SnapManager® 6.0.4 for Microsoft® Exchange Installation and Administration Guide

NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089 USA
Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 463-8277
Web: www.netapp.com
Feedback: doccomments@netapp.com

Part number: 215-07021_A0
September 5 2012

Contents

SnapManager overview 15
SnapManager for Exchange limitations 15
Where you install and run SnapManager 15
How SnapManager works with other system components 16
How Snapshot technology is used by SnapManager 17
How SnapManager and SnapDrive work together 17
How SnapManager uses VSS 17
Management of Snapshot copies created with VSS 18
Data ONTAP VSS Hardware Provider requirement 18
Comparison of methods for creating restorable backups 18
SnapManager and conventional backup processes 19
Types of external backup to use with SnapManager 19
SnapManager snap-in 19
The SnapManager graphical user interface 20
The SnapManager command-line interface 20
Preparation for installing or upgrading SnapManager 21
Configuring the Windows firewall 22
Where to install SnapManager 23
SnapManager requirements in a DAG environment 23
Backing up system resources and data 23
SnapManager licensing options 23
Applying a SnapManager license to the Exchange server 24
Applying a per-storage-system license to the storage system 25
Exchange permission level required for SnapManager 25
SnapManager for Exchange Service identity account requirements 26
Account permissions for the report directory share 26
Windows host system requirements 27
If your storage system has multiple IP addresses 28
Storage system requirements 29
Remote administration server requirements 29
Remote verification server requirements 30
Exchange Server 2003 files to back up with a Windows backup utility 31

Exchange Server 2007 files to back up with a Windows backup utility 31
SnapManager installation and upgrade 33
Compatibility of SnapManager builds with Windows Server 33
Installation of SnapManager on a stand-alone Windows host system 33
Comparison of installation modes for a stand-alone Windows host system 33
Installing SnapManager in interactive mode 34
Installing SnapManager in unattended mode 35
Displaying the SnapManager software license agreement 37
Examples of unattended installation 38
SnapManager installation in a Windows cluster 38
Requirements for installing SnapManager in a Windows cluster 38
Installing SnapManager after creating a new Windows cluster 39
Installing SnapManager in an existing Windows cluster 40
Preparing to upgrade SnapManager 41
Comparison of interactive and unattended upgrade modes 42
SnapManager upgrade paths 42
Upgrade considerations for fractional space reserve monitoring 42
Upgrading SnapManager in interactive mode 43
Upgrading SnapManager in unattended mode 43
Example of an unattended upgrade 44
Updating legacy scheduled jobs 45
Prerequisites for uninstalling SnapManager 45
Comparison of interactive versus unattended uninstallation 46
Moving Exchange data to a local disk 46
Uninstalling SnapManager in interactive mode 47
Uninstalling SnapManager in unattended mode 47
Examples of unattended uninstallation 49
SnapManager reinstallation with or without uninstallation 49
Comparison of interactive and unattended reinstallation 50
Reinstalling SnapManager in interactive mode 50
Reinstalling SnapManager in unattended mode 51
Example of an unattended reinstallation 52
When you start SnapManager for the first time 53
What SnapManager verifies at startup 53
Starting SnapManager and connecting to the default server 53

Starting SnapManager and connecting to the CCR server 55
SnapManager Dashboard view 55
Scheduled jobs 57
Database Availability Group 57
Database migration considerations 58
Exchange configuration requirements 58
Deletion of Storage Groups from clustered systems 60
Rules for Exchange Server storage groups and databases enforced by the Configuration wizard 60
Recommended Exchange Server storage groups and databases configurations 61
Exchange Server storage groups and databases configurations to avoid 61
Exchange message tracking in an MSCS configuration 62
NTFS volume mount points 62
Limitations of NTFS volume mountpoints 62
SnapManager support for volume mountpoints 63
Drive letter limitations and individual database restoration 63
Mounted volume restrictions with SnapManager 64
Mounted volume naming conventions with SnapManager 64
How mounted volumes are shown in SnapManager 64
Transaction log archiving 66
NTFS hard links 66
Why SnapManager uses NTFS hard links for transaction log archiving 66
Support for multiple SnapInfo directories 67
Example Exchange configurations supported with SnapManager 68
Example: single Storage Group and individual database restoration not required 68
Example: single Storage Group and individual database restoration required 69
Example for multiple storage groups and individual database restoration required 69
Example: multiple Storage Groups and individual database restoration not required 70
Creation of LUNs on qtrees 71
Storage system volume and LUN planning 71
Configuration and migration of Exchange data using SnapManager ... 72

The SnapManager Configuration wizard 73
What the SnapManager Configuration wizard does 73
When to use the SnapManager Configuration wizard 74
Settings configurable only with the SnapManager Configuration wizard 75
Placement of Exchange and SnapManager components 76
Viewing the placement of Exchange and SnapManager components 77
Exchange Server storage groups and databases migration and configuration considerations 78
Migrating and configuring your Exchange Server storage groups and databases using the SnapManager Configuration wizard 78
Moving a Storage Group to a LUN 80
Moving an individual database to a LUN 81
Changing the location of a Storage Group or database 81
Viewing the full path for database files 82
Dataset protection policy 83
Migration of transport database paths for Exchange Server 2010 83
How to move transaction logs using the SnapManager Configuration wizard 84
Moving transaction logs to a LUN 84
Changing the location of transaction log files 84
Viewing the full path for transaction log files 85
Configuration using the SnapManager control file 86
Configuration of SnapManager in a DAG using a control file 86
Importing Exchange Server configuration information using the control file 87
Exporting Exchange Server configuration information using the control file 88
Migration of LCR-enabled databases 89
Prerequisites for configuring an LCR-enabled Exchange server 89
Comparison of moving production and LCR replica storage groups 90
Database seeding in an LCR 90
Migration of CCR-enabled databases 91
Prerequisites for configuring a CCR-enabled Storage Group 91
Considerations for configuring SnapManager in a CCR-enabled Storage Group 92
Configuring SnapManager in a CCR-enabled Storage Group 92
Migration of Exchange Server 2010 mailbox databases 94
Prerequisites for migrating Exchange Server 2010 mailbox databases 94
Available LUNs in a Database Availability Group 94
Migrating Exchange Server 2010 mailbox databases 94

Guidelines for migrating to mount points for LUN mapping in Exchange 2007 ... 96
Scenario: migrating an existing configuration from using drive letters to using mountpoints for LUN mapping ... 96
SnapManager Backup overview ... 99
How SnapManager Backup works ... 99
Copy backup support ... 99
The SnapInfo directory ... 99
SnapManager backup sets ... 100
Exchange Storage Group/database sets ... 101
SnapManager minimum unit of backup ... 101
SnapManager naming-convention options ... 101
Backup process in a Windows Server 2003 or Windows Server 2008 environment ... 102
Typical VSS backup process ... 102
Why a transaction log backup might contain two Snapshot copies ... 103
SnapManager Snapshot copy naming conventions ... 104
When to run a SnapManager backup ... 105
How SnapManager checks database integrity in backup sets ... 106
LUN requirements for verifying databases in a backup set ... 106
Database verification load management ... 107
Backup verification status reporting ... 109
Where to run database and transaction log integrity verification ... 109
When to verify the databases in a backup set ... 109
Backup set retention ... 110
Maximum number of Snapshot copies per volume ... 110
Ways to delete Snapshot copies ... 110
Automatic deletion of Snapshot copies ... 111
Transaction log management ... 111
Option to back up transaction logs that Exchange will truncate ... 114
Exchange System Manager in a SnapManager environment ... 114
Displaying the time of the last full backup ... 115
Exchange page zeroing and deleted item retention in Exchange 2007 and earlier ... 115
Database backup using SnapManager ... 117
Exchange storage groups and databases display ... 117

Exchange databases display in a DAG 117
Decisions to make before performing a SnapManager backup 117
Backup planning for databases in Database Availability Groups 120
Backing up using the Backup wizard 121
Backing up using the Backup and Verify window 122
Using Database Filter to display the databases to back up in a Database Availability Group 124
Activation Preference Number of a mailbox database in Exchange Server 2010 126
LCR-enabled database backups 126
CCR-enabled database backups 126
Considerations before backing up a CCR replica database and production database 127
Backing up and verifying a CCR replica database and production database 127
Creating a secondary backup on a remote CCR node using the Backup wizard 129
Creating a secondary backup on a remote CCR node using the Backup and Verify window 130
How SnapManager creates a secondary backup on a remote CCR node 130
Reasons that a SnapManager backup might fail 131
Problem: cluster failover during backup 131
Problem: Snapshot copy limit reached 131
Problem: SnapInfo directory being accessed 131
Problem: SnapInfo directory out of space 132
Problem: data does not match 132
Problem: busy Snapshot copy 132
Problem: Snapshot copy already exists 132
Problem: out of disk space 132
Problem: SnapManager server initialization failed 132
Verifying database backups 132
Decisions to make before database verification 133
Starting or scheduling database verification 134
Backup management groups 135
Backup management group assignments 135
Example using backup management groups 135

Assigning a backup set to a different backup management group 136
Frequent Recovery Point backup operation 137
How the Frequent Recovery Point feature works 137
Frequent Recovery Point backup operations 137
Frequent Recovery Point backup operation on clustered configurations 137
Frequent Recovery Point backup operation in a DAG 138
Verification of Frequent Recovery Point backup copies 138
Deletion of Frequent Recovery Point backup copies 138
Frequent Recovery Point backup reports 138
Performing a Frequent Recovery Point backup operation 138
Advantages of using the remote additional copy backup feature 140
Configuring gapless DAG backups using remote additional copy 141
PowerShell new-backup cmdlet parameters for DAG backups 141
control.xml file settings for remote additional copy backup 143
Types of database restore operations using SnapManager 145
When to choose SnapManager Restore to recover Exchange 2010 mailbox databases 145
How SnapManager Restore works 146
How to choose the type of restore operation to perform 146
Types of SnapManager Restore operations 147
Snapshot copies created during a restore process 147
Methods that can decrease restore process time 148
Transaction log sequence verification options 148
LUN Clone Split Restore method 149
SnapManager Restore in a Windows cluster 149
SnapManager Restore in a live Exchange virtual server cluster 150
Guidelines for using SnapManager Restore 150
How to choose the type of restore operation to perform 150
Guidelines for restoring from a SnapManager backup copy 151
Restore from a SnapManager backup copy 151
Decisions to make before restoring from a SnapManager backup copy 152
Restoring databases using the Restore wizard 153
Restoring databases using the Restore window 154
Restoring data to a specified Frequent Recovery Point 156
Restoring Exchange 2010 databases in a DAG 158
Backups available for restore in a DAG 158

Restoring a backup copy from one Exchange server to another in a DAG 158
Reseed after a restore operation in a DAG 159
Reseeding a database using a PowerShell script 159
Restore from an LCR-enabled Storage Group 161
Restoring a backup copy from an LCR replica Storage Group 162
What to do if corruption occurs in an LCR-enabled Storage Group 163
Restore from a CCR-enabled Storage Group 163
Recovery Storage Groups 167
Limitations of using a Recovery Storage Group 167
Restoring a database to a Recovery Storage Group in Exchange Server 2007 167
Restoring a database to a Recovery Storage Group in Exchange Server 2003 172
Restoring multiple databases to a Recovery Storage Group 175
Restoring an unverified backup copy to a Recovery Storage Group 176
Restore of backups created at different Exchange server locations 176
Restoring backup copies that were created on other Exchange servers in Exchange 2007 176
Restoring backup sets from unmanaged media 177
Up-to-the-minute restore from an archive backup 180
Recovery Database 180
Limitations of using a Recovery Database 180
Restoring a mailbox database to the Recovery Database in Exchange Server 2010 181
When to delete a Recovery Database 182
Troubleshooting restore operations 182
Mailbox restore using Single Mailbox Recovery 184
Recovering mailbox data 184
Cleaning up after Single Mailbox Recovery 185
Deletion of Snapshot copies 186
Criteria for deleting backups 186
Automatic deletion of Snapshot copies 187
Explicit deletion of Snapshot copies 188
Option to retain up-to-the-minute restore ability 188
Explicitly deleting individual backup copies 189
Explicitly deleting backup sets or SnapInfo Snapshot copies 190

Explicitly deleting Snapshot copies created by SnapManager Restore 192
Problem deleting backups due to busy Snapshot copy error 193
How SnapManager uses SnapMirror 194
Volume replication using SnapMirror 194
Where to find more information about configuring and using SnapMirror 194
Requirements for using SnapMirror with SnapManager 195
How SnapManager uses SnapMirror and SnapDrive 195
How SnapMirror replication works 195
Integrity verification on SnapMirror destination volumes 196
Selecting the SnapMirror destination volumes for verification 196
Requirements to run destination volume integrity verification 197
Troubleshooting integrity verification failure on SnapMirror destination volumes 197
Types of destination volume integrity verification 198
Backup with verification 199
Integrity verification for test restore operations 199
Integrity verification for a restore process 200
Remote destination volume integrity verification 200
Deferred integrity verification 200
Concurrent backup verification 201
Managing integrity verification jobs 202
Disaster recovery using SnapManager 204
SnapManager backup archiving 205
Why organizations archive data 205
Guidelines for archiving SnapManager backups 205
Methods of archiving SnapManager backups 206
Archives created with NDMP or the dump command 206
Evaluation of the NDMP and dump command method of archiving 207
Example using NDMP or dump to archive SnapManager backups 207
Archives created using a Windows backup utility 209
Evaluation of the Windows backup utility method of archiving 210
Example: Using a Windows backup utility to archive SnapManager backups 210
How to use the Exchange Backup Agent to archive Exchange backups 211
Evaluation of the Exchange Backup Agent method of archiving 212

Example: Using Exchange Backup Agent to archive Exchange backup copies ... 212
If you use a centralized backup model ... 212
Automatic backup archiving using the Run Command After Operation feature ... 213
Command arguments supported by the Run Command After Operation feature ... 213
Specifying the command to be run by the Run Command After Operation feature ... 214
Enabling the launch of SnapManager scripts from a UNC path ... 215
SnapManager reports and the report directory ... 217
SnapManager reports in a Database Availability Group (DAG) ... 217
Reasons to change the report directory location ... 217
Changing the SnapManager report directory ... 218
Locating the report directory in a Windows cluster ... 218
Viewing SnapManager reports ... 219
Printing SnapManager reports ... 219
Deleting SnapManager reports ... 219
Dataset and SnapVault integration ... 221
Dataset concepts ... 221
Available functionalities of dataset and SnapVault integration with SnapManager ... 222
Dataset and SnapVault integration with SnapManager ... 223
Prerequisites for dataset and SnapVault integration with SnapManager ... 224
Limitations of dataset and SnapVault integration with SnapManager ... 224
Dataset configuration ... 225
Creating a dataset using SnapManager ... 226
Editing a dataset using Protection Manager ... 227
SnapVault relationships ... 227
Local backup protection using dataset and SnapVault integration ... 228
Information used to create remote backups ... 228
Remote backup retention ... 229
Deferred database integrity verification with SnapVault ... 229
Restoring from a remote backup ... 230
SnapManager application settings configuration ... 231
Where to access SnapManager application settings ... 231
Adding Exchange servers to be managed ... 232

Enabling database migration back to local disks 233
Disabling database migration back to local disks 233
Considerations for selecting the database verification server 233
Configuring the verification server 234
Remote verification prerequisites 235
Selecting the Snapshot copy access method for database verification 236
Database verification throttling 237
How database verification throttling works 237
Database verification throttling options 238
Calculating the verification throttling sleep interval 238
Configuring database verification throttling 239
Throttling entries in the SnapManager backup and verification report 240
Verification override entry in the SnapManager restore report 240
Impact of database verification on performance 240
Database verification override during restore operation 241
Configuring the database verification override option 241
Verification override entry in the SnapManager restore report 242
Verification override entry in the SnapManager restore report 242
Configuring default settings for the Run Command After Operation option 242
Fractional space reservation 243
What can happen with a fractional-space-reserved volume 244
Fractional space reservation policies 244
Fractional space reservation policies to manage Exchange data 247
Viewing current fractional space reservation data for a LUN 249
Event notification options 251
Configuring automatic event notification settings 252
SnapManager control file XML schema 255
Storage layout settings XML schema 255
Notification settings XML schema 257
Verification settings XML schema 258
Report directory settings XML schema 259
Backup settings XML schema 260
SnapMirror relationship settings XML schema 260
SnapManager command-line reference 262
Guidelines for using the SnapManager for Exchange PowerShell command-line tool 262

Launching SnapManager for Exchange PowerShell 262
new-backup 262
verify-backup 271
delete-backup 276
get-backup 278
get-mirrors 280
resync-mirrors 282
release-mirrors 283
break-mirrors 285
restore-backup 286
Get-JobStatus 293
Change-JobPriority 294
Cancel-Job 295
exec-bc 296
Export-config 299
Import-config 300
Copyright information 303
Trademark information 304
How to send your comments 305
Index 306

SnapManager overview

SnapManager provides you an integrated data management solution for Microsoft Exchange that enhances the availability, scalability, and reliability of Exchange databases. SnapManager provides rapid online backup and restoration of databases, along with local or remote backup set mirroring for disaster recovery.

SnapManager uses online Snapshot technology that is part of Data ONTAP and integrates Exchange backup and restore APIs and Volume Shadow Copy Service (VSS). SnapManager uses SnapMirror to support disaster recovery.

SnapManager provides the following data management capabilities:

- Migrating Exchange databases and transaction logs to LUNs on storage systems
- Backing up Exchange databases and transaction logs from LUNs on storage systems
- Verifying the backed-up Exchange databases and transaction logs
- Managing backup sets
- Archiving backup sets
- Restoring Exchange databases and transaction logs from previously created backup sets

SnapManager for Exchange limitations

SnapManager for Exchange cannot be used with some software versions, features, databases, and server editions.

- SnapDrive versions prior to 6.1
- Microsoft Windows 2000
- Creation or restoration of backups of Exchange databases on third-party storage except for databases on third-party storage arrays that provide storage for a V-Series system
- Backup and restoration of Microsoft Exchange 5.5 and Exchange 2000 databases
- Windows Server 2003 or the Windows Server 2008 IA-64 editions

Where you install and run SnapManager

You can use SnapManager with configurations having multiple servers. You can perform local administration, remote administration, and remote verification.

SnapManager provides the following capabilities:

- Local administration
  You can install SnapManager on the same Windows host system as your Exchange server.
- Remote administration

You can install and run SnapManager on a remote computer and perform any task that you can perform on a locally installed SnapManager system.

- Remote verification

  You can perform remote database verification from a remote administration server that is configured with SnapDrive and Exchange server. Remote verification offloads the CPU-intensive database verification operations that can affect the performance of your production Exchange server.

How SnapManager works with other system components

SnapManager coordinates with SnapDrive and Volume Shadow Copy Service (VSS) and uses Snapshot technology to create database backups.

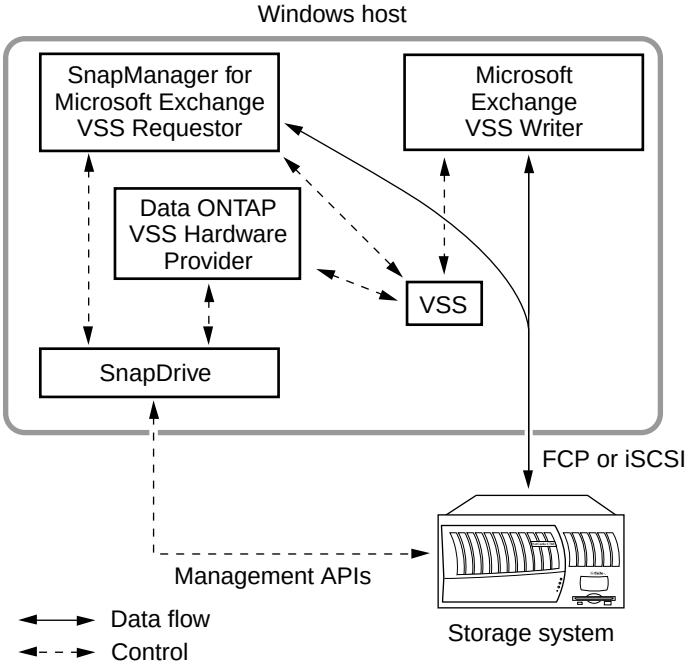

SnapManager on Windows Server 2003 and 2008

VSS coordinates between servers, backup applications, and storage management software to help SnapManager to create and manage Snapshot copies. SnapDrive provides the underlying layer of support for SnapManager to help you to manage resources on the storage system in the Windows environment.

How Snapshot technology is used by SnapManager

A Snapshot copy is a point-in-time, real-time, online, and read-only copy of a LUN stored on a volume. SnapManager Backup uses the Snapshot technology to create copies of Exchange databases that reside in these LUNs.

Data ONTAP software allows a maximum of 255 Snapshot copies per volume. To avoid reaching the limit, you should delete SnapManager backups that are no longer needed, because SnapManager backups automatically create Snapshot copies. Note that the number of Snapshot copies on a volume can be greater than the number of SnapManager backups being retained. For example, if a single volume contains both the SnapInfo directory and the Exchange databases, each SnapManager backup generates two Snapshot copies on that volume. Therefore, the number of Snapshot copies is double the number of backups on the same volume.

Note:

For more information about Snapshot technology, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

How SnapManager and SnapDrive work together

SnapDrive provides the underlying layer of support for SnapManager by working with the NTFS Windows file system, Windows Volume Manager, and LUNs to help you to manage resources on the storage system in the Windows environment.

Note: Use SnapManager for all backup-related operations. Use SnapDrive only to create and manage the LUNs and storage system volumes that contain your Exchange data.

How SnapManager uses VSS

SnapManager is a Microsoft Volume Shadow Copy Service (VSS) requestor, which means that it uses the VSS subsystem to initiate backups.

The Data ONTAP VSS Hardware Provider integrates the SnapDrive service and storage systems running Data ONTAP into the VSS framework. For more information about VSS, see your SnapDrive documentation.

Management of Snapshot copies created with VSS

You must manage the Snapshot copies that SnapManager creates using SnapManager or SnapDrive. Do not use the VSS administration tool `vssadmin` to manage Snapshot copies that are created by SnapManager using VSS.

Data ONTAP VSS Hardware Provider requirement

You must have the Data ONTAP VSS Hardware Provider installed for SnapManager to function properly. Data ONTAP VSS Hardware Provider integrates the SnapDrive service and storage systems running Data ONTAP into the VSS framework.

The Data ONTAP VSS Hardware Provider is now included with SnapDrive 6.0 and later and does not need to be installed separately.

Comparison of methods for creating restorable backups

You can create Snapshot copy backups in a variety of ways. It is important to understand when each of these methods can produce a restorable image and when each cannot. For example, the SnapManager backup method might create several Snapshot copies, depending on how you configure your Exchange database.

Note: Always use SnapManager to manage SnapManager backups. Do not use the storage system console, OnCommand System Manager, FilerView, SnapDrive, a Windows backup utility that is not Exchange aware, or VSS to manage SnapManager backups, to avoid an inconsistent image of the database that is unusable for a restore process.

For some system states, other Snapshot copy backup tools can create restorable backups. The following table illustrates when each of the Snapshot copy backup tools can create restorable backups.

Tool used to create Snapshot copy or backup	**System state**		
	Exchange online	**All Exchange databases unmounted**	**All LUNs disconnected**
SnapManager Backup	OK	NA	NA
SnapDrive	Invalid	OK	NA
Exchange-aware Windows backup utility	OK (copy backup only)	NA	NA
Non-Exchange-aware backup utility	Invalid	OK (copy backup only)	NA
Storage system tools	Invalid	Invalid	OK

SnapManager and conventional backup processes

SnapManager complements conventional backup processes. However, some backup processes interfere with SnapManager's ability to back up and restore correctly. It is important that you understand which of your current backup tools and processes you need to modify or eliminate when you implement SnapManager.

SnapManager backups reside on disks. To archive SnapManager backups to another storage medium, such as tape, you can use the same tools and processes you currently employ.

Note: SnapManager backups cannot replace periodic disaster recovery measures, such as archiving to tape or other offline media.

Types of external backup to use with SnapManager

SnapManager backups replace the backups that are created by Exchange-aware backup tools (such as Windows NTBackup). If you need to create a backup using an Exchange-aware backup tool on a system with SnapManager installed, you must use only copy or differential backup.

Attention: If you create a backup of a type other than copy or differential backup using an Exchange-aware backup tool, your previous backups become unusable for up-to-the-minute restore operations.

The following table describes the backup types.

