required and should not be included in the command.

Note: If you include -cahost in the command, you must also specify the *hostname* of the system (local or remote) hosting the operation.

ca_backup Oracle RMAN Specific Syntax

To back up Oracle RMAN database objects, use the following syntax:

```
ca_backup -source [<hostname>] [node options] -database ORACLERMAN <oracle_sid> [  
[<tablespaces ...> | "ARCHIVE LOG" | "CONTROL FILE" | "PARAMETER FILE"] | [-  
tablespace <tablespace name>| "ARCHIVE LOG" | "CONTROL FILE" | "PARAMETER FILE" |  
[<tablespace file>] ...] ] [dbase options]
```

ca_backup Usage

The ca_backup commands allow you to set global options and filters, and specify the source, destination, the rotation schedule, and the run schedule for your backup operation. To build a backup operation, you must set one category of options at a time, in the order specified in the preceding Syntax section.

ca_backup Miscellaneous Options

The ca_backup command supports the following miscellaneous options:

-list

Displays a list of sources and destinations available for the backup.

Note: To execute this option on the local machine, you do not need to include the -cahost *hostname* switch.

-f <filename>

Used to specify a file name that contains the switches and parameters for the command. This switch overcomes the shell limitation of 1024 character input from command line. You can also use this switch to hide passwords by saving them in a file.

-usage

Displays a list of all basic ca_backup commands.

allusage

Displays a list of all ca_backup commands and their switches.

ca_backup Staging Options

To execute a staging backup job using the ca_backup command line utility, use the following syntax:

```
ca_backup -diskstage [Miscellaneous Options] [Full Backup Policy]
[Differential/Incremental Backup POLICY]
```

Full Backup Policy Syntax

```
ca_backup -diskstage [-fullbackup [-copyDataToDestination [afterjobstarts <weeks>
<days> <hours> <minutes>]|[afterjobends <weeks> <days> <hours>
<minutes>]|[aftersessionends <weeks> <days> <hours> <minutes>]|[at <hh:mm:ss>
[afterjobends]]]
```

Differential/Incremental Backup Policy Syntax

```
ca_backup -diskstage [-incdiffbackup [-copyDataToDestination [afterjobstarts
<weeks> <days> <hours> <minutes>]|[afterjobends <weeks> <days> <hours>
<minutes>]|[aftersessionends <weeks> <days> <hours> <minutes>]|[at <hh:mm:ss>
[afterjobends]]]
```

Miscellaneous Staging Policy Syntax

```
ca_backup -diskstage <GROUP NAME> [-maxstreams <Max # Streams>] [-chunksize
<#chunkSize>] [-purgefailedsessions] [-purgecancelledsessions] [-makeupjobtotape]
[-createdmymakeupjobonhold] [-leaveCatalogsOnDisk]
```

The ca_backup command line utility supports the following staging backup options:

Option	Description
-diskstage <GROUP NAME>	Specifies that the backup job will use staging functionality and the name of the staging device group.
-maxStreams <Max # Streams>	Specifies how many streams BrightStor ARCserve Backup will use as it runs the backup job to the staging device.
-fullbackup	Specifies that the staging backup job consists of full backups.
-incdiffbackup	Specifies that the staging backup job consists of either incremental or differential backups.
-copyDataToDestination [afterjobstarts <weeks> <days> <hours> <minutes>] [afterjobends <weeks> <days> <hours> <minutes>] [aftersessionends <weeks> <days> <hours> <minutes>] [at <hh:mm:ss> [afterjobends]]	Specifies when the copy to final destination operation should commence. For -afterjobstarts, -afterjobends, and -aftersessionends, you need to input the desired length of time.
-DONOTCOPY	Directs the Disk Staging Option to back up data to a staging device, but do not copy the media to a final destination after the retention period elapses.
-purgeData [afterjobstarts <weeks> <days> <hours> <minutes>] [afterjobends <weeks> <days> <hours> <minutes>] [at <hh:mm:ss>]	Specifies when the purge data from disk operation should commence. For -afterjobstarts and -afterjobends, input the desired length of time that must elapse before the purge operation commences.
-ENABLESNAPLOCK	Directs backup job to use SnapLock security on the backup job.
-chunksize <#chunkSize>	During a staging backup, directs BrightStor ARCserve Backup to divide the data into smaller chunks (or sub jobs) to be written to the disk. You must specify the maximum size or amount of data (in KB) contained in each chunk.

Option	Description
-purgefailedsessions	If a session fails during the backup to disk (staging) process, directs BrightStor ARCserve Backup to mark this session for deletion (purged from disk) immediately. This helps to reclaim disk space as soon as possible.
-purgecancelledsessions	If a session is cancelled during the backup to disk (staging) process, directs BrightStor ARCserve Backup to mark this session for deletion (purged from disk) immediately. This helps to reclaim disk space as soon as possible.
-makeupJobToTape	During the backup to disk (staging) process, if an error occurs because the disk is full, directs a makeup job to be created, which when run will directly backup to the final destination media (tape). This increases the chances of a successful backup even though the disk is full.
-createDMJMakeupJobOnHold	During a data migration job (DMJ), if a media or a tape drive error occurs, a makeup job would automatically be created on Hold. As a result, you do not have to create a tapecopy job. After fixing the drive or media error, you would then just need to change the status of the makeup job from Hold to Ready to execute the migration process (disk to tape).

ca_backup Global Options

The ca_backup command line utility supports the following global options.

Option	Description
-scan	Verify the integrity of the backup. Scans backup media and checks the header of each file. If the header is readable, the data is assumed to be reliable.

Option	Description
-compare	Verify the integrity of the backup. Reads blocks of data from backup media and compares the data byte for byte against the source files on the source machine. Note: The Compare Media to Disk option is not supported with agents for databases and applications.
-partialdbupdate	Record only Job and Session information into the BrightStor ARCserve Backup database.
-sessionpassword <i>session password</i>	Apply a password to each session backed up to media. To restore data from one of these sessions, the password must be provided. See ca_restore.
-encryption <i>encryption key</i>	Encrypt files before the backup. To restore encrypted files, the encryption password must be provided. See ca_restore.
-compression	Compress files before backup.
-logfile <i>filename</i> [summary allactivity]	Record activities during the running of the backup job to the specified filename. The user can specify to record all activity or a summary of the activity.
-snmp	Enable SNMP Alert.
-tng	Enable Unicenter NSM Alert.
-email <i>email address</i>	Send a copy of the Activity log to a specified email address.
-printer <i>printer name</i>	Send a copy of the Activity log to a specified printer. The printer must be set up in the configuration file \$BAB_HOME/config/caloggerd.cfg.
-preexec <i>command</i>	Run the specified command before the job starts. The entire path of the command should be included.
-postexec <i>command</i>	Run the specified command after the job finishes. The entire path of the command should be included.

Option	Description
-preexectimeout <i>minutes</i>	The time to wait, in minutes, before the backup job starts, to allow time for the pre-execute command to finish.
-prepostpassword <i>user password</i>	The password of the user submitting this backup job.
-exitcode <i>exit code</i>	Specify the exit code of the pre-execute command. Used with the -skip_delay, -skip_job, and -skip_post switches.
-condition <i>equalto greaterthan lessthan notequalto</i>	Specify the condition for pre execute command exit codes. Used with -exitcode. For greater than and less than conditions, only one exit code should be specified.
-skip_delay	Run the backup job immediately if a specified exit code is received. This option will be activated only if BrightStor ARCserve Backup detects that the exit codes meet the specified condition (Equal To, Greater Than, Less Than, or Not Equal To).
-skip_job	Skip the backup job completely if a specified exit code is received. This option will be activated only if BrightStor ARCserve Backup detects that the exit codes meet the specified condition (Equal To, Greater Than, Less Than, or Not Equal To).
-skip_post	Skip the post-execute command if a specified exit code is received. This option will be activated only if BrightStor ARCserve Backup detects that the exit codes meet the specified condition (Equal To, Greater Than, Less Than, or Not Equal To).
-retry now later off	Specify to retry backing up open files missed during the initial backup.
-retrycount <i>count</i>	Specify the number of retry attempts.
-retryinterval <i>seconds</i>	Specify the interval in seconds between retry attempts.
-virus skip delete rename cure	Virus scan option; specify the action to take when an infected file is encountered.
-deletefiles	Delete files after backup.

Option	Description
-noestimation	Disable file estimation.
-clearconn	Clear user connections before backup. (NetWare only)
-cleararchbit	Clear archive bit after backup. (DOS and Windows only)
-savescript <i>script name</i>	Backup job saved as a script rather than submitting it to Job Queue. The script can be loaded into the Queue later.
-accessmethod usedenynoneifdenywritefails uselockmodeifdenywritefails usedenywrite usedenynone	Specifies the actions BrightStor ARCserve Backup will take if any files are not available during the backup job. These are file sharing options: ■ denynoneifdenywritefails--BrightStor ARCserve Backup attempts to place the file in deny write mode. If this is not possible because file is already open, it will be placed into deny none mode. This is default setting. ■ lockifdenywritefails--BrightStor ARCserve Backup attempts to place the file in deny write mode. If this is not possible because file is already open, the file will be locked completely so that no user can open it or write to it. This option ensures that the most recent version of file is backed up. ■ denywrite--Select this option to prevent another process from writing to the file while BrightStor ARCserve Backup has it open. If another process opens the file before BrightStor ARCserve Backup opens it, BrightStor ARCserve Backup will not backup the file, unless you selected one of the open File Retry options. ■ denynone--Select this option to allow other processes to read or write to the file, regardless of whether BrightStor ARCserve Backup opens it first or opens it after another process already has it open. Although the backed up file may not be the most recent version, this option ensures that the file is up-to-date.

Option	Description
-waitForJobStatus <polling interval (secs)>	The ca_backup command 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 ca_backup checks the jobs' status with the Queue services. The default interval is 60 seconds. This is useful for Unicenter NSM (formerly known as TNG) scheduling. For more information on how to use this to integrate with Unicenter NSM, see Unicenter NSM Integration in the chapter "Introducing BrightStor ARCserve Backup."

ca_backup Global Filters

Apply a filter to the backup job. Filters can be global (applied to the entire job), Node-level (applied to a specific node), or Volume-level (applied to a specific file system). The position of the -filter switch in the ca_backup command determines the filter level applied.

When you apply filters to your job, remember that not all file systems record the creation date or access date. Therefore, these filters may not be available for your job.

Important! *Incorrect use of filters can result in data being missed during backup. Exercise care when specifying or applying filters!*

Option	Description
-filter include exclude file dir pattern	Specify to include or exclude files or directories based on the specified pattern. Note: If you select the include directory pattern filter and do not specify an absolute path, empty directories for all the directories that do not match the user provided criteria will be backed up. To avoid creating these empty directories during restore, disable the global restore option Create Empty Directories when creating your restore job. For more information on this option, see Create Empty Directories Option in the chapter "Restoring Data."

Option	Description
-filter include exclude attribute [hidden readonly system archive]	Specify to include or exclude files with the specified file attribute.
-filter include exclude date modify create access onorbefore onorafter <i>mm/dd/yyyy</i>	Specify to include or exclude files changed, last modified, or accessed before, on, or after the specified date.
-filter include exclude date modify create access between <i>mm/dd/yyyy</i> <i>mm/dd/yyyy</i>	Specify to include or exclude files changed, last modified, or accessed between the specified dates. Note: For UNIX and Linux servers, BrightStor ARCserve Backup automatically interprets the create option as specifying the file changed date.
-filter include exclude date modify create access within count days months years	Specify to include or exclude files changed, last modified, or accessed within the specified number of days, months, or years. Note: For UNIX and Linux servers, BrightStor ARCserve Backup automatically interprets the create option as specifying the file changed date.
-filter include exclude size equalto greaterthan lessthan <size val> Bytes KBytes MBytes GBytes	Specifies the size of the files that you want to include or exclude when you run the filter. Note: For UNIX and Linux servers, BrightStor ARCserve Backup automatically interprets the create option as specifying the file changed date.
-filter include exclude size between <low size val> Bytes KBytes MBytes GBytes <high size val> Bytes KBytes MBytes GBytes	Specifies a range of sizes that you want to include or exclude when you run the filter.

ca_backup Source Arguments

The ca_backup command line utility supports the following source arguments.

Option	Description
-source [<i>hostname</i>]	Specify source machines to back up. The default, if <i>hostname</i> is not provided, is the local machine. This switch can appear multiple times in a ca_backup command, and must appear for each source to be backed up. If used without additional switches, the entire source machine is backed up by default.
-filesystem <i>filesystem name</i> [<i>relative directory</i>]	Specify the file system to back up and, optionally, the directory or directories under the file system. This switch can appear multiple times in a ca_backup command and must appear for each file system to be backed up.
-filelist <i>list of files</i>	Specify individual files to back up. Use with the -filesystem switch.
-fsfile <i>filename</i>	Used to specify the path and name of an external text file that lists the file systems you want to back up. When you use this option, the ca_backup command opens the file and reads its contents to create the backup job. In this external text file, you can define the following information, depending on the level of granularity you want for your backup job: ■ The file systems you want to back up ■ The relative directories of the file systems you want to back up ■ The -filelist option and file names to specify the files to use within the targeted file system ■ The -inputfile option and file name to add files from another external file. To do this, use the following syntax: <pre>[<i>filesystem name</i>] [<i>relative_dir</i>][-filelist <file1><file2>][-inputfile <filename>]</pre>
-raw <i>raw device</i>	Specify raw device to backup.

Option	Description
-database <i>database type</i> <i>database name</i> [<i>tablespaces</i> <i>dbspaces</i>]	Back up supported databases using BrightStor ARCserve Backup database agents. Specify a database type, name, and, optionally, a list of tablespaces or dbspaces to back up. Supported, valid database types are—INFORMIX, ORACLE, ORACLE8, ORACLE_AS66, SAP, SYBASE, and INGRES. Note: ORACLE_AS66 is used to support the Oracle Database agent.
-table <i>tablespace name</i> <i>tablespace files</i>	Allow the user to back up specific tablespace files under a tablespace. Oracle 8 specific.
-NDS <i>NDS tree name</i>	Specify the NetWare NDS tree name.
-NDSserver <i>NDS server name</i>	Specify the NetWare NDS server name.
-NDSaddress <i>NDS server address</i>	Specify the NetWare NDS server address.
-username <i>user name</i>	Specify the username of the source machine to back up. This is the <i>user</i> used to log into the source machine.
-password <i>password</i>	Specify the password for the user to be used to log into the source machine.
-traversesymlink	Traverse symbolic links during the backup, and back up the actual file the link points to, not simply the link itself.
-traversenfs	Traverse mounted NFS file systems during the backup. By default, mounted file systems are skipped during the backup.
-resetaccesstime on off	Specify whether to reset the file access time, changed when BrightStor ARCserve Backup accesses a file to perform a backup.
-noestimation	Disable file estimation prior to backup.
-acrossfs	Traverse across the file system during backup.
-tapeformat tar cpio	Specify the tape format of the backup job. Both tar and cpio tape formats are supported, as well as BrightStor ARCserve Backup's own tape format.
-priority <i>priority level</i>	Assign a backup priority to the Nodes/Volumes in a job. Priority level ranges from 1 (highest priority) to 255 (lowest priority).