Backup type	Description	Result
Copy	Backs up all selected files; does not clear archive bit	OK. SnapManager backups unaffected
Normal (full)	Backs up all selected files and clears archive bit	Do not use. Truncates transaction logs; disables SnapManager up-to-the- minute restores
Differential	Backs up all selected files with archive bit set	OK. SnapManager backups unaffected
Incremental	Backs up all selected files with archive bit set and clears archive bit	Do not use. Truncates transaction logs; disables SnapManager up-to-the- minute restores

SnapManager snap-in

The SnapManager snap-in is based on Microsoft Management Console 3.0 (MMC). The SnapManager snap-in allows you to manage the SnapManager application from Microsoft Management Console.

The following components of SnapManager snap-in help you to perform all operations:

- Scope pane
  This is the left pane. It lists SnapManager for Exchange instances.
- Actions pane
  This is the right pane. It displays all of the actions that you can perform, based on the instance that you select in the Scope pane.
- Results pane
  This is the center pane. It displays details of the type of instance that you select in the Scope pane.

The Dashboard view appears in the center pane when you select any server or server instance. The Dashboard view enables you to:

- View the Exchange server configuration
- Add new Exchange servers

The SnapManager graphical user interface

SnapManager 6.0 for Microsoft Exchange has a graphical user interface based on Microsoft Management Console (MMC). You can use it with other MMC snap-ins to create a single console for managing Exchange and storage system components.

The new graphical interface enables you to do the following tasks:

- Manage and administer multiple instances of SnapManager for Exchange successfully.
- Manage backup and restore operations of Exchange database files and transaction log files on LUNs.
- Schedule backups and verify the integrity of databases in SnapManager backup sets.
- Administer SnapManager for Exchange on another Exchange server computer on the network.
- Configure Exchange database, transaction logs, and Simple Mail Transfer Protocol (SMTP) queue locations for SnapManager backup and restore operations.

The SnapManager command-line interface

SnapManager Microsoft Exchange supports a SnapManager command-line functionality called cmdlet, through SnapManager for Exchange PowerShell. Cmdlets enable you to perform almost all of the jobs that you can perform using the GUI.

Preparation for installing or upgrading SnapManager

Before you install or upgrade to SnapManager 6.0 for Microsoft Exchange, you might need to configure Windows host systems, hardware, and software components; back up your databases; and check licensing and account requirements.

For the most current list of system requirements, see the SnapManager for Microsoft Exchange Description page on the NetApp Support Site on the Web site.

For information about compatible versions of SnapManager, SnapDrive, and Data ONTAP, see the SnapManager and SnapDrive Compatibility Matrix.

You might need to install, upgrade, or configure the following components:

- Data ONTAP
- iSCSI and FC protocols
- SnapManager license
- SnapRestore license
- SnapMirror license
- FlexClone license
- Microsoft Windows operating system
- Microsoft Windows hotfixes
- Microsoft Exchange
- SnapDrive
- SnapDrive preferred IP address (if your storage system has multiple IP addresses)

  Note: If you do not configure a SnapDrive preferred IP address for a storage system that has multiple IP addresses, SnapDrive times out when attempting to create Snapshot copies simultaneously for more than one LUN on a storage system. For details, see the *SnapDrive for Windows Installation and Administration Guide* for your version of SnapDrive.

Installing or upgrading SnapManager involves performing the following tasks:

1. Backing up system resources and databases
2. Determining whether you want to use per server SnapManager licensing or per storage system SnapManager licensing
3. Ensuring that the SnapManager COM server identity account on each Exchange server to be used by SnapManager has Exchange Administrator or higher permissions
4. Configuring or upgrading your storage system according to the requirements of SnapManager and SnapDrive

5. If you upgrade SnapManager and underlying SnapDrive or Microsoft iSCSI initiator versions, removing the iSCSI dependency with respect to SnapManager
6. Noting whether your storage system has multiple IP addresses
7. Configuring or upgrading your Windows host systems to meet the requirements for SnapDrive and SnapManager
8. Ensuring that the TCP port 810 is open and the firewall is correctly configured

Related information

NetApp Support Site on the Web
SnapManager and SnapDrive Compatibility Matrix

Configuring the Windows firewall

If you are using the Windows firewall and Exchange 2010 Database Availability Groups, you might need to adjust the configuration of the Windows firewall.

Steps

1. Open the Windows Control Panel and choose **Firewall.**
2. Click **Advanced Settings** and then select **Inbound Rules**.
3. Select **New Rule**.
4. Select **Program** and click **Next**.
5. Browse to `%ProgramFiles%\NetApp\SnapManager for Exchange \SnapMgrService.exe` and click **Next**.
6. Choose **Allow the connection** and click **Next**.
7. Select **Domain only** and click **Next**.
8. Enter a value for **Name for rule** (for example, **`SME Inbound connection`**) and click **Finish**.
9. After you are back in **Inbound rules**, click **New Rule**.
10. Select **Port** and click **Next**.
11. Select **TCP** and enter 810 for **Specific Port**; click **Next**.
12. Choose **Allow the Connection** and click **Next**.
13. Select **Domain only** and click **Next**.
14. Enter **Name for rule** (for example,

    `SME DAG Port`
) and click **Finish**.

15. Repeat Steps 1 through 14 on all nodes in the Exchange 2010 DAG.

Where to install SnapManager

In a basic configuration, SnapManager is installed on the same Windows host system as Exchange. You can also install SnapManager on one or more remote Windows hosts for remote administration of the Exchange computer or for remote verification of the databases in backup sets.

SnapManager requirements in a DAG environment

You must install SnapManager and SnapDrive for Windows on all member servers of the Database Availability Group (DAG) to ensure proper working of SnapManager.

For Exchange Server 2010, you should use the same drive letters or mount points for the Exchange data LUNs on all the nodes of a DAG.

You should install SnapDrive for Windows on all nodes, and note that if you do not install SnapManager for Exchange on all nodes, you are not able to use it to migrate databases using the configuration wizard.

You can use SnapManager for Exchange only on the Exchange server that has it installed.

Backing up system resources and data

Before you install SnapManager, back up your system resources and data using Windows NTBackup or another industry standard backup utility to prevent any data loss during configuration.

Steps

1. Back up the operating system installation on the Exchange server. This includes backing up all of the server system state which consists of the registry, the boot files, and the COM+ class registry.
2. Back up the data on the local drives on the Exchange server.
3. Back up the boot and system drives.
4. Use your backup utility to create and maintain a current emergency repair disk (ERD).

SnapManager licensing options

SnapManager supports two licensing options: Per Server and Per Storage system licensing.

A Per Server SnapManager license is a 14-character license code for a specific Exchange server. When you use Per Server licensing, you do not require a SnapManager license on the storage system. Instead, you apply a SnapManager server-side license on every Exchange server. You can specify a server-side license while you are installing SnapManager, or you can defer this activity until you

have completed SnapManager installation. After you install SnapManager, you can apply the license from **Help > About SnapManager**.

If a SnapManager license is not enabled on the Exchange server, you must enable a 7-character SnapManager Per Storage system license code directly on the storage system, using the `license add` command. If no server license is detected, SnapManager checks the storage system for a SnapManager license whenever a SnapManager operation starts. If the SnapManager license is not enabled on the storage system either, the SnapManager operation fails and logs an error in the Windows event log.

Note: With storage system SnapManager licensing, if you use SnapMirror with SnapManager, both SnapMirror and SnapManager must be licensed on both the source and target storage systems.

Applying a SnapManager license to the Exchange server

You can add a Per Server license to your Exchange server. When you use Per Server licensing, you do not require a SnapManager license on the storage system. Instead, you apply a SnapManager server-side license to the Exchange server.

Before you begin

You must have a SnapManager Per Server license key.

Steps

1. Run the SnapManager .exe file.

   The SnapManager InstallShield wizard appears.

2. In the **Welcome** window, click **Next**.

   The License Agreement window appears.

3. If you accept the terms of the agreement, select **I accept the terms in the license agreement** and click **Next**.

   The Customer Information window appears.

4. In the **Customer Information** window, specify the user name and the name of your organization.
5. Under **License Type**, select **Per Server**.
6. In the **License Key** box, enter the license key for your server-side license.
7. If you do not have a license key, you can leave the License Key box empty for now and enter your server-side license key later from the SnapManager interface. In the **Actions** pane, click **License settings**.
8. In the **Per Server License Key** window, specify your license key.
9. Click **OK**.

Applying a per-storage-system license to the storage system

You can add per-storage-system license to your storage system. If a SnapManager operation starts and no server license is detected, SnapManager checks the storage system for a SnapManager license. If the SnapManager license is not enabled on the storage system either, the SnapManager operation fails and logs an error in the Windows event log.

Before you begin

- You must have a SnapManager per-storage-system license key.
- If you use SnapMirror with SnapManager, both SnapMirror and SnapManager must be licensed on both the source and target storage systems.

Steps

1. Run the SnapManager .exe file.
2. In the **Welcome** window, click **Next**.
3. If you accept the terms of the agreement, select **I accept the terms in the license agreement** and then click **Next**.
4. In the **Customer Information** window, specify your user name and the name of your organization.
5. Under **License Type**, select **Per Storage System**.
6. Continue with the instructions in the **InstallShield** wizard to apply the license.

Exchange permission level required for SnapManager

If you need a granular control for the SnapManager for Exchange Service, you can apply Exchange-specific permissions to the object.

The minimum Exchange-specific permissions required for SnapManager are as follows:

- Read
- Execute
- Read Permissions
- List Contents
- Read Properties
- Administer Information Store
- View Information Store Status

To launch or access the SnapManager for Exchange Service from a different domain the user account, or the group that you are configuring must also be a member of the Exchange server's local administrators group.

The local administrator group permissions enable you to launch the SnapManager for Exchange Service account that needs to be part of the Exchange server administrator group.

Ensure that there are no firewall restrictions between the remote and local servers.

SnapManager for Exchange Service identity account requirements

The SnapManager for Exchange Service identity account must have proper Exchange permissions for SnapManager to function. SnapManager uses a Windows NT Service (SnapManagerService) that hosts a number of services, such as SME Service, FSR Service, and SME Business Continuance Service.

There are three predetermined permission levels for Exchange:

- Exchange View Only
- Exchange Administrator
- Exchange Full Administrator

The minimum level required for SnapManager is Exchange Administrator.

Use one of these methods to meet the requirement:

- Assign Exchange Administrator permissions to the organization object.
- Create an Exchange administrative group and assign the correct permissions to that group.

  In Exchange Server 2003, right-clicking an object in the Exchange System Manager and selecting **Delegate Control** opens the Exchange Administration Delegation wizard, which enables you to change that object's permissions.

  In Exchange Server 2007, select **Organization node** in the Exchange Management Console and then click **Add Exchange Administrator wizard** in the Actions pane to change the object's permissions.

Exchange Server 2010 uses the Role Based Access Control (RBAC) permissions model. You can add users or members to a role group through the Exchange Management Console. From the **Toolbox** of the Exchange Management Console, open the **Role Based Access Control (RBAC) User Editor** and add users to management role groups. Add accounts to **Organization Management**.

Account permissions for the report directory share

The first time SnapManager is launched, it creates a shared report directory, `SMEReportFolder`, to allow remote administration. It grants the logged-on user on the local computer, and the administrators group on the local computer full control on the shared directory

If you recently upgraded SnapManager to SnapManager 6.0, other accounts might already have permissions on the share. If only the group Everyone has permissions, SnapManager removes the full control permissions to the report directory for that group. If multiple accounts (including Everyone) have permissions, SnapManager does not modify those accounts.

Windows host system requirements

Your system must meet the operating system, hotfix, Exchange server, hardware, protocol, and licensing requirements to run SnapManager.

The most current system requirements can be found using the Interoperability Matrix.

SnapManager 3.2 for Microsoft Exchange has fractional reserve monitoring enabled by default. When you upgrade from SnapManager 3.2 to this version of SnapManager, fractional space reservation remains enabled.

For SnapManager versions prior to 3.2, fractional space reservation is disabled by default. If you are not using fractional space reservation on the volumes that contain LUNs that are used for Exchange, you can disable the monitoring and thereby improve the time required for a backup operation.

Operating system	One of the following: • Windows Server 2003 Standard Edition with Service Pack 2 • Windows Server 2003 Enterprise Edition with Service Pack 2 • Windows Server 2003 Standard Edition R2 with Service Pack 2 • Windows Server 2003 Enterprise Edition R2 with Service Pack 2 • Windows Server 2008 Standard Edition with SP2 • Windows Server 2008 Enterprise Edition with SP2 • Windows Server 2008 R2 SP1
Microsoft Windows hotfixes	See the SnapDrive section, later in this topic.
Microsoft Exchange Server	One of the following versions of Microsoft Exchange (Standard or Enterprise Edition): • Exchange Server 2010 with SP2 • Exchange Server 2007 with SP3 • Exchange Server 2003 and with Service Pack 2 This requirement does not apply to a remote administration server. By default, Exchange Server 2007 uses `ChkSgFiles` for the integrity verification of databases and transaction logs. The option to throttle `Eseutil.exe` database checksum verification requires Exchange Server 2003 with SP2 or later.
FCP and iSCSI Multipath I/O for Windows	See the following SnapDrive requirements.
SnapDrive	SnapDrive 6.4.1.

	If you need to install or upgrade SnapDrive, see the *SnapDrive Installation and Administration Guide* for detailed instructions. For a remote administration server, SnapDrive is optional unless you intend to use the remote administration server to remotely administer SnapDrive. For a remote verification server, SnapDrive is required.
SnapDrive preferred IP address	If your storage system has multiple IP addresses, configure the SnapDrive preferred IP address. See the *SnapDrive Installation and Administration Guide* for your version of SnapDrive. If you do not configure a SnapDrive preferred IP address for a storage system that has multiple IP addresses, SnapDrive times out when attempting to create Snapshot copies simultaneously for more than one LUN on a storage system.
iSCSI protocol and FC protocol	The appropriate LUN access protocol software must be installed on the Windows host that runs the Exchange server. See your SnapDrive documentation for complete details about the system requirements. This requirement does not apply to a remote administration server.
SnapManager license	If SnapManager is licensed per server, a SnapManager license is required on the Windows host system. **Note:** For per-server SnapManager licensing, you can install SnapManager without specifying a server-side license; after SnapManager has been installed, you can apply the license from the License Settings dialog box.
Microsoft .NET 3.5 framework	SnapManager installation package installs Microsoft .NET framework if it is not present in the host system.
Microsoft Management Console (MMC)	You need MMC 3.0 x64 Edition or MMC 3.0 x86 Edition (required for Exchange 2007 and Exchange 2003) to launch the SnapManager snap-in console. This is a prerequisite to starting the SnapManager installation.
Windows PowerShell	PowerShell 1.0 RTM x86 Edition or PowerShell 1.0 RTM x64 Edition (required for Exchange 2007 and Exchange 2003) or PowerShell 2.0 (required for Exchange 2010). This is a prerequisite to running the SnapManager installation.

If your storage system has multiple IP addresses

If your storage system has multiple IP addresses, you must configure a SnapDrive preferred IP address; otherwise, SnapDrive times out when attempting to create Snapshot copies simultaneously

for more than one LUN on a storage system. See the *SnapDrive Installation and Administration Guide* for your version of SnapDrive.

Storage system requirements

Your storage system must meet some hardware, protocol, and licensing requirements before you use it with SnapManager.

Data ONTAP	See the SnapDrive software requirements described in the *SnapDrive Installation and Administration Guide* for your version of SnapDrive.
iSCSI protocol or FC protocol	You must install the appropriate LUN access protocol software on the storage system that stores the Exchange databases. For more information, see the *SnapDrive Installation and Administration Guide* for your version of SnapDrive.
SnapManager licenses	If SnapManager is licensed per-storage system, you need a SnapManager license on the storage system. **Note:** With storage system SnapManager licensing, if you use SnapMirror with SnapManager, SnapManager must be licensed on both the source and target storage systems.
SnapRestore license	For restore operations to succeed, the SnapRestore feature of SnapDrive must be licensed on the storage system that stores the Exchange databases. For more information, see the *SnapDrive Installation and Administration Guide* for your version of SnapDrive.
FlexClone license	If you plan to use the integrity verification feature on the destination volume, you must have a FlexClone license on that volume. See the *SnapDrive Installation and Administration Guide* for your version of SnapDrive.
SnapMirror license	If you plan to use the SnapMirror software with SnapManager, you require a SnapMirror license on both the source and target storage systems. For more information, see the *SnapDrive Installation and Administration Guide* for your version of SnapDrive.

Remote administration server requirements

You can remotely administer an Exchange server running SnapManager from another Windows system that is configured as a remote administration server. Your remote administration server must meet some installation requirements.

- You do not need to install an iSCSI driver or an HBA driver on this system.
- If you want to use the remote administration server to remotely administer SnapDrive, you need to install SnapDrive.

- You must install SnapManager.
- If you will be running SnapManager Services, you must install Microsoft Exchange System Management Tools.

Remote verification server requirements

Remote verification is performed using the same mechanisms that are used for local verification, except that the verification is performed on a host that is different from the Exchange server that initiated the backup set. To run remote verification, your remote system must meet some configuration requirements.

- SnapDrive must be installed.

  Note: Do not try to connect the Exchange server's LUNs to the remote verification server.
- SnapManager must be installed, but it does not need to be configured.

  You must specify the user account that you use for the production Exchange server.
- The appropriate LUN driver (iSCSI or FC protocol) must be installed.
- The remote Windows system must have connectivity (iSCSI or FC protocol) to the storage system.

  Note: If you are using iSCSI to connect to the storage system on the remote verification server, you must create an iSCSI connection.
- Exchange System Management Tools must be installed on the server (whether remote or local) for SnapManager to perform integrity verification.

  For more information about the installation of Exchange System Management Tools, see Microsoft documentation.

The connectivity required to perform remote verification is shown in the following illustration.

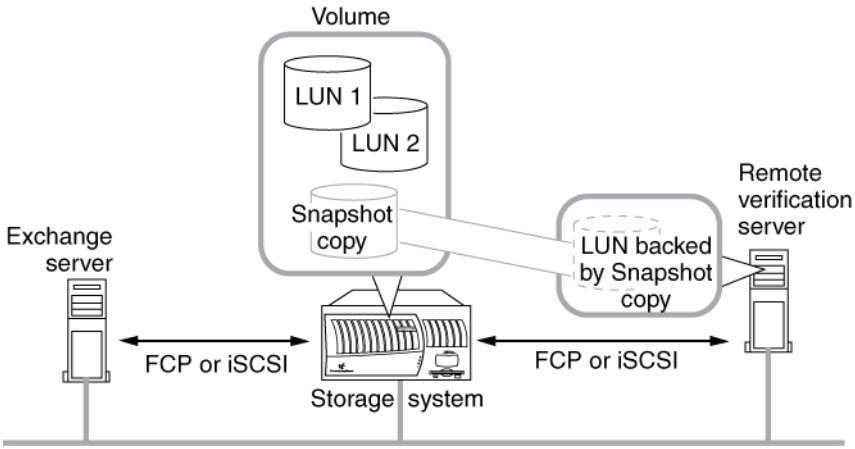

Exchange Server 2003 files to back up with a Windows backup utility

SnapManager does not back up all the files commonly used by Exchange Server 2003. You should use a standard Windows backup utility, such as Windows NTBackup.

Back up the following files using Windows NTBackup:

- File systems on hard disks connected to the Exchange server
- Windows registry
- Windows Active Directory
- Internet Information Service (IIS) metabase
- Key Management Service (KMS) database
- Site Replication Services (SRS) database
- System state

  Note: Instead of using tape, you can use Windows NTBackup to back up to a file and write or store that file directly on a storage system.

Exchange Server 2007 files to back up with a Windows backup utility

SnapManager does not back up all the files commonly used by Exchange Server 2007. You should use a standard Windows backup utility, such as Windows NTBackup.

Back up the following files using Windows NTBackup:

- File systems on hard disks connected to the Exchange server
- Windows registry
- Windows Active Directory
- Internet Information Service (IIS) metabase
- System state

  Note: Instead of using tape, you can use Windows NTBackup to back up to a file and write or store that file directly on a storage system.

- Mailbox server data−Offline Address Book that uses Web distribution
- Hub Transport server data−message tracking and protocol logs
- Edge Transport data−ADAM (Config Clone for backup and restore)
- Content Filter ESE database (Config Clone for backup and restore)
- Message tracking and protocol logs
- Client Access server data
- Microsoft Office Outlook Web Access web site, and Web.config file

- IMAP4 and POP3 protocol settings
- IIS metabase
- Exchange ActiveSync
- Outlook Web Access virtual directories
- Unified Messaging server data
- Custom audio prompts
- Incoming calls

SnapManager installation and upgrade

You can install SnapManager for Exchange on a stand-alone Windows host system or a Windows cluster using either the interactive or the unattended modes of installation. SnapManager also enables you to upgrade an older version and reinstall SnapManager on your configuration.

Unless it is specified for a particular upgrade path or for a particular troubleshooting situation, you do not need to uninstall SnapManager before reinstalling it or upgrading to a newer version.

Note: If you plan to reinstall SnapManager later, ensure that you record the drive letter and path of your current SnapInfo directories. After you reinstall SnapManager, ensure that you reconfigure SnapManager to use those same SnapInfo directory locations. The reconfiguration preserves SnapManager's records of backups created before you uninstalled and reinstalled SnapManager.

Compatibility of SnapManager builds with Windows Server

SnapManager 6.0 for Microsoft Exchange is available in x86 and x64 builds. The build that you use depends on your system configuration.

Build	Compatibility	WOW64
x86	Windows Server 2003 or Windows Server 2008 x86 edition	Not supported
x64	Windows Server 2003 or Windows Server 2008 x64 edition	Supported

Note: Ensure that you install MMC 3.0 and Windows PowerShell 1.0 (x86 or x64) on the Windows system before installing SnapManager.

Installation of SnapManager on a stand-alone Windows host system

You can install SnapManager on a stand-alone Windows host system that is used for a production Exchange server, a remote administration server, or a remote verification server.

Comparison of installation modes for a stand-alone Windows host system

You can run the software installation utility for SnapManager in either the interactive mode or the unattended mode. SnapManager guides you through the interactive mode; the unattended mode requires that you type certain commands and then installation takes place on its own.

The minimum required input for both modes is as follows:

- Acceptance of the terms of the license agreement
- SnapManager server account
  - User name and password

The optional input for both modes is as follows:

- User name
- Organization name
- SnapManager server-side license key
- SnapManager installation directory

The following table lists and describes the differences between interactive mode and unattended mode:

Installation activity	Interactive mode	Unattended mode
Access	Requires user interaction and access to the graphical user interface	Allows automated installation by executing a script or command-line command
SnapManager software license agreement	Is displayed in the installation utility	Is displayed in the command-line interface if you pass a specific parameter to the installation utility
After the installation finishes	If a system reboot is required to activate new software, a dialog box appears and prompts you to select whether you want to reboot the target system.	If a system reboot is required to activate new software, a dialog box appears and prompts you to select whether you want to reboot the target system. You can override this default behavior by including an optional command-line parameter.

Installing SnapManager in interactive mode

You can install SnapManager using the software installation utility in the interactive mode. The InstallShield wizard guides you through the installation.

Before you begin

The SnapManager Server Identity account must meet these requirements: *Preparation for installing or upgrading SnapManager* on page 21.

About this task

You do not need to stop Exchange services while you install SnapManager. Exchange can continue to run while you install SnapManager and afterward.

The SnapManager software installation program does not allow you to continue with the installation process if Microsoft .NET 3.5 is not installed. If .NET is not installed on a 32-bit system, the installation setup automatically installs it; on a 64-bit system, SnapManager prompts you to install it.

The default installation directory for SnapManager 6.0 for Microsoft Exchange installation is `C:\Program Files\NetApp\SnapManager for Exchange\`.

Steps

1. Download the software, browse to the SnapManager installation package, and double-click `setup.exe`.
2. In the **License Agreement** window, accept the license agreement.
3. In the **Customer Information** window, specify the user name, the organization name, and the SnapManager license type.
4. Note the full path of the folder in which you want to install SnapManager.

   You can change this default directory by clicking **Change**.
5. In the **SnapManager Server Identity** window, specify the user account you want to use to run SnapManager.

   If SnapDrive is installed and configured, the text box is prepopulated with the account that SnapDrive is configured for. Otherwise, browse to locate and select the account name. The user account name is specified in either of the following formats:

   - *`DomainName\UserName`*
   - *`UserName@DomainName`*
6. Type and confirm the password.
7. Click **Install**.
8. Wait until the **InstallShield Wizard Completed** window appears; then click **Finish** to exit the software installation utility.

Installing SnapManager in unattended mode

You can install SnapManager by running the software installation utility from the command-line interface. You can also install the SnapManager software under the control of a script for an unattended installation.

About this task

The command name that you specify depends on the SnapManager installation media that you want to access. If you are using a software package downloaded from the network, run `SME6.0.x.exe`.

You do not need to stop the Exchange service when you install SnapManager 6.0 for Microsoft Exchange. Exchange can continue to run while you install SnapManager and afterward.

Steps

1. Access the command line of the target host system and enter the following command, either directly at the command line or through a script:

   Example

   ***CommandName*/v"[AGREETOLICENSE={Yes|No}] [USERNAME=*UserName*] [COMPANYNAME=*CompanyName*] [ISX_SERIALNUM=*LicenseKey*] [INSTALLDIR=*InstallationDirectory*] SVCUSERNAME=*Domain**UserName*SVCUSERPASSWORD=*Password*SVCCONFIRMUSERPASSWORD=*PassWord*[REBOOT=0] [/L* *TempDirPath**LogFileName*] /qb"**

Command or parameter	Description
CommandName	Either `setup.exe` or `SME6.0.x.exe`, depending on the SnapManager installation media that you use.
`AGREETOLICENSE=`{Yes|No}	Set this parameter to Yes only if you have read and accept the terms of the SnapManager software license agreement. The installation does not proceed unless you agree.
`USERNAME=`*UserName*	Optional. If you do not specify a user name, the default value is retrieved from the registry.
`COMPANYNAME=`*CompanyName*	Optional. If you do not specify a company name, the default value is retrieved from the registry.
`ISX_SERIALNUM=`*LicenseKey*	Optional. Used only to specify an Exchange server-side license for SnapManager.
`INSTALLDIR=`*InstallationDirectory*	Optional. If you do not specify a directory, the default installation directory is used: `C:\Program Files\NetApp\SnapManager for Exchange\`
`SVCUSERNAME=`*Domain*`\`*UserName*	The account from which SnapManager is to be run.
`SVCUSERPASSWORD=`*Password*	
`SVCCONFIRMUSERPASSWORD=`*Password*	

Command or parameter	Description
`REBOOT=0`	Optional. After the installation finishes, the installation utility automatically reboots the Windows host system if that is required to activate updated software. If you specify this option, however, the system does not reboot.
`/L* TempDirPath\LogFileName`	Optional. If you specify this option, detailed information about the installation is written to the specified log file. This information can be used to investigate details about how a particular instance of SnapManager for Microsoft Exchange is installed. The asterisk (“*”) is a wildcard character that specifies that all the installation information (such as status messages, nonfatal warnings, and error messages) is to be logged. *`TempDirPath`* is the fully qualified name of the directory in which the installation log is to be created or overwritten. *`LogFileName`* is the name of the file to which the installation information is to be written.

The software installation utility installs the SnapManager for Microsoft Exchange software in the specified directory. If the `INSTALLDIR` parameter is not specified, the utility installs the software in the default directory `C:\Program Files\NetApp\SnapManager for Exchange\`.

2. Start SnapManager.

Displaying the SnapManager software license agreement

You can view the SnapManager software license agreement by running a command from the command-line or a script.

Step

1. Enter the following command, either directly at the command-line or through a script:

   Example

   ***CommandName*/v”SHOWEULA=Yes /qb”**
   outputclass="nolinebreak">

CommandName
is either `setup.exe` or `SME6.0.x.exe`, depending on the SnapManager installation media being used.

Note: The *`SHOWEULA=Yes`* command argument cannot be used in combination with other command-line arguments.

Examples of unattended installation

You can perform an unattended installation by using the downloaded installation package.

Scenario 1

Enter the following at the command-line or from a script:

```
C:\NetApp\downloads\SME6.0.x.exe /v"AGREETOLICENSE=Yes
SVCUSERNAME=MKTG2\Administrator SVCUSERPASSWORD=STeeL
SVCCONFIRMUSERPASSWORD=STeeL ISX_SERIALNUM=ABCDEFGHIJKLMN /qb"
```

Installation details	
Installation media	Installation package downloaded to `C:\NetApp\downloads\`
Server account	Account user name: MKTG2\Administrator Account password: STeeL
SnapManager license	Server-side (license key ABCDEFGHIJKLMN)

SnapManager installation in a Windows cluster

You can install and configure SnapManager for Microsoft Exchange either in an existing Windows cluster, or on a new Windows cluster.

Requirements for installing SnapManager in a Windows cluster

SnapManager is supported in Windows clusters with a maximum of eight nodes. Exchange Server 2003 requires a Windows cluster that uses shared LUN resources. Exchange Server 2007 supports two types of Windows cluster configurations--SCC and CCR.

- The SCC, which uses LUN resources that are shared among all nodes in the cluster.
- The CCR, which uses LUN resources that are dedicated to each of the two cluster nodes and are not shared among the nodes

  For more information, see Microsoft Exchange server documentation.

Installing SnapManager after creating a new Windows cluster

You can create a new cluster, and install and configure SnapManager for Microsoft Exchange, as part of the cluster.

Before you begin

When you install SnapManager for Microsoft Exchange in a new Windows cluster, consider the following points:

- You have two to five systems that you are ready to configure into a Windows cluster.
- Microsoft Exchange is not installed on either system.
- You are using a LUN as the quorum disk.

Create Exchange System Attendant Resource in the cluster group from the node that owns the cluster group containing the LUN physical disk resources.

Microsoft Management Console 3.0, and Windows PowerShell 1.0 (x86 or x64) must be installed. If Microsoft .NET 3.5 Framework is not already installed, SnapManager automatically installs it.

You do not have to use the default cluster group, but whatever group you use, must contain an IP address resource, a network name resource, and at least one physical disk cluster resource (LUN).

Steps

1. Create the cluster by following the procedure in your SnapDrive documentation for creating a Windows cluster with a virtual quorum disk.
2. Install and configure Exchange on one node.

   Then install Exchange on the remaining nodes.

   Note: Exchange virtual servers are not created during this step. See the Microsoft documentation for information about installing Exchange in a cluster.
3. Using **Cluster Administrator**, create an Exchange System Attendant resource in the cluster group, on which the LUN physical disk resources are to reside.

   That LUN should store the Exchange databases and transaction logs.

   During the process you are prompted for an Exchange data directory. This directory is also where the Exchange system files are placed. You can accept the default location now; these system files can be moved in a later step by using the Configuration wizard.
4. Specify the LUN physical disk resources and the Network Name resource as dependencies for the Exchange System Attendant Resource.
5. If it is not online, bring the cluster group online, and verify that the Exchange cluster is functioning correctly by moving the cluster group containing the newly created Exchange virtual server to the other node and back.

6. Install SnapManager on all nodes, starting with the node that currently owns the cluster resources.

   You can install SnapManager either in interactive mode, or in unattended mode.

7. Start SnapManager.

Installing SnapManager in an existing Windows cluster

You can install and configure SnapManager for Microsoft Exchange in an existing Windows cluster.

Before you begin

The shared LUNs must be physical disk resources in the cluster group that contains the Exchange virtual server that uses them.

You need to install Microsoft Management Console 3.0, Windows PowerShell 1.0 (x64 or x86) and Microsoft .NET 3.5 Framework (x64 or x86) on 64-bit and 32-bit operating system respectively, for SnapManager to work.

Prepare your SnapManager data configuration plan.

About this task

SnapManager adds the dependencies automatically for all the LUNs that Exchange uses, if the Configuration wizard detects that it is being run on a cluster. There are no dependencies set on the LUN physical disk resource in a Continuous Cluster Replication (CCR) configuration.

Steps

1. Verify that the Exchange virtual servers and the cluster services are functioning, by moving the Exchange virtual server from one cluster node to the other, and back.

   If any errors occur, or if any of the cluster resources do not start correctly, resolve the issue before continuing.

2. Install or upgrade SnapDrive as required.

   For details, see the SnapDrive documentation.

3. Create the shared LUNs to hold the Exchange databases and transaction log files, according to your SnapManager data configuration plan.

   Create the shared LUNs from the node that owns the cluster group, which contains the Exchange virtual server.

   For details, see your SnapDrive documentation.

4. Verify that dependencies for the System Attendant Resource are set correctly.

   The System Attendant Resource depends on all the LUN physical disk resources, and the Network Name resource in this cluster group, for Exchange Server 2003 clusters and Exchange Server 2007 single copy cluster.

5. Verify that the Exchange virtual servers and the cluster services are functioning correctly by moving the cluster group containing the newly created Exchange virtual server to the other node and back.
6. Install SnapManager on all nodes, starting with the node that currently owns the cluster resources.

   Use either the interactive installation procedure or the unattended installation procedure for a stand-alone Windows host system.
7. Start SnapManager.

Preparing to upgrade SnapManager

Ensure that you prepare the target system, the environment, and the existing SnapManager application before you begin the upgrade process.

Before you begin

If you need to interrupt connectivity to the Exchange data when upgrading other components, such as storage systems or connection protocols, take Exchange offline.

Close SnapManager before you begin upgrading to the new SnapManager version.

Uninstall SnapManager versions earlier than 3.2 before you install SnapManager 6.0.

About this task

When you upgrade SnapManager on a 32-bit system, set the SnapMgrService account to a Windows domain account that has the required rights to manage Exchange databases. Ensure that you use the same account, as the one used by the SnapManager server. The SnapManager installation program prompts for this account during installation.

You do not need to stop Exchange services while you upgrade SnapManager. Exchange can continue to run while you upgrade SnapManager.

You do not need to stop Exchange services while you upgrade SnapManager; Exchange can continue to run while you upgrade.

> **Note:** When you upgrade, all previous backup jobs will run correctly. The new functionality will not be available for existing jobs.

Steps

1. Prepare your target system and environment.
2. If you have upgraded the underlying SnapDrive or Microsoft iSCSI initiator versions to prepare your storage system for upgrading SnapManager, remove the iSCSI dependency with respect to SnapManager during the SnapManager upgrade process.

   There are two ways to do this:

- Run the Configuration wizard and use the Add Microsoft iSCSI Service Dependency page to remove this dependency.
- Stop the Microsoft Exchange System Attendant resource service.

After you finish upgrading the SnapManager application, you can restore the dependency.

3. Prepare the existing SnapManager application.

Comparison of interactive and unattended upgrade modes

You can run the software upgrade utility for SnapManager in either interactive mode or unattended mode. SnapManager guides you through interactive mode; in unattended mode, you type certain commands and then the upgrade occurs without you doing anything else.

Upgrade activity	Interactive mode	Unattended mode
Access	Requires user interaction and access to the graphical user interface	Enables automated uninstallation by executing a script or command-line command
SnapManager software license agreement	Is displayed in the software installation utility	Is displayed at the command-line interface if you pass a specific parameter to the installation utility

SnapManager upgrade paths

Unless specified for a particular upgrade path, or for a particular troubleshooting situation, you do not need to uninstall SnapManager before reinstalling it, or upgrading to a newer version.

You can upgrade to SnapManager 6.0 for Microsoft Exchange Server 2003 from the following older versions of SnapManager:

- SnapManager 5.0 for Microsoft Exchange
- SnapManager 4.0 for Microsoft Exchange
- SnapManager 3.2 for Microsoft Exchange

Upgrade considerations for fractional space reserve monitoring

SnapManager has fractional space reserve monitoring enabled by default. When you upgrade from SnapManager 4.0 or SnapManager 5.0, fractional space reservation remains enabled. If you upgrade from versions earlier than SnapManager 4.0, it is disabled by default.

If you are not using fractional space reservation on the volumes that contain Exchange LUNs, you can disable the monitoring to improve the backup completion time.

Upgrading SnapManager in interactive mode

You can reinstall SnapManager on a Windows host system to repair missing or corrupt files, shortcuts, and registry entries. SnapManager guides you through the interactive mode.

About this task

When you run the SnapManager installation utility as part of an upgrade, you are asked whether you want to retain your current installation directory. Because changing your installation directory might cause existing scripts that rely on the SnapManager installation directory path to fail, you should probably keep what you have.

Steps

1. Upgrade to SnapManager 6.0 for Microsoft Exchange by installing the new product.
2. After upgrading to SnapManager by following one of the installation procedures, proceed to start SnapManager.

   Note: Upgrading to SnapManager 6.0 removes all of the previously managed servers. You need to select the servers that you want to manage, using the **Add Servers to be Managed** dialog box.

Upgrading SnapManager in unattended mode

You can upgrade the SnapManager software under the control of a script for an unattended installation. You can also upgrade SnapManager by running the software installation utility from the command-line interface.

Before you begin

SnapManager operations must be stopped.

About this task

The command that you specify depends on the SnapManager installation media that you want to access. If you are using a software package downloaded from the network, run `SME6.0.x.exe`.

You do not need to stop Exchange services before or during the SnapManager software upgrade process.

Steps

1. From the command line of the target host system, enter the name of the script that executes the upgrade:

Example

***CommandName* /v"REINSTALLMODE=vomus REINSTALL=ALL [/L* *TempDirPath* *LogFileName*] /qb"**

Command or parameter	Description
CommandName	Either `setup.exe` or `SME6.0.x.exe`, depending on the SnapManager installation media being used
`REINSTALLMODE=vomus`	Causes the software installation utility to re-cache the new software package
`REINSTALL=ALL`	Causes the software installation utility to update all the software components
`/L*`*`TempDirPath`* `\`*`LogFileName`*	Optional If you specify this option, detailed information about the installation is written to the specified log file. This information can be used to investigate details about how a particular instance of SnapManager for Microsoft Exchange is installed. The asterisk (*) is a wildcard character that specifies that all the installation information (such as status messages, nonfatal warnings, and error messages) should be logged. *`TempDirPath`* is the fully qualified name of the directory in which the installation log is created or overwritten. *`LogFileName`* is the name of the file to which the installation information is written.

The software installation utility installs the SnapManager for Microsoft Exchange software.

2. Start SnapManager.

Example of an unattended upgrade

You can perform an unattended upgrade by using the downloaded installation package.

Assume that you are upgrading from the installation package downloaded to `C:\NetApp\downloads\`. Enter the following at the command-line interface:

```
C:\NetApp\downloads\SME6.0.x.exe /v"REINSTALLMODE=vomus
REINSTALL=ALL /qb"
```

Updating legacy scheduled jobs

When you upgrade from SnapManager 3.2 or 4.0, you must update the scheduled legacy jobs so that they are compatible with SnapManager 6.0.

About this task

In an Microsoft Cluster Services Cluster environment, the SnapManager upgrade tool shows all the nodes in a list. Although you can select a specific node to migrate a legacy job to that particular node, you should schedule the job on all nodes in the cluster, to achieve fault tolerance.

In SnapManager 3.2 and 4.0, the jobs that are scheduled to run against a server are not required to reside there. In SnapManager 6.0, the jobs must be scheduled and run in the same server.

By default, SnapManager enables the **Delete legacy job** and **Replace the job if the job exists** check boxes if the target server is different from the server on which the SnapManager 4.0 jobs exist or if the name of the specified job is different from the name of the SnapManager 4.0 job.

Note: When you upgrade from SnapManager 5.0 to SnapManager 6.0, you do not have to update the SnapManager 5.0 jobs as they are compatible with SnapManager 6.0.

Steps

1. Launch `SMEUpgradeJobs.exe` from the SnapManager Installation directory.

   The Update SnapManager for Exchange legacy scheduled jobs window appears, listing all the SnapManager 3.2 and 4.0 jobs.

2. To see the jobs in a different server, click **Browse** to select a different server and then click **Refresh**.

   SnapManager lists the 3.2 and 4.0 jobs for the selected server.

3. To update the SnapManager 3.2 and 4.0 jobs for SnapManager 6.0, click **Update**.

   A **Scheduling** dialog box appears that you can use to migrate the 3.2 and 4.0 to 6.0 jobs.

4. If you do not want to use SnapManager 3.2 and 4.0 jobs for SnapManager 6.0, click **Actions > Delete**.

Prerequisites for uninstalling SnapManager

If you have used SnapManager to manage your Exchange databases, and you plan to reinstall SnapManager later, ensure that you record the drive letter and the path of the SnapInfo directory locations before proceeding.

If you have set up a single SnapInfo directory for all databases on your host, then record the drive letter and the path of the LUN that contains a single SnapInfo directory for all Exchange servers, and their associated databases.

If you have set up multiple SnapInfo directories, then record the drive letter and path of each LUN that contains a SnapInfo directory.

SnapManager reports record the current SnapInfo directory locations in the most recent logs contained in the `Backup` folder and in the `Config` folder.

After you reinstall SnapManager, ensure that you reconfigure SnapManager with the same SnapInfo directory locations that SnapManager had used earlier.

Attention: If you configure SnapManager with different SnapInfo directory locations than used previously, then SnapManager no longer has records of any backups made before the reinstallation of SnapManager occurred. As a result, your prior backup sets could be invalidated or deleted the next time you perform a backup.

Comparison of interactive versus unattended uninstallation

You can run the software uninstallation utility for SnapManager in either the interactive mode or the unattended mode. SnapManager guides you through the interactive mode; the unattended mode requires that you type certain commands and then uninstallation takes place on its own.

Uninstall activity	Interactive mode	Unattended mode
Access	Requires user interaction and access to the graphical user interface	Allows automated uninstallation by executing a script or command-line command
Method used	Can be implemented by using **Add or Remove Programs** in the Control Panel	Can be implemented by using the software installation utility for SnapManager

You can remove the report directory using both the modes.

Moving Exchange data to a local disk

Before you uninstall SnapManager, you might want to move your Exchange data store from the storage system to a local disk, so that there is no data loss if you face a problem in the uninstallation operation.

Before you begin

Ensure that you have enough space on your local disk before you move your database back to it.

Steps

1. Use Exchange System Manager to move your databases, transaction logs, and system files from the LUNs onto a local disk.
2. Confirm that the data was moved correctly and that your Exchange server is functioning normally.

3. Disconnect or delete your LUNs from your SnapManager host.

   For information about disconnecting or deleting LUNs, see the SnapDrive documentation.

Uninstalling SnapManager in interactive mode

You can uninstall SnapManager and all its components by using the Windows Add or Remove Programs utility. You can also remove the SnapManager report directory. Unless it is specified for a particular upgrade path, or for a particular troubleshooting situation, you do not need to uninstall SnapManager before reinstalling it or upgrading to a newer version.

About this task

Do not attempt to remove a currently installed version of SnapManager using an interactive method other than the approach described here. By doing so, you might cause system problems that result from unknown remaining files.

Ensure that you uninstall SnapManager from all the nodes of the cluster in a clustered configuration.

You do not need to stop the Exchange service or remove the Exchange databases from the LUNs before you uninstall SnapManager.

Steps

1. Close SnapManager.
2. In the **Control Panel**, select **Add or Remove Programs**, and then select the entry for SnapManager.
3. Click the following buttons, depending to what you want to remove:

If you want to...	Then...
Remove only the SnapManager software and leave the report directory	Click **Remove**.
Remove both the SnapManager software and the report directory	**a.** Click **Change**. **b.** Click **Remove Reports**. **c.** Click **Remove**.

4. Click **Yes**.

Uninstalling SnapManager in unattended mode

You can uninstall SnapManager by running the software installation utility from the command-line interface. You can also uninstall the SnapManager software under the control of a script for an unattended uninstallation. Unless it is specified for a particular upgrade path or for a particular

troubleshooting situation, you do not need to uninstall SnapManager before reinstalling it or upgrading to a newer version.

About this task

Do not attempt to remove a currently installed version of SnapManager using an unattended method other than that described here. Doing so might cause system problems that result from unknown remaining files.

The command that you use to initiate the uninstallation depends on the SnapManager installation media that you access. If you are using the software package that you downloaded from the network, then use `SME6.0.x.exe` (the name of the software package itself).

Ensure that you uninstall SnapManager from all the nodes of the cluster in a clustered configuration.

Step

1. Access the command line of the target host system, and enter the following command either directly at the command line or through a script:

   ***CommandName* /v"REMOVE=ALL [REMOVEREPORTFOLDER=1] [/L* *TempDirPath**LogFileName*] /qb"**

Command or parameter	Description
CommandName	Either `setup.exe` or `SME6.0.x.exe`, depending on the SnapManager installation media being used
`REMOVE=ALL`	Causes the software installation utility to remove SnapManager (performs the same function as the **Remove** option in the Program Maintenance window)
`REMOVEREPORTFOLDER=1`	Optional Causes the software installation utility to remove the Report directory (performs the same function as the **Remove Report Folder** option in the Remove the Program window)

Command or parameter	Description
`[/L* TempDirPath\LogFileName`	Optional If you specify this option, detailed information about the installation is written to the specified log file. You can use this information to investigate details about how a particular instance of SnapManager is installed. The asterisk `*` is a wildcard character that specifies that all the installation information (such as status messages, nonfatal warnings, and error messages) should be logged *`TempDirPath`* is the fully qualified name of the directory in which the installation logs are created or overwritten *`LogFileName`* is the name of the file to which the installation information is written

Examples of unattended uninstallation

You can perform an unattended uninstallation by using the downloaded installation package.

Scenario 1

Enter the following at the command-line:

```
C:\NetApp\downloads\SME6.0.x.exe /v"REMOVE=ALL
REMOVEREPORTFOLDER=1 /qb"
```

Installation details	
Installation media	Installation package downloaded to `C:\NetApp\downloads\`
SnapManager license	Server-side (key ABCDEFGHJIJKLMN)

SnapManager reinstallation with or without uninstallation

You can reinstall the version of SnapManager that you are currently using to repair missing or corrupt files, shortcuts, and registry entries.

Unless it is specified for a particular upgrade path or for a particular troubleshooting situation, you do not need to uninstall SnapManager before reinstalling it or upgrading to a newer version.

You do not need to stop Exchange services before, or during the SnapManager software reinstallation process.

If you have already uninstalled SnapManager

If you have uninstalled SnapManager, then the reinstallation procedure is identical to a new installation of the software.

If you used SnapManager to manage your Exchange databases before you uninstalled the SnapManager application, then ensure that you configure SnapManager with the same SnapInfo directory locations that SnapManager used before the reinstallation. If you configure SnapManager with different SnapInfo directory locations than used previously, then SnapManager no longer has records of any backups made before the reinstallation of SnapManager occurred.

If you did not uninstall SnapManager

You can upgrade to this version of SnapManager.

Comparison of interactive and unattended reinstallation

You can run the software reinstallation utility for SnapManager in either the interactive mode or the unattended mode. SnapManager guides you through the interactive mode; the unattended mode requires that you type certain preliminary commands, after which reinstallation takes place independently.

Reinstallation activity	Interactive mode	Unattended mode
Access	Requires user interaction and access to the graphical user interface	Allows automated uninstallation by executing a script or command-line command
SnapManager software license agreement	Is displayed in the installation utility	Is displayed in the command-line interface if you pass a specific parameter to the installation utility

Reinstalling SnapManager in interactive mode

You can reinstall the same version of SnapManager on a Windows host system. This option repairs missing or corrupt files, shortcuts, and registry entries. SnapManager guides you through this process in the interactive mode.

Before you begin

- Exit SnapManager.
- Back up system resources using an industry-standard backup utility.
- Back up the operating system, including the system state, the boot and system drives, and the registry.
- Back up your Exchange databases and transaction log files.

- Use a backup utility that is part of Windows to create and maintain a current emergency repair disk (ERD).
- Verify that your host system meets the minimum requirements.
- Verify that your storage system meets the minimum requirements.

About this task

You do not need to stop Exchange services before or during the SnapManager software reinstallation process.

Steps

1. Download the software, browse to the SnapManager installation package, and double-click `setup.exe`.
2. In the **Program Maintenance** page, select **Repair/Upgrade**.
3. In the **Ready to Install** page, click **Install**.
4. Click **Finish**.
5. Start SnapManager.

Reinstalling SnapManager in unattended mode

You can reinstall SnapManager by running the software installation utility from the command-line interface. You can also reinstall the SnapManager software under the control of a script for an unattended reinstallation. SnapManager also provides you with the option to reinstall using an interactive mode in which the InstallShield wizard guides you through the reinstallation.

Before you begin

Exit SnapManager.

Steps

1. Access the command line of the target host system, and then enter the following command either directly at the command line, or through a script:*`CommandName`* `/v"REINSTALLMODE` =vomus `REINSTALL`=ALL `[/L*`*`TempDirPath\LogFileName`*`] /qb"`.

Command or parameter	Description
`CommandName`	Either `setup.exe` or `SME6.0.x.exe`, depending on the SnapManager installation media that you use.
`REINSTALLMODE`=vomus	Causes the software installation utility to recache the new software package.