Option	Description
-volscan	Verify the integrity of the file system (volume) backup. Scans the backup media and checks the header of each file. If the header is readable, the data is assumed to be reliable.
-volcompare	Verify the integrity of the file system (volume) backup. Reads blocks of data from the backup media and compares the data byte for byte against the source files on the source machine.
-volsessionpw <i>session password</i>	Apply a session password to the session on tape containing the file system (volume) backed up.
-volencryption <i>encryption key</i>	Encrypt files before the backup. To restore the encrypted files in this session, the password must be provided.
-volcompression	Compress files before the backup, for this file system (volume) only.
-volgroomdisable	Disables the volume groom option.
-dbusername <i>database username</i>	Specify the <i>database username</i> to use to log into the database to be backed up.
-dbpassword <i>database password</i>	Specify the password for the database user to use to log into the database to be backed up.
-db2_prunedays <i>number of days</i>	DB2 database backup specific. Specifies the number of days to retain backup information in the DB2 history file.
-db2_multiplexing <i>number of parallel tablespaces</i>	DB2 database backup specific. Specifies the number of tablespaces which can be read in parallel by the DB2 backup utility. The default value is 1.
-db_offline	DB2 database backup specific. Use this option to perform an offline backup of your DB2 database. Note: If the target DB2 database does not have LOGRETAIN or USEREXIT enabled, you must use this option to back up the database.
-inf_level 1 2	IBM Informix database backup specific. Specify the backup level. By default, the level is 0.
-inf_currLogOnly	IBM Informix database backup specific. Back up current log only.

Option	Description
-inf_salvageLogs	IBM Informix database backup specific. Salvage logs.
-inf_scan	Informix database backup specific. Set Informax scan option.
-oracle_sid	Oracle database backup specific. Specify the Oracle SID of the Oracle database to back up.
-oracle_offline	Oracle database backup specific. Back up the Oracle database in off-line mode (no tablespace backup).
-oracle_purgelog	Oracle database backup specific. Purge the log after it has been backed up.
-oracle_timefinder	Oracle database backup specific. Specifies that you want to use the Symmetrix Timefinder technology option for database backups. This option creates a temporary mirror image of the database, which the agent then backs up.
-transactionlog	Sybase database backup specific. Set transaction log option. The database user to use to log into the database to be backed up.
-multiplextape <i>Max # of Streams</i>	Use this to submit a multiplexing job. Max # Streams sets the maximum number of streams that can write to a tape at the same time.
-muxChunkSize	Used with -multiplextape. Sets the performance of restore operations and memory usage. The chunk size value determines the amount of contiguous data written for one session before data from another session is multiplexed. The higher the value, the faster the restore on some drives, but at the cost of memory size during backup.

ca_backup Destination Arguments

The ca_backup command line utility supports the following destination arguments.

Option	Description
-eject	Eject media at backup completion.

Option	Description
-export all duplicate	Export all media or duplicate media after backup job completion.
-group <i>groupname</i>	Specify media group to use for the backup job.
-tape <i>tape name</i>	Specify name of media to use for the backup job.
-mediapool <i>pool name</i>	Specify media pool to use for the backup job.
-multistream max # streams	Specify the maximum number of subjobs to split the backup job into. Splitting a job into multiple subjobs can result in better usage of system resources and increased performance.
-streamlevel node volume	Specify whether to enable multi-streaming at the node level or at the file system level.
-firsttapeopt owritesameblank owritesameblankany owrites ameanyblank	Specify media options for the first media used in the backup job. By default, set to Append to Media. Blank media and Any media are not the same. Any indicates a formatted media with a different media name than that provided in the job. If option owritesameblankany is specified, BrightStor ARCserve Backup first searches for a media with the same name as the job. If one is found and is usable, the media is formatted using the same name, and used for the backup. If not, BrightStor ARCserve Backup searches for a Blank media to use. If no Blank media is available, BrightStor ARCserve Backup searches for Any usable media to format and use for the backup.
-spantapeopt owritesameblank owritesameblankany owritesameanyblank	Specify media options for any span media used in the backup job. By default, set to Overwrite Same or Blank. During spanning of tape, if the default is specified, BrightStor ARCserve Backup first searches for a media with the same name and a higher sequence than the original tape. If a tape is found and is usable, the media is formatted and used as the next tape. If not, BrightStor ARCserve Backup searches for a Blank media to use.
-firsttapetimeout <i>minutes</i>	Specify the time, in minutes, to wait for a usable media to be made available for a backup job. By default, this value is 5 minutes. If a usable media is not made available within this time period, the job times out and fails.

Option	Description
-spantapetimeout <i>minutes</i>	Specify the time, in minutes, to wait for a usable span media to be made available for a backup job. By default, this value is infinite, and the job continues to wait and prompt until a usable media is loaded or the user cancels the job.
-multiplextape [<max # streams>] [-muxChunkSize <MuxChunkSize>]	Specifies a multiplexing (MUX) job. Note: BrightStor ARCserve Backup for Mainframe Linux does not support multiplexing. ■ Max # Streams--the maximum number of streams into which a multiplexing job can be split ■ MuxChunkSize--During a multiplexing job, directs BrightStor ARCserve Backup to divide the data into smaller chunks (or sub jobs) to be written to the disk. You must specify the maximum size or amount of data contained in each chunk.

ca_backup Schedule Arguments

The ca_backup command line utility supports the following schedule arguments.

Option	Description
-custom	Specify the schedule type of the backup job as a custom schedule. By default, this is the schedule type used for backup jobs.
-repeat <i>months days hours minutes</i>	Use with -custom. Specify a repeating interval for a backup job. By default, there is no repeating interval and a job only runs once. Specify a repeating interval so your job runs every <i>X</i> minutes/hours/days/months. The syntax of the command requires a value for each field of months, days, hours, and minutes. For example, To schedule a repeating job every 1 day and 2 hours enter ca_backup -custom -repeat 0 1 2 0.

Option	Description
-exclueday <i>Sun Mon Tue Wed Thu Fri Sat *</i>	Use with -custom to exclude specific days from a repeating backup job.
-rotation	Specify the schedule type of a backup job as a rotation schedule.
-mediapool <i>pool name</i>	Use with -custom or -rotation switches. With -rotation, it specifies the mediapool to be created and associated with the rotation job. With -custom, the mediapool specified must have been previously created.
-method incr diff [full backup job ID]	Use with -custom to create a repeating incremental or differential backup job. Use the <i>Full Backup Job ID</i> variable to specify the full backup job on which you want to base the incremental or differential backups. Note: If no full backup job is specified, then the first backup of this job will be a full backup. All subsequent backups will then be either incremental or differential, depending on the option you specified.
-saveset <i>number of tapes</i>	Use with -rotation. Specify the minimum number of media to keep in the created media pool's save set.
-retention <i>days</i>	Specify the media retention period, in days, for the created media pool.
-jobunit full diff incr off append overwrite <i>media name</i>	Each job unit represents a day in the rotation scheme, from Sunday to Saturday. You can customize each day, although there are certain restrictions, such as not combining differential and incremental backups in the same rotation schedule. The first -jobunit switch is for Sunday, the next for Monday, and so on. You must account for each day of the week, so seven (7) -jobunit switches are required. By default, any days not represented by a -jobunit switch are set to Off day, and no backup takes place on that day. You have the option of not specifying the -jobunit switch to set the rotation schedule to the default 5-day incremental with a full backup on Friday. This is the same schedule seen from the front-end backup manager.

Option	Description
-reschedule <i>hh:mm</i>	Use with the -custom, -rotation, or -gfsrotation switch. Specify the time to run a Makeup job if there are missed targets in the original backup.
-exception full diff incr off append overwrite <i>mm/dd/yyyy</i>	Specify an exception condition from the regular rotation schedule. This feature is useful in the case of a holiday or other event when a different behavior for the backup job is needed on that date.
-gfsrotation	Specify the schedule type of the backup job as a GFS (Grandfather, Father, Son) rotation schedule.
-mpoolprefix <i>mediapool prefix</i>	Use with -gfsrotation as a prefix for naming three mediapools (Daily, Weekly, and Monthly) to be created and associated with this GFS rotation job. For example, if the prefix is GFSJOB1, the 3 pools created are: GFSJOB1_DLY, GFSJOB1_WLY, GFSJOB1_MLY.
-preservedaily <i>number of tapes</i>	Use with -gfsrotation. Specify the minimum number of media to preserve in the daily media pool save set.
-preserveweekly <i>number of tapes</i>	Use with -gfsrotation. Specify the minimum number of media to preserve in the weekly media pool save set.
-preservemonthly <i>number of tapes</i>	Use with -gfsrotation. Specify the minimum number of media to preserve in the monthly media pool save set.
-jobunit full diff incr off	Same as the description for rotation jobs, except that, for GFS rotation, the arguments are limited to specifying the type of backup to occur on the selected day.
-exception full diff incr off <i>mm/dd/yyyy</i>	Same as the description for rotation jobs, except that, for GFS rotation, the arguments are limited to specifying the type of backup to occur on the exception date.

ca_backup Run Job Arguments

The ca_backup command line utility supports the following run job arguments.

Option	Description
-at <i>hh:mm</i>	Specify the execution time of the backup job.
-on <i>mm/dd/yyyy</i>	Specify the execution date of the backup job.
-hold	Submit the backup job on hold.
-runjobnow	Submit and execute the backup job immediately.
-description <i>description string</i>	Add comments to the job. You must use double quotes "" to enclose the string and handle blank spaces.

ca_dbadmin Command

The Database Administration and Maintenance command is the command line interface for the database used by BrightStor ARCserve Backup, and residing on the Advantage™ Ingres® relational database management system (RDBMS). The ca_dbadmin command allows you to maintain the database and perform administrative and maintenance tasks. You can perform verifications on the Ingres database, and generate reports that describe maintenance tasks.

Using the ca_dbadmin command, you can:

- Generate reports describing maintenance requirements
- Perform maintenance on specified tables and indexes
- Update Ingres database statistics
- Export the database
- Manage the overall size of the database

ca_dbadmin Syntax

```
ca_dbadmin -report -type -maintain
```

ca_dbadmin Options

The ca_dbadmin command line utility supports the following options.

Option	Description
-allusage	Displays a list of all ca_dbadmin commands and their switches.
-usage	Displays a list of basic ca_dbadmin commands.
-report	Generate a report.
-type <i>maintenance</i> <i>env_info</i>	Create a maintenance report (<i>maintenance</i>). Creates a description of all Ingres table information in the system, and creates the database information output (<i>env_info</i>).
-filename <i>filename</i>	Specify the file name. Defaults to <i>maintenance_report.txt.ddmmyy</i> or to <i>env_info.txt</i> . All reports are generated in <i>\$BAB_HOME/dbase/maintenance/reports</i> .
-entity details filename filepath session job tape marked staging	Performs maintenance on the specified or marked entities.
-overwrite	Overwrite the existing report, and recreates a new report using the same file name. If not specified, the system prompts you if to overwrite an existing report.
-delete <i>nbdays</i>	Specify an alternate retention period to the value in the configuration file. When ca_dbadmin is used to generate reports, it automatically purges all older reports based on the retention period specified in the configuration file.
-severity warning critical	Mark the table for requiring maintenance, if a warning or critical condition is met. Used with the -mark option.
-mark	Mark the table for requiring maintenance, if a warning or critical condition is met.
-maintain	Perform maintenance tasks.

Option	Description
-entity details filename filepath session job tape mmo marked	Perform maintenance on the marked entities.
-structure keep btree heap	Support details entity.
-index_only	Perform maintenance tasks only on external indexes.
-index clean rebuild	Shrinks btree indexes for the entity specified (clean). Recreate all indexes for the entity specified (rebuild).
-missing_index	Rebuild the missing or dropped indexes. Supports the details entity.
-mark	Mark the selected entity for maintenance only. Do not run the maintenance now.
-gen_script	Create the SQL script file, and allows you to do the maintenance later. Default location for script file: \$BAB_HOME/dbase/maintenance. Script file name: maintain_script.sql. For example: sqlingres <maintain_script.sql>
-use_all_locations	Force use of all available database locations for that table. Default is to use the same location(s) as defined for the entity.
-statistics no yes only	Perform updates for the maintained entity statistics (yes). Performs updates only for the statistics, but does not perform maintenance on the table. Performs no updates for the maintained entity statistics.
-force	Force the maintenance program to run, even if it detects that cadbd is running. The maintenance program also attempts to pause the MergeCat and dbclean programs, if, running prior to starting the maintenance program. Default condition for the Maintenance program is not to run if cadbd is running.
-type statistics	Collect statistics on all known tables (optimizedb).

Option	Description
-type system	Run sysmod.

Compress the Advantage Ingres Database using the ca_dbadmin Command

When you prune or purge data from the Ingres database, the overall size of the database itself may not decrease.

You should check the size of the Ingres database periodically (once every seven days is recommended) to verify that the overall size does not exceed the recommended levels. If the recommended levels are exceeded, you can use the ca_dbadmin command line utility to compress and reduce the size of the Ingres database.

To compress the size of the Ingres database, perform the following steps:

1. Run the ca_dbadmin command to generate a Maintenance Report as follows:

```
ca_dbadmin -report -type maintenance
```

Note: The default location for the generated Maintenance Report is \$BAB_HOME/dbase/maintenance/reports. The report is labeled maintenance_report.txt.MMDDYY. for example, maintenance_report.txt.010505.

2. Review the Maintenance Report and check the Diskspace Usage Report category and the Database/Diskspace Usage Report category.

Verify that the actual Diskspace Usage and Database/Diskspace Ratio values do not exceed the recommended values in the report.

3. If either of the actual values exceed the corresponding recommended values, you should run the ca_dbadmin command to perform maintenance and compress the databases as follows:

```
ca_dbadmin -maintenance -entity details
```

4. After the database has been compressed, generate another Maintenance Report and verify that the actual Diskspace Usage and Database/Diskspace Ratio values have been reduced to within the recommended values.

ca_dbadmin Hidden Options

The ca_dbadmin command line utility supports the following hidden options.

Option	Description
-version	Returns the Ingres version information.
-exportdb	Exports the complete Ingres database system.
-compress <yes no>	Compresses the tar file (default). Uses gzip to compress the file with the file extension .gz for Linux platform. Uses compress with the file extension .Z for all other UNIX platforms.
-exportdir <i>dir_name</i>	Exports to the specified directory name. Default export directory is \$BAB_HOME/dbase/exports.

ca_dbmgr Command

This command is the command line interface to the Database Manager and the Media Pool Manager. It allows you to maintain the BrightStor ARCserve Backup database, including configuring media pools. Using this command, you can query database information and set database options. This powerful utility allows other programs to interact easily with backup events. All of the features available from the Database Manager and Media Pool Manager are available from the command line.

ca_dbmgr Syntax

```
ca_dbmgr[-cahost <hostname>]-show[display options]-mediapool[media pool  
management options][database management options]
```

The [-cahost <hostname>] switch is an optional switch that identifies the name of the system hosting the operation. If you want to execute the operation on a remote system, this switch must be included in the command. If you want to execute this operation on your local system, this switch is not required and should not be included in the command.

Note: If you include -cahost in the command, you must also specify the *hostname* of the system (local or remote) hosting the operation.

ca_dbmgr Options

The options and switches available for the ca_dbmgr command allow you to display a variety of information, manage and control media pools, and manage the database.

Miscellaneous Options

Option	Description
-allusage	Displays a list of all ca_dbmgr commands and their switches.
-usage	Displays a list of basic ca_dbmgr commands.

ca_dbmgr Display Options

The ca_dbmgr command line utility supports the following display options.

Option	Description
-show prune purge	Display pruning or purging status and settings.
-show summary	Show database size and limits. Display status of pruning and purging, space information, and database information.
-show jobs [-completed -cancelled -failed -incomplete] [-last <i>no_of</i> days weeks months]	Display information for all jobs of a particular type. For example, enter [-completed] to view all completed jobs. Alternatively, you can enter a time limit, for example: [-last 2 weeks].
-show jobsessions <i>jobID</i>	Show all the sessions contained in the specified job ID.
-show tapes tapesessions <i>tapeID[:seqNo]</i>	Show information about the specified tape or tape sessions.
-show pools poolmedia <i>poolName</i>	Show information about the specified pools or pool media.
-show scratchmedia	Shows information about all medias current in the scratch set of a media pool. The information includes the tape name, serial number, tape ID, sequence number, format date, expiration date, and the media pool it belongs to.

Option	Description
-show savemedia	Shows information about all medias current in the save set of a media pool. The information includes the tape name, serial number, tape ID, sequence number, format date, expiration date, and the media pool it belongs to.
-show clients	Display information about clients.

ca_dbmgr Media Pool Management Options

The ca_dbmgr command line utility supports the following media pool management options.

Option	Description
-mediapool add <i>poolName</i> <i>saveTapes</i> [-b <i>baseSerial</i>] [-i <i>serialIncr</i>] [-m <i>maxSerial</i>] [-retention <i>retentionDays</i>]	Create a new media pool. Add media to the specified media pool. Switches allow you to specify information about the serial number and the retention period.
-mediapool modify <i>poolName</i> [-save <i>saveTapes</i>] [-retention <i>retentionDays</i>]	Modify information for the media pool.
-mediapool delete [-f] <i>poolName</i> [<i>tapeID[:seqNo]</i>] 	Delete a specified tape from a media pool.
-mediapool move <i>tapeID[:seqNo]</i> <i>fromPoolName toPoolName</i> DEFAULT SCRATCH SAVE	Move tapes from media pool to media pool. You can also move tapes from the Scratch Set to the Save Set, or back. This command also has the same function as Assign Media in the UI. The DEFAULT media pool refers to tapes that are not assigned to a media pool.
-mediapool applyRetention	Scan all available media pools for media that have expired their save set retention periods and moves them to their scratch sets. If you do not use this switch, this process occurs only when you perform a backup job to a media pool, and it occurs for that particular media pool only.

ca_dbmgr Database Management Options

The ca_dbmgr command line utility supports the following database management options.

Option	Description
-tape delete <i>tapeID[:seqNo]</i>	Delete a specified tape from the database.
-prune on off set <i>count</i> <i>days</i> <i>months</i> <i>years hh:mm</i>	Set database pruning to on or off, set pruning period and scheduled time for pruning job. Pruning removes detail records older than the specified age, but retains job and session records.
-purge on off set <i>count days</i> <i>months</i> <i>years hh:mm</i>	Set database purging to on or off, set purging period and scheduled time for purging job. Purging removes detail records older than the specified age, as well as job and session records.
-client add <i>hostname</i> [-ip <i>xxx.xxx.xxx.xxx</i>] [-os <i>type</i>]	Add a client to the database so it can be backed up. You will need to enter the IP address and Operating System type.
-client delete <i>hostname</i>	Delete a client from the database.
-setdbsize <size [G M]>	Set the database size and limit in either GB or MB.
-help	Get help.

Note: BrightStor ARCserve Backup uses the Advantage Ingres embedded database Version 2.6, Service Pack 2 with UNIX and Mainframe Linux platforms, and Version 3.0.2 with Linux platforms. For more information about how to maintain the Advantage Ingres embedded database, see the Advantage Ingres Embedded Edition Administrator's Guide included with your documentation.

ca_devmgr Command

The `ca_devmgr` command, the command line device management program, allows you to perform various device management commands in BrightStor ARCserve Backup without interfacing with the Device Manager. These device management commands are used to obtain information, or manipulate the tape or library device. This command allows you to control storage devices, and to format and erase media in drives or changers. All of the features available from the Device Manager are available from the command line.

While these operations can be performed using the Device Manager, the commands are useful if you do not have access to a browser. BrightStor ARCserve Backup must be running, and your database must be authorized. If the database is not authorized, run `cauth_setup` to authorize it.

ca_devmgr Syntax

```
ca_devmgr [-cahost <hostname> ] [-command parameters ]
```

The `[-cahost <hostname>]` switch is an optional switch that identifies the name of the system hosting the operation. If you want to execute the operation on a remote system, this switch must be included in the command. If you want to execute this operation on your local system, this switch is not required and should not be included in the command.

Note: If you include `-cahost` in the command, you must also specify the *hostname* of the system (local or remote) hosting the operation.

ca_devmgr Usage

The following are commands for `ca_devmgr`. Some commands detect tape drives only, while some detect tape libraries only. There are also some commands that pertain to both a tape drive and a tape library.

The `ca_devmgr` commands are:

- **Common Commands**—These commands can be used with either tape drives or tape libraries attached to your system. You can use these commands to obtain information about adapters, devices, groups, or media, and specify that one media be copied to another (the source and destination device must be the same type of device).
- **Tape Drive Commands**—These commands apply only to tape drives. You can use these commands to format, erase, retension, or eject tapes, enable compression, or lock or unlock tape drives.

- **Tape Library Commands**—These commands apply only to tape libraries. You can use these commands to load and unload, mount and dismount, import and export, bring a changer online or take a changer offline, enable, or disable drives, and get or set the cleaning time.
- **Miscellaneous Commands**—These commands do not interact with tape drives or tape libraries. You can use these commands to interact with BrightStor ARCserve Backup, such as setdebug, or usage.
- **Staging Commands**—These commands let you query and purge sessions from file systems devices that you are using for staging backup operations.

ca_devmgr Common Commands

The ca_devmgr command line utility supports the following common commands.

Command	Description
-v <-adapterinfo> <-deviceinfo ...> <-groupinfo > <-mediainfo ...>	More of a verbose command, which can only be used with adapterinfo, deviceinfo, groupinfo, and mediainfo. The only difference in using the -v switch is that it prints additional information on the 4 commands.
-adapterinfo	Cycle through all the SCSI adapters attached to the system and print out the adapter name, adapter number, SCSI ID, vendor ID, product ID, and firmware of any tape drive or tape library connected to the system.
-deviceinfo [<i>adapter#</i> <i>scsi ID</i>]	Print the type of device, SCSI ID, vendor ID, product ID, firmware, status, and device sharing (tape libraries only) information for any tape drive or tape library. The adapter number and SCSI ID options are required.
-groupinfo	Print the adapter number, SCSI ID, vendor ID, product ID, firmware, and status (tape libraries only) information on all groups that have been configured in BrightStor ARCserve Backup.

Command	Description
<code>-mediainfo [adapter# scsi ID scsi LUN][ieinfo]</code>	Print the tape name, tape ID, sequence number, serial number, and expiration date for any tape device. For tape libraries, the same information displays, including slot number, and if the tape is loaded and write protected. For tape libraries, each slot is displayed. Adding the optional ieinfo switch displays information on the library's import/export slot, including whether or not it is full and, if full, the barcode number of the tape in the slot. The adapter number, SCSI ID, and SCSI LUN are required.
<code>-regenerate <adapter #> <scsi ID> <scsi lun> <tape name> [<tape ID> <mm/dd/yyyy>]</code>	This option is for use with file system devices if you accidentally delete the tape header. It lets you regenerate or rebuild a tape header with a specified tape name (<tape name>) for a file system device. After you generate a new tape header, you can then merge all of the sessions on the file system device into the BrightStor ARCserve Backup database, which enables the capability for point-and-select restores. If you know the original tape name (<tapename>) and its tape ID (<tapeID>), you can reuse them so that the session records in the BrightStor ARCserve Backup database can reconnect to the volume without having to merge the session records again. (You can check the original tape record in the BrightStor ARCserve Backup database for the tape name and tape ID). You can use the <mm/dd/yyyy> parameter to specify a different tape expiration date from the default date.

Command	Description
<code>-copy [source adapter # source scsi ID source scsi LUN dest adapter # dest scsi ID dest scsi LUN]</code>	Copy the entire contents of one tape to another. Can be performed between two standalone tape drives, two tape drives in a tape library, or one drive in a tape library and a standalone tape drive. Both the source and destination tape device must be the same type of device. The source media cannot be blank. The destination media must be blank. The source adapter number, source SCSI ID, source SCSI LUN, destination adapter number, destination SCSI ID, and the destination SCSI LUN are required.

ca_devmgr Tape Drive Commands

The ca_devmgr command line utility supports the following tape drive commands.

Command	Description
<code>-format [<adapter #><scsi ID><scsi LUN>< tape name>] [<mm/dd/yyyy><serial no.>]</code>	Format a tape in a tape drive. The adapter number, SCSI ID, SCSI LUN, and the new name of the tape are required. The date and serial number are optional.
<code>-erase [<adapter #><scsi ID><scsi LUN>] [q qz l alex]</code>	Erase a tape in a tape drive. The adapter number, SCSI ID, and SCSI LUN are required. The options q (quick erase, which destroys media label), qz (quick erase plus, which destroys media label and serial number), and l (long erase, which destroys all data on the media), and alex (quick erase Alexandria tape) are optional. The default is quick erase. Important! <i>Using long erase erases the entire media from the beginning of the tape to end of the tape and may take a long time.</i>
<code>-retension [<adapter #><scsi ID><scsi LUN>]</code>	Retension a tape in a tape drive. Both adapter number and SCSI ID are required.

Command	Description
-compression [<i><adapter #></i> <i><scsi ID></i> <i><scsi LUN></i>] on off	Enable or disable compression on a tape drive. The tape device must support compression, and there must be a blank tape in the drive, for this command to work. The adapter number, SCSI ID, SCSI LUN, and an on or off flag are required.
-eject [<i><adapter #></i> <i><scsi ID></i> <i><scsi LUN></i>]	Eject a tape from the tape drive. The tape drive must support the eject command for this command to work. The adapter number, SCSI ID, and SCSI LUN options are required.
-reserve [<i><drive adapter #></i> <i><drive scsi ID></i> <i><drive scsi LUN></i>]	BrightStor ARCserve Backup takes control of the specified tape device, and locks it, not allowing any other programs to use that tape device. This command works only if device sharing has been enabled. The tape drive adapter number, drive SCSI ID, and drive SCSI LUN options are required.
-unreserve [<i><drive adapter #></i> <i><drive scsi ID></i> <i><drive scsi LUN></i>]	BrightStor ARCserve Backup relinquishes control of the specified tape device, so other programs can use the device. This command works only if device sharing is enabled. The tape drive adapter number, drive SCSI ID, and drive SCSI LUN options are required.

ca_devmgr Tape Library Commands

The ca_devmgr command line utility supports the following tape library commands.

Command	Description
<code>-chformat [<i>changer adapter # changer scsi ID scsi LUN groupname</i>] [<i>groupname</i>] SLOT <i>slot # tape name</i> [SLOT <i>slot # tape name</i> [<i>mm/dd/yyyy</i> def [<i>serial no.</i>]] ...]</code>	Format a single tape or multiple tapes in a tape library. There are two ways to use this command. In the first method, the changer adapter number, changer SCSI ID/LUN, and the group name are required. The word SLOT is required before each slot number specified. The slot number of the tape to be formatted and the tape name are required. The expiration date is optional. The def switch stands for the default expiration date. You cannot use both the expiration date and the def switch at the same time. The serial number is optional. Specify as many slot numbers as required. In the second method, the group name, slot number, and tape name are required. The expiration date and the serial number are optional. The word SLOT must be used before each slot number you specify.
<code>-cherase [<i>adapter # changer scsi ID scsi LUN groupname</i>] [<i>groupname</i>] slot # [<i>q qz l</i>] [, slot # [<i>q qz l</i>],...]</code>	Erase a single tape or multiple tapes in a tape library. There are two ways to use this command. In the first method, the adapter number, changer SCSI ID/LUN, the group name, and the slot are required. The q, qz, and l are optional. The q, qz, and l options perform quick erase, quick erase plus, and long erase procedures. Quick erase (q) is the default. In the second method, erase a tape using only the group name and the slot number. Quick erase is the default setting. You can specify multiple slots to be erased.
<code>-cherase [<i>adapter # changer scsi ID groupname</i>] [<i>groupname</i>] slot_range all [<i>q qz l</i>] [, slot_range all [<i>q qz l</i>] ,...]</code>	Erase all slots that have media, and it has no effect to slots without media. The command line will print an message indicating some slots are empty.

Command	Description
-load [<i>changer adapter #</i> <i>changer scsi ID changer scsi</i> <i>LUN drive adapter # drive</i> <i>scsi ID scsi LUN groupname</i>] [<i>drive adapter # drive scsi</i> <i>ID drive scsi LUN</i>] <i>[groupname] slot #</i>	Load a tape from a specified slot into a tape drive. There are two ways to use this command. In the first method, the changer adapter number, changer SCSI ID, changer SCSI LUN, drive adapter number, drive SCSI ID, drive SCSI LUN, group name, and slot number are required. The second method provides an easier way to perform a load command. Only the group name and slot number are required.
-unload [<i>changer adapter #</i> <i>changer scsi ID changer scsi</i> <i>LUN groupname</i>] <i>[groupname] slot #</i>	Unload a tape from a tape drive and put it back in the specified slot. There are two ways to use this command. In the first method, the changer adapter number, changer SCSI ID, changer SCSI LUN, group name, and the slot number are required. In the second method, only the name of the group and the slot number to put the tape in are required.
-mount [<i>changer adapter #</i> <i>changer scsi ID scsi LUN</i> <i>groupname</i>] [<i>groupname</i>] <i>beg. slot end slot [q l]</i>	Inventory your entire tape library. If your tape library does not have a bar code reader, BrightStor ARCserve Backup puts all the tapes into the tape drive and reads them. If your tape library does have a bar code reader, specify whether to obtain the information from the database or have BrightStor ARCserve Backup read in all the tapes. There are two ways to use this command. Using the first method, the changer adapter number, changer SCSI ID/LUN, group name, a beginning slot, and an end slot are required. The beginning slot and end slot can be any valid slot number in the group you are mounting. The end slot cannot be smaller than the beginning slot. The q (Quick Mount) or l (Long Mount) are optional. Quick Mount, the default for tape libraries with bar code readers, obtains information from the BrightStor ARCserve Backup database. Long Mount forces BrightStor ARCserve Backup to put each tape into the tape drive and read the information on it. If your tape library has no bar code reader, the q switch is disabled. The second method requires only the group name, the beginning slot, and end slot. The q and l options are not required.

Command	Description
<code>-dismount [<i>changer adapter # changer scsi ID scsi LUN groupname</i>] [<i>groupname</i>] <i>beg. slot end slot</i> [export]</code>	Dismount the slots in the specified group in the range provided in the options. The slots are then renamed to Dismounted Slot. Unless BrightStor ARCserve Backup is restarted or a mount command is issued to the dismounted group, no other command can be issued to the dismounted slots. There are two ways to use this command. In the first method, the changer adapter number, changer SCSI/LUN ID, group name, beginning slot, and the end slot are required. The export option exports the tapes from the specified slot range after the dismount is completed. The second method requires only the group name, beginning slot, and end slot. Export is optional.
<code>-import [<i>changer adapter # changer scsi ID scsi LUN groupname</i>] [<i>groupname</i>] <i>slot # [slot # ...]</i></code>	Move a tape from the import/export slot and place it into its destination slot. If your tape library has a bar code reader enabled in BrightStor ARCserve Backup, the information is taken from the database and the drive does not read the tape. If your tape library does not have a bar code reader, it then places the tape into the drive to be read. There are two ways to use this command. In the first method, the changer adapter number, changer SCSI ID/LUN, group name, and slot number are required. The only optional switch is the extra <i>slot #</i> for tape libraries with more than one import/export slot. You can import multiple tapes at a time unless your tape library has only one import/export slot. The second method requires only the group name and slot number. As in the first method, the extra <i>slot #</i> allows you to import multiple tapes at a time.