Command or parameter	Description
`REINSTALL=ALL`	Causes the software installation utility to update all the software components.
`SVCUSERNAME=`*`Domain\UserName`* `[/L*`*`TempDirPath\LogFileName`*`] /qb”`	Optional. If you specify this option, detailed information about the installation is written to the specified log file. This information can be used to investigate details about how a particular instance of SnapManager is installed. The asterisk `*` is a wildcard character that specifies that all of the installation information (such as status messages, nonfatal warnings, and error messages) should be logged. *`TempDirPath\LogFileName`* is the fully qualified name of the directory in which the installation log will be created or overwritten. *`LogFileName`* is the name of the file to which the installation information will be written.

2. Start SnapManager.

Example of an unattended reinstallation

You can perform an unattended reinstallation by using the downloaded installation package.

Enter the following at the command line or from a script:

```
C:\Netapp\downloads\SME6.0.x.exe /v"REINSTALLMODE=vomus
REINSTALL=ALL /qb"
```

Installation details	
Installation media	Installation package downloaded to `C:\NetApp\downloads\`
Installed version of SnapManager	SnapManager 6.0 for Microsoft Exchange

When you start SnapManager for the first time

After you install or upgrade SnapManager and start the application for the first time, SnapManager verifies the Exchange server, SnapDrive version, and databases for compatibility. You must follow certain guidelines to start SnapManager for the first time.

What SnapManager verifies at startup

When SnapManager starts for the first time, it verifies the Exchange server first, checks the SnapDrive version, and finally checks for database configuration, ensuring that the basic prerequisites to run the Configuration wizard have been met.

1. SnapManager checks if there is a supported version of Microsoft Exchange installed and tells you if there is not.

   SnapManager does not check the Service Pack level.

   If you have another available Exchange server that has the minimum Service Pack required by SnapManager, you can connect to that server instance instead. Otherwise, you must close the SnapManager application, upgrade Microsoft Exchange as needed, and then restart SnapManager.

2. Next, SnapManager checks the SnapDrive version.

   If SnapManager detects that SnapDrive is installed, but that it is a version not supported with this version of SnapManager, a warning appears. Close the SnapManager application, upgrade SnapDrive as needed, and then restart SnapManager.

3. Finally, if SnapManager detects that no databases have been configured for use with SnapManager, a message appears, prompting you to run the Configuration wizard.

Starting SnapManager and connecting to the default server

When you start SnapManager for the first time after either an installation or upgrade, you have to select one or more Exchange servers that you want to manage.

Before you begin

You must know the name of the Exchange server that you want to manage, the computer that you want to use for database integrity verification, and the default mountpoint path for integrity verification.

You can also decide if you want to enable verification throttling and to override the existing database verification.

Steps

1. To start SnapManager, go to the **Start** menu and select **Program Files > NetApp > SnapManager for Exchange Management Console**.
2. Click **SnapManager for Exchange**.

   SnapManager prompts you to specify an Exchange server that you want to manage.
3. Click **OK**.

   The **Add Exchange Servers to be Managed** window appears.
4. Select an Exchange server from the displayed list, or click **Browse** to select an Exchange server from the network.
5. Click **Add**.

   SnapManager prompts you to specify the computer that you want to use for database integrity verification.
6. Click **OK**.

   The **Database Verification Settings** window appears.
7. Specify the verification server that you want to use for integrity verification.

   SnapManager prompts you to specify the default mountpoint directory.
8. In the **Access LUN in Snapshot** window, specify the default mountpoint directory.

   Note: You can also select **Automatically assign available drive letter**.
9. Click the **Override Verification** tab if you want to override the existing database verification, and then click **Override database verification requirement for restore**.
10. If you want to enable verification throttling, in the **Verification Throttling** tab, click **Throttle database checksum verification**.
11. Click **OK**.

After you finish

You can run the Configuration wizard to start SnapManager configuration.

Starting SnapManager and connecting to the CCR server

After starting and connecting to the default server, you can connect to a CCR-enabled server by using SnapManager for Exchange. You need to add the CCR-enabled server if you are using a clustered configuration.

Before you begin

You must know the name of the Exchange server that you want to manage.

Steps

1. To start SnapManager, go to the **Start** menu and select **Program Files > NetApp > SnapManager for Exchange Management Console**.

   The SnapManager 6.0 for Microsoft Exchange console appears.
2. Click **Add Servers to be managed** in the **Actions** pane.
3. Select a CCR-enabled Exchange server from the displayed list.

   Note: You can use the **Browse** button to select a different CCR-enabled Exchange server from the network.
4. Click **Add**.

   SnapManager connects to the CCR-enabled Exchange server.

After you finish

You can run the Configuration wizard to start SnapManager configuration.

SnapManager Dashboard view

You can view the status of different SnapManager servers connected to the SnapManager for Exchange network. This is a dynamic view that is refreshed frequently. You can view your Exchange server configuration and add new Exchange servers.

Exchange server configuration

You can select any Exchange server in the Scope pane to view the server configuration details. The following details are displayed in the Results pane:

- Name of the Exchange server
- Name of the host

Note: In the case of clustered configurations, SnapManager displays the host name of the node to which SnapManager is connected.

- Exchange server version
- SnapManager version
- Exchange server configuration

  This value can be any one of the following: Standalone, Single Copy Cluster, or Cluster Continuous Replication.
- Verification server
- Next job that is scheduled
- Last backup operation, including a timestamp and a hyperlink to the corresponding report
- Last restore operation, including a timestamp and a hyperlink to the corresponding report
- Last configuration operation, including a timestamp and a hyperlink to the corresponding report

Current Job Status view

This view displays the jobs that are running and those in queue, for local as well as remote verification. The following list describes the elements of the pane:

View Status	Specifies a link to the SnapManager report associated with the running job.
Priority	Specifies the queue position of the job.
Job ID	Displays the Job ID of the job.
State	Displays whether the job is running, cancelling, or in queue.
Job Type	Displays whether the job is a backup job or backup with verification job. This element can have any one of the following values: • Full Backup • Frequent Recovery Point Backup • Backup Verification • Restore • Test Restore
Backup Source	Displays the name of the server that creates the backup.
Submission time	Displays the time the job was submitted.
Start time	Displays the time the backup job started.

Add new Exchange servers

You can add a new Exchange server from the Actions pane and manage it through SnapManager.

Scheduled jobs

When you select **Scheduled Jobs** in the Scope pane, the details of all scheduled jobs such as backup and integrity verification jobs are displayed in the Results pane. You can view, edit, and delete these jobs.

When you select **Scheduled Jobs** in the Scope pane, the details of each scheduled job like job name, host name, last run time of the job, last return code, next run time of the job, and so on are displayed in the Results pane.

You can view and edit the settings for a job such as its schedule, password, and so on. You can also delete scheduled jobs.

Database Availability Group

The Database Availability Group (DAG), a high availability feature of Exchange Server 2010, provides database-level recovery from failures and data corruption.

The DAG is a group of up to 16 servers in which some or all of the mailbox databases are replicated to one or more servers within the DAG. Thus there can be up to 16 copies of a database in the DAG. A database is not associated with any server. The current active database and its copies use the same path on each server. All database names in a DAG have to be unique.

The DAG is for mailbox databases and not for public folder databases.

Database migration considerations

You must migrate and store your database in correctly sized and appropriately configured LUNs. Use the minimum possible number of LUNs for faster backup operations and faster restore operations. You must group your databases if you do not need to restore them individually.

Determine how many LUNs you need for your SnapManager configuration and what data those LUNs hold. Then create your data configuration plan by calculating the required sizes for each LUN and the volume that contains it.

Exchange configuration requirements

Microsoft Exchange Server 2010, Exchange Server 2007, and Exchange Server 2003 have data configuration requirements with respect to the number of Storage Groups per server, the number of databases per Storage Group, the number of databases per server, and the number of Recovery Storage Groups, Recovery Databases, and Storage Groups per virtual server.

Member servers in a Database Availability Group (DAG)

A DAG can have a maximum of 16 Exchange 2010 member servers.

Storage Groups per server

If you use Exchange Server 2003, each server can have 1 to 4 Storage Groups per server. If you use Exchange Server 2007, each server can have 1 to 50 Storage Groups per server. Storage Groups are not present in Exchange Server 2010.

Databases per Storage Group

Exchange Server 2003 Enterprise Edition supports a total of four Storage Groups, each containing a maximum of five databases. Exchange Server 2007 Standard Edition, supports a total of five Storage Groups and five databases. Exchange Server 2007 Enterprise Edition supports up to 50 Storage Groups and a maximum of 50 databases per server. Exchange 2007 permits up to five databases in each Storage Group, but best practice is to put only one database in a Storage Group.

A Continuous Cluster Replication configuration (CCR) can have only one database in a Storage Group.

Note: Storage Groups are not present in Exchange Server 2010.

Databases per Exchange Server 2010

The maximum number of databases that can be associated with each Exchange Server 2010 is 100.

Recovery Storage Group

A Recovery Storage Group enables you to mount a copy of a mailbox store on a production server, and recover the data within the restored mailbox store when the current mailbox store is still running.

If you are using Exchange Server 2003, you must use Exchange System Manager (ESM) to manually create the Recovery Storage Group.

If you are using Exchange 2007, SnapManager creates a Recovery Storage Group automatically when you restore a backup to a Recovery Storage Group. You can also create and manage Recovery Storage Groups by using the Exchange Management Shell.

There can be only one Recovery Storage Group per Exchange Server.

Recovery Database

A Recovery Database enables you to restore, recover, and mount your database to a different Exchange server. It is a feature available with Exchange Server 2010. Each Exchange Server 2010 can have only one such Recovery Database mounted at any time.

Storage Groups per virtual server in Exchange Server 2003

An Exchange Server 2003 cluster has one or more Exchange virtual servers. In a two-node cluster, you can have a maximum of four Storage Groups on the complete cluster.

An active/passive Exchange Server cluster can have more than two cluster nodes. The maximum number of Exchange virtual servers in an active/passive cluster is one less than the number of nodes in the cluster. An active/passive Exchange cluster can have more than four Storage Groups, because only one Exchange virtual server runs on a cluster node at a time.

Maximum configurations supported by Exchange servers

The following table shows the maximum number of Storage Groups (including the Recovery Storage Group) and databases that Exchange Server 2003 and Exchange Server 2007 support:

Exchange Server	Configurations supported for Standard Edition		Configurations supported for Enterprise Edition	
	Storage Group support	Number of databases per Storage Group	Storage Group support	Number of databases per Storage Group
Exchange Server 2003	1	1 Mailbox store and 1 public folder store	4	5
Exchange Server 2007	5	5	50	50

The following table shows the maximum number of databases (including the Recovery Database) that Exchange Server 2010 supports:

Exchange Server 2010	Standard Edition	Enterprise Edition
Maximum number of databases	5	100

Deletion of Storage Groups from clustered systems

You can delete a Storage Group from a clustered system by deleting it from the Exchange Management Console and the Cluster Administrator. If you do not delete the Storage Group from the Exchange Management Console, the Configuration wizard and the Backup wizard detect the Storage Group even after you delete it.

Rules for Exchange Server storage groups and databases enforced by the Configuration wizard

When you configure your Exchange Server storage groups and databases for SnapManager, the Configuration wizard creates some rules for databases, transaction logs, SnapInfo directories, for sharing of LUNs or occupation of a SAN boot LUN, and Storage Group system files.

- Databases

  Databases from different Storage Groups cannot share a single LUN. Also, databases cannot share a single LUN with transaction logs, even if both belong to the same Storage Group.

  Note: On Exchange 2003, SnapManager restores all databases on a LUN together. To restore an individual database without restoring the rest of its Storage Group, move the database to a separate LUN.

  Note: In a Database Availability Group (DAG) configuration, it you have a drive or LUN existing on all the nodes and this drive is used by a database only on some nodes, then other databases cannot use this drive on other nodes of the DAG.

  Note: On Exchange 2010, SnapManager restores 1 database per LUN.
- Transaction logs

  Transaction logs from different Storage Groups can share a single LUN.
- SnapInfo directory

  A SnapInfo directory can have its own LUN, or it can share a LUN with transaction logs if you use NTFS hard links. The SnapInfo directory cannot share the LUN with databases.
- SAN boot LUN

  You cannot place databases, transaction logs, or a SnapInfo directory on a SAN boot LUN (a LUN configured as a boot device for a SAN host).
- Storage Group system files

  Storage Group system files must be on a LUN that contains the transaction logs for that Storage Group. If you modify the database or move the transaction logs for a Storage Group, the

Configuration wizard automatically moves the Storage Group system files to the same LUN as the transaction logs.

Recommended Exchange Server storage groups and databases configurations

Ensure that you place your Exchange Server storage groups and databases on dedicated volumes, place your LUNs correctly, and use the minimum number of LUNs possible, for fast and successful backup and restore operations.

- Use dedicated volumes for your data.
- You must always group your databases unless you need to perform an individual database restore operation.

You should ensure that for a Continuous Cluster Replication (CCR) configuration, all Exchange servers hosting a mailbox database and its copies have same storage level configuration layout. For example, if two LUNs share the same storage volume (vol1) on host A, the same two LUNs on host B should share same storage volume (vol2).

Exchange Server storage groups and databases configurations to avoid

Before you migrate your Exchange database and transaction logs to LUNs during configuration, ensure that you do not manually store any files or directories on the LUNs that contain databases, SnapInfo directories, or host a mountpoint.

Note: After adding or moving Exchange Server storage groups and databases, always run the Configuration wizard. If you add more databases or move databases to different LUNs without using the Configuration wizard, you create an invalid configuration that causes backup or restore operations to fail.

Before you migrate your database and transaction logs to LUNs, ensure that you follow these points:

- Do not manually store any directories or files (including system paging files) on the LUNs that are used by databases and transaction logs.
- Do not manually store any directories or files on the LUNs that are used by the SnapInfo directory.
- Do not manually store databases on a LUN that hosts NTFS volume mountpoints.
- Do not place LUNs containing data from different Exchange servers in the same volume, for the following reasons:
  - If you create backups for a different server on a volume that contains data from your current server, SnapManager uses all the Snapshot copies for the restore operation, but you can restore only the backups of the other server and not the backups of your current server.
  - Conflicting backup schedules between two Exchange servers can lead to busy Snapshot copies.
- Do not place LUNs in the storage system’s root volume, which is reserved for use by Data ONTAP.

Exchange message tracking in an MSCS configuration

When you use Exchange message tracking in an Microsoft Cluster Services (MSCS) configuration, place the message tracking directory on a shared LUN that does not store a database. If you specify a shared LUN with stored database, SnapManager restores an old copy of the message tracking log during the restore operation.

To create an Exchange virtual server in an MSCS configuration, install the Exchange binaries on both nodes and then create a new System Attendant cluster resource. As part of creating the System Attendant cluster resource, Exchange prompts you as follows:

```
Enter [a] path to the data directory.
```

Exchange places the new databases and other related files (including the SMTP and Message Transfer Agent (MTA) queues, message tracking logs, and Microsoft Search files, if configured) on this path on a shared disk.

Suppose that the path that you specify for the data directory at the System Attendant cluster resource is a shared LUN that is also specified for a database using the SnapManager Configuration wizard later. In this case, the Exchange message tracking log and the Exchange database are on the same LUN.

SnapManager does not prevent this configuration, nor does this configuration pose an immediate problem. During a subsequent restore operation, SnapManager restores the entire LUN from the Snapshot copy, with the desired database and an old copy of the message tracking log.

If you already created a System Attendant cluster resource with the data directory on a shared LUN with a stored database, you can move the message tracking location to a different location on a shared disk. For a description of how to move the message tracking location in Exchange Server 2003, see Microsoft Knowledge Base article 841089.

NTFS volume mount points

A volume mount point is a drive or volume in Windows that is mounted to a folder in the NTFS file system.

A mounted drive is assigned a drive path instead of a drive letter. Volume mount points enable you to exceed the limitation imposed by naming drives using one of the 26 letters of the alphabet. By using volume mount points, you can mount a target partition into a folder on another physical disk. For more information about volume mount points, see the Microsoft Web site.

Limitations of NTFS volume mountpoints

Volume mountpoints in an NTFS environment, or on a server cluster have limitations with their creation, the supported clustered configuration, and placement with respect to the Quorum disk.

The NTFS volume mountpoints feature imposes the following limitations:

- You can create volume mountpoints only on a dedicated disk or a shared cluster resource disk.

- You cannot use volume mountpoints as storage locations for the Exchange Server 2003 binaries.
- You cannot use volume mountpoints on a server cluster.
- You cannot place volume mountpoints between clustered and nonclustered disks.
- You cannot create mountpoints that refer to the Quorum disk.

Using server cluster mountpoints includes the following limitations:

- The mounted volume must be of the same type as its root; that is, if the root volume is a shared cluster resource, the mounted volume must also be shared, and if the root volume is dedicated, the mounted volume must also be dedicated.
- You cannot create mountpoints to the Quorum disk.
- If you have a mountpoint from one shared cluster resource disk to another, ensure that the disks are in the same group and that the mounted disk resource is dependent on the root disk source.

  For more details, see the full text of Microsoft TechNet article 280297.

SnapManager support for volume mountpoints

SnapManager supports volumes mounted on Data ONTAP LUNs in certain environments. You must accommodate particular drive letter limitations, restrictions, naming conventions, and procedures for using volume mount points. Exchange Server 2003 supports only a stand-alone Windows Server 2003 configuration.

Drive letter limitations and individual database restoration

By using NTFS volume mountpoint support, SnapManager can manage databases that are stored on mounted volumes, in addition to those stored on standard Windows volumes. With SnapManager, your configuration is not limited to the 26 letters of the alphabet to name drives in Windows.

SnapManager needs one LUN and one available drive letter to perform backup, verification, or restore operations.

The following SnapManager operations require additional drive letters:

- Individually restored databases require their own LUN and drive letter.

  If you need to restore many databases individually, you likely need more than the total limit of 26 drives.
- Performing a remote verification of more than one backup set requires a second LUN and therefore another available drive letter or mountpoint.

Windows supports up to 26 drives, and you can allocate a maximum of 25 drive letters with releases of SnapManager prior to 5.0. You can avoid running out of available drive letters by determining which databases must be restored individually and which can be combined on a LUN with one or two other databases and all restored together. With this approach, you must still reserve at least one free drive letter for SnapManager.

Mounted volume restrictions with SnapManager

In addition to the limitations inherent in the NTFS volume mount point feature of Windows, SnapManager does not allow you either to store files or to back up databases on an NTFS volume that has mount points.

You cannot put databases or transaction logs on a mount point root volume. Consider the following points:

- The SnapManager Configuration wizard does not allow you to place Exchange databases on an NTFS volume that has mount points.
- SnapManager Backup does not allow you to back up a database on an NTFS volume that has mount points.

Note: If databases reside on a LUN, do not add mount points to that LUN, even after you finish a backup operation. A subsequent LUN-level restore operation removes those volume mount points. This disrupts access to the data on the mounted volumes, to which the volume mount points refer.

Mounted volume naming conventions with SnapManager

You can refer to a mounted volume by the path of a volume mountpoint that is created in an empty directory. The SnapManager user interface represents a mounted volume based on the path of a volume mountpoint that references that mounted volume, as follows: *`DriveLetter:`* *`MountPointName`*.

Here *`DriveLetter`* represents the drive letter to which the mounted volume is assigned and *`MountPointName`* represents the mountpoint name assigned to the empty directory used to reference the mounted volume.

An example of a path name is as follows:

```
F:\ My_Mount_Point
```

How mounted volumes are shown in SnapManager

A mounted volume is a volume that is mapped to a folder and referenced by a mount point that is created in an empty directory. You can access the mounted volumes in the SnapManager user interface either from the Configuration wizard or from the Backup and Verification window.

The path-style representation of a mounted volume can appear in any part of the SnapManager user interface that refers to LUNs accessed by SnapManager.

Path-style representation in the Configuration Wizard

The Configuration wizard displays disk list selection in the following different ways:

- Select a Storage Group/Database to move to a LUN.
- Select a set of logs from one or more Storage Groups/Databases to move to a LUN.
- Configure the SnapInfo directory to store the backup information.

The representation depends upon the following types of LUNs:

LUNs that are referenced more than once	If the storage system is configured with multiple references to the same LUN, each such LUN reference that has a label that includes any other references to the same LUN. For example, suppose the drive letter H: and the mountpoint `G:\DB\` reference the same LUN. In this case, the Disk List selection contains two entries for the one LUN: LUN H: <`G:\DB\`> LUN `G:\DB\` <H>
LUNs that have mounted volumes	If SnapManager accesses a LUN with an NTFS volume that is referenced by a mountpoint, that LUN has a label that indicates that it was accessed by SnapManager. For example, suppose the drive letter F: references a LUN that hosts a mountpoint. In this case, the Disk List selection shows that LUN as follows: LUN F: (MPRoot)

The Configuration wizard does not allow you to store Exchange database files on LUNs that host NTFS volume mountpoints.

Path-style representation in the Backup wizard and the Backup and Verify window

The left navigation pane of the Backup and Verify window and the Backup wizard lists the location of the Storage Group/Database components as follows:

- SnapInfo directory
- EDB path
- Storage Group/Database log path

Path-style representation in the Restore wizard and Restore window

The right display pane of the Restore window and the Restore wizard displays the location of the Storage Group components as follows:

- SnapInfo directory
- EDB path
- Storage Group/Database log path

Access to Snapshot copies during database integrity verification

To specify the method to be used to access Snapshot copies during database integrity verification, use the Access LUN in Snapshot tab to assign either a drive letter or directory path to the Snapshot copy as a mounted LUN.

You can access this setting from the following locations within the SnapManager user interface:

- Configuration wizard

- Actions pane
- Backup wizard
- Restore wizard

Transaction log archiving

SnapManager archives transaction log files using NTFS hard links, which enables conservation of disk space in storage systems and improves the system performance.

During a backup operation, SnapManager archives transaction logs on the live file system to the SnapInfo directory. To optimize backup operations in which the Storage Group transaction logs and the SnapInfo directory are on the same NTFS volume, SnapManager archives transaction logs to the SnapInfo directory by creating NTFS hard links instead of performing a file copy operation. This conserves the disk space of the storage system and improves system performance during the transaction log archival phase of the backup operation.

NTFS hard links

A hard link is a directory entry for a file, and it serves as a file system-level shortcut to the file. Unlike an application-level link, any updates to the file contents are seen by all applications that access the file using a hard link. Unlike a symbolic link, a hard link created on one NTFS volume cannot point to a file in a different NTFS volume.

You can consider every file to have at least one hard link. Within a particular NTFS volume, a single file can have multiple hard links. Having multiple hard links enables a single file to appear in multiple directories or to appear multiple times (under different names) within the same directory. Because all the links reference the same file, multiple applications can open any of the hard links and modify the same file.

Note: The NTFS file system does not delete a file until all hard links to the file are deleted.

When you create a hard link for a file on an NTFS volume, NTFS adds hard link information and a reference count to the file's directory entry at the NTFS level. Creating the hard link does not create a duplicate original file or duplicate file-based reference (the application-level link).

While the physical file remains intact in its original location, you can access the same content by two or more names. Because the file system is responsible for managing the various path names to a single physical content, using a hard link can conserve system resources and storage system disk space compared to performing a standard file copy operation.

Why SnapManager uses NTFS hard links for transaction log archiving

By creating NTFS hard links, SnapManager archives transaction log files without physically copying the log file. During a backup operation, SnapManager archives transaction log files on the live file

system to the SnapInfo directory, from which the files later can be retrieved, during a restore operation.

When SnapManager uses a file copy operation to copy the transaction log files from the live file system to the SnapInfo directory, the following resources are consumed:

- Management of path names and the copy operation incurs file system overhead and storage system overhead, resulting in slower host and storage system performance.
- Each additional copy of a transaction log file not only consumes storage system disk space but also increases the use of overwrite reserve for the LUN.

SnapManager 5.0 and later supports an optimization feature for cases in which the backup operation archives transaction logs to a SnapInfo directory that resides on the same volume as the transaction logs. To achieve this optimization, SnapManager archives transaction logs to the SnapInfo directory by creating NTFS hard links to the live transaction log files, without physically copying the log file.

After the backup operation completes, Exchange truncates the transaction logs. However, the physical log files on the original location are not deleted, because the NTFS file system detects that other hard links to the files still exist. Instead, the original NTFS hard links to the live transaction log files are removed, and the links to the SnapInfo directory remain intact and available for access during SnapManager restore operations.

SnapManager automatically uses NTFS hard links to archive transaction logs if they reside on the same LUN as the SnapInfo directory.

The entries that are logged in the backup report depend on the archival method that you use. If a transaction log file is archived by performing a file copy operation, for example, the backup report entry is similar to the following:

```
Log file copying from:
L:\Program Files\Exchsrvr\mdbdata\E000000A.log
Log file copying to:
S:\SME_SnapInfo\EXCH__CLPUBS-WINSRVR3\SG__First Storage Group\CLPUBS-
WINSRVR3__recent\Logs\E000000A.log
```

If the same transaction log file is archived by creating an NTFS hard link instead of performing a file copy operation, the backup report entry is similar to the following:

```
Log file:
L:\Program Files\Exchsrvr\mdbdata\E000000A.log
Archived to:
S:\SME_SnapInfo\EXCH__CLPUBS-WINSRVR3\SG__First Storage Group\CLPUBS-
WINSRVR3__recent\Logs\E000000A.log
```

Support for multiple SnapInfo directories

The main advantage of having multiple SnapInfo directories is the support of NTFS hard links for archiving transaction logs, which helps to free the system resources and disk space. The use of

multiple SnapInfo directories also minimizes your exposure to the loss of backups if a SnapInfo directory is lost.

Multiple SnapInfo directories enable copy-less transaction log archiving with the use of NTFS hard links. The transaction logs must reside on the same NTFS volume as the SnapInfo directory. If you configure an environment with a single SnapInfo directory for all Storage Groups, SnapManager can still use NTFS hard links to archive transaction logs.

In most SnapManager environments, however, the transaction log files are stored on multiple LUNs in multiple volumes. In this configuration, the use of NTFS hard links to archive transaction logs requires a separate SnapInfo directory for each Storage Group. Each SnapInfo directory should be placed on the same NTFS volume as the transaction logs for that Storage Group.

Example Exchange configurations supported with SnapManager

Learn about the sample configurations for single and multiple storage group and individual database restore operation that satisfy the configuration rules for SnapManager.

Example: single Storage Group and individual database restoration not required

When you do not need to restore each database in your configuration individually, you can configure your system to contain a Storage Group; a LUN to contain all of your databases; another LUN to contain transaction logs and SnapInfo directory; and an Exchange host. A LUN can contain more than one database.

The following illustrated example configuration has one Storage Group containing up to five databases. You require two LUNs for this configuration, as shown in the following diagram:

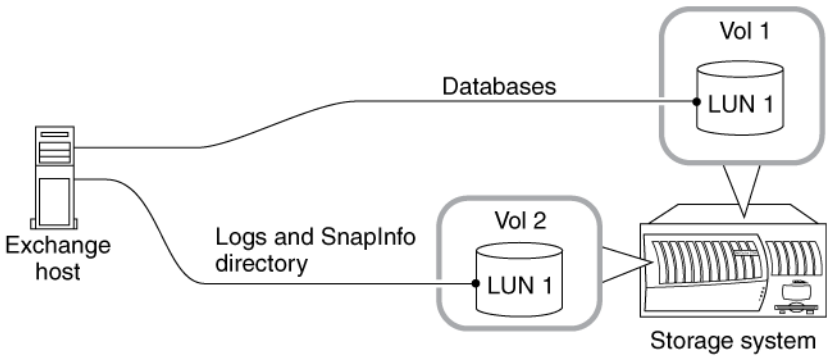

Sample configuration 1

Example: single Storage Group and individual database restoration required

When you need to restore each database in your configuration individually, you can configure your system to contain a Storage Group; a LUN for each of the individual databases; another LUN to contain transaction logs and SnapInfo directory; and an Exchange host.

The following illustrated example configuration has one Storage Group with two databases. Because you require the ability to restore either of the databases individually, place each database on its own LUN. Therefore, you require three LUNs for this configuration, as shown in the following diagram.

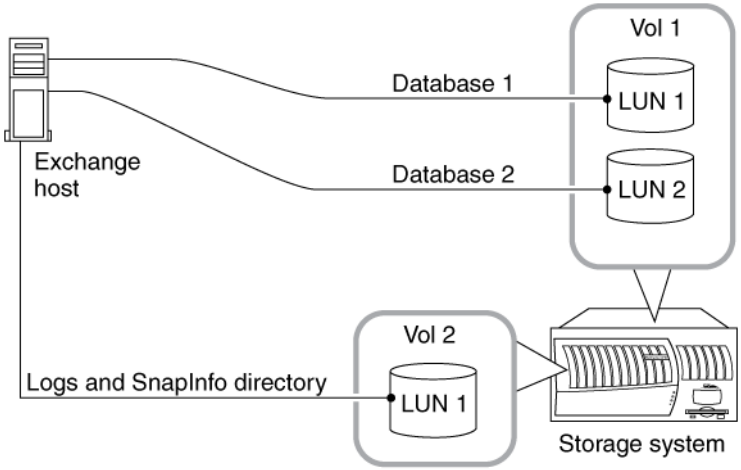

Sample configuration 2

Example for multiple storage groups and individual database restoration required

When you need to restore individual databases in a configuration having multiple storage groups, you can configure your system to contain LUNs for several of your databases; separate LUNs for any databases that you need to restore individually; another LUN to contain transaction logs and SnapInfo directory; and an Exchange host.

The following illustrated example configuration has two storage groups with three databases each. In this configuration, you can restore Database C in Storage Group 1 without restoring Databases A or B, because Database C has its own LUN. However, you have to restore Databases A and B at the same time, because they share a LUN. You require a total of five LUNs for this configuration, as shown in the following diagram.

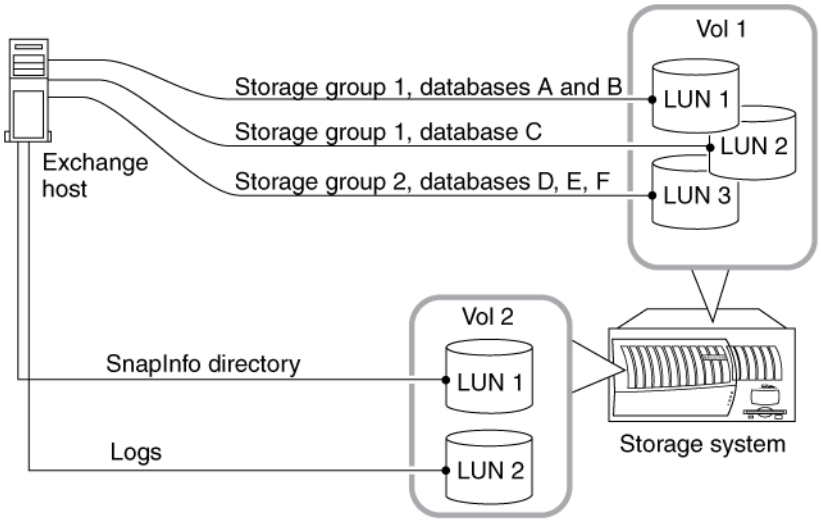

Sample configuration 4

Example: multiple Storage Groups and individual database restoration not required

When you do not need to restore individual database in a configuration having multiple Storage Groups, you can configure your system to contain multiple Storage Groups; LUNs for all of your databases; another LUN to contain transaction logs and SnapInfo directory; and an Exchange host.

The following illustrated example configuration has three Storage Groups with up to five databases each. You require four LUNs for this configuration, as shown in the following diagram. This applies only to Exchange Server 2003.

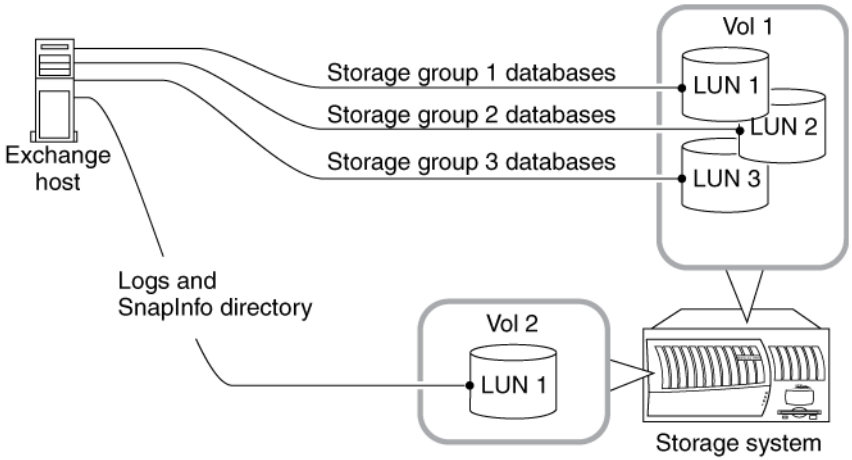

Sample configuration 3

Creation of LUNs on qtrees

To avoid performance setbacks, you should create your LUNs only on the qtrees that contain volumes that have UNIX as the default security type. You can check if your qtree is a UNIX qtree by using the Data ONTAP command `qtree status`. For more information, see Data ONTAP documentation.

Storage system volume and LUN planning

After you determine how many LUNs you need for your SnapManager configuration and what data those LUNs hold, you can develop your SnapManager data configuration plan and prepare storage system volumes and LUNs for use with SnapManager.

For the most current information about storage system volume and LUN planning, see *TR article 4033*, *Microsoft Exchange Server 2010 and SnapManager for Exchange Best Practices Guide*.

Configuration and migration of Exchange data using SnapManager

The SnapManager Configuration wizard enables you to select database verification servers, move Exchange databases and transaction logs to Data ONTAP LUNs, and configure automatic event notification.

You can perform the following configuration and migration tasks using SnapManager for Exchange:

- Implement the planned storage system layout
- Configure Microsoft Exchange data
- Migrate LCR-enabled Exchange data
- Migrate CCR-enabled Exchange data
- Migrate Exchange 2010 databases from local disks/non N series storage LUNs to N series storage LUNs

Attention: You must run SnapManager from the system console.

The SnapManager Configuration wizard

You can use the Configuration wizard to migrate Exchange databases from your local disk to Data ONTAP LUNs, so that you can make a backup copy and restore the databases by using SnapManager.

Use the SnapManager Configuration wizard to move Exchange databases, transaction logs, or system files for a Storage Group. SnapManager ensures that these files are placed in locations that meet SnapManager configuration requirements. Incorrectly located Exchange databases, transaction logs, or system files for a Storage Group impair SnapManager operation. If Exchange System Manager is used to move Exchange databases, transaction logs, or system files for a Storage Group, run the SnapManager Configuration wizard after the operation, to ensure that these files are in correct locations.

Note: Protection Manager integration can be enabled without reinstalling SnapManager for Exchange anytime by rerunning the Configuration wizard. You must first ensure that you have configured SnapDrive for Protection Manager integration. If you have not already done so, you can set up SnapDrive Protection Manager integration via the `sdcli` command, and then restart the SnapDrive service, or you can rerun the SnapDrive setup. (See the SnapDrive help for more information on Protection Manager integration.)

What the SnapManager Configuration wizard does

The Configuration wizard enables you to move your Exchange data in the following ways—from local disk to LUN, from LUN to LUN, or from LUN to local disk.

From local disk to LUN

If you want to move Exchange databases, you can use the wizard to unmount the Exchange databases in a Storage Group, move the Exchange databases and transaction log files to the selected LUN, and remount the databases.

Note: SnapManager takes Exchange databases offline during the transfer operation.

The wizard creates a SnapInfo directory that SnapManager uses to store information about the backup sets and the backed up transaction logs.

The wizard disables circular logging for all Storage Groups that are moved to a LUN, to enable up-to-the-minute restoring of those Storage Groups.

The wizard also guides you through several application settings. These settings include enabling notification of SnapManager events using e-mail, Syslog, or the AutoSupport feature.

From LUN to LUN

Resource management issues might require LUN-to-LUN transfer of Exchange data. An example of a situation that requires a LUN-to-LUN transfer is when you want to consolidate the Exchange data on another storage system.

From LUN to local disk

You need to move Exchange data from a LUN to a local disk if the data is no longer managed by using SnapManager. However, Microsoft Cluster Services configurations do not support Exchange data transfer from LUN to local disk.

Note: Protection Manager integration can be enabled without reinstalling SnapManager for Exchange anytime by rerunning the Configuration wizard. You must first ensure that you have configured SnapDrive for Protection Manager integration. If you have not already done so, you can set up SnapDrive Protection Manager integration via the `sdcli` command, and then restart the SnapDrive service, or you can rerun the SnapDrive setup. (See the SnapDrive help for more information on Protection Manager integration.)

Note: Ensure that all databases in a Storage Group are migrated to LUNs from a local disk. SnapManager does not support heterogeneous configuration. An example of a heterogeneous configuration is Storage Groups with some databases that reside on LUNs, and some on a local disk. Exchange databases must reside on NetApp LUNs to enable SnapManager to back up.

When to use the SnapManager Configuration wizard

You can use the Configuration wizard to move the databases and transaction logs to the newly configured LUN, to modify data configuration settings, or to validate the databases that you moved or added. You must re-run Configuration wizard if you change the Exchange configuration manually.

For initial configuration

Before using SnapManager to create a backup and to restore an Exchange database, you must use the SnapManager Configuration wizard to migrate the databases and transaction logs from your Exchange servers to the LUNs you configured on your storage system with SnapDrive.

To view or change the database configuration

After the initial configuration, you can re-run the Configuration wizard at any time to review or make changes to your Exchange database configuration.

To validate the database configuration

If you add databases, or move databases to different LUNs without using SnapManager, run the Configuration wizard to ensure that the databases are stored in valid locations and to map those databases to their respective SnapInfo subdirectories.

If you changed the Exchange configuration manually

Whenever you change the Exchange configuration manually using SnapDrive or the Exchange System Manager, you need to re-run the Configuration wizard to inform SnapManager that the Exchange configuration has been modified. For example, when you manually change the SnapManager configuration from using volume mountpoints to using drive letters, or from using drive letters to using volume mountpoints, you need to re-run the Configuration wizard before the changes take effect. Another reason to re-run the Configuration wizard is when a new database is added by using the Exchange System Manager.

Settings configurable only with the SnapManager Configuration wizard

You can configure certain settings of LUN mappings, SnapInfo directory, iSCSI service, and automatic event notification by using only the Configuration wizard, and not other wizards or windows.

SnapManager setting	Configuration wizard page
Mapping your database to LUNs	Select a Storage Group/Database to move to a LUN
Mapping your transaction log files to LUNs	Select a set of logs from one or more Storage Groups/Databases to move to a LUN
SMTP and MTA file-LUN mappings	Configure SMTP and MTA Services **Note:** SMTP and MTA system files are applicable to Exchange Server 2003. The Hub Transport server role in Exchange Server 2007 and Exchange Server 2010 uses the SMTP system file alone. **Note:** This setting is skipped when you connect to Exchange Server 2010.
Specifying the SnapInfo directory location	Configure the SnapInfo directory to store the backup information
Adding Microsoft iSCSI Service as a dependency	Add Microsoft iSCSI Service Dependency
Configuring SnapManager event notification enabling for SMTP e-mail, storage system Syslog, and AutoSupport	Configure Automatic Event Notification (Recommended)

Placement of Exchange and SnapManager components

When you use the Configuration wizard, you have to specify the placement of several components of Exchange and SnapManager like databases, transaction logs, SnapInfo directories, SMTP, and MTA queues.

Storage Group system path	The path to the directory that contains the Exchange system files for the Storage Group or database SnapManager places system files for a Storage Group on the same LUN that stores the transaction logs for that Storage Group. The system files for a Storage Group must remain where the Configuration wizard places them.
Storage Group databases	The databases in a Storage Group or data store
Database	The paths for the `.edb` file for the Exchange database The `.edb` file stores message and attachment content that is generated by Microsoft Messaging Application Program Interface (MAPI) clients, and message properties such as author, recipients, subject, and priority. The `.stm` file holds native Internet content (for example, Multipurpose Internet Mail Extensions (MIME)). **Note:** The `.stm` file is supported only in Exchange Server 2003.
Transaction log files	The location of the Exchange transaction logs The transaction logs contain changes made to the Exchange databases since the last backup, enabling an up-to-the-minute restore operation.
SMTP queue path	The directories that are used by the Simple Mail Transfer Protocol (SMTP) queue
MTA queue path	The directories that are used by the Message Transfer Agent (MTA) queue **Note:** The MTA queue path is applicable only to Exchange Server 2003. Exchange Server 2007 uses the SMTP queue path for its Hub Transport server role. SnapManager copies the queue database, and updates the database path in a configuration file, to move the queue databases.
SnapInfo directory	Contains SnapManager Backup information, copies of transaction log files, and other data critical to the backup set

Viewing the placement of Exchange and SnapManager components

Viewing the placement of Exchange and SnapManager components helps you in your configuration planning.

About this task

You can also see database placement information from within the Configuration wizard if you highlight a database displayed in a wizard page, and click **Details**.

Exchange Server 2003 supports both SMTP and MTA system files, whereas Exchange Server 2007 and Exchange Server 2010 support only SMTP data files.

Steps

1. In the Scope pane, click **Backup**.
2. Select the component in the Results pane.

   SnapManager displays the component properties in the Results pane.

Exchange Server storage groups and databases migration and configuration considerations

Before you migrate your database using the Configuration wizard, consider how the different requirements for stand-alone systems and Windows clusters can impact where the SMTP and MTA queues and the report directories should reside.

The SMTP and MTA queue location

The Simple Mail Transfer Protocol (SMTP) and Message Transfer Agent (MTA) queues are used by Exchange. Ensure that these files do not reside on the same LUN as your Exchange databases.

For stand-alone systems, the SMTP and MTA system files can be placed on any disk, or on a LUN that does not contain Exchange databases (the default location of the original local installation disk is acceptable).

For Windows clusters, the SMTP and MTA system files must be moved to a shared LUN that does not contain Exchange databases. This is the data directory that Exchange prompts you to specify when you create a new Exchange virtual server on a cluster. The data directory becomes the default system path. The Configuration wizard places the MTA and SMTP queues there by default.

> **Note:** In Exchange Server 2007, the SMTP queue is valid only if the bridgehead, or the Hub Transport server role is installed. The MTA queue is not applicable to Exchange Server 2007 and Exchange Server 2010.

Report directories in a Windows cluster

In a Windows cluster you might want to have one set of report directories, no matter which node SnapManager runs on. Move your report directory to a shared disk resource that does not contain Exchange data or SnapManager data.

Migrating and configuring your Exchange Server storage groups and databases using the SnapManager Configuration wizard

To manage your database using SnapManager, you need to migrate it from your local disk and configure it on the storage systems assigned for use by SnapManager.

Before you begin

If the computer that you want to use for database verification does not have Exchange installed, you must configure the necessary files on it according to Windows host system requirements.

Ensure that you go through documentation on automatic e-mail notification settings before you start the configuration.

About this task

For configuring archived database by using Protection Manager with SnapManager, SnapManager currently allows only the Backup and Remote backups only policies to be implemented. Back up is selected as the default option. You can change the policy by using the NetApp Management Console.

You can configure the following automatic e-mail notification settings using the Configuration wizard:

- Automatic SMTP e-mail notification for event notification
- Enable Syslog (storage system event log) support
- Enable AutoSupport for SnapManager
- Limit event logging to failure events

After you complete the Configuration wizard, you can also view or change these settings by selecting **Notification Settings** from the Actions pane to access the Auto Notification Settings dialog box. You should enable all of the automatic event notification options; this ensures that if any problems with your SnapManager backups occur, you are alerted and can take corrective action in a timely manner.

The list of available drives changes as you make choices that rule out other options.

Any time you use the Configuration wizard to move data from one drive to another, backups performed before the migration become invalid. After you migrate your databases to a new LUN, always perform an immediate backup.

Steps

1. In the **Actions** pane, click **Configuration wizard**.
2. If you want to use a control file for your configuration, click **Use Control File**, and then specify the path to the control file that you want to use.
3. In the **Database Verification Server** page, specify the Exchange server to be used to perform Exchange database verification, or select the option **Select a verification server later using the Options menu**.

   If you select the option **Select a verification server later using the Options menu**, SnapManager for Exchange will use the existing verification server if you have already specified one.
4. In the **Select a Storage Group/Database to move to a LUN** page, associate your CCR-enabled Exchange Storage Groups or databases with the correct LUNs, according to your SnapManager data configuration plan.
5. In the **Select a set of logs from one or more Storage Groups/Databases to move to a LUN** page, associate your transaction logs with the correct LUNs, according to your SnapManager data configuration plan.

6. In the **Configure the SnapInfo directory to store the backup information** page, select the SnapInfo files, and the LUN to which you want to move them and click **<=>**.
7. If you have datasets configured on your system, in the **Configure Dataset for Backup Archival** page, select the Storage Groups that you want to add to your dataset, and click **Next**.
8. Choose the default protection policy as **Back up** or **Remote backups only** as per your requirement.
9. In the **Add Microsoft iSCSI Service Dependency** window, select **Yes, add the..** if you want to add iSCSI service as a dependency for the MSExchangeSA service.
10. Use the **Configure Automatic Event Notification (Recommended)** page to configure the automatic event notification options for SnapManager.
11. Review the configuration summary in the **Completing the Configuration Wizard** page, and click **Finish**.
12. Click **Start Now** to migrate your databases and their transaction logs and SnapInfo files to the LUNs you specified.

    Note: If you are moving the location of Exchange databases in this step, it could take some time to complete.

13. Click **OK**.

After you finish

Create a SnapManager backup immediately.

Moving a Storage Group to a LUN

The Configuration wizard enables you to select the LUNs in which you want to store your Exchange databases when you complete the wizard. You use the Select a Storage Group/Database to move to a LUN window to select the LUNs.

About this task

The list of available drives changes as you make choices that rule out other choices.

Steps

1. In the **Actions** pane, click **Configuration wizard**.
2. In the **Select a Storage Group/Database to move to a LUN** page, select the Storage Group that you want to move.
3. In the **Disk** list, select the drive letter for the LUN, or the local disk that you want to move the Storage Group or database to.
4. Click **<=>**.

The selected Storage Group or database and disk gets associated, and the Storage Group files or database are then moved to that disk when the Configuration wizard moves your data.

5. Step through the wizard, and review the configuration summary in the **Completing the Configuration Wizard** page, and click **Finish**.
6. Click **Start Now**.

   Note: If you are moving the location of Exchange databases in this step, it could take some time to complete.

7. Click **OK**.

Moving an individual database to a LUN

The Configuration wizard enables you to move an individual database to a LUN. You use the Select a Storage Group/Database to move to a LUN page to move an individual database to a LUN.

About this task

The list of available drives changes as you make choices that rule out other choices.

Steps

1. In the **Actions** pane, click **Configuration wizard**.
2. In the **Disk** list, select the drive letter for the disk that you want to move the database to.
3. Click **<=>**.

   The selected database and the disk gets associated, and the database files are then moved to that disk, when the Configuration wizard moves your data.

4. Step through the wizard, and review the configuration summary in the **Completing the Configuration Wizard** page, and click **Finish**.
5. Click **Start Now**.

   Note: If you are moving the location of Exchange databases in this step, it could take some time to complete.

6. Click **OK**.

Changing the location of a Storage Group or database

The Configuration wizard enables you to change the location of a Storage Group that you had associated with a LUN, or local disk, if you want to change your configuration plan.

About this task

The list of available drives changes as you make choices that rule out other choices.

If the Exchange system file path for a Storage Group is on the database LUN (placed for a previous version of SnapManager), and then you use the Configuration wizard to reconfigure a database path

of the Storage Group or file path of the transaction log, then the Configuration wizard automatically moves the system file path to the same LUN as the file path of the transaction log.

Steps

1. In the **Actions** pane, click **Configuration wizard**.
2. Under **Database Location Results**, highlight the Storage Group you want to move, and click **Reconfigure**.

   The selected Storage Group or database is returned to the list of unconfigured databases, and you can re-associate it with any available disk.
3. Step through the wizard, and review the configuration summary in the **Completing the Configuration Wizard** page, and click **Finish**.
4. Click **Start Now**.

   Note: If you are moving the location of Exchange databases in this step, it could take some time to complete.
5. Click **OK**.

Viewing the full path for database files

The Configuration wizard enables you to view the full path for your database, or Storage Group files.

Steps

1. In the **Actions** pane, click **Configuration wizard**.
2. In the Configuration wizard, highlight a Storage Group or a database under **Database Location Results**, and click **Details**.

   The Storage Group/Database Settings window appears for the selected Storage Group or database.
3. To see the file path for the databases in the selected Storage Group, highlight the database and click **Details**.

   The current path and file name for the selected database are displayed, along with the new path and file name if the location of the selected database has changed.
4. Step through the wizard, and review the configuration summary in the **Completing the Configuration Wizard** page, and click **Finish**.
5. Click **Start Now**.

   Note: If you are moving the location of Exchange databases in this step, it could take some time to complete.
6. Click **OK**.

Dataset protection policy

You have to specify the protection policy for your datasets using the Configuration wizard. The protection policy assigned to a dataset determines how replications are done. It also determines how long backups are retained on the secondary storage SnapVault target.