Command	Description
-export [<i>changer adapter # changer scsi ID scsi LUN groupname</i>] [<i>groupname</i>] <i>slot #</i> [<i>slot #...</i>]	Take a tape from a slot in the tape library and put it in the import/export slot. There are two ways to use this command. In the first method, the changer adapter number, changer SCSI ID/LUN, group name, and the slot number are required. The <i>slot # ...</i> is for multiple exporting of tapes, which can be done only if your tape library has more than one import/export slot. The second method requires only the group name and slot number. As in the first method, the second <i>slot #</i> allows you to export multiple tapes.
-onlinechanger [<i>changer adapter# changer scsi ID scsi LUN</i>]	Set the tape library online, making it available for use. The changer adapter number and changer SCSI ID/LUN are required.
-offlinechanger [<i>changer adapter # changer scsi ID scsi LUN</i>]	Mark the specified tape library offline. The changer adapter number and changer SCSI ID/LUN are required.
-getclntime	Get the clean time, in hours, and display the information.
-setclntime [<i>hours</i>]	Set the auto clean time, in hours. Every number of hours that you specify, BrightStor ARCserve Backup cleans the tape drive in the tape library. The hours switch is required.
-clean [<i>changer adapter # changer scsi ID scsi LUN drive adapter # drive scsi ID scsi LUN groupname</i>] [<i>groupname</i>] <i>slot #</i>	Put a cleaning tape, if one is installed in the tape library, into the specified drive and clean the tape drive. There are two ways to use this command. In the first, the changer adapter number, changer SCSI ID/LUN, tape drive adapter number, tape drive SCSI ID/LUN, the group name, and the slot number are required. The second method requires only the group name and slot number.
-enabledrv [<i>drive adapter # drive scsi ID scsi LUN groupname</i>]	Enable a tape device inside a tape library. The tape drive adapter number, tape drive SCSI ID/LUN, and the group name are required. This command cannot be used on a standalone tape device.
-disabledrv [<i>drive adapter # drive scsi ID scsi LUN groupname</i>]	Disable a tape device inside a tape library. The tape drive adapter number, tape drive SCSI ID/LUN, and the group name are required. This command cannot be used on a standalone tape device.

ca_devmgr Miscellaneous Commands

The ca_devmgr command line utility supports the following miscellaneous commands.

Command	Description
-setdebug <i>debug level</i>	Set the debug level of BrightStor ARCserve Backup. This parameter must be set to a value of 0 or a positive number, otherwise BrightStor ARCserve Backup will not put any debugging information into the file. If you specify a negative value, the command will fail and BrightStor ARCserve Backup presents an error message. For more information about debugging values, see the next section.
-usage	Display a list of all ca_devmgr commands and switches.

Debugging Values

BrightStor ARCserve Backup supports the following debugging values:

0—Disable debugging (camediad will no longer write debug messages to the camediad.dbg file.

1—SCSI information (without the detailed SCSI READ/WRITE commands when doing the job).

2—Timeouts, Semaphore information.

4—Configuration information.

8—Callback information.

16—Media information.

32—Read/Write information.

64—Changer Command information.

128—Shared changer (ASCLS) information

256—Detailed SCSI READ/WRITE commands when doing the job.

512—Multiplexing information.

Important! *The sum of individual debug levels (1 through 512) can be combined to generate a new overall combined debug level. For example, a debug level of 3 is the combined sum of debug level 1 (SCSI information) and debug level 2 (Timeouts, Semaphore information). As a result, if all the individual debug levels are combined (1 + 2 + 4 + 8 + 16 + 32 + 64 + 128 + 256 + 512), the overall result will be the maximum debug level of 1023.*

1023—Highest level of debug information (includes all debug levels). 1023 is the sum of debug level 1 through debug level 512).

Note: For information about the configuring the DEBUG level using the camediad.cfg file, see the topic *Device Configuration File - camediad.cfg* in the online help.

ca_devmgr Environment Variables

The ca_devmgr command line utility supports the following environment variable.

Variable	Description
BAB_HOME	Location of the BrightStor ARCserve Backup product.

Staging Command Line Query Tool

If a copy to media operation fails, BrightStor ARCserve Backup provides you with the capability to query the staging device designated to perform the copy to media operations. The query tool can help you identify problems and take corrective actions.

The query tool, `ca_devmgr`, provides you with a description of the following staging information:

- All copied and uncopied sessions
- All sessions with SnapLock security
- Retention period for each session
- All sessions that can be purged

ca_devmgr Syntax

To run the query tool, use the following syntax:

```
ca_devmgr [-cahost <hostname>]
           -query <adapter #> <scsi id> <lun>
           -all | -copied | -uncopied | -snaplocked | -purgable
           -sessions <list of space separated sessions> | all
```

The staging query tool supports the following command line options:

Option	Description
-copied	Displays a list of all copied sessions.
-uncopied	Displays a list of all sessions that were not copied.
-snaplocked	Displays a list of all sessions with SnapLock security enabled and the retention period for each session.
-purgable	Displays a list of all sessions that can be purged.

Option	Description
-sessions <list of space separated sessions> all	Specifies the session number, the group of session numbers, or all session numbers that you want to query. ■ To specify a session number or group of session numbers, you must provide the session number or a list of space separated session numbers using the following syntax for the -sessions argument: -sessions 1 2 3 ■ To specify all sessions, use the following syntax for the <session range> argument: -sessions all

Staging Command Line Purge Tool

If the staging device becomes full or exceeds its storage threshold, backup jobs will fail. To remedy this situation, BrightStor ARCserve Backup provides you with capability to purge sessions forcefully from a file system device using the ca_devmgr command line utility.

Note: Use the query tool to analyze the sessions on the disk device before running the purge tool.

Using the purge tool, you can perform the following tasks:

- Purge copied sessions
- Purge uncopied sessions

ca_devmgr Syntax:

To run the purge tool, use the following syntax:

```
ca_devmgr [-cahost <hostname>]
          [-force]
          -purge <adapter #> <scsi id> <lun>
          -sessions <list of space separated sessions> | <session range>
```

For example:

```
ca_devmgr [-cahost <hostname>]
          [-force]
          -purge <adapter #> <scsi id> <lun>
          -sessions <session001 session002 session003 ...> | -sessions <start session
number>-<end session number> | -sessions all
```

The purge tool supports the following command line options:

Option	Description
-purge	Purges the specified copied sessions (only).
-force	Purges the specified uncopied sessions (only).
-sessions <list of space separated sessions> <session range>	Specifies the session number, the group of session numbers, the range of session numbers, or all session numbers that you want to purge. ■ To specify a session number or group of session numbers, you must provide the session number or a list of space separated session numbers using the following syntax for the -sessions argument: -sessions 1 2 3 ■ To specify a range of sessions, use the following syntax for the <session range> argument: -sessions 1-99 ■ To specify all sessions, use the following syntax for the <session range> argument: -sessions all

ca_generic Command

The ca_generic command is the command line interface to the Generic Job Manager. It is used to create and submit generic jobs to the Job Queue. All features available from the Generic Job Manager are also available from the this command line utility.

ca_generic Syntax

```
ca_generic [-cahost <hostname>][-command <command>][global options][schedule arguments][run job arguments]
```

The [-cahost <hostname>] switch is an optional switch that identifies the name of the system hosting the operation. If you want to execute the operation on a remote system, this switch must be included in the command. If you want to execute this operation on your local system, this switch is not required and should not be included in the command.

Note: If you include -cahost in the command, you must also specify the *hostname* of the system (local or remote) hosting the operation.

ca_generic Global Options

The ca_generic command line utility supports the following global options.

Option	Description
-command <i>command</i>	Specify the generic command to be executed.
-preexec <i>command</i>	Run the specified command before the job starts. The entire path of the command should be included.
-preexec timeout <i>minutes</i>	The time to wait, in minutes, before the backup job starts, to allow time for the pre-execute command to finish.
-postexec <i>command</i>	Run the specified command after the job finishes. Include the entire path of the command.
-prepostpassword <i>user password</i>	The password of the user who is submitting this generic job. Note: There is no -prepostusername option. The username of the user who is submitting this job is used as prepostusername.
-exitcode <i>exit code</i>	Specify the exit code of the pre-execute command. Used with the -skip_delay, -skip_job, and -skip_post switches.

Option	Description
-condition <i>equalto greaterthan lessthan notequalto</i>	Specify the condition for pre execute command exit codes. Used with -exitcode. For greater than and less than conditions, only one exit code should be specified. Note: To enable multiple exit codes, you must use the equalto or notequalto conditions.
-skip_delay	Run the generic job immediately if the specified exit code is received. This option will be activated only if BrightStor ARCserve Backup detects that the exit codes meet the specified condition (Equal To, Greater Than, Less Than, or Not Equal To).
-skip_job	Skip the generic job completely if the specified exit code is received. This option will be activated only if BrightStor ARCserve Backup detects that the exit codes meet the specified condition (Equal To, Greater Than, Less Than, or Not Equal To).
-skip_post	Skip the post-execute command if the specified exit code is received. This option will be activated only if BrightStor ARCserve Backup detects that the exit codes meet the specified condition (Equal To, Greater Than, Less Than, or Not Equal To).
-savescript <i>script name</i>	Save the generic job as a script rather than submit it to the Job Queue. The script can be loaded into the Job Queue at a later time. See ca_qmgr.
-username <i>user name</i>	Specify the user name of the machine where the generic job is to be run. This is the user name that was used to log in to the preferred machine.
-password <i>password</i>	Specify the password to log in to the preferred machine.

ca_generic Scheduling Options

The ca_generic command line utility supports the following scheduling options.

Option	Description
-repeat <i>months days hours minutes</i>	Specify a repeating interval for a generic job. By default, there is no repeating interval and a job only runs once. Specify a repeating interval so your job runs every <i>X</i> minutes/hours/days/months. The syntax of the command requires a value for each field of months, days, hours, and minutes. To schedule a repeating job every 1 day and 2 hours enter <code>ca_backup -custom -repeat 0 1 2 0</code> .
-exclueday < Sun Mon Tue Wed Thu Fri Sat >	Exclude specific days from a repeating generic job. For example, if a generic job has been schedule to repeat ever day, but you want to skip Tuesday and Thursday, issue the following: <code>ca_generic -repeat 0 1 0 0 -exclueday Tue Thu</code>

ca_generic Run Job Options

The ca_generic command line utility supports the following run job options.

Option	Description
-at <i>hh:mm</i>	Specify the execution time of the generic job.
-on <i>mm/dd/yyyy</i>	Specify the execution date of the generic job.
-hold	Submit the generic job on hold. Cannot be used with -runjobnow.
-runjobnow	Submit and execute the generic job immediately. Cannot be used with -hold.
-description <i>description string</i>	Add comments to the job. You must use double quotes "" to enclose the string and handle blank spaces.

ca_jobstat Command

The `ca_jobstat` command can be used to generate real time, backup job status information. The command captures data from logs stored in the `$BAB_HOME/logs/` directory and database. The report's fields include job ID, job description given during backup, current status, client details, and error messages. You can also use this command to generate custom information. For example, a list of all canceled backup jobs that started today between 5:00 PM and 10:00 PM.

You can print and send the report via email messaging.

ca_jobstat Syntax

The `ca_jobstat` command line utility supports the following syntax:

```
ca_jobstat [-d <char>] [[-s "mm/dd/yyyy hh:mm:ss" [-e "mm/dd/yyyy hh:mm:ss" ]]] [-jobid <job id1,job id2,job id3,...>][[-all]] [-status < F | C | L | I> ] [-m <mail id> ] [-p] [-usage]
```

ca_jobstat Options

The `ca_jobstat` command line utility supports the following options.

Option	Description
<code>[-d <char>]</code>	Specify a delimiter that will be used to separate the columns of the output.
<code>[-s mm/dd/yyyy hh:mm:ss]</code>	Specify the backup job start time. All backup jobs that start after this time will be included in the report.
<code>[-e mm/dd/yyyy hh:mm:ss]</code>	Specify the backup job end time. All backup jobs that end before this time will be included in the report.
<code>[-jobid <job id1,job id2,job id3,...>]</code>	Specifies the backup job ID for which information is required. To specify more than one job, enter multiple job IDs separated by a comma.
<code>[-all]</code>	Specifies all backup job IDs to be included in the report.
<code>[-status < F C L I>]</code>	Specify the backup job status. All backup jobs that have this job status will be included in the report.

Option	Description
[-m <mail id>]	Specify the email address where the report should be sent via email.
[-p]	Specify the name of the printer where the report should be printed.
-usage	Display the usage for the command.

Notes

- If you specify the -m option and do not specify the -d option, the default delimiter is set to comma and the attachment file becomes a .csv file. A .csv file can be opened with spreadsheet software. If you specify the -d option, the delimiter is set to the specified delimiter.
- If you specify the -s option and not the -e option, the end time is considered to be the current system time.
- For the email feature to function, the shell PATH environment variable must contain the path for the mail and uuencode/mutt command. The email feature uses mutt for sending emails. In the event the mutt utility is not available, it uses the mail utility with uuencode to send the file attachment. If uuencode is not available, you must download and install a copy of mutt.
- For the printing feature to function, your system must have a default printer specified.

ca_log Command

The ca_log command is the command line interface to the Report Manager. This command allows you to view and maintain BrightStor ARCserve Backup logs. All of the features available from the Report Manager are available from the command line.

Note: Because ca_log communicates with the caloggerd daemon to perform these operations, the caloggerd daemon must be running for ca_log to function properly.

ca_log Syntax

```
ca_log [-cahost <hostname>] [-clear filename] [-delete filename] [-browse] [-view  
<filename> <view options>] [-purge <filename> <purge options>] [-schedprune  
<schedprune options>]
```

The [-cahost <hostname>] switch is an optional switch that identifies the name of the system hosting the operation. If you want to execute the operation on a remote system, this switch must be included in the command. If you want to execute this operation on your local system, this switch is not required and should not be included in the command.

Note: If you include -cahost in the command, you must also specify the *hostname* of the system (local or remote) hosting the operation.

ca_log Log File Manipulation Options

The ca_log command line utility supports the following log file manipulation options.

Option	Description
-usage	Displays a list of basic ca_log commands.
-browse	Show all available log files.
-clear <i>filename</i>	Clear the log file and index files associated with a log file. All of the information in the log file is discarded and the resulting log file will have zero length. Note: All of the information in the log file will be lost.
-delete <i>filename</i>	Delete the specified log file. Note: This command deletes the specified log. This command is generally used for deleting User logs and GFS. Some of the daemon log files, such as activity log, camediad.log, cadbd.log, cauth.log, are essential to BrightStor ARCserve Backup and these log files cannot be deleted by this command. You can clear daemon log files only using the clear command.
-view <i>filename</i>	Show the specified log file. Supports viewing BrightStor.log (Activity Log) only.
-view <i>filename</i> [-jobID <i>ID</i>]	Show the specified log file by job ID.
-view <i>filename</i> [-before <i>mm/dd/yyyy</i>]	Report all entries in the specified log file before the specified date.

Option	Description
-view <i>filename</i> [-after <i>mm/dd/yyyy</i>]	Report all entries in the specified log file after the specified date. Note: You can use the -before and -after options together to display logs across a period of time.
-view <i>filename</i> [-monitor]	Keep the ca_log command from terminating after displaying the last specified log and continues to read and display additional logs from the log file as they become available.
-view <i>filename</i> -sev	Display the severity level of each log message— I (Information), W (Warning), or E (Error). The severity levels are displayed after the date column.
-purge <i>filename</i> [-olderthan <i>num</i> , < day[s] week[s] months[s] year{s}>]	Purge the information of specified log file based on age criteria from daemon log files, such as BrightStor.log, camediad.log, cadbd.log, cauth.log, User log, GFS log and so on.
-schedprune [-prunesize <i>kilobytes</i> -olderthan <i>num</i> <day[s] week[s] month[s] year[s] > -prunetime <i>HH:MM</i>]	Schedule the pruning of a log file based on the file size or older than criteria. If you specify both size and older than criteria, the pruning removes the oldest part of the log file keeping the latest information intact. Note: You can schedule pruning for activity log files only.
-prunesize <i>kilobytes</i>	Specify the maximum log file size above which the pruning will occur. The size should be specified in kilobyte units. For example, -prunesize 3200 allows activity log file to grow a maximum of 32MB.
-olderthan <num> <day[s] week[s] months[s] year{s}>	Specify that logs older than given number of days, weeks, months, or years will be removed from the log file.
-prunetime <HH:MM>	Specify the time at which the prune will happen on a daily basis.

ca_mediarep Command

The `ca_mediarep` command can be used to generate information about media used during a backup. This report captures data from the `astape` table in the BrightStor ARCserve Backup database.

You can print and send the report via email messaging.

ca_mediarep Syntax

The `ca_mediarep` command line utility supports the following syntax:

```
ca_mediarep [-d <char>] [-p ] [-m <mail id>] [<-t tape id>|<-n tape name>] [-usage]
```

ca_mediarep Options

The `ca_mediarep` command line utility supports the following options.

Option	Description
[-d <char>]	Specify a delimiter that will be used to separate the columns of the output.
[-p]	Specify the name of the printer where the report should be printed.
[-m <mail id>]	Specify the email address where the report should be sent via email.
[<-t tape id> <-n tape name>]	Specify the tape name and tape ID from which you want to filter the output.
-usage	Display the usage for the command.

Notes

- For the email feature to function, the shell PATH environment variable must contain the path for the mail and uuencode/mutt command. The email feature uses mutt for sending emails. In the event the mutt utility is not available, it uses the mail utility with uuencode to send the file attachment. If uuencode is not available, you must download and install a copy of mutt.
- For the printing feature to function, your system must have a default printer specified.
- If you specify the -m option and do not specify the -d option, the default delimiter is set to comma and the attachment file becomes a .csv file. A .csv file can be opened with spreadsheet software. If you specify the -d option, the delimiter is set to the specified delimiter.

ca_merge Command

The `ca_merge` command is the command line interface to the BrightStor ARCserve Backup Merge Manager utility. Use this command to create and submit merge jobs to the Job Queue. You can merge database information from backup media into your BrightStor ARCserve Backup database. All of the features available from the Merge Manager are available from the command line.

ca_merge Syntax

```
ca_merge [-cahost <hostname>] <source args> <run-job args> <options>
```

The `[-cahost <hostname>]` switch is an optional switch that identifies the name of the system hosting the operation. If you want to execute the operation on a remote system, this switch must be included in the command. If you want to execute this operation on your local system, this switch is not required and should not be included in the command.

Note: If you include `-cahost` in the command, you must also specify the *hostname* of the system (local or remote) hosting the operation.

ca_merge Usage

The commands, arguments, and switches available for the `ca_merge` command allow you to specify the data to be merged, allow you to submit the merge job to be run immediately, to submit the job on Hold, or to schedule the job for a later date and time.

Additionally, you can specify Pre/Post command options (including passwords), log options, output devices, first and span media options, and to save the merge job as a script.

ca_merge Source Arguments

The `ca_merge` command line utility supports the following source arguments.

Option	Description
-cahost <hostname>	Specify the BrightStor ARCserve Backup server to use for the merge by providing the hostname where the desired server is running. If this switch is not used, <code>cahost</code> is set to the local machine by default.
-group <group name>	Specify the tape group to use for the merge job.
-tape <tape name> [tape ID]	Specify the tape to use for the merge job. The tape ID is optional and is used if there are multiple tapes with the same name.
-currenttapeseq [-allsessions -session <session #>]	For Windows platforms, this option is used to specify to use the current tape sequence for the merge job.
-currenttapeseq [-allsessions -session <session range>]	For UNIX and Linux platforms, this option is used to specify to use the current tape sequence for the merge job.
-allsessions	Specify to merge all the sessions of the tape for the merge job.
-session session range	Specify to merge a single session or multiple sessions of the tape. Specify a session range to merge multiple sessions.

ca_merge Run Job Arguments

The ca_merge command line utility supports the following run job arguments.

Option	Description
-at <hh:mm>	Specify the execution time of the merge job.
-on <mm/dd/yyyy>	Specify the execution date of the merge job.
-hold	Submit the merge job on hold. Cannot be used with -runjobnow.
-runjobnow	Submit and execute the merge job immediately. Cannot be used with -hold.
-description <description string>	Add comments to the job. You must use double quotes "" to enclose the string and handle blank spaces.

ca_merge Miscellaneous Options

Media Options

The ca_merge command line utility supports the following media options.

Option	Description
-firsttapetimeout <i>minutes</i>	Specify the time, in minutes, to wait for a usable media to be made available for the merge job. The default is 5 minutes. If a usable media is not made available within this time, the job times out and fails.
-spantapetimeout <i>minutes</i>	Specify the time, in minutes, to wait for a usable span media to be made available for the merge job. The default is infinite; the job continues to wait and prompt until a usable media is loaded or the user cancels the job.

Job Status Options

The ca_merge command line utility supports the following job status options.

Option	Description
-waitForJobStatus <polling interval (secs)>	The ca_merge command 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 ca_merge checks the jobs' status with the Queue services. The default interval is 60 seconds. This is useful for Unicenter NSM (formerly known as TNG) scheduling. For more information on how to use this to integrate with Unicenter NSM, see Unicenter NSM Integration in the chapter "Introducing BrightStor ARCserve Backup."

Miscellaneous Options

The ca_merge command line utility supports the following miscellaneous options.

Option	Description
-list	Used to display a list of tapes available for the merge job.
-savescript <scriptname>	Used to save a job script (for use on UNIX and Linux platforms only).
-f <filename>	Used to specify a file name that contains the switches and parameters for the command. This switch overcomes the shell limitation of 1024 character input from command line. You can also use this switch to hide passwords by saving them in a file.
-help	Opens the ca_merge Help topic.
-usage	Displays a list of basic ca_merge commands.
-allusage	Displays a list of all ca_merge commands and their switches.
-examples	Opens a Help topic with ca_merge examples.

ca_mmomgr Command

The `ca_mmomgr` command is the command line interface to the BrightStor ARCserve Backup Media Management Admin from the command prompt. Use this command to control and monitor vaulting operations and reports. Many of the features available from the Media Management Admin are available from the command line.

ca_mmomgr Syntax

```
ca_mmomgr [-vault cycle options] [report options] [report by date options]
```

ca_mmomgr Options

`ca_mmomgr` has three types of commands:

- Vault Cycle commands—These commands provide control of the vault cycle.
- Report commands—These commands allow you to monitor the current Media Management reports.
- Report by Date commands—These commands allow you to view Media Management reports for a particular date.

You can automate Media Management operations by saving any of these commands as scripts.

Miscellaneous Options

Option	Description
-allusage	Displays a list of all <code>ca_mmomgr</code> commands and their switches.
-usage	Displays a list of basic <code>ca_mmomgr</code> commands.

ca_mmomgr Vault Cycle Options

The ca_mmomgr command line utility supports the following vault cycle options.

Option	Description
-start_vault_cycle [-export] [-email]	Start the vault cycle to assign media to slots, and update location information for your media. Use the -export option if you want export a tape after vaulting. Use the -email option if you want to email vault cycle reports to a customizable list of recipients that you can define in the cadbd.cfg file. To define the list of recipients, go to \$BAB_HOME/config/cadbd.cfg and enter the email addresses next to the EMAIL_MMOREPORTS setting. Note: Recipient information is read from ca_mmomgr at run time. You do not have to stop and restart any process.
-simulate_cycle [<# of days from today>][-email]	Simulates the vault cycle to produce five reports, including a Vault Selection Report, and to predict how many tape volumes will be moved in the next vault cycle without updating location information. Use the -email option if you want to email simulated reports to a customizable list of recipients that you can define in the cadbd.cfg file. To define the list of recipients, go to \$BAB_HOME/config/cadbd.cfg and enter the email addresses next to the EMAIL_MMOREPORTS setting. Note: Recipient information is read from ca_mmomgr at run time. You do not have to stop and restart any process.
-reset	Manually reset the current status of Vault Processing if something went wrong during the vault cycle, such as the Media Management database became corrupted while the vault cycle process was running. After the status is reset, you should be able to restart another vault cycle.
-exportAll	Export all tapes. Use this option independently if -export is not used with the -start option. This is useful if you do not want to export every time you run a vault cycle.

ca_mmomgr Report Options

The ca_mmomgr command line utility supports the following report options.

Option	Description
-show shipping	Display the current Shipping report. By default, the output device is the Media Management Admin window.
-show receiving	Display the current Receiving Reports. By default, the output device is the Media Management Admin window.
-show inventory_media	Display the current Media Inventory Report. By default, the output device is the Media Management Admin window.
-show inventory_vault	Display the current Vault Inventory Report. By default, the output device is the Media Management Admin window.
-show vault_selection	Display the current Vault Selection Report. By default, the output device is the Media Management Admin window.

ca_mmomgr Report Options by Date

The ca_mmomgr command line utility supports the following report options by date.

Option	Description
-display shipping <i>yyyy-mm-dd</i>	Display the Shipping Report for the specified date in the Media Management Admin window.
-display shipping_content <i>yyyy-mm-dd</i>	Display the Shipping Content Report for the specified date in the Media Management Admin window.
-display shipping_simulate <i>yyyy-mm-dd</i>	Display the Shipping Simulate Report for the specified future date in the Media Management Admin window.
-display shipping_simulate_content <i>yyyy-mm-dd</i>	Display the Shipping Simulate Content Report for the specified date in the Media Management Admin window.

Option	Description
-display receiving yyyy-mm-dd	Display the Show Receiving Report for the specified date in the Media Management Admin window.
-display receiving_content yyyy-mm-dd	Display the Receiving Content Report for the specified date in the Media Management Admin window.
-display receiving_simulate yyyy-mm-dd	Display the Show Receiving Simulate Report for the specified future date in the Media Management Admin window.
-display receiving_simulatecontent yyyy-mm-dd	Display the Receiving Simulate Content Report in the Media Management Admin window.
-display inventory_media yyyy-mm-dd	Display the Media Inventory Report for the specified date in the Media Management Admin window.
-display inventory_vault yyyy-mm-dd	Display the Vault Inventory Report for the specified date in the Media Management Admin window.
-display vault_selection yyyy-mm-dd	Display the Vault Selection Report for the specified date in the Media Management Admin window.
-display vault_selection_simulate yyyy-mm-dd	Display the Vault Selection Report for the specified future date in the Media Management Admin window.

ca_qmgr Command

This command, the command line interface with the Job Status Manager, allows you to monitor and control jobs submitted to the BrightStor ARCserve Backup job queue. All of the features available from the Job Status and Activity Log Manager are available from the command line.

ca_qmgr Syntax

ca_qmgr [-cahost <hostname>]

The [-cahost <hostname>] switch is an optional switch that identifies the name of the system hosting the operation. If you want to execute the operation on a remote system, this switch must be included in the command. If you want to execute this operation on your local system, this switch is not required and should not be included in the command.

Note: If you include -cahost in the command, you must also specify the *hostname* of the system (local or remote) hosting the operation.

ca_qmgr Commands

ca_qmgr has three kinds of commands:

- Job Queue commands—These commands allow you to view and control the Job Queue.
- Job Script commands—These commands allow you to control and use job scripts.
- Job-Specific commands—These commands allow you to monitor and control individual jobs.

ca_qmgr Job Queue Commands

The ca_qmgr command line utility supports the following job queue commands.

Command	Description
-list	Display current job queue.
-listprevid	Shows current job queue, including previous job IDs.
-holdq [on off]	Activate and deactivate the job queue. When on, the queue is deactivated and no jobs are run. When off, the queue is activated, and jobs run normally.
-usage	Display a list of all ca_qmgr commands and their switches.

ca_qmgr Job Script Commands

The ca_qmgr command line utility supports the following job script commands.

Command	Description
-listscripts	Display available job scripts in \$BAB_HOME/jobscrip
-load <i>job script</i> [<i>script owner</i>]	Load and run previously saved job scripts.
-addscript <i>job script</i>	Import and register a "foreign" job script. For a job script to be available for BrightStor ARCserve Backup, it must be registered for use on the backup server.
-removescript <i>job script</i> [<i>script owner</i>]	Remove and de-register a job script.

ca_qmgr Job-Specific Commands

The ca_qmgr command line utility supports the following job-specific commands.

Command	Description
-changestatus <i>job ID</i> ready hold	Change the job status to ready or put a job on hold.
-changedate <i>job ID</i> mm/dd/yyyy	Change the date a job will run.
-changetime <i>job ID</i> hh:mm	Change the time a job will run.
-stop <i>job ID</i>	Stop a currently running job. If it is a repeating job, the next job in the sequence is queued. If it is a run-once job, it is stopped and deleted. If it is a job on hold, no action is taken. Important! No confirmation is asked! The job is stopped without asking if you are sure!

Command	Description
-delete <i>job ID</i>	Delete an inactive job. Deleting inactive jobs completely removes them from the job queue. If you delete a makeup job, a new backup job is scheduled. Note: If you want to delete an active job, you must first stop the job before you can delete it.
-view <i>job ID</i>	View details of job ID (Job Summary).

ca_recoveryrep Command

The `ca_recoveryrep` command can be used to generate reports listing media that is needed to recover files from a failed machine. This report captures the data from the BrightStor ARCserve Backup database. The `ca_recoveryrep` command eliminates the need to search for job and corresponding media information stored in the Database Manager.

You can print and send the report via email messaging.

ca_recoveryrep Syntax

The `ca_recoveryrep` command line utility supports the following syntax:

```
ca_recoveryrep [ -d <char> ] [ -p ] [ -m <mail id> ] [ -h <node name> ] [ -usage ]
```

ca_recoveryrep Options

The `ca_recoveryrep` command line utility supports the following options.

Option	Description
[-d <char>]	Specify a delimiter that will be used to separate the columns of the output report.
[-h <node name>]	Specify the node name from which to capture media details.
[-m <mail id>]	Specify the email address where the report should be mailed.

Option	Description
[-p]	Specifies the name of the printer where the report should be printed.
-usage	Display the usage for the command.

Notes

- For the email feature to function, the shell PATH environment variable must contain the path for the mail and uuencode/mutt command. The email feature uses mutt for sending emails. In the event the mutt utility is not available, it uses the mail utility with uuencode to send the file attachment. If uuencode is not available, you must download and install a copy of mutt.
- For the printing feature to function, your system must have a default printer specified.
- If you specify the -m option and do not specify the -d option, the default delimiter is set to comma and the attachment file becomes a .csv file. A .csv file can be opened with spreadsheet software. If you specify the -d option, the delimiter is set to the specified delimiter.

ca_restore Command

This command, the command line interface to the Restore Manager, allows you to create and submit restore jobs to the BrightStor ARCserve Backup Job queue, and to set all associated options. All of the features available from the Restore Manager are available from the command line.

ca_restore Syntax

```
ca_restore [-cahost <hostname>] [global options] [global filters] - source  
[source arguments] - dest [destination arguments] [schedule arguments] [run job  
arguments]
```

The [-cahost <hostname>] switch is an optional switch that identifies the name of the system hosting the operation. If you want to execute the operation on a remote system, this switch must be included in the command. If you want to execute this operation on your local system, this switch is not required and should not be included in the command.