You can assign any these protection policies to the dataset:

- Back up
- Remote backups only
- Any customized policy based on the above two policy templates

Both the **Back up** policy and the **Remote backups only** policy have scheduled SnapVault replication that happens in the background along with "on demand" SnapVault replication that happens when you do a backup or verify operation with SnapManager.

The "backup retention type determines the duration for which the archived backups are retained and which backups are verified during a "verification only" operation.

Migration of transport database paths for Exchange Server 2010

The Configuration wizard cannot be used to migrate transport database paths for Exchange Server 2010. The Configuration wizard does not support migration of transport database paths of Exchange 2010 transport server role.

When SnapManager runs on Exchange Server 2010, the option to Configure SMTP and MTA Services is skipped on the Configuration wizard. You cannot access this option when you connect to the Exchange 2010 server.

To migrate transport database or temporary storage to NetApp LUNs, use the Exchange script `move-transportdatabase.ps1` normally found under `C:\Program Files\Microsoft\Exchange Server\V14\Scripts`.

The script can change the paths of the following:

- QueueDatabasePath <path>
- QueueDatabaseLoggingPath <path>
- IPFilterDatabasePath <path>
- IPFilterDatabaseLoggingPath <path>
- TemporaryStoragePath <path>

How to move transaction logs using the SnapManager Configuration wizard

You can use the Configuration wizard to move transaction logs to a LUN, to change the location of the transaction logs, and to view their full path. You use the Where to Move Transaction Logs window to perform the tasks.

Note: The list of available drives changes as you make choices that rule out other choices.

Moving transaction logs to a LUN

You can move the transaction logs to a LUN by using the Configuration wizard.

Steps

1. In the **Actions** pane, click **Configuration wizard**.
2. From the **Storage Group/Database** list in the Configuration wizard, select the Storage Group or database whose transaction logs you want to move.
3. From the **Disk List**, select the drive letter for the LUN or local disk you want to move the transaction logs to.
4. Click **<=>**.

   The transaction logs for the selected Storage Group or database and LUN gets associated, and the transaction log files are then moved to that LUN when the Configuration wizard moves your data.
5. Step through the wizard, and review the configuration summary in the **Completing the Configuration Wizard** page, and click **Finish**.
6. Click **Start Now**.

   Note: If you are moving the location of Exchange databases in this step, it could take some time to complete.
7. Click **OK**.

Changing the location of transaction log files

You can reconfigure the location of the transaction log files by using the Configuration wizard.

Steps

1. In the **Actions** pane, click **Configuration wizard**.
2. In the Configuration wizard, go to **Transaction Logs Location Results**, and highlight the Storage Group or database for the transaction files that you want to move, and click **Reconfigure**.

The selected Storage Group or database is returned to the list of unconfigured transaction logs, and you can re-associate it with any available LUN or local disk.

3. Step through the wizard, and review the configuration summary in the **Completing the Configuration Wizard** page, and click **Finish**.
4. Click **Start Now**.

   Note: If you are moving the location of Exchange databases in this step, it could take some time to complete.

5. Click **OK**.

Viewing the full path for transaction log files

You can view the full path of the transaction log files from the Configuration wizard.

Steps

1. In the **Actions** pane, click **Configuration wizard**.
2. In the Configuration wizard, highlight a Storage Group or database in **Transaction Logs Location Results**, and then click **Details**.

   The **Transaction Logs Path** window appears for the selected Storage Group or database. If the location of the transaction log files has changed, the current paths and file names for the transaction logs, for the selected database are displayed, along with the new paths and file names.

3. Step through the wizard, and review the configuration summary in the **Completing the Configuration Wizard** page, and click **Finish**.
4. Click **Start Now**.

   Note: If you are moving the location of Exchange databases in this step, it could take some time to complete.

5. Click **OK**.

Configuration using the SnapManager control file

SnapManager creates a control file that contains configuration information about the Exchange server. You can use the control file either to import or to export the Exchange server configuration.

The configuration data in the control file is represented in XML format. You can edit the file manually using an XML editor.

The configuration data in the control file is grouped into the following sections, so that you can manage it more easily:

- Storage layout settings
- Notification settings
- Verification settings
- Report directory settings
- Backup settings
- Scheduled jobs
- SnapMirror relationship settings

Instead of running the Configuration wizard every time you want to migrate databases and transaction logs to LUNs, you should use the control file in the following scenarios:

- Large configurations
- Business continuance or Disaster Recovery
- Mass deployment

Configuration of SnapManager in a DAG using a control file

You can use a control file instead of the configuration wizard to configure SnapManager. The control file support for Exchange Server 2010 mailbox databases is based at the Database Availability Group (DAG) level and not at a member server level.

You can export the configuration settings of a DAG to another DAG or import the configuration settings from another DAG. This is possible by using either a SnapManager snap-in or a SnapManager for Exchange PowerShell cmdlet.

Because all database instances in all member servers of a DAG have to be migrated at the same time to keep database replication operating properly, the control file support for Exchange Server 2010 mailbox databases is DAG based, not member server based. It is not possible to export the configuration of a member server in a DAG to a control file and then import that control file to another member server in the same DAG.

Importing Exchange Server configuration information using the control file

You can configure an Exchange server by using a control file containing configuration settings exported from another Exchange server. You can import configuration details and specific sections of the Exchange configuration from a control file.

Steps

1. In the **Actions** pane, click **Configuration wizard**.
2. Select **Use Control File**.
3. To import configuration details for the server, click **Next**.

   The Import or Export Selection window appears.
4. Select **Import**.
5. You can review your configuration settings by selecting **Review current settings in the Configuration wizard**.
6. In **Use control file**, either type the complete path of the control file or use the Browse feature to locate and select the file path.
7. Click **Advanced**.
8. In the **Configuration Import/Export Advanced Options** window, specify the configuration settings that need to be imported, or exported.
9. Select **OK**.
10. Click **Next** and then, to load the control file and validate the imported configuration and settings, click **Apply**.

    If you do not want to apply the new configuration, click **Cancel**.

    If you click **Apply**, SnapManager loads the control file, and validates the imported configuration and settings.

Exporting Exchange Server configuration information using the control file

You can export specific settings of an Exchange server configuration to a control file that you can use to configure other Exchange servers. You can export the current configuration details to a control file, and export a specific section of the current configuration to a control file.

Steps

1. In the **Actions** pane, click **Configuration wizard**.
2. Select **Use Control File**.
3. To export configuration details for the server, click **Next**.

   The Import or Export Selection window appears.
4. Select **Export**.
5. In **Use control file**, either type the complete path of the control file or use the Browse feature to locate and select the file path.
6. Click **Advanced**.
7. In the **Configuration Import/Export Advanced Options** window, specify the configuration settings that need to be exported.
8. Select **OK**.
9. Click **Next** and then, to load the control file and validate the imported configuration and settings, click **Apply**.

   If you do not want to apply the new configuration, click **Cancel**.

   If you click **Apply**, SnapManager loads the control file, and validates the imported configuration and settings.

Migration of LCR-enabled databases

You can move the databases, transaction logs, and system files with Local Continuous Replication (LCR) replicas to LUNs. You can follow the same procedure as you use for the production Storage Group for migrating the databases, transaction logs, and system files.

In an LCR-enabled database, you can perform the following actions by using SnapManager:

- Move Storage Groups to a LUN
- Move databases to a LUN
- Change the location of a Storage Group or database
- View the full path of the database files
- Change the location of a transaction log file
- View the full path of transaction log files
- Seed the databases in an LCR copy
  See your version of Microsoft Exchange Server 2007 Installation Guide.

Prerequisites for configuring an LCR-enabled Exchange server

Before you use SnapManager to enable Local Continuous Replication (LCR) on Exchange, ensure that the Exchange server is LCR-enabled, the databases in the Storage Group are seeded, and that you put the database, transaction logs and replica database on a different volume from the LCR environment.

Ensure that you meet the following criteria before you enable the LCR feature on an Exchange server:

- The Exchange server on which you want to perform SnapManager configuration is LCR enabled.
  Use the Microsoft Exchange cmdlet or the Exchange Management Console to enable LCR on a Storage Group. SnapManager does not automatically enable LCR on a Storage Group if it is not enabled already.
- The databases in the Storage Groups are seeded.
  If the replica database in an LCR is not seeded or broken, SnapManager seeds it during migration.
- In an LCR environment, you should put the database, transaction log files of the production database, and replica database on a different volume than that which contains the LCR relationship.

Comparison of moving production and LCR replica storage groups

The main difference between the production Storage Group and an Local Continuous Replication (LCR) replica Storage Group is the way that it is represented on the SnapManager user interface. A Storage Group with an LCR replica appears as "Storage Group (LCR)."

Database seeding in an LCR

A database (either blank, or a copy of a production database) is added to the Storage Group copy by seeding. SnapManager performs database seeding in an LCR copy during a SnapManager migration operation, if a Storage Group is configured for LCR, and if a replica database is not seeded.

The added database becomes the baseline database for the copy. In Exchange Server 2007, after an LCR replica is enabled on the Storage Group, the LCR replica is automatically seeded.

When seeding is complete, the copy is given a status of good health. To reduce the I/O burden on the active Storage Group volumes, seed the database during an off-peak time. The existing database needs to be deleted before the seeding process.

SnapManager uses the Exchange cmdlet `Update-StorageGroupCopy` to seed the LCR replica. The `Update-StorageGroupCopy` cmdlet initiates or re-synchronizes the copying for a specified Exchange Server 2007 LCR.

Migration of CCR-enabled databases

You can migrate Cluster Continuous Replication (CCR) databases and transaction logs to a specified location to enable SnapManager to manage it. The configuration is similar to that of the production Storage Group; SnapManager internally migrates databases to the specified location on both the nodes.

CCR uses three computers or nodes combined in a single cluster. Two nodes host a clustered mailbox server. CCR uses the third node, referred as "voter," to avoid an occurrence of network partition within the cluster. A node that is currently running a clustered mailbox server is an active node, and a node that is not running a clustered mailbox server is a passive node.

You can connect to both the CCR active node and passive node at the same time. You can also migrate Storage Groups from either node.

When you connect to a CCR-enabled Exchange server, SnapManager connects to both the active and the passive nodes by default.

Note: The Storage Groups must be in a healthy state before you can run the Configuration wizard to migrate the CCR-enabled database.

Prerequisites for configuring a CCR-enabled Storage Group

Ensure that you enable the Continuous Cluster Replica (CCR) feature on the Exchange server, the cluster continuous replication Storage Group copy is seeded and is healthy, SnapManager is installed on both nodes of the CCR, and the CCR active and passive nodes have identical dedicated LUNs before configuring a CCR-enabled Storage Group.

Before you plan to configure and migrate the CCR enabled Storage Group, ensure that you meet the following system requirements:

- The CCR feature is enabled on Exchange Server 2007.

  For more information about how to enable CCR on your version of Exchange Server 2007, see the Microsoft documentation for Exchange Server 2007.
- SnapManager is installed on both nodes of the CCR.
- The cluster continuous replication Storage Group copy is seeded, and it is in a healthy state.

  For more information about how to seed the databases in a CCR environment, see the Microsoft documentation for Exchange Server 2007.
- The CCR active and passive nodes have identical dedicated LUNs, with the same drive letter or mountpoint path.

  If a LUN exists only on one node, it cannot be used by CCR databases.

Note: Try to give the same LUN on each node the identical volume layout for the storage system. If LUNs E, F, and G on Node A are sharing a volume, LUNs E, F, and G on Node B also must share an identical volume.

Considerations for configuring SnapManager in a CCR-enabled Storage Group

Before configuring SnapManager in a CCR-enabled Storage Group, ensure that database verification server is configured and enable all of the automatic event notification options.

If the computer used for database verification does not have Exchange installed, you must put the necessary files on it.

You need to set up the database verification server on both active and passive nodes by performing the Backup Verification Settings action from both Cluster Continuous Replication (CCR) nodes, or by rerunning the Configuration wizard on both the nodes. The Configuration wizard does not set up the verification server on the remote CCR node.

Enable all of the automatic event notification options. Enabling automatic event notification ensures that if any problems occur with your SnapManager backups, you are alerted and can take immediate corrective action.

Configuring SnapManager in a CCR-enabled Storage Group

You can configure SnapManager in a Storage Group that has Cluster Continuous Replication (CCR) enabled.

Before you begin

Keep the name of the verification server ready.

About this task

SnapManager migrates a Storage Group or database from one location to another on both CCR nodes, because both nodes use the same database path.

SnapManager migrates transaction logs from one location to another on both CCR nodes. The transaction log path is same for both the active and the passive nodes.

The SnapInfo directory path is similar for both active and passive nodes.

You should configure your Microsoft iSCSI service as a dependency for the Microsoft Exchange System Attendant so that Microsoft Exchange cannot start before the Microsoft iSCSI service is ready to accept connections.

After you complete the Configuration wizard, you can also view or change the Notification settings using the **Auto Notification Settings** in the **Actions** pane. You should enable all of the automatic

event notification options. When you enable automatic event notification, if any problems occur with the SnapManager backups, you are alerted and can take immediate corrective action.

If you are moving the location of Exchange databases, it could take some time to complete.

Steps

1. In the Scope pane, double-click **SnapManager for Exchange**.
2. Click the Exchange server that is CCR-enabled, and for which the database seeding is complete.
3. In the **Actions** pane, click **Configuration Wizard**.
4. In the **Database Verification Server** window, specify the Exchange server to be used to perform Exchange database verification.
   a) If you want to specify the database verification server now, type the name of the Exchange server or browse to an Exchange server in the **Verification Server** window, or select **Select a verification server later**.
5. In the **Select a Storage Group/Database to move to a LUN** window, associate your CCR-enabled Exchange Storage Groups or databases with the correct LUNs, according to your SnapManager data configuration plan.
6. In the **Select a set of logs from one or more Storage Groups/Databases to move to a LUN** window, associate your transaction logs with the correct LUNs, according to your SnapManager data configuration plan.
7. In the **Configure the SnapInfo directory to store the backup information** window, select the SnapInfo files, and the LUN to which you want to move them and click **<=>**.
8. If you are installing SnapManager on a system that has the Microsoft iSCSI initiator installed, select whether you want to configure your Microsoft iSCSI service as a dependency for the Microsoft Exchange System Attendant.
9. In the **Configure Automatic Event Notification** (**Recommended**) window, configure the automatic Simple Mail Transfer Protocol (SMTP) e-mail notification for event notification, enable Syslog (storage system event log) support, enable AutoSupport for SnapManager, and limit event logging to failure events.
10. Review the configuration summary in the **Completing the Configuration Wizard** window, and click **Finish**.
11. To migrate your databases, and their transaction logs and SnapInfo files to the LUNs you specified, click **Start Now**.
12. Click **OK**.

After you finish

Create a SnapManager Backup immediately.

Migration of Exchange Server 2010 mailbox databases

You can migrate the Exchange Server 2010 mailbox databases and transaction logs to a specified location to enable SnapManager to manage them.

When the SnapManager MMC snap-in connects to a Database Availability Group (DAG), all mailbox databases in the DAG are available for migration. When the SnapManager MMC snap-in connects to a DAG member server, only the mailbox databases in that server are available for migration.

Prerequisites for migrating Exchange Server 2010 mailbox databases

Before you migrate the Exchange Server 2010 mailbox databases, ensure that an active database and its copy databases use the same path in all servers. You have to install SnapManager on all member servers of the Database Availability Group (DAG) before migrating the mailbox databases.

If SnapManager is not installed on a member server of the DAG, all databases hosted by that member server, including both the active and passive databases, will not be migrated by SnapManager, and SnapManager will not be able to back up databases on that server.

Ensure that the database replication status is healthy before the migration.

Available LUNs in a Database Availability Group

When the SnapManager MMC snap-in connects to the Database Availability Group (DAG), the available disks include all LUNs from all member servers of the DAG.

Migrating Exchange Server 2010 mailbox databases

To manage your database using SnapManager, you need to migrate it from your local disk and configure it on the storage systems assigned for use by SnapManager. You can migrate Exchange Server 2010 mailbox databases by using the SnapManager Configuration wizard.

About this task

SnapManager enables you to migrate the mailbox databases and transaction logs to LUNs, specify the database verification server, and configure automatic event notification.

Steps

1. In the **Scope** pane, select the Database Availability Group (DAG) or a member server of the DAG.
2. In the **Actions** pane, click **Configuration Wizard**.
3. In the **Database Verification Server** window, either type or browse to the name of the Exchange server that you want to be the database verification server, or select **Select a verification server later using the Options menu**.
4. In the **Select a Storage Group/Database to move to a LUN** window, associate your databases with the correct LUNs, according to your SnapManager data configuration plan.
5. In the **Select a set of logs from one or more Storage Groups/Databases to move to a LUN** window, associate your transaction logs with the correct LUNs, according to your SnapManager data configuration plan.
6. In the **Configure the SnapInfo directory to store the backup information** window, select the SnapInfo files, and the LUN to which you want to move them and click **<=>**.
7. If you have datasets configured in your system, in the **Configure Dataset for Backup Archival** window of the **Configuration wizard**, select the databases that you want to add to your dataset.

   Note: Dataset configuration is not available when you connect to the DAG.

8. Choose the default protection policy as **Back up** or **Remote backups only**, according to your SnapManager data configuration plan.
9. If you are installing SnapManager on a system that has the Microsoft iSCSI initiator installed, select whether you want to configure your Microsoft iSCSI service as a dependency for Microsoft Exchange System Attendant.
10. In the **Configure Automatic Event Notification** (**Recommended**) window, configure the automatic Simple Mail Transfer Protocol (SMTP) e-mail notification for event notification, enable Syslog (storage system event log) support, enable AutoSupport for SnapManager, and limit event logging to failure events.
11. Review the configuration summary in the **Completing the Configuration Wizard** window, and click **Finish**.

Guidelines for migrating to mount points for LUN mapping in Exchange 2007

In an Exchange 2007 environment, you can migrate an existing configuration that uses drive letters to map to Data ONTAP LUNs to a new configuration that uses volume mount points instead. The migration process uses SnapDrive and SnapManager.

Changing from drive-letter mapping to mount-point mapping briefly takes the affected Storage Groups offline as their paths are changed. You do not need to copy the databases themselves from one LUN to another. You save time, and Storage Groups go offline for less time.

Note: A drive letter is still used for the mount point root volume and not all drive letters are eliminated.

The mount point root can be any of the following:

- The local C:\ drive
- The LUN that stores the transaction log files
- A separate dedicated LUN
  The best practice is to use the LUN that stores the transaction log files. Using a separate LUN as the mount point root adds an additional point of failure to the Exchange server.

  Note: In clustered configurations, the mount point root must reside on a shared storage device, such as the transaction logs LUN.

For configurations using more than one Storage Group, use one LUN with a drive letter for the transaction logs of each Storage Group.

Because the System Attendant cluster resources require dependencies on the physical disk cluster resources in clustered configurations, the entire Exchange virtual server is brought offline during a migration, even if only a single Storage Group is migrated. Taking the physical disk resource for the transaction logs offline for one Storage Group takes the System Attendant resource offline and with all other Storage Groups.

Scenario: migrating an existing configuration from using drive letters to using mountpoints for LUN mapping

You can migrate an existing configuration that uses drive letters to map to Data ONTAP LUNs, to a new configuration that uses volume mountpoints instead. In this scenario, the migration process is described using a sample configuration.

About this task

The sample configuration contains the following components:

- A single Storage Group
- A single database located on LUN `I:\`

This scenario assumes the following points:

- The Exchange transaction logs, MTA and SMTP system files, and the SnapManager SnapInfo directory are collocated on LUN `L:\`.
- The mountpoint root LUN requires a drive letter, and the LUN containing the transaction logs is used as the mountpoint root.
- Double-clicking the mountpoint (directory) name takes you to the root of the mounted LUN rather than to an empty directory on the mountpoint root LUN.
- The original drive letter mapping, `I:\`, should not be removed until after you verify that the LUN mappings have been created successfully.
- For the Configuration wizard to function properly, SnapManager requires that the original drive letter mappings be in place along with the newly added mountpoints. Unneeded mappings are removed later.

Steps

1. Launch Windows Explorer, and then browse to the root of the transaction logs, LUN `L:\`.
2. For example, create a directory named database under `L:\`.

   Use a directory `database`.
3. Launch the SnapDrive MMC, right-click the database LUN `I:\`, and then select **Change Drive Letter and Paths**.
4. In **Change Drive Letter and Paths**, click **Add**.
5. When prompted, enter the path to the appropriate directory on your mountpoint root LUN, `L:\database`
6. Verify that the additional LUN mappings are successfully created and reflected in the SnapDrive MMC.

   In this example, the LUNs show drive letters of `L:\` and `L:\database\`.

   You can also browse to the root of the mountpoint root LUN, drive `L:\`. The icon for the folder database changes from the standard Windows folder icon to the Windows icon representing a mountpoint.
7. Launch SnapManager, and then run the Configuration wizard.
8. In the **Welcome** window, click **Next**.
9. In the **Database Verification Server** window, click **Next**.

   Although changes can be made here if desired, doing so is not part of this example.
10. In the **Exchange Server to Configure** window, verify that you are about to configure the correct server, and then click **Next**.

11. In the **Select a Storage Group/Database to move to a LUN** window, do the following: a. b. c.

a) Select a Storage Group listed in the Database Location Results field, and then click **Reconfigure**.

The Storage Group appears in the list of unconfigured databases, and you can reassociate it with any available disk.

b) Select LUN `L:\database\`, and then click **<=>** to associate the database with the new path.

The Storage Group appears in the Database Location Results field again. The old drive letter LUN I appears under the From column. The new mountpoint path `L:\database\:`appears under the To column.

c) Click **Next**.

12. Keep clicking **Next** to proceed through the next few windows, making additional configuration changes if desired. At the final window, click **Finish**.

The Configurator Status window appears.

13. Click **Start Now**.

The databases are taken offline, and the migration proceeds.

14. After the SnapManager Configuration wizard completes successfully, return to the SnapDrive MMC.

15. Using the SnapDrive MMC, right-click the database `LUN I:\`, and then select **Change Drive Letter and Paths**.

16. In Change Drive Letter and Paths, highlight the original drive letter for the database LUN (I:\) and click **Remove**.

The SnapDrive MMC displays the two LUNs `L:\` and `L:\database`.

SnapManager Backup overview

SnapManager creates backups of your databases that are used to further restore and recover your databases. Your database is stored as a backup set that also contains transaction logs and metadata for your database. SnapManager also performs integrity verification of your database.

You can create or schedule a SnapManager backup by using the SnapManager Backup wizard or the Backup and Verify window. You can also manage the transaction logs and the number of backups you need to retain for your database.

Note: To ensure that your database is backed up by SnapManager for Exchange, you must place the LUNs for your database, logs, and snapinfo in a qtree.

How SnapManager Backup works

SnapManager Backup uses the Snapshot technology to create online, read-only copies of databases. After you back up the selected Storage Groups, SnapManager deletes the transaction logs that are committed to the databases.

Copy backup support

Although previous releases included only full backup and remote copy backup in a Cluster Continuous Replication configuration, starting with the 6.0.2 release, SnapManager for Microsoft Exchange contains new backup features available as advanced options of the backup dialog box and Backup wizard.

You can perform the following two types of backup:

Full backup Choosing this backup type produces a backup of the selected databases and truncates their transaction logs.

Copy backup Choosing this backup type produces a backup of the selected databases but does not truncate their transaction logs.

Copy backup information is not logged into AutoSupport and EMS. You can back up truncated logs with a copy backup, but it is not necessary to do so because these logs are not required for a restore operation.

The same backup operations as full backup also apply to copy backup: for example, backup verification SnapMirror updates, and archive backup using SnapVault with a dataset.

The SnapInfo directory

SnapInfo is the directory storing backup metadata. A snapshot will not be a backup if there is no backup metadata associated with it. The SnapInfo directory is the location storing transaction log

backups. SnapManager for Exchange uses snapshot to protect SnapInfo directory. SnapManager uses NTFS hardlink for log backups in SnapInfo.

You can change the default name of the SnapInfo directory (`SME_SnapInfo`), as well as its default location using the Configuration wizard. If you rename the directory (or its Snapshot copies), you must follow strict naming standards.

Specify the location of this directory is when you run the Configuration wizard. By default, the directory name is `SME_SnapInfo`. However, you can specify a different directory name.

By default, the SnapInfo directory is on the LUN that stores the transaction log files, but this is not a requirement. The SnapInfo directory cannot reside on the same LUN that stores the database files.

The SnapInfo directory name is `EXCH__` followed by the Exchange server host name as `EXCH__`*`ExchServerName`*.

For example, you would name the subdirectory for databases that belong to the Exchange server on the Windows host system CLPUBS-WINSRVR3 as `EXCH__CLPUBS-WINSRVR3`.

For SnapInfo directory backups, Snapshot copy names begin with the string `eloginfo__`.

Backup management group	**Format of the SnapInfo directory Snapshot copy name**
Standard	Depending on unique or generic naming convention: • `eloginfo__`*`ExchServerName_date_time`* • `eloginfo__`*`ExchServerName`*`__recent`
Weekly or Daily	Depending on unique or generic naming convention: • `eloginfo__`*`ExchServerName_date_time__BkpMgmtGrp`* • `eloginfo__`*`ExchServerName`*`__recent`

Every time a backup set is created, SnapManager creates a new backup set subdirectory under the SnapInfo directory. The contents of this subdirectory are the backed up transaction logs, and the Snapshot copy recovery information.

A complete backup set consists of this SnapInfo subdirectory and the corresponding Snapshot copies of the LUNs that store the databases and transaction logs.

SnapManager backup sets

SnapManager stores backup data in backup sets. A backup set consists of all data that you need to be able to perform a restore process.

A backup set contains the following items:

- Exchange databases
- Exchange transaction logs
- SnapInfo directory

A backup set contains these components regardless of whether the data exists on the same LUNs and volumes or not.

SnapManager backup set names are displayed in the SnapManager Restore window and in the SnapManager Restore wizard. Backup set names include the name of the server and management group, as well as a timestamp.

Exchange Storage Group/database sets

A set of Storage Groups on a single volume that belong to the same server is known as a Storage Group/database set. When you select a Storage Group for backup, Snap Manager backs up the complete Storage Group set to which the Storage Group belongs.

All Storage Groups of a Storage Group/database set share the same backup name and timestamp.

Storage Groups do not exist in Exchange Server 2010.

Note: You can restore any Storage Group of a Storage Group set individually. You can also restore a single database.

SnapManager minimum unit of backup

The smallest object you can back up is a Storage Group. In Exchange 2010, the smallest object you can back up is a database. For better administration, back up all the Storage Groups in the same Exchange server at the same time.

SnapManager naming-convention options

SnapManager offers both unique and generic naming conventions for naming Snapshot copies. The unique naming convention contains the variable *`date_time`* in the name, and the generic naming convention includes the string `recent` in the name of the most recent Snapshot copy.

When you use the unique naming convention, the most recent Snapshot copy is identified by the most recent date and time. SnapManager does not rename this Snapshot copy when the next Snapshot copy is created.

When you use the generic backup naming convention, the most recent Snapshot copy is identified by the Snapshot copy name that includes the string `recent`. This is the naming convention older versions of SnapManager use and is the default setting to enable backward compatibility.

When you have datasets configured in your system and you choose the generic naming convention, no archives are created. To create archives, apply the unique naming convention with the archival process enabled. If you archive the backups using PowerShell, SnapManager changes the generic naming convention to the unique naming convention. In the GUI, you can select the backup naming convention in the Backup Settings dialog box.

Note: Select the unique naming convention unless you have legacy scripts that need a Snapshot copy with `recent` in its name.

Backup process in a Windows Server 2003 or Windows Server 2008 environment

SnapManager performs backup individually for each volume you choose to back up, using the following sequence. For example, if two separate volumes are involved, SnapManager performs the sequence of steps once for the first volume, then again for the second volume.

To create a backup, SnapManager performs the following sequence of steps:

1. Checks the SnapManager license
2. Renames the most recent SnapInfo directory, if required
3. Renames the most recent Snapshot copy, if you choose the generic naming convention
4. Creates a SnapInfo directory for the backup
5. Initiates a VSS backup
   When you back up data from an LCR or a CCR location, SnapManager uses VSS replication writer to back up the database.
6. Gathers Exchange metadata and activates the VSS Snapshot copy set
7. Adds all required volumes to the VSS Snapshot copy set
8. Creates a volume shadow copy
   The transaction logs are backed up and truncated during this step.
9. Verifies the physical integrity of the database
10. Creates a SnapInfo Snapshot copy
11. Deletes the SnapInfo Snapshot copies for the deleted backups
12. Deletes any specified older backups

Typical VSS backup process

A typical backup using SnapManager for Microsoft Exchange, Microsoft Exchange 2003 or 2007, and the Data ONTAP VSS Hardware Provider is outlined in the following process.

1. SnapManager determines which LUNs it wants to capture and makes sure that Exchange 2003 or Exchange 2007 is present as a valid writer.
2. SnapManager initiates the shadow copy process.
3. VSS informs Exchange and the Data ONTAP VSS Hardware Provider that a shadow copy is starting. Exchange stops writing to disk.
4. VSS ensures that NTFS is in a consistent state.
5. VSS requests the Data ONTAP VSS Hardware Provider to create a shadow copy.

6. The Data ONTAP VSS Hardware Provider requests SnapDrive to create a Snapshot copy of the storage system volume that contains the specified LUN.
7. SnapDrive requests the storage system to create a Snapshot copy of the specified volume.
8. When the shadow copy is complete, VSS returns NTFS to a normal state and informs Exchange that it can resume disk writes.
9. VSS manages the shadow copy of the LUN based on the attributes specified by the requestor. For example, VSS could mount the LUN in a Snapshot copy. In a case, however, in which SnapManager is the requestor, SnapManager tells VSS to forget about the shadow copy it just created. This enables SnapManager to have complete control of the Snapshot copy.

Why a transaction log backup might contain two Snapshot copies

When you use SnapManager with Exchange Server 2003 in a Windows Server 2003 environment, a single backup creates two Snapshot copies instead of one if the SnapInfo directory is configured on the same volume as the transaction log.

The following Snapshot copies are created:

- A VSS Snapshot copy in the transaction log
  The `exchsnap__*` Snapshot copy is created at the same time as the corresponding database Snapshot copy of the same name.
- The SnapInfo backup Snapshot copy for protecting the SnapInfo directory
  The `eloginfo__*` Snapshot copy is created at the end of the backup operation. It is used for transaction log verification and for backup protection. If the backup set is archived, the `eloginfo__*` Snapshot copy, along with the Exchange Snapshot copy, can be used for a SnapManager restore operation.

  Note: For VSS backups, the transaction log LUN Snapshot copy is used instead of the SnapInfo directory LUN Snapshot copy to verify the integrity of the transaction logs during remote database verification.

When SnapManager uses VSS to back up and restore data, the SnapInfo directory and transaction log are on the same volume.

Suppose the SnapInfo directory is on the same LUN as the transaction log in an Exchange Server 2003 environment with Windows Server 2003. The console output shows that the transaction log volume contains Snapshot copies with the same name as that of the Exchange database Snapshot copy.

```
Volume netappnj1_logs
working...
date name
____________ ________
Jan 30 01:04 eloginfo__mailnj1__recent
Jan 30 01:00 exchsnap__mailnj1__recent
Jan 29 13:03 eloginfo__mailnj1_08-29-2006_13.00.00__daily
Jan 29 13:00 exchsnap__mailnj1_08-29-2006_13.00.00__daily
Jan 29 01:04 eloginfo__mailnj1_08-29-2006_01.00.00
```

```
Jan 29 01:00 exchsnap__mailnj1_08-29-2006_01.00.00
Jan 28 13:03 eloginfo__mailnj1_08-28-2006_13.00.00__daily
Jan 28 13:00 exchsnap__mailnj1_08-28-2006_13.00.00__daily
Jan 28 01:04 eloginfo__mailnj1_08-28-2006_01.00.00
Jan 28 01:00 exchsnap__mailnj1_08-28-2006_01.00.00
Jan 27 15:36 eloginfo__mailnj1_08-27-2006_15.30.15
Jan 27 15:30 exchsnap__mailnj1_08-27-2006_15.30.15
Jan 20 15:35 eloginfo__mailnj1_08-20-2006_15.27.35
Jan 20 15:27 exchsnap__mailnj1_08-20-2006_15.27.35
```

If you are not aware of the configuration details, you might expect to see only the `eloginfo__` Snapshot copies and not the `exchsnap__` Snapshot copies.

Note: Volume size is unchanged, because the backup set still captures the same total number of changed blocks on the storage system, even though two Snapshot copies are used instead of one.

SnapManager Snapshot copy naming conventions

Data ONTAP automatically names SnapManager Snapshot copies. The name of each Snapshot copy created during a SnapManager backup operation includes information like the server name, the backup management group, and whether the backup is the most recent backup.

SnapManager Snapshot copy names and SnapInfo directory Snapshot copy names include the name of the server for which the backup was taken.

SnapManager Snapshot copy names and SnapInfo directory Snapshot copy names also include the backup management group to which you assigned the full database backup. SnapManager provides backup management groups for designating various levels of backup retention: Standard, Daily, and Weekly.

- If you assign a full database backup to the Standard backup management group, the Snapshot copy names for the databases and SnapInfo directory do not include a backup management group name.
- If you assign a full database backup to the Daily or Weekly management groups, the Snapshot copy names for the databases and SnapInfo directory include the name of the backup management group.

For database backups, Snapshot copy names begin with the string `exchsnap__`.

Backup management group	**Format of the database Snapshot copy name**
Standard	Depending on unique or generic naming convention: • `exchsnap__`*`ExchServerName_date_time`* • `exchsnap__`*`ExchServerName__recent`*

Backup management group	Format of the database Snapshot copy name
Daily or Weekly	Depending on unique or generic naming convention: • `exchsnap__`*`ExchServerName_date_time__BkpMgmtGrp`* • `exchsnap__`*`ExchServerName__recent`*

When to run a SnapManager backup

You need to balance the frequency of backups against any overhead incurred by the database verification process. In addition, you must ensure that no SnapManager operations overlap with each other.

Observe the following recommendations for scheduling backups and verifications:

- Do not schedule verifications on the Exchange server during peak usage hours.

  The verification process is CPU-intensive and can degrade Exchange performance if you run it on the Exchange server during peak usage hours. To minimize the impact of backups on the client response time, run integrity verification during off-peak Exchange usage hours, or from a remote computer. Typical off-peak times are between 6:00 p.m and 7:00 a.m.
- Balance the frequency of backups against any process of verification.
- Do not schedule a backup when SnapManager performs database verification, even if you perform the verification on a remote verification computer.

  This can result in a backup that you cannot delete easily.
- You must schedule your backup operations such that they cannot be deleted before their SnapVault transfer to the secondary storage is complete.

  For example, you configure a backup operation such that only two backup copies can be retained on the primary storage. You schedule them at 12 PM, 1 PM, 2 PM, and 3 PM. The backup operations take 3 hours before their SnapVault transfer to the secondary storage is complete. When the 3 PM backup operation is running, the 12 PM backup operation is still transferring. Hence SnapManager deletes the 1 PM backup, and the Protection Manager job fails as it fails to find the deleted 1 PM backup to transfer. You can avoid this by using number of days for local backup retention instead of using number of backup copies for local backup retention.
- To upgrade the performance of your production server, run your database verification operations on a remote server.
- More backup operations cause fewer transaction logs to be played forward during a restore operation.

  Perform a minimum of one SnapManager full database backup every 24 hours.
- Do not schedule any operation to overlap any other operation.

  Only one SnapManager operation can be running on the same computer at the same time.
- Back up the databases at the end of a migration process.

  Any previous non-SnapManager backups are no longer valid after the migration process.

- If you schedule backups in a cluster, schedule them on only one node.
  If a virtual server fails over to another node, that performs the scheduled backup. If the entire node fails, reschedule the backups.
- Be sure to reschedule the scheduled jobs that are affected by any changes you make in the configuration. Update the cmdlet parameters manually because they are hard coded in the scheduler.

How SnapManager checks database integrity in backup sets

SnapManager uses Eseutil, a Microsoft Exchange utility, to verify the Exchange Server 2003 database files; it uses `ChkSgFiles.dll`, an integrity verification library, to verify the Exchange Server 2007 databases and transaction logs; it uses either ChkSgFiles or Eseutil to verify the page-level integrity of the databases. If the databases are in a Database Availability Group (DAG), consistency checking might be unnecessary.

The Eseutil utility is installed automatically with binaries. SnapManager checks the database and the transaction logs for physical-level corruption. For more information about physical-level corruption, see your Microsoft Exchange documentation.

The `ChkSgFiles.dll` library checks the databases and transaction log files for physical and logical corruption, by using checksum verification. You must install Exchange Server 2007 Management Tools on the server that performs integrity verification.

Note: By default, SnapManager Restore requires that you restore only from verified backups, but you can override this requirement.

Exchange 2010 database verification is not a support requirement for databases with at least two copies in a Database Availability Group. Microsoft recommends verification of the transaction logs. By default, when performing a DAG backup with SnapManager for Exchange, verification is off. You can monitor both locally running and remote verification jobs through the main SnapManager for Exchange management console. A single Exchange mailbox server can run only one verification process at a time on a particular verification server. A verification server can simultaneously run one verification job from each Exchange mailbox server. More than one verification server can be used to simultaneously verify more than one backup job on a single Exchange mailbox server. A verification server can utilize a virtual machine to off-load the verification.

See *http://msdn.microsoft.com/en-us/library/aa579091(v=exchg.140)* for further information.

LUN requirements for verifying databases in a backup set

SnapManager requires that all databases be mounted before verification. The SnapManager verification server must have enough drive letters or mountpoints to mount the LUNs that store the backup sets that you verify, depending on your verification scenario.

SnapManager mounts the LUNs that contain the backup sets that you select for verification with SnapDrive commands. Each mounted LUN requires one available drive letter or mountpoint.

To verify backups that are stored on a single LUN or across multiple LUNS, the verification server must have at least one drive letter or mountpoint available.

If you are using a remote verification server, you need two unassigned drive letters or mountpoints to verify multiple backup sets in a single job.

In other situations, you need one unassigned drive letter or mountpoint.

The one unassigned drive letter or mountpoint is to mount the LUN in a Snapshot copy. The Snapshot copy stores the transaction log directory for multiple Storage Groups for transaction log verification. When SnapManager verifies the transaction logs for the first backup set, it does not unmount the LUN or release the assigned drive letter. SnapManager reuses the same LUN for the transaction log verification for the second backup set.

If a second drive letter or mountpoint is not available, schedule the backup or verification jobs to verify one backup set at a time.

Note: If you verify the destination volumes that contain more than one Storage Group sharing the same LUN for the transaction logs, ensure that a second drive letter or mountpoint is available for the Exchange server.

Database verification load management

Because running database verification on a production server can place a significant load on both the server and the storage systems, you must manage this workload to avoid degraded Exchange response, particularly during peak work hours.

You can use any of the following methods to manage your database verification workload:

- Deferred database verification
- Remote database verification
- Remote verification on a SnapVault destination (Data ONTAP operating in 7-Mode only)
- Verification throttling
- Integrity verification on destination SnapMirror volume

Deferred database verification

You can distribute your system workload by disabling automatic integrity verification, which takes place immediate after your backup is created, and then performing a separate verification later.

Remote database verification

You can distribute your system workload by disabling automatic integrity verification, which takes place immediate after your backup is created, and then run verification from another Exchange server.

Verification throttling

You can slow down the integrity verification to decrease the load on the Windows host and the storage systems by using verification throttling if you are using Microsoft Exchange Server 2003 SP2, or higher, or Microsoft Exchange Server 2007.

SnapManager can throttle integrity verification (the Microsoft Exchange consistency checker utility) in checksum verification mode to reduce the load on the host CPU, and on the storage subsystem. This .dll file checks the databases and transaction log files for integrity verification.

Integrity verification on the destination SnapMirror volume

SnapManager 5.0 and later supports integrity verification on the destination SnapMirror volume. SnapDrive uses FlexClone to access data in the Snapshot copy on destination SnapMirror volumes.

Integrity verification on the destination SnapMirror volume is available through the following:

- Backup and Verification window
- Backup wizard
- Deferred verification
- Test Restore button and Restore wizard

  Note: When you perform a restore operation from the Restore window, SnapManager prompts for integrity verification on the source or destination volume, if there are any SnapMirror volume relationships associated with it.

The Choose SnapMirror Destination Volumes for Integrity Verification window displays each SnapMirror volume as a tree showing the relationship among the volume, the LUN, and the databases contained in it. For each source volume, there is a list of destination volumes, and each destination volume displays a SnapMirror state and a FlexClone state. You can select a SnapMirror destination volume for each SnapMirror source volume for which you want verification.

Backup verification status reporting

You can use SnapManager Restore to determine the verification status of each backup and of the transaction logs for each backup.

SnapManager Restore shows you a list of the backups that have been created and indicates the verification status of each backup. For each listed backup, the date and time of the backup operation is displayed, as well as an icon that indicates the verification status of the backup.

You can also determine whether the transaction logs for a specific backup have been verified by selecting the backup in the Restore window. The transaction log verification status is included with the backup information shown in the Result pane.

Backup verification status icons

The following verification status icons represent the verification status of the databases.

Verification status icon	Database verification status
	Verified databases.
	Unverified databases.
	Databases failed verification and cannot be restored.
	Verified databases, but unverified transaction logs. If you need to restore from backups with this symbol, contact technical support.

Where to run database and transaction log integrity verification

Regardless of when you verify the databases in a backup set, you can perform the verification either on the production server or on a remote verification system.

In the simplest SnapManager configuration, verification is run from your production server; however, this type of verification is CPU and disk I/O intensive, so you might want to avoid performing verification on the production server during peak usage as it can affect Exchange performance.

Performing integrity verification on a remote system minimizes the negative performance impact on Exchange system resources and the backup schedule.

When to verify the databases in a backup set

You can verify the databases in your SnapManager backup sets immediately after creation, at a scheduled time later, or when you restore them.

By default, SnapManager automatically verifies full backup sets at the time the backup is created. This is simple and ensures that each database in the backup set is verified. However, this method significantly increases the time required to complete the backup.

You can start an operation to verify the databases contained in one or more backup sets that have already been created. You can start the verification immediately, or you can schedule the verification to occur later, when it does not affect performance or delay later backups.

If you attempt to restore from a backup set on which a database consistency check has not been run successfully, SnapManager prompts (but does not require) you to first verify the databases in that backup set.

Note: When you perform verification on a LUN clone, and you make a Snapshot copy of the volume while a LUN clone exists, a "busy Snapshot copy" is created, which might cause problems when you attempt to delete Snapshot copies. To avoid this, do not schedule backups while a verification is in progress, to enable the use of FlexClone for Snapshot copy access. If the FlexClone volume is licensed on the storage system, it is used without SnapMirror verification.

Backup set retention

When you plan your SnapManager backup schedules, you also need to manage the number of Snapshot copy-based backup sets that are retained online.

You can do the following to manage the number of backup sets kept online:

- Regulate the number of Snapshot copies per volume.
- Delete old and unnecessary backups.
- Specify a backup deletion criteria that deletes the oldest backups for each Storage Group and backup management group.
- Explicitly delete any backups or Snapshot copies of LUNs created during a restore operation.

Maximum number of Snapshot copies per volume

The Data ONTAP software used with SnapManager supports a maximum of 255 Snapshot copies per volume, including copies not created by SnapManager. Because each SnapManager backup operation creates Snapshot copies, a SnapManager backup operation fails if the volume that contains the database LUN exceeds the 255 Snapshot copy capacity.

Note: The total number of Snapshot copies on a volume might exceed the number of retained backups. For example, if a single volume contains both the SnapInfo directory and the databases, each backup operation generates two Snapshot copies on the volume.

Ways to delete Snapshot copies

To avoid reaching the limit of 255 Snapshot copies per volume, delete your old SnapManager backups when they are no longer needed.

SnapManager provides two ways to delete backups:

- Automatic deletion of older backups in the management group
- Explicit deletion of backups or SnapInfo Snapshot copies only

Note: If a storage group is no longer available in the Exchange server, then backups associated with it cannot be removed by the delete backup module.

Attention: Do not use SnapDrive or the storage system administration tools to delete Snapshot copies created by SnapManager. Doing so will leave behind unwanted data that cannot be removed.

Automatic deletion of Snapshot copies

You can manage the number of Snapshot copies you store by configuring SnapManager to delete backups automatically, based on how old the backups are or based on how many of them are stored.

Automatic deletion deletes a backup only if the backup has the following characteristics:

- The backup is in the same management group as the management group of the backups that you just created.
- The backup is the oldest backup of the Storage Group
- The number of backups exceeds the backup retention level that you specified in the “Delete backups older than” option or the “Delete backups in excess of” option.

Note: If a storage group is no longer available in the Exchange server, then backups associated with it cannot be removed by the delete backup module.

If you do not select automatic backup deletion, backups that are created after the current backup are retained. This would require manual removal of backups, or enough storage capacity for all backups and transaction logs. You can delete the retained backups by selecting automatic backup deletion in the next backup that you take.

Transaction log management

SnapManager provides advanced options to manage the transaction logs that belong to a backup management group.

You can choose not to back up transaction logs that Exchange truncates after the backup operation finishes.

If deleting Snapshot copies of LUNs that contain transaction logs related to the selected backup management group breaks the continuity of transaction logs between the previous backup and the present time, you can skip the deletion.

In releases earlier than Snap Manager 6.0.2 for Microsoft Exchange, the oldest backups were deleted based on a user-specified number of backups to keep, per management group. Once an older backup was deleted, the capability of up-to-the-minute restore on other backup management groups was lost if the option “Retain up-to-minute restore ability for older backups in other backup management groups” was disabled.

If the previously mentioned option was enabled, only the Snapshot copy was deleted, and the related transaction log files were retained in the SnapInfo directory.

However, in the prior scenario, a large amount of storage was consumed to retain the large number of transaction logs.

Starting in SnapManager 6.0.2 for Microsoft Exchange, the retention management algorithm more efficiently limits the number of saved transaction logs based on administrator-specified options: the number of days or number of backups to keep to ensure up-to-the-minute restore ability.

New backup options

The specific “Retain up-to-minute restore ability” backup choices are as follows:

- Backups generated in the last (days)
- The most recent (number of backups)
- All the older backups (do not delete any logs)

Where to specify the new options

You can specify the new options in one of three locations:

- **Backup wizard**, **Retain up-to-the-minute restore ability for older backups**
- **Backup and Verify**, **Up-to-minute Restore Options**
- **Delete Backup**, **Advanced Options**

After you set the new options, other backups outside the specified backups in the same management group become point-in-time only. For example, assume you have the following backup set:

Bkup_std_5	Oldest backup
Bkup_std_4	
Bkup_std_3	
Bkup_std_2	
Bkup_std_1	
Bkup_std_0	Latest backup

If backup retention is set to five, and up-to-the-minute retention is set to three, the sixth backup, Bkup_std_5, is deleted and the retained backups above the third become point-in-time backups:

Bkup_std_4	Point-in-time
Bkup_std_3	Point-in-time
Bkup_std_2	Up-to-the-minute
Bkup_std_1	Up-to-the-minute
Bkup_std_0	Up-to-the-minute

The SnapManager for Microsoft Exchange graphical user interface (GUI) contains options only for the new backup retention management algorithm; the option in releases prior to SnapManager 6.0.2 for Microsoft Exchange to retain up-to-the-minute restore ability has been removed. If you run

backup or delete operations from the SnapManager for Microsoft Exchange GUI, the new backup retention method is used.

PowerShell continues to support the `NoUtmRestore` switch; however, the switch is disabled as soon as you specify values for `RetainUtmBackups` or `RetainUtmDays`.

Valid input for UTM-related switches

`RetainUtmBackups` can have any integer value between 1 and 255.

A float value >0 for `RetainUtmDays` is allowed. However, a negative number generates an error message.

If you specify values for both `RetainUtmBackups` and `RetainUtmDays`, an error message is generated. However, there is one exception: in PowerShell, if you specify -1 for either `RetainUtmBackups` or `RetainUtmDays`, SnapManager ignores the switch. In that case, SnapManager does not generate an error. Instead, it uses the prior backup retention method and takes the default value of `NoUtmRestore`.

If none of the three UTM-related switches is specified, the prior backup retention method is used with the default value of `NoUtmRestore`=0 applied.

New registry keys

Three new registry keys have been added under `HKEY_LOCAL_MACHINE->Software->Network Appliance->SnapManager for Exchange->Client`:

`KeepLastNSnapshotUtm REG_DWORD`

`KeepLastNSnapshotOlderDayUtm REG_SZ`

`KeepLastNSnapshotOptUtm REG_DWORD`

When both `KeepLastNSnapshotUtm` and `KeepLastNSnapshotOlderDayUtm` have a valid value, SnapManager uses the value of `KeepLastNSnapshotOptUtm` to determine if the operation is either day-based or number-based.

`KeepLastNSnapshotOptUtm` =

0: Number-based

1: Day-based

2: Retain up-to-the-minute restore ability to all the older backups

Three new entries are added to the control file under `<BACKUP_CLIENT_SETTING>`:

`<BACKUP_SET_TO_KEEP_UTM>`

`<BACKUP_SET_TO_KEEP_IN_DAYS_UTM>`

`<DELETE_BACKUPS_OPTION_UTM>`

When both `BACKUP_SET_TO_KEEP_UTM` and `BACKUP_SET_TO_KEEP_IN_DAYS_UTM` have a valid value, SnapManager uses the value of `DELETE_BACKUPS_OPTION_UTM` to determine if the operation is either day-based or number-based.

`DELETE_BACKUPS_OPTION_UTM` =

0: Number-based

1: Day-based

2: Retain up-to-the-minute restore ability to all the older backups

Option to back up transaction logs that Exchange will truncate

If you do not need to retain up-to-the-minute restore ability from a backup that is not the most recent, you can reduce the amount of disk space required by omitting the backup of transaction logs that Exchange truncates after the SnapManager Backup operation finishes.

By default, SnapManager backs up database and the transaction logs. SnapManager creates a database Snapshot copy and a SnapInfo Snapshot copy, and truncates transaction logs by removing any entries already committed to the database.

Manage this feature using the **Back up transaction logs that will be truncated by Exchange at the end of the backup** option in the Advanced Options dialog box.

The option to back up all transaction logs (including those that Exchange truncates after the backup operation finishes) is selected by default. When you clear the option to back up all transaction logs, the database and transaction log verification cannot be deferred, and the option to back up all transaction logs is automatically enabled the next time you start SnapManager.

Exchange System Manager in a SnapManager environment

Exchange System Manager (ESM) is an Exchange Server 2003 component that includes a timestamp that tells you the last time that a full backup of a database took place.

The timestamp is displayed on the database Properties page. For an Exchange Server 2003 database, ESM 2003 displays a date and timestamp, and whether you made the backup with SnapManager using VSS Snapshot copy or with a streaming method such as Windows NTBackup.

Attention: Do not use the Exchange System Manager to move databases or transaction logs for a system that you configure using SnapManager. Doing so prevents SnapManager from functioning correctly.

Displaying the time of the last full backup

To display the time of the last full backup of a database, use the SnapManager Restore window rather than Exchange System Manager.

About this task

To display the time of the last full backup of a given database, use the SnapManager Configuration wizard rather than Exchange System Manager.

Steps

1. In the SnapManager console, double-click **Restore** in the Scope pane.

   SnapManager displays the backup sets in the Result pane.

2. In the Result pane, click the Storage Group name to get the available backup sets.
3. Double-click the name of the backup set whose properties you want to view.

   Note the timestamp information included in the Result pane.

   Name displays the backup set name. If you configure SnapManager to use unique backup naming (the default setting) as opposed to generic naming, the backup name includes the date and timestamp. Backup Date and Time displays the date and time at which the backup was made.

Exchange page zeroing and deleted item retention in Exchange 2007 and earlier

Because SnapManager relies on the Snapshot technology to create a backup, SnapManager does not read every page of data. Hence, the page zeroing and deleted item retention features are not triggered when you use only SnapManager to perform backups. This concern does not apply to Exchange releases after 2007.

The page zeroing feature of Exchange increases your database security by enabling you to find deleted database pages and is used by high-security installations. The deleted item retention settings feature enables you to mark items and mailboxes to delete after they are backed up.

> **Attention:** SnapManager archives items before you select them for deletion. To ensure that deleted items are archived into at least one backup set, select the **Do not permanently delete mailboxes and items until the store has been backed up** option. However, this option overrides the "Deleted item retention" intervals for items and mailboxes. Therefore, to recover deleted items without restoring and extracting Exchange data from a tape, do not permanently delete mailboxes and items until you back up the store.

You can cause these Exchange features to be triggered by scheduling a periodic backup using a standard backup utility, such as NTBackup.

Ensure that the backup that you use for this purpose is a copy or differential backup. Performing any other type of backup causes the existing SnapManager backups to become unusable for up-to-the-

minute restore operations. Also ensure that you perform no extra backup operation at the same time as a SnapManager backup.

Database backup using SnapManager

You can back up Exchange database storage sets using SnapManager. You can also back up LCR, CCR and DAG configurations using SnapManager.

Exchange storage groups and databases display

The Result pane in the SnapManager console lists the Exchange Server storage groups and databases that are managed from the current Exchange server.

- Storage groups or databases that reside on the same volume are shown with disk icons of the same color.
- Storage groups or databases that span multiple volumes are shown with disk icons of three colors.
- Storage groups or databases that cannot be backed up by SnapManager are shown with the label Invalid next to the LUN drive letter or mountpoint.

For a CCR-enabled Exchange Server storage groups and databases, the server's navigation tree in the Scope pane lists the Backup, Restore, Scheduled jobs, and the Reports directories for both the active and the passive nodes.

Exchange databases display in a DAG

When the SnapManager MMC snap-in connects to a Database Availability Group (DAG), all mailbox databases in the DAG, including both the active and passive databases, are displayed. If the SnapManager MMC snap-in connects to a member server of the DAG, only the mailbox databases on that server are displayed.

When you connect to the DAG in the Scope pane, you can view all the databases in the DAG that are available for backup. When you connect to a member server of the DAG, you can view and back up the databases on that server only.

Decisions to make before performing a SnapManager backup

Before you perform a SnapManager backup, you need to gather the information required to complete the backup operation.

Feature used for backup

- Do you want to use the Backup and Verify window or the Backup wizard?

  You can create a backup more quickly using the Backup and Verify window on your own. The SnapManager Backup wizard helps you create and schedule your backups using default settings.

Storage Group

- Which Storage Groups do you want to back up to?
  The types of Storage Groups you can back up to are:
  - Production Storage Groups
  - LCR replica Storage Groups
  - Continuous Cluster Replica (CCR) replica Storage Groups

Databases in a Database Availability Group (DAG)

- Which databases do you want to back up?
  The types of databases you can back up are:
  - Active databases
  - Passive or copy databases

Type of backup

- What type of backup do you want to perform?
  You can do a full backup or a copy backup.

  Note: Copy backup is available only for CCR configuration created along with full backup on the remote CCR node.

Backup management group

- Which backup management group you want to assign to this backup?
  You can assign your backup to a standard, daily, or weekly management group.

Integrity verification

- Do you want to verify databases and transaction logs on the SnapMirror destination volume?
- Do you want to verify the backup databases and transaction logs after the backup is created for a restore operation?
  Although it is possible to restore from an unverified backup, it is advisable to restore only from verified backups.

Backup retention

- Do you want to automatically delete the oldest backups of this backup management group?
  If you do not select automatic backup deletion, backups that are created after the current backup are retained.
- Do you want to back up transaction logs that Exchange truncates at the end of the backup?

If you do not need to retain up-to-the-minute restore ability from a backup that is not the most recent, you can reduce the amount of disk space required by omitting the backup of transaction logs that Exchange truncates after the SnapManager Backup operation finishes.

Up-to-the-minute restore ability

- Do you want to retain up-to-the-minute restore ability for any backups from other backup management groups that are older than backups you would delete?

  By default, automatic deletion of older backups within a backup management group is done selectively, in such a way that it does not cause a break in the continuity of transaction logs between the previous backup and the present time.

  However, if you do not need to retain up-to-the-minute restore ability from a backup that is not the most recent, you can allow the automatic deletion feature to delete all older backups within the backup management group. The storage system space consumption reduces.

  If you choose not to retain up-to-the-minute restore ability, it reduces the amount of disk space consumed by transaction log backups that Exchange truncates after the backup finishes. If all backups have the same backup management group designation, not retaining up-to-the-minute restore ability has no effect.

Naming convention

- Do you want to use the unique naming convention or the generic naming convention?

  When you use the unique naming convention, the most recent Snapshot copy is identified by the most recent date and time. When you use the generic backup naming convention, the most recent Snapshot copy is identified by the Snapshot copy name that includes the string `recent`.
- Do you want to rename a Storage Group?

  When you rename a Storage Group, all the databases in that Storage Group need to be unmounted. After you unmount all the databases, you can remount the databases. At this point Exchange recognizes the new Storage Group name, and you can create new SnapManager backups of this Storage Group. Refresh the backup window before the backup operation.

Backup schedule

- Do you want to run the backup now or schedule it for later?

  If you want to schedule the backup to run later, you must have the job scheduling information.
- Do you want to run a command after the backup is complete?

SnapMirror

- Do you want to update the SnapMirror destination volume after the backup?

Backup planning for databases in Database Availability Groups

By planning the backup strategy for your Exchange databases, you can minimize the chances of losing data if a restore operation is necessary, while still controlling the resources needed to create and maintain the backups. The strategy you choose depends, in part, on the number of Database Availability Group (DAG) nodes, the number of databases, the size of the databases, the network links between the DAG nodes, the service level agreements, and how quickly you must be able to reseed a database.

If your DAG resides on three nodes and consists of no more than eight databases per node, you can choose a gapless backup strategy based on the SnapManager for Exchange remote additional copy backup feature. Remote additional copy backup greatly simplifies the administration of backups. If the DAG is larger or has more than eight databases on a node, choose a server-oriented backup strategy; depending on your recovery requirements, you might add frequent recovery point (FRP) backups to the server-oriented backups.

Planning backups for three-node DAGs

For a Database Availability Group consisting of three nodes and no more than eight databases per node, you can use SnapManager remote additional copy backup to simplify backup planning. These plans still permit up-to-the-minute database restore and reseed operations for any database in the DAG, independent of the availability of any specific node in the group.

For more information about using remote additional copy backup for a DAG, see *Remote additional backup copy* on page 140.

Planning backups for DAGs with more than three nodes

The basic approach for DAGs with more than three nodes is to perform one full backup of any of the Exchange copies of the database (the active or any of the passive versions) and to copy backups on other. Because of the size of the full backup, you should allow additional time for it to complete, compared to the copy backups. If a database has four copies (Copy 1 through Copy 4), and you decide to run the full backup on Copy 4, you might choose a schedule such as the following:

1. 00:30 to 02:00 full backup of Copy 4
2. 02:00 to 03:00 copy backup of Copy 3
3. 03:00 to 04:00 copy backup of Copy 1 and, simultaneously, copy backup of Copy 2

Depending on your service level requirements, this schedule might be sufficient. If it is not, rather than repeat this schedule, you can add FRP backups. FRP backups are very quick, because they are Snapshot copies of the transaction logs. Adding the FRP backups, the modified schedule might be something like the following:

1. 00:30 to 02:00 full backup of Copy 4

2. 02:00 to 03:00 copy backup of Copy 3
3. 03:00 to 04:00 copy backup of Copy 1 and copy backup of Copy 2
4. 03:00 to 04:00 FRP backup of Copy 3 and FRP backup of Copy 4
5. 05:00 to 00:00 FRP backups of all copies every 30 minutes

For more information about configuring frequent recovery point backups, see *Frequent Recovery backup operations* on page 137.

Backing up using the Backup wizard

You can back up databases by using the SnapManager Backup wizard. Creating a backup enables you to store and recover your database.

Steps

1. In the Scope pane, select the Exchange server node you want to back up.
2. Click **Backup**.

   SnapManager displays the list of Storage Groups or Databases in the Backup view in the Result pane and the corresponding actions that you can perform for SnapManager backup in the Actions pane.

   Note: If the SnapManager MMC snap-in is connected to a Database Availability Group (DAG), you have to use the Database Filter to specify the criteria to display the Exchange 2010 databases for backup.

3. Select the Storage Groups or Databases you want to back up.

   The **Result** pane shows whether the Storage Group or Database is dataset enabled, the name of the enabled dataset, the SnapMirror status, the SnapVault status, and other details about the Storage Group or Database.

4. Click **Backup Wizard** in the **Actions** pane.

   The Welcome window appears.

5. Follow the instructions in the **Backup wizard** to initiate a backup process.
6. At the **Completing the Backup wizard** dialog box, click **Finish** after you verify that all the settings in the window are what you want.

   The **Backup Status** window appears.

7. In the **Backup Status** window, click **Start Now** to start the backup.

   The backup is performed and the Snapshot copy is written to the volume. SnapManager Backup completes each task and checks it off on the list shown in the **Backup Task List** view. You can

alternate between the task check-off list and the progress report. If the backup is successful, the Task view shows the check-off list with the tasks completed.

Note: If you enable the Notification option, an e-mail message is sent and the event is posted to the Windows event log.

Backing up using the Backup and Verify window

Use the Backup and Verify window to create a backup more quickly than using the Backup wizard.

About this task

You can choose to archive your database to a secondary storage system and when to verify the archived backup.

Steps

1. In the **Actions** pane, select **Backup**.
2. In the **Backup and Verify** window, select the Storage Groups or Databases that you want to back up.

   Note: If you decide to back up a Storage Group with an invalid status, SnapManager does not allow you to proceed with the backup operation.

3. In the **Backup Management Group** list, select a management group for the backup you want to create.
4. If you have datasets configured in your system, select a backup archiving option under **Backup archiving options** to archive the database at the secondary storage system.
5. If you have datasets configured in your system, select a backup retention group under **Backup archiving options** to determine the retention time of the dataset on the archived secondary storage system.
6. If you want to delete older backups of this backup management group automatically, select one of the **Delete Backups** options.

   Note: Ensure to select one of the **Delete Backups** options to manage your Snapshot copies.

7. To view or change the setting to retain up-to-the-minute restore, click **Advanced**.

If you want to...	Then do this...
Avoid creating a break in the continuity of transaction logs (between the previous backup and the present time)	Select the **Retain up-to-the-minute restore ability for older backups in other management groups** option. Snapshot copies of LUNs that contain transaction logs related to any management group that is not selected are not deleted from the SnapInfo directory.

If you want to...	Then do this...
Reduce the space consumption of the storage system by transaction logs (by allowing more transaction logs to be deleted)	Clear the **Retain up-to-the-minute restore ability for older backups in other management groups** option.

8. To view or change backup transaction logs settings, do the following.

If you want to...	Then do this...
Retain up-to-the-minute restore ability from a backup that is not the most recent	Select **Back up transaction logs that will be truncated by Exchange at the end of the backup**. This option is automatically selected each time you launch SnapManager.
Reduce the amount of disk space consumed by transaction log backups that Exchange truncates after the backup finishes	Clear the **Back up transaction logs that will be truncated by Exchange at the end of the backup** option. The option remains disabled only until you exit SnapManager. **Note:** Clearing this option causes SnapManager Backup to also remove the option to defer database verification.

9. To verify the backup immediately after the backup is complete, select **Verify Backed Up Databases and Transaction Logs**.

   Do not select the option if you want to verify your backup later.

   Note: Although it is possible to restore from an unverified backup, you should not do so.

10. If you want to run a command after the backup process is complete, select the **Run Command After Operation** check box.

    Selecting the option **Run Command After Operation** archives backups after the SnapVault process completes.

11. Depending on your requirement of a SnapMirror update, do either of the following:

If...	Then do this...
Your volume is a SnapMirror source volume and you do not want the destination volume to be updated after this backup process is complete	Clear the **Update SnapMirror After Operation** check box.
Your volume is a SnapMirror destination volume and you want the destination volume to be updated after the backup is complete	Select the **Update SnapMirror After Operation** check box.

12. To run integrity verification on the Exchange databases and transaction logs that are stored on the destination volume, select **Verify on available SnapMirror destination volumes**.

    The **Verify on available SnapMirror destination volumes** is available and checked by default, only when the selected Storage Group or database have at least one SnapMirror relationship.

The **Verify on available SnapMirror destination volumes** option is grayed out if the FlexClone license is not installed on the destination volume.

13. Depending on when you want to run the backup process, do either of the following:

If you want to...	Then do this...
Run your backup process immediately	Click **Backup Now**, and then proceed to step to back up the selected Storage Groups or databases. The Backup Status window appears, showing the Task list.
Defer your backup process	Click **Schedule** to use the Windows Task Manager to schedule your backup process, and then proceed to the next step.

14. In the **Schedule Jobs** window, name your backup job, provide the user ID and password for the job, then click **OK**.

    Note: If this name already exists as a Windows scheduled task and you want to replace it with a new job, select the **Replace if it Exists** check box, then click **OK**.

    The Schedule jobs window appears.

15. In the **Schedule Jobs** window, use the **Schedule** tab to specify the following parameters of your job schedule—When the job is to run, and if you want the job to repeat, at what frequency.

16. After you schedule your job, click **OK**.

    You can use Control Panel to modify the schedule or cancel the scheduled job.

    The backup job runs at the times you specified in the **Schedule Task** view.

17. In the **Backup Status** window, click **Start Now** to back up the selected Storage Groups or Databases.

    Clicking **Start Now** completes the backup operation and the Snapshot copy is written to the volume. SnapManager Backup completes each task and checks it off on the list shown in the **Backup Task List** view.

    You can alternate between the task check-off list and the progress report.

    If the backup process is successful, the Task view shows the check off list with the tasks completed.

    Note: If you enable Notification, an e-mail message is sent and the event is posted to the Windows Application event log.

Using Database Filter to display the databases to back up in a Database Availability Group

You can use Database Filter to specify the criteria to display the Exchange Server 2010 databases for backup in a Database Availability Group (DAG). Databases that satisfy the filter criteria are

displayed, enabling you to select them for backup and verification. The Database Filter also affects how the backup job is scheduled.

About this task

A DAG is a group of up to 16 servers, each with many databases. The Database Filter enables you to narrow the number of databases that are displayed for selection. You can access the Database Filter when you connect to the DAG from the following locations within the SnapManager user interface:

- Backup wizard
- Actions pane for backup

Note: When you specify criteria that is different from the default criteria "any," the word "Filtered" is displayed in italics in the **Results** pane of the Backup window.

Steps

1. To display the databases on a member server, select one of the following:

If you want to...	Then do this...
Display databases on a member server of the DAG	Select the server from the **Show databases located on an Exchange member server** list.
Display all databases in the DAG	Select **Any** from the **Show databases owned by a particular server** list.

2. To display the databases based on database activation attributes, select one of the following from the **Show databases based on the copy specified** box:

If you want to...	Then do this...
Narrow the display to active primary databases on the member server selected in Step 1 or in the DAG	Select the **Active Primary Database** option from the **Copy** list.
Narrow the display to passive copy databases on the member server selected in Step 1 or in the DAG	Select the **Passive Copy Database** option from the **Copy** list.
Display both active and passive databases on the member server selected in Step 1 or in the DAG	Select the **Any** option from the **Copy** list.

3. To narrow the display of databases based on their Exchange activation order, select one of the following from **Show databases by their Exchange activation order** box:

If you want to...	Then do this...
Display the databases irrespective of their Activation Preference numbers	Select **Any** from the **Activation Preference** list.
Display the databases based on their Activation Preference number	Select the Activation Preference number from the **Activation Preference** list.

Activation Preference Number of a mailbox database in Exchange Server 2010

The Activation Preference Number of a mailbox database in Exchange Server 2010 is a number specifying the preference value of that database copy. You can select the database that has to be backed up by specifying its Activation Preference Number through the Database Filter.

The original database in a Database Availability Group (DAG) always has a Activation Preference Number of 1. During database activation, when multiple database copies satisfy the criteria for activating, the Activation Preference Number is used to decide which database copy is to be activated.

LCR-enabled database backups

SnapManager allows you to back up and schedule a database from either its Local Continuous Replication (LCR) copy location or its original location. Ensure that you do not create a backup on the replica database and its production database at the same time.

Using either the SnapManager Backup and Verify window or the Backup wizard, you can view the replica Storage Group simultaneously with its production Storage Group. The replica Storage Group is represented as an LCR Storage Group on the SnapManager graphical user interface.

Note: You must use different volumes to store the database files, the transaction log files of the replica Storage Group, and the production Storage Group.

When SnapManager creates a backup on the production database and its replica database at the same time, SnapManager creates two successive sets of timestamps, one for each backup operation. Only one of the backups has the `recent` backup name, if you select the generic naming convention in Backup Settings. Two copies of the backup set are created that consume extra storage space in the LUN.

Note: SnapManager does not perform any special operation to delete the transaction log files on a replica database after making the backup on the production database. When you create a backup on an LCR-enabled database, Exchange Writer handles the truncation of the log files automatically, on both the production and the replica databases.

CCR-enabled database backups

You can back up both the active and the passive nodes in a Continuous Cluster Replica (CCR) environment in Exchange Server 2007. Before starting the backup operation, ensure that SnapManager is installed on both nodes of the CCR configuration.

Considerations before backing up a CCR replica database and production database

Before you back up a CCR replica database and production database, you need to make some considerations about the use of active and passive nodes in your CCR configuration.

- When you perform an operation such as adding, reconnecting, or removing a CCR-enabled Exchange server, the operation applies to both the active and passive nodes together.
- You can back up your Storage Group from either the active node or the passive node.
  - You cannot connect to the active nodes and passive nodes at the same time and perform backup operations on both the nodes simultaneously.
  - If you do not select the option **Create a copy database on another CCR node along with current backup**, then you only back up the CCR node that you are connected to, which is either an active node or a passive node.
  - When you rename a Storage Group, you must unmount all of the databases in that Storage Group. After you unmount all of the databases, you can remount them.

    At this point Exchange recognizes the new Storage Group name, and you can create new SnapManager backups of this Storage Group.
- When you back up a CCR configuration, SnapManager creates full database backups either on a CCR active node, using Exchange Store Writer, or on a CCR passive node, using Exchange Replication Writer.

Note: SnapManager does not perform any special operation to delete the log files on a replica database after the production database backup in a CCR environment. When a backup is created on a CCR replica database, Exchange Writer handles the truncation of the log files automatically, on both the production and the replica databases.

Backing up and verifying a CCR replica database and production database

You can back up and verify a CCR replica database and production database using SnapManager Backup.

Steps

1. In the SnapManager console, select the Continuous Cluster Replica (CCR) node on which you want to create the primary backup.
2. Click **Backup** either in the CCR Active node or passive node in the **Scope** pane.
3. Click **Backup Wizard** in the **Actions** pane.
4. Select the storage group that you want to back up or verify.
5. In the **Backup or Verify Databases and Transaction Logs** window, do either of the following:

If you want to...	Then do this...
Back up databases and transaction logs in the selected Storage Group	Select **Back up databases and transaction logs**.
Verify the databases and transaction logs in an unverified backup	Select **Verify databases and transaction logs in the most unverified backups**.

6. In the **Backup Management Group** window, select the management group for this backup operation.
7. In the **Select the naming convention for the newest backup** window, select **Use generic ("_recent")** or **Use unique ("time stamp")** according to your requirement.
8. To delete older backups, enter the number of backups to be retained or the number of days to retain the backups in the **Delete Older Backups** window.
9. In the **Retain Up-to-The-Minute Restore ability for Older backups** window, do either of the following:

If you want to...	Then do this...
You do not want to delete the transaction logs that are saved in the SnapInfo directory	Select **Yes, the retain up-to-the-minute restore capability for older backups**.
You want to delete the transaction logs that are saved in the SnapInfo directory	Select **No, allow older backups to become point-in-time only**.

10. To back up the truncated logs, select **Yes, backup truncated transaction logs** in the **Truncated log backup** window.
11. In the **Verify the Databases and Transaction Logs in this Backup** window, do either of the following:

If you want to...	Then do this...
Verify the physical integrity of the databases and transaction logs	Select **Yes, I want to verify the databases and transaction logs after this backup is complete**. **Note:** Click **Verification Settings** button to launch the **Databases Verification Settings** dialog box to select the SnapMirror relationship for the integrity verification.
Verify the physical integrity of the databases and transaction logs later	Select **No, I want to verify the databases and transaction logs in this backup later**.

12. In the **Update SnapMirror After Operation** window, do either of the following:

If you want to...	Then do this...
You want to request a SnapMirror update	Select **Yes, request a SnapMirror update after operation**.

If you want to...	Then do this...
You do not want to request a SnapMirror update	Select **No at this time**.
You want to run integrity verification on the Exchange databases and transaction logs that are stored on the destination volume	Select **Verify on available SnapMirror destination volumes**.
You do not want to run integrity verification on the Exchange databases and transaction logs that are stored on the destination volume	Clear **Verify on available SnapMirror destination volumes**.

13. To run a command after the backup is complete, select **Yes, run a command after this operation** in the **Run Command After Operation** window.
14. In the **Complete the Backup wizard** window, verify the backup settings and click **Finish**.
15. You can also schedule a backup process by selecting the **Schedule** button in the **Complete the Backup wizard** screen.
16. In the **Backup Status** window, click the **Start Now** button to initiate a backup process.

Creating a secondary backup on a remote CCR node using the Backup wizard

You can create a copy of the database backup and verify it on another CCR node before the full database backup is created on the target node. This backup is called a secondary backup. You can use the Backup wizard to create secondary backup copies. This option is enabled by default in the wizard.

About this task

Create the primary backup considering whether the remote