Note: If you include -cahost in the command, you must also specify the *hostname* of the system (local or remote) hosting the operation.

ca_restore Oracle RMAN Specific Syntax

To restore Oracle RMAN database objects, use the following syntax:

```
ca_restore- source <hostname> < [<tablespace> | "ARCHIVE LOG" | "PARAMETER FILE"
| "CONTROL FILE"] [<absolute path of the datafile>]
```

ca_restore Usage

The options and switches for the `ca_restore` command allow you to set global options and filters, select your source and destination for the restore job, and submit the restore job to run immediately or at a scheduled time. You must specify them in the above order.

ca_restore Global Options and Global Filters

The `ca_restore` command line utility supports the following global options and global filters.

Option	Description
<code>-firsttapetimeout <i>minutes</i></code>	Specify the time, in minutes, to wait for a usable media to be made available for the restore job. The default is 5 minutes. If a usable media is not made available within this time period, the job times out and fails.
<code>-spantapetimeout <i>minutes</i></code>	Specify the time, in minutes, to wait for a usable span media to be made available for the restore job. The default is infinite; the job continues to wait and prompt until a usable media is loaded or the user cancels the job.
<code>-logfile <i>filename</i></code> <code>[summary allactivity]</code>	Record activities during the running of the restore job to the specified file name. The user can specify to record all activity or only a summary of the activity.
<code>-snmp</code>	Enable SNMP Alert.
<code>-tng</code>	Enable Unicenter NSM Alert.
<code>-email <i>email address</i></code>	Send a copy of the Activity log to the specified email address.

Option	Description
-printer <i>printer name</i>	Send a copy of the Activity log to the specified printer. Note: The printer must be set up in the configuration file: \$BAB_HOME/config/caloggerd.cfg.
-preexec <i>command</i>	Run the specified command before the job starts. The entire path of the command should be included.
-postexec <i>command</i>	Run the specified command after the job finishes. The entire path of the command should be included.
-preexec timeout <i>minutes</i>	Specify the time to wait, in minutes, before the restore job starts to allow time for the pre-execute command to finish.
-prepost username <i>username</i>	Specify the user name of the user submitting this restore job.
-preserve userspace off	Turn off preserve user space restrictions (applies to NetWare agents).
-preserve dirs space off	Turn off preserve directory space restrictions (applies to NetWare agents).
-exitcode <i>exit code</i>	Specify the exit code of the pre-execute command. Used with the -skip_delay, -skip_job, and -skip_post switches.
-condition <i>equalto greaterthan lessthan notequalto</i>	Specify the condition for pre execute command exit codes. Used with -exitcode. For greater than and less than conditions, only one exit code should be specified.
-skip_delay	Run the restore job immediately if the specified exit code is received.
-skip_job	Skip the restore job completely if the specified exit code is received.
-skip_post	Skip the post-execute command if the specified exit code is received.
-session password <i>session password</i>	Specify the session password needed to restore data from tape. Required only if a session password was applied during the backup job.
-create empty dir off	Do not create empty directories during restore.

Option	Description
-entirepath	Create the entire path from root during restore.
-base	Create directories from the base during restore.
-nobase	Do not create directories from the base during the restore.
-onconflict rename skip overwriteold	Specify the action to take when a file with the same name is encountered during the restore.
-filter <i>include exclude</i> [<i>file dir attribute date</i>]	Apply a filter to the job. User must specify to include or exclude the desired pattern. Filter types are files or directories, attributes, or dates. Note: If you select the include directory pattern filter and do not specify an absolute path, empty directories for all the directories that do not match the user provided criteria will be restored. To avoid creating these empty directories during restore, disable the global restore option Create Empty Directories when creating your restore job. For more information on this option, see Create Empty Directories Option.
- f <script filename>	Used to specify a file name that contains the switches and parameters for the command. This switch overcomes the shell limitation of 1024 character input from command line. You can also use this switch to hide passwords by saving them in a file.

ca_restore Source Arguments

The ca_restore command line utility supports the following source arguments.

Option	Description
-source [<i>hostname</i>] <i>filelist</i>	Specify the files/directories to restore. If the -source switch is used alone, without other switches; the Restore is treated as a File system view restore. BrightStor ARCserve Backup determines the version of the file to restore. For example, if a file was backed up several times, each time to a different session or even a different tape, and the user does not specify the tape or session of this file to restore, the database finds the most recent version and restores this file.
-tape <i>tape name</i> [<i>tape ID</i>]	Specify the tape for the restore job. The tape ID is optional and is used if there are multiple tapes with the same name. If the -tape switch is used with the -source switch, the restore is treated as a Tape Session view restore and the BrightStor ARCserve Backup database is used. BrightStor ARCserve Backup checks if it has a record of the file and tape specified for the restore. If not, the restore job is not submitted, even if all of the information provided is actually correct. The tape and session in question must be merged into the BrightStor ARCserve Backup database before this restore job can be submitted. If the -tape switch is not used with the -source switch, the restore is treated as a Tape view restore, and the BrightStor ARCserve Backup database is not used. If the specified tape name or session number is invalid, the restore job fails at run-time. The -tape switch must be used with the -session switch.
-session <i>session number</i>	Specify the tape session number to use for the restore job. Must be used with the -tape switch.
-group <i>group name</i>	Specify the tape group to use for the restore job.

Option	Description
<code>-tapesessionpw <i>session password</i> /<i>encryption key</i></code>	Specify the session password or encryption key needed to restore data from tape. Required only if a session password or encryption key was applied during the backup job.
<code>-listgroups</code>	Display a list of groups available for the restore job.
<code>-listtapes</code>	Display a list of tapes available for the restore job.
<code>-listsessions <i>tape name</i> [<i>tape ID</i>]</code>	Display a list of tape sessions backed up to the specified tape and available for restore.
<code>-version [<i>host name</i>] <i>path</i></code>	Display a version history of the specified file/directory that has been backed up. The host name is optional and defaults to the local machine if not provided.
<code>-findfile <i>file name</i> ignorecase casesensitive <i>hostname</i> any <i>search path</i> inclsubdir noinclsubdir <i>mm/dd/yyyy</i> today within # days months years [timeout <timeoutvalue>] [maxrecord <maxrecordvalue>]</code>	Search the BrightStor ARCserve Backup database to determine if a file has been backed up. The user provides the file name, whether the name is case-sensitive, the host name (or any if any host name applies), the path to search for the file (use / to search at the top-most level), whether to include sub-directories in the search, the starting modification date, and the number of days, months or years from the starting modification date to search from. Optionally, you can specify a time out value (the job times out if no matching files are found within the specified time; the default value is 600 seconds), and the maximum number of matches to return (by default, all matches are returned).

ca_restore Destination Arguments

The `ca_restore` command line utility supports the following destination arguments.

Option	Description
<code>-dest [hostname] path</code>	Specify the destination machine and directory path to restore files to. The <i>hostname</i> is optional; if not provided, the default is the local machine.
<code>-orglocation</code>	Restore files to their original machine and path. Note: The <code>-orglocation</code> switch is supported only for UNIX file system, Linux file system, NT file system, and NT System State restores. It is not supported when restoring a database from the command line.
<code>-username user name</code>	Specify the <i>user name</i> of the destination machine to which to restore. This is the user used to log in to the desired machine.
<code>-password password</code>	Specify the password to use to log into the destination machine.
<code>-database database type database name</code>	Specify the database type and name to restore to. Supported, valid database types are—INFORMIX, ORACLE, ORACLE8, ORACLE_AS66, SAP, SYBASE, and INGRES. ORACLE_AS66 supports the Oracle Database agent.
<code>-dbusername database username</code>	Specify the <i>database username</i> to use to log in to the destination database.
<code>-dbpassword database password</code>	Specify the password for the database user to use to log in to the destination database.
<code>-oracle_controlfile</code>	Specify to restore the Oracle control file. Oracle database restore specific.
<code>-oracle_overwritelog</code>	Specify to overwrite the existing log file. Oracle database restore specific.

Option	Description
-preserveuserspaceoff	Disables the preserve user space option. Does not restore the user space restrictions along with the files. By default, the preserve user space option is applied and the same user space restrictions assigned during the backup will also be applied during the restore. Note: This option is only applicable when files or sessions will be restored to a machines running the NetWare Agent.
-preservedirspaceoff	Disables the preserve directory space option. Does not restore the directory space restrictions along with the files. By default, the preserve directory space option is applied and the same directory space restrictions assigned during the backup will also be applied during the restore. Note: This option is only applicable when files or sessions will be restored to a machines running the NetWare Agent.

ca_restore Schedule and Run Job Arguments

The ca_restore command line utility supports the following schedule and run job arguments.

Option	Description
-at <i>hh:mm</i>	Specify the execution time of the restore job.
-on <i>mm/dd/yyyy</i>	Specify the execution date of the restore job.
-hold	Submit the restore job on hold. Cannot be used with -runjobnow.
-runjobnow	Submit and execute the restore job immediately. Cannot be used with -hold.
-description <i>description string</i>	Add comments to the job. You must use double quotes "" to enclose the string and handle blank spaces.

Option	Description
-savescript <i>script name</i>	Save the restore job as a script rather than submit it to the Job Queue. The script can be loaded into the Job Queue at a later time. See <i>ca_qmgr</i> .

ca_scan Command

This command is the command line interface to the Scan Manager utility, and allows you to create and submit scan jobs to the Job Queue. All of the features available from the Scan Manager are available from the command line. Reports information about one or more backup sessions on media.

ca_scan Syntax

```
ca_scan [-cahost <hostname>] <source args> <run-job args> <options>
```

The [-cahost <hostname>] switch is an optional switch that identifies the name of the system hosting the operation. If you want to execute the operation on a remote system, this switch must be included in the command. If you want to execute this operation on your local system, this switch is not required and should not be included in the command.

Note: If you include -cahost in the command, you must also specify the *hostname* of the system (local or remote) hosting the operation.

ca_scan Usage

The commands, arguments, and switches available for the *ca_scan* command allow you to specify the data to be scanned, allow you to submit the scan job to be run immediately, to submit the job on Hold, or to schedule the job for a later date and time.

Additionally, you can specify Pre/Post command options (including passwords), log options, output devices, first and span media options, and to save the scan job as a script.

ca_scan Source Arguments

The ca_scan command line utility supports the following source arguments.

Option	Description
-cahost <hostname>	Specify the BrightStor ARCserve Backup server to use for the scan by providing the hostname where the desired server is running. If this switch is not used, cahost is set to the local machine by default.
-group <group name>	Specify the tape group to use for the scan job. If you do not know the name of the group, you can use the wildcard character '*' as in the following example: Group *. However, when you use the wildcard character, ca_scan will only scan media that corresponds to the first available tape group in the list of tape groups to be scanned, for example, "Group0."
-tape <tape name> [tape ID]	Specify the tape to use for the scan job. The tape ID is optional and is used in the event that there are multiple tapes with the same name.
-currenttapeseq [-allsessions -session <session #>]	For Windows platforms, this option is used to specify to use the current tape sequence for the scan job.
-currenttapeseq [-allsessions -session <session range>]	For UNIX and Linux platforms, this option is used to specify to use the current tape sequence for the scan job.
-allsessions	Specify to scan all the sessions of the tape for the scan job.
-session session range	Specify to scan a single session or multiple sessions of the tape. Specify a session range to scan multiple sessions.

ca_scan Run Job Arguments

The ca_scan command line utility supports the following run job arguments.

Option	Description
-at <hh:mm>	Specify the execution time of the scan job.
-on <mm/dd/yyyy>	Specify the execution date of the scan job.

Option	Description
-hold	Submit the scan job on hold. Cannot be used with -runjobnow.
-runjobnow	Submit and execute the scan job immediately. Cannot be used with -hold.
-description <description string>	Add comments to the job. You must use double quotes "" to enclose the string and handle blank spaces.

ca_scan Logging Options

The ca_scan command line utility supports the following logging options. These options are for use on UNIX and Linux hosts.

Option	Description
-logfile <filename> [summary allactivity]	Record activities during the running of the scan job to the specified filename. The user can specify to record all activity or only a summary of the activity.
-snmp	Enable SNMP Alert.
-tng	Enable Unicenter NSM Alert.
-email <email address>	Send a copy of the Activity log to the specified email address.
-printer <printer name>	Send a copy of the Activity log to the specified printer. Note: The printer must be set up in the configuration file: \$BAB_HOME/config/asloggerd.cfg.

ca_scan Pre/Post Options

The ca_scan command line utility supports the following Pre/Post options.

Command	Description
-preexec <command>	Run the specified command before the job starts. The entire path of the command should be included. Note: To use this option you must also specify the -prepostuser option.

Command	Description
-postexec <command>	Run the specified command after the job finishes. The entire path of the command should be included. Note: To use this option you must also specify the -prepostuser option.
-preexec timeout <minutes>	Specify the time to wait, in minutes, before the scan job starts to allow time for the pre-execute command to finish.
-prepostuser <username>	Specify the user name of the user submitting this scan job.
-prepostpassword <user password>	Specify the password of the user submitting this scan job.

ca_scan Exit Options

The ca_scan command line utility supports the following exit options.

Command	Description
-exitcode <exit code>	Specify the exit code of the pre-execute command. Used with the -skip_delay, -skip_job, and -skip_post switches.
-condition equalto greaterthan lessthan n_ notequalto	Specify the condition for pre execute command exit codes. Used with -exitcode. For greater than and less than conditions, only one exit code should be specified.
-skip_delay	Run the scan job immediately if the specified exit code is received. This option will be activated only if BrightStor ARCserve Backup detects that the exit codes meet the specified condition (Equal To, Greater Than, Less Than, or Not Equal To).
-skip_job	Skip the scan job completely if the specified exit code is received. This option will be activated only if BrightStor ARCserve Backup detects that the exit codes meet the specified condition (Equal To, Greater Than, Less Than, or Not Equal To).

Command	Description
-skip_post	Skip the post-execute command if the specified exit code is received. This option will be activated only if BrightStor ARCserve Backup detects that the exit codes meet the specified condition (Equal To, Greater Than, Less Than, or Not Equal To).

ca_scan Miscellaneous Options

Media Options

The ca_scan command line utility supports the following media options.

Option	Description
-firsttapetimeout <i>minutes</i>	Specify the time, in minutes, to wait for a usable media to be made available for the scan job. The default is 5 minutes. If a usable media is not made available within this time, the job times out and fails.
-spantapetimeout <i>minutes</i>	Specify the time, in minutes, to wait for a usable span media to be made available for the scan job. The default is infinite; the job continues to wait and prompt until a usable media is loaded or the user cancels the job.

Job Status Options

The ca_scan command line utility supports the following job status options.

Option	Description
-waitForJobStatus <polling interval (secs)>	The ca_scan command 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 ca_scan checks the jobs' status with the Queue services. The default interval is 60 seconds. This is useful for Unicenter NSM (formerly known as TNG) scheduling. For more information on how to use this to integrate with Unicenter NSM, see Unicenter NSM Integration in the chapter "Introducing BrightStor ARCserve Backup."

Miscellaneous Options

The `ca_scan` command line utility supports the following miscellaneous options.

Option	Description
-list	Used to display a list of tapes available for the scan job.
-savescript <scriptname>	Used to save a job script (for use on UNIX and Linux platforms only).
-f <filename>	Used to specify a file name that contains the switches and parameters for the command. This switch overcomes the shell limitation of 1024 character input from command line. You can also use this switch to hide passwords by saving them in a file.
-help	Opens the <code>ca_scan</code> Help topic.
-usage	Displays a list of basic <code>ca_scan</code> commands.
-allusage	Displays a list of all <code>ca_scan</code> commands and their switches.
-examples	Opens a Help topic with <code>ca_scan</code> examples.

ca_stagingrep Command

The `ca_stagingrep` command line utility can be used to query the BrightStor ARCserve Backup database to obtain information and create various reports about staging session data. You can select the type of report to be generated as either a complete report based on a specified function (migration, purge, or snaplock) or a summary report of all migration, purge, and snaplock jobs based on a specified job number or within a specified time period.

You can print and send the report via e-mail messaging.

ca_stagingrep Syntax

The ca_stagingrep command supports the following syntax:

```
ca_stagingrep
  -r <migration | purge | snaplock>
  -j <jobid> [-s <yyyy-mm-dd>] [-e <yyyy-mm-dd>]
  [-d <char>]
  [-m <email id>]
  [-p ]
```

Note: Use the -r option to generate static data reports about migration, purge, and snaplock sessions. Use the -j option to generate summary reports based upon a date range or job number. The Summary Report captures the details for all session categories (migration, purge, and snaplock).

ca_stagingrep Options

The ca_stagingrep command supports the following options:

Option	Description
-r [migration purge snaplock]	Generates a complete report for the specified function (migration, purge, or snaplock), including all job numbers and all dates.
migration	Directs BrightStor ARCserve Backup to capture information about sessions that migrated (Migrated), have not yet migrated (Not Migrated), migration failed (Migration Failed) , or migration is not applicable (No Migration).
purge	Directs BrightStor ARCserve Backup to capture information and create a report about sessions that are prior to or past the purge date.
snaplock	Directs BrightStor ARCserve Backup to capture information about sessions with and without SnapLock, WORM protection.
-j	Generates a summary report of all migration, purge, and snaplock jobs for a specified master job ID or within a specified time period (start and end date).
<jobid>	Directs BrightStor ARCserve Backup to capture information about sessions with the same master job ID. For example, GFS jobs.

Option	Description
[-s <yyyy-mm-dd>]	Directs BrightStor ARCserve Backup to capture information about sessions from the start time.
[-e <yyyy-mm-dd>]	Directs BrightStor ARCserve Backup to capture information about sessions to the end time.
-d <char>	Specifies the delimiter that will be used to separate the columns of the output.
-m <mail id>	Specifies the email address where the report should be sent via email.
-p	Specifies the name of the printer where the report should be printed.
-usage	Displays the usage information.

Notes:

- For the email feature to function, the shell PATH environment variable must contain the path for the mail and uuencode/mutt command. The email feature uses mutt for sending emails. In the event the mutt utility is not available, it uses the mail utility with uuencode to send the file attachment. If uuencode is not available, you must download and install a copy of mutt.
- For the printing feature to function, your system must have a default printer specified.
- If you specify the -m option and do not specify the -d option, the default delimiter is set to comma and the attachment file becomes a .csv file. A .csv file can be opened with spreadsheet software. If you specify the -d option, the delimiter is set to the specified delimiter.

ca_summaryrep Command

The `ca_summaryrep` command can be used to generate reports that summarize a client's start time, end time, backup type, backup job status and the name of the host server. The command captures data from logs stored in the `$BAB_HOME/logs/` directory and the BrightStor ARCserve Backup database.

You can print and send the report via email messaging.

ca_summaryrep Syntax

The ca_summaryrep command line utility supports the following syntax:

```
ca_summaryrep
  [-d <char>]
  [-j <job id1, ,job id2, ,job id3,...>] | [ -h host names ] | [-f ] | [-s
mm/dd/yyyy hh:mm:ss [-e mm/dd/yyyy hh:mm:ss]]
  [-v]
  [-p]
  [-m <mail id>]
  [-usage]
```

ca_summaryrep Options

The ca_summaryrep command line utility supports the following options.

Option	Description
[-d <char>]	Specify a delimiter that will be used to separate the columns of the output report.
[-j <job id1,job id2,job id3,...>]	Filters the report based on job ID. You can specify more than one job ID by separating each job ID with a comma.
[-h host names]	Specify the source host from which you want to filter the output. You can specify more than one host by separating the host names with a comma.
[-f]	Displays the backup summary report on file system level.
[-s mm/dd/yyyy hh:mm:ss]	Specify the backup job start time. All backup jobs that start after this time will be included in the report.
[-e mm/dd/yyyy hh:mm:ss]	Specify the backup job end time. All backup jobs that end before this time will be included in the report.
[-v]	Specifies time elapsed, throughput, total data in mb (verbose mode).
[-p]	Specify the name of the printer where the report should be printed.
[-m <mail id>]	Specify the email address where the report should be mailed.

Option	Description
-usage	Display the usage for the command.

Notes

- For the email feature to function, the shell PATH environment variable must contain the path for the mail and uuencode/mutt command. The email feature uses mutt for sending emails. In the event the mutt utility is not available, it uses the mail utility with uuencode to send the file attachment. If uuencode is not available, you must download and install a copy of mutt.
- For the printing feature to function, your system must have a default printer specified.
- If you specify the -s option and not the -e option, the end time is considered to be the current system time.
- If you specify the -m option and do not specify the -d option, the default delimiter is set to comma and the attachment file becomes a .csv file. A .csv file can be opened with spreadsheet software. If you specify the -d option, the delimiter is set to the specified delimiter.

ca_utilizationrep Command

The `ca_utilizationrep` command can be used to generate reports that provide media usage information. The command captures data from the BrightStor ARCserve Backup database and the library file for the media type. The report's fields include the media information, (tape name, sequence number, serial number, tape ID), percentage of media used, amount of data stored on the media, and job information (job ID, description, and type).

You can print and send the report via email messaging.

ca_utilizationrep Syntax

The `ca_utilizationrep` command line utility supports the following syntax:

```
ca_utilizationrep [-d <char>] [-p] [-m <mail id>] [-l <location status>] | -e
<expire date> | -s <data in MB>] [-usage]
```

ca_utilizationrep Options

The ca_utilizationrep command line utility supports the following options.

Option	Description
[-d <char>]	Specify a delimiter that will be used to separate the columns of the output.
[-p]	Specify the name of the printer where the report should be printed.
[-m <mail id>]	Specify the email address where the report should be sent via email.
[-l <location status> -e <expire date> -s <data in MB>]	Specify the information that you want to view. ■ -l <location status>—Location status of the media. The status will display as 0 (Online), 1 (Offline), or 2 (Offsite). ■ -e <expire date>The expiration date of the media. ■ -s <data in MB>—Size of the data written to the tape.
-usage	Display the usage for the command.

Notes

- For the email feature to function, the shell PATH environment variable must contain the path for the mail and uuencode/mutt command. The email feature uses mutt for sending emails. In the event the mutt utility is not available, it uses the mail utility with uuencode to send the file attachment. If uuencode is not available, you must download and install a copy of mutt.
- For the printing feature to function, your system must have a default printer specified.
- If you specify the -m option and do not specify the -d option, the default delimiter is set to comma and the attachment file becomes a .csv file. A .csv file can be opened with spreadsheet software. If you specify the -d option, the delimiter is set to the specified delimiter.

ca_vaultrep Command

The `ca_vaultrep` command can be used to generate reports that list media to be exported to a vault or imported from a library. The command captures data from the `$BAB_HOME/bin/ca_mmomgr`, and the `vslot`, `astape`, and `vault` tables in the BrightStor ARCserve Backup database. The report's fields include the media name, sequence number, serial number, and tape ID.

You can print and send the report via email messaging.

ca_vaultrep Syntax

The `ca_vaultrep` command line utility supports the following syntax:

```
ca_vaultrep [-p] [-m <mail id>] [-d <char>] [-usage]
```

ca_vaultrep Options

The `ca_vaultrep` command line utility supports the following options.

Option	Description
[-p]	Specify the name of the printer where the report should be printed.
[-m <mail id>]	Specify the email address where the report should be sent via email.
[-d <char>]	Specify a delimiter that will be used to separate the columns of the output.
[-usage]	Display the usage for the command.

Notes

- For the email feature to function, the shell PATH environment variable must contain the path for the mail and uuencode/mutt command. The email feature uses mutt for sending emails. In the event the mutt utility is not available, it uses the mail utility with uuencode to send the file attachment. If uuencode is not available, you must download and install a copy of mutt.
- For the printing feature to function, your system must have a default printer specified.
- If you specify the -m option and do not specify the -d option, the default delimiter is set to comma and the attachment file becomes a .csv file. A .csv file can be opened with spreadsheet software. If you specify the -d option, the delimiter is set to the specified delimiter.

cadiag Command

The BrightStor ARCserve Backup Diagnostic Utility (cadiag) is a convenient tool for gathering and compressing (packing) various BrightStor ARCserve Backup data and system logs, which may be valuable tools for troubleshooting problems. Using this utility you can start, stop, check the status of the diagnostic daemons on the local or a remote machine, and unpack collected diagnostic information to a specific directory or location.

cadiag Syntax

```
cadiag start|stop|status|unpack
```

cadiag Options

The cadiag command line utility supports the following options.

Option	Description
-usage	Displays a list of basic cadiag commands.
start stop status unpack	Starts, stops, and checks the status of the diagnostic daemons. Use the unpack option to unpack and view the collected diagnostic information.

pfc Command

This command allows you to run vital checks on the BrightStor ARCserve Backup server and agents to detect conditions that may cause backup jobs to fail. The checks performed by pfc fall into four categories— system checks, BrightStor checks, agent checks, and media checks:

- **System Checks**—These include checking system requirements for the server, available disk space for the database, and RPC service registration.
- **BrightStor Checks**—These include checking the BrightStor system account and its privileges, the status of the BrightStor engines, SAN server connectivity (if the 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.
- **Media Checks**—These include checking the availability of media in the scratch set (if a media pool is specified for the job), the media expiration dates, and for source and destination conflicts for file system devices.

When you run the pfc utility, it creates the following log:

```
PFC_SERVERNAME_#####.LOG
```

This log includes the same information that appears in the output generated in the Command Prompt windows when you run pfc and is located in the BrightStor LOG directory. You can change this directory by using the -logpath option.

pfc Syntax

```
pfc [<hostname>] [options] [filename(s)]
```

Examples:

Use the following syntax to perform all checks, in non-interactive mode, on all READY jobs in the job queue:

```
pfc -allchecks
```

Use the following syntax to perform system checks in verbose and non-interactive mode:

```
pfc -syschecks -v -n
```

Use the following syntax to perform BrightStor checks and to start any BrightStor ARCserve Backup engines that are not running:

```
pfc -bchecks -s
```

Use the following syntax to perform agent checks for all READY jobs in the queue:

```
pfc -agentchecks -a
```

Use the following syntax to perform agent checks for job103:

```
pfc -agentchecks job103
```

Use the following syntax to perform agent checks for a job which is on HOLD:

```
pfc -agentchecks job105
```

Use the following syntax to perform media checks for job103 and job104:

```
pfc -mediachecks job103 job104
```

Use the following syntax to perform media checks for job103, display the output on the console, and also log the output in a file in the /tmp directory:

```
pfc -mediachecks -logpath /tmp/ job103
```

pfc Options

The pfc command line utility supports the following options.

Option	Description
-allchecks	Performs all checks (system checks, BrightStor checks, agent checks, and media checks), in non-interactive mode, on all Ready jobs in the job queue. You cannot specify filenames when you use this switch.
-syschecks	Performs system checks.
-bchecks	Performs BrightStor checks.
-agentchecks <filenames>	Performs Agent checks. When you use this, you must specify one or more job script file names. For more information, see the filenames option. You can also use the -a switch with this option to run Agent checks for all jobs in the queue.

Option	Description
-mediachecks <filenames>	Performs media checks. When you use this, you must specify one or more job script file names. For more information, see the filenames option. You can also use the -a switch with this option to run media checks for all jobs in the queue.
-a	Specifies all Ready jobs in the job queue. You cannot specify filenames when you use this switch.
-n	Runs in non-interactive mode. When you use this, pfc does not stop during execution to prompt for input.
-s	Attempts to start any BrightStor ARCserve Backup engines that are not running.
-v	Runs in verbose mode. When you use this, pfc provides detailed information in its output to the Command Prompt window and log about the checks being performed. This includes information used for debugging, such as the name of the failing function and the error code returned when an API call fails.
-logpath <pathname>	Sets the path for log files. The default path is the BrightStor ARCserve Backup LOG directory.
filename(s)	Specify a job script file name if you want to perform a check on a specific job. For example, job105. These files are located in the \$BAB_HOME/queue directory.

tapecopy Command

This command allows you to make logical, media to media copies at session level, or of two different type media.

Note: You cannot use the tapecopy command to copy data to VM:Tape media.

tapecopy Syntax

```
tapecopy <-s[source group] [Database Query options]>< -d[destination group]>
[source group options] [destination group options] [Span tape option][Media Pool
option] [Script option]
```

tapecopy Options

The default is \$BAB_HOME/bin/tapecopy -s[Source Group] -d[Destination Group]

This command copies all non-blank tapes in the source group to a tape in the destination group. The tape copy looks for a blank tape in the destination group and formats it, giving it the same name as the source tape.

Miscellaneous Options

Option	Description
-activitybardisable	When tapecopy is running, a rotating activity bar displays indicating that tapecopy is active and copying sessions. This option lets you disable the activity bar.
-allusage	Displays a list of all tapecopy commands and their switches.
-usage	Displays a list of basic tapecopy commands.

tapecopy Database Query Options

These options let you select source sessions based on specific attributes. When you specify a database option, the database is queried and all of the sessions that meet the search criteria become source sessions for tapecopy. One or more of these switches can be used to specify a complex query.

The tapecopy command line utility supports the following database query options.

Options	Description
-qMediaPool <media pool name>	Used to include tapes belonging to the specified Media Pool. It supports searches based on wildcard characters like '*' and '?'.

Options	Description
-qType < <i>backup session type</i> > TYPE: INFORMIX ORACLE ORACLE8 NTORACLE ORACLERMAN SYBASE LOTUS NTDB BABDATABAS UNIX UNIXRAW MSWIN NETWARE TAR CPIO DRTAR	Used to include only sessions of the selected type for copy.
-qMethod < <i>backup session method</i> > METHOD: CUSTOM GFS_FULL GFS_DIFF GFS_INCR ROTATE_FULL ROTATE_DIFF ROTATE_INCR	Used to include only those session that were backed up using the specified backup method.
-qNode < <i>node name</i> >	Used to include sessions backed up from the specified node only.
-qMID < <i>master job ID</i> >	Used to include sessions with the specified Master Job ID only.
-qOnOrBefore < <i>mm/dd/yy</i> > [<i><hh:mm></i>]	Used to include sessions that were backed up on or before the specified date and time. It is optional to specify the time. Data and time should be separated by space.
-qOnOrAfter < <i>mm/dd/yy</i> > [<i><hh:mm></i>]	Used to include sessions that were backed up on or after the specified date and time. It is optional to specify the time. Data and time can be separated by a space.
-qPastTime < <i>number of days</i> >	Used to include sessions that were backed up during the last specified number of days. Days are counted in duration of 24 hours starting from the time the tape copy operation is run. The difference in number of days in each month is taken into consideration.
-qIgnoreRep	Used to ignore the replication flag so that sessions that were already copied by tape copy are included. If you do not ignore the flag, previously copied sessions are ignored by tape copy.

Options	Description
-qPreview	Must be used with one or more of previous query switches. This switch puts tape copy in Preview mode so that tapecopy only displays a list of sessions that satisfy the query criteria; the actual tape copy operation is not performed.

tapecopy Source Group Options

The tapecopy command line utility supports the following source group options.

Options	Description
-s[<i>Source Group</i>]	Specify the name of the group from which to copy.
-sTapeName [<i>TapeName</i>] -sTapeID [<i>TapeID</i>] {-sSeq [<i>Tape Sequence</i>]}	Tape name of the source tape to be copied. This copies all tapes with the specified tape name and tape ID into the destination group.
-n	Beginning session number to copy from.
-N	Number of sessions on the source tape to be copied
-erase	Erase source tape when entire tape is copied successfully.

tapecopy Destination Group Options

The tapecopy command line utility supports the following destination group options.

Option	Description
-d[<i>Destination Group</i>]	Specify the name of the group to copy to. If you omit this option, any available group is used.
-dTapeName [<i>TapeName</i>] -dTapeID [<i>TapeID</i>] -dSeq [<i>TapeSeq</i>]	Append to the specified tape only when the source tape <i>TapeName</i> , <i>TapeID</i> and <i>Seq</i> are provided. If not provided, the option is invalid.
-f [<i>Tapename</i>]	Format a tape with the specified name.

Option	Description
-appendSameBlank	With this switch, Tapecopy first looks for the tape with tapename same as the source tape. If such tape is found, copied sessions are appended to that tape. Otherwise tapecopy looks for blank tape to use as destination.
-overwriteSameBlank -overwriteSameAnyBlank -overwriteSameBlankAny	These switches set the order of destination tape selection as per the individual switch.
-merge	After Tapecopy is complete, submits a merge job so that the details of copied tape can be available in the database.
[-waitForMerge]	Used with -merge, only. Use this if you want tapecopy to wait until the merge job is finished.
-MediaPool[media pool name]	Moves the copied tape to the specified media pool after Tapecopy is finished.
-export	Export destination tape if it is in a changer.
-offline	Take the destination tapes offline at the end of the tapecopy operation if it is in a changer.
-f [Tapename]	Use the specified tape name for format when formatting a blank tape for destination.

tapecopy Span Tape Options

By default, there is no timeout for span tape for either reading from a tape in the source group or writing to a tape in the destination group. There is, however, one span tape option:

-t[span tape timeout]

tapecopy Media Pool Options

[-mediapool <media pool name>]

This switch specifies an existing media pool name to which the destination tape must be moved after the tapecopy operation is completed.

This switch also enables tapecopy to select a tape from the scratch set to use for the source. When looking for scratch tapes, tapecopy looks for the tape in the scratch set of the specified media pool (or scratch set of all the media pools if global scratch set mode is enabled) before looking for a blank tape.

tapecopy Script Options

The script option is as follows:

-SCRIPT <script file name>

The script file format is as follows:

[QUERY]

QueryPreview=NO ;default is NO
QuerySessionType=session type
QuerySessionMethod=session method
QueryMediaPool=media pool
QueryBackupNode=backup node
QueryMasterJobID=job ID
QueryBackupOnOrBefore=mm/dd/yy hh:mm
QueryBackupOnOrAfter=mm/dd/yy hh:mm

[Source]

Group= ;specify
MediaName= ;default is empty
MediaID= ;default is empty
MediaSequence= ;default is empty
StartSessionNumber= ;default is empty
NumberOfSessionToBeCopied=;default is empty

[Destination]

Group= ;specify
MediaName= ;default is empty
MediaID= ;default is empty
MediaSequence= ;default is empty
MediaNameForFormat= ;default is empty
AppendSameBlank=NO ;default is NO
OverWriteSameBlank=NO ;default is NO
OverWriteSameBlankAny=NO ;default is NO
OverWriteSameAnyBlank=NO ;default is NO

[Option]

EraseMediaInSource=NO ;default is NO
Merge=NO ;default is NO
ExportDestinationMedia=NO ;default is NO
OfflineDestinationMedia=NO ;default is NO
SpanMediaTimeOut=9999 ;default is no time out

Appendix B: Acknowledgements

Portions of this product include software developed by third-party software providers. The following appendix provides information regarding this third-party software.

RSA Data Security, Inc. Acknowledgement

MD5C.C - RSA Data Security, Inc., MD5 message-digest algorithm.

Copyright (C) 1991-2, RSA Data Security, Inc. Created 1991. All rights reserved.

License to copy and use this software is granted provided that it is identified as the "RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing this software or this function.

License is also granted to make and use derivative works provided that such works are identified as "derived from the RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing the derived work.

RSA Data Security, Inc. makes no representations concerning either the merchantability of this software or the suitability of this software for any particular purpose. It is provided "as is" without express or implied warranty of any kind.

These notices must be retained in any copies of any part of this documentation and/or software.

Apache Acknowledgement

Portions of this product include software developed by the Apache Software Foundation (<http://www.apache.org/>). The Apache software is distributed in accordance with the following license agreement.

Version 2.0, January 2004

<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. **Grant of Copyright License.** Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.
3. **Grant of Patent License.** Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.
4. **Redistribution.** You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:
 - a. You must give any other recipients of the Work or Derivative Works a copy of this License; and
 - b. You must cause any modified files to carry prominent notices stating that You changed the files; and

- c. You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and
- d. If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

- 5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.
- 6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.
- 7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.

8. Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.
9. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

Copyright 2004 Computer Associates International, Inc.

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at:

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

OpenSSL Acknowledgement

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>). This product also includes libraries from an SSL implementation written by Eric Young (eay@cryptsoft.com)."

This product includes software written by Tim Hudson (tjh@cryptsoft.com).

LICENSE ISSUES

The OpenSSL toolkit stays under a dual license, i.e. both the conditions of the OpenSSL License and the original SSLeay license apply to the toolkit. See below for the actual license texts. Actually both licenses are BSD-style Open Source licenses. In case of any license issues related to OpenSSL please contact openssl-core@openssl.org.

OpenSSL License

Copyright (c) 1998-2003 The OpenSSL Project. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment:

"This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)"

4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact openssl-core@openssl.org.
5. Products derived from this software may not be called "OpenSSL" nor may "OpenSSL" appear in their names without prior written permission of the OpenSSL Project.
6. Redistributions of any form whatsoever must retain the following acknowledgment:

"This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>)"

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT ``AS IS'' AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com). This product includes software written by Tim Hudson (tjh@cryptsoft.com).

Original SSLeay License

Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com) All rights reserved.

This package is an SSL implementation writte by Eric Young (eay@cryptsoft.com). The implementation was written so as to conform with Netscapes SSL.

This library is free for commercial and non-commercial use as long as the following conditions are aheared to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, lhash, DES, etc., code; not just the SSL code. The SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson (tjh@cryptsoft.com).

Copyright remains Eric Young's, and as such any Copyright notices in the code are not to be removed. If this package is used in a product, Eric Young should be given attribution as the author of the parts of the library used. This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement:

"This product includes cryptographic software written by Eric Young (eay@cryptsoft.com)" The word 'cryptographic' can be left out if the routines from the library being used are not cryptographic related :-).

4. If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgement:

"This product includes software written by Tim Hudson (tjh@cryptsoft.com)"

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The licence and distribution terms for any publically available version or derivative of this code cannot be changed. i.e. this code cannot simply be copied and put under another distribution licence [including the GNU Public Licence.]

Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com) All rights reserved.

This package is an SSL implementation written by Eric Young (eay@cryptsoft.com). The implementation was written so as to conform with Netscapes SSL.

This library is free for commercial and non-commercial use as long as the following conditions are adhered to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, lhash, DES, etc., code; not just the SSL code. The SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson (tjh@cryptsoft.com).

Copyright remains Eric Young's, and as such any Copyright notices in the code are not to be removed. If this package is used in a product, Eric Young should be given attribution as the author of the parts of the library used. This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement:

"This product includes cryptographic software written by Eric Young (eay@cryptsoft.com)" The word 'cryptographic' can be left out if the routines from the library being used are not cryptographic related :-).

4. If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgement:

"This product includes software written by Tim Hudson (tjh@cryptsoft.com)"

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The licence and distribution terms for any publically available version or derivative of this code cannot be changed. i.e. this code cannot simply be copied and put under another distribution licence [including the GNU Public Licence.]

Index

A

- access tab • 64
- add a job • 152
- agent view tab • 116
- Apache Acknowledgement • 382
- assign media to a media pool • 190
- authentication services • 14
- automated media spanning • 171
- automatic performance tuning • 181

B

- bab command • 288
- back-end services • 20
- backup
 - actions • 89
 - custom • 84
 - of volumes • 62
 - options • 88
 - types • 83
 - verifying • 65

C

- ca_auth command • 289
- ca_backup command • 99, 292
- ca_dbmgr command • 191, 233, 314
- ca_devmgr command • 182, 318
- ca_log command • 268, 337
- ca_merge command • 256, 340
- ca_mmomgr command • 344
- ca_qmgr command • 347
- ca_restore command • 351
- ca_scan command • 262, 359
- cabd • 14
- cadiag • 269, 371
- cadiscovd • 14
- calendar views • 139
- caloggerd • 14
- camediad • 14
- caqd • 14
- caroot administrator profile • 21
- caroot user profile • 21
- caservd • 14
- CASVTUTL utility • 159
- cauthd • 14

- Circular Logging • 166
- clean tape heads • 180
- client tab • 233
- consoles • 151, 160
- control services • 288
- current stats tab • 232
- custom
 - backup schedule • 84
 - jobs • 125
 - schedules • 136
- cycle table • 137

D

- database
 - extension • 247
 - management • 225
 - manager • 226
 - optimization • 243
 - options • 97, 226
 - recording • 255
 - recovery • 234, 236
 - reports • 267
 - server • 14
 - service • 15
- databaseBackup log • 225, 273
- date tab • 231
- delete a job • 154
- destination tab • 115, 121
- detail tab • 162, 163
- device
 - groups • 169
 - manager • 159
 - toolbar • 169
- device drivers
 - Linux/390 • 159
- differential backups • 83
- directory
 - filter • 130
 - structure options • 122
- disaster recovery • 273
- discovery services • 14
- Disk Staging Option • 22, 23
 - architecture • 22
 - command line options • 293

- purge tool • 330
- query tool • 329
- configuration tasks • 68
- features • 24
- modify a schedule • 76
- multistreaming • 66
- overview • 24
- pause migration • 77
- policies • 72
 - copy policies • 70
 - policy configuration • 70, 71
 - purge policies • 71
- reports • 365
- restore data • 113
- run a backup job • 74, 75
- Smart Restore • 24
- SnapLock • 24
- dismount devices • 178, 179, 180
- domains • 16
- Duplicate Sessions dialog • 113

E

- eject media • 178
- Enabling devices • 178
- Equivalence list • 21
- Event Logger • 14
- Exception view • 139
- Expiration dates • 173
- Exporting media • 92

F

- File
 - filters • 132
 - menu • 20, 151
 - pattern filter • 128
 - retry options • 98
 - sharing options • 98
- Filter tab • 59, 62
- Filters • 88, 117, 125, 155
- Firewall communication
 - modifying the configuration file • 222
 - rpc.cfg • 222
- Format tab • 60
- Full backups • 83
- Functions • 18

G

- generic job type • 144
- GFS rotation • 87, 137, 186
- Global scratch set • 184
- Grouping devices • 169

H

- History of jobs • 154
- Home page • 18
- httpd daemon • 14

I

- Importing media • 92
- Incremental backups • 83

J

- Java • 14
- Job
 - add a • 152
 - delete a • 154
 - history • 154
 - level filters • 125
 - log • 156
 - modify a • 152
 - monitor a • 154
 - queue • 150
 - report • 267
 - reschedule • 153
 - scheduler • 14
 - scripts • 142
 - status • 155
 - Status Manager • 149
 - stop a • 153
 - summary • 156

L

- Large library support • 181
- Last result field • 156
- Loader • 14
- Loading media • 179
- Log
 - options • 96, 120, 254
 - overview • 265
 - tab • 261
- Logger E-mail Alert Messages • 28

M

- Machine-level filters • 88, 125

main product components • 18

Manager

- Backup • 53
- console • 151
- overview • 14
- selecting • 18

Manual disaster recovery • 274

Media

- importing and exporting • 92
- management • 201
- merging • 249
- pools
 - assign media • 190
 - maintenance • 188
 - manager • 187
- report • 267
- rules • 84, 121, 138, 255, 261
- server • 14
- service • 15
- session report • 267
- tab • 231
- view • 229

Media Management

- Back up and restore the primary MMO server • 213
- Back up the primary server • 215
- Demote an MMO primary server • 219
- Reinstitute the MMO primary server • 219
- Restore the MMO primary server data • 216

Merge

- manager • 250
- options • 252

Merge media • 249

Method tab • 82

Minimum save set • 187

Missed targets • 137

MMO administrator • 191

Modifying jobs • 152

Monitoring services • 288

mount and dismount • 178

Mounting devices • 178, 179, 180

Multiple code pages

- about • 57
- applying • 58, 114

N

NDS option tab • 61

Node

level options • 58

option tab • 59

O

Object information tab • 58, 62, 117, 252, 259

object status • 201

Offline library • 180

OpenSSL Acknowledgement • 386

Options • 290

P

Password • 63, 90

Priority • 60, 64

Property tab • 231

Q

Queue service • 15

R

Repeating interval • 84

Report

- manager • 263
- object • 199

Reports • 200, 225, 267

Rescheduling jobs • 153

restore

- data • 105
- methods • 112
- options • 117

Retention media option • 177

Retention periods • 187

Retry • 98, 137

Rotation • 87, 137, 186, 198

rpc.cfg

- about • 222
- modification • 222

Run command • 95, 119, 253, 260

Run Now option • 141

S

Save set • 183, 186

Scan Manager • 257

Schedule

- jobs • 140
- methods • 83
- object • 196
- option • 141
- tab • 82

- Scheduler • 14
- Scratch set • 184, 186
- Scripts • 142
- Security • 59, 117
- Serial number information • 185
- serial numbers and bar codes • 185
- server and group configuration • 149
- Session password • 63, 90
- Sharing
 - options • 98
 - slots • 171
- Simple rotation • 87, 186
- Slot
 - integrity • 210
 - sharing • 171
- Source tab • 54
- Span media options • 86
- Stopping jobs • 153
- Submit
 - button • 256
 - jobs on hold • 142
- Summary tab • 161, 229
- Switches • 290

T

- Tape
 - clean heads option • 180
 - create a VCD • 206
 - special tape movement • 211
- Tape Management
 - tasks
 - establish vault storage system • 192
- tapecopy command • 175, 375
- Targets, missed • 137
- Total tab • 233

U

- Unload media • 179
- Use Rotation Scheme • 137
- User Profile Manager • 22

V

- Vault Management overview • 192
- Vault object • 200
- Vaults • 192, 196, 197, 201, 204, 206
 - Simulating the Vault Cycle • 202
 - Starting the vault cycle process • 202
 - Vault cycle reports • 203

- VCD • 192, 197, 206
- Verification • 65, 92
- Views • 229
- Volume
 - backup options • 62
 - option tab • 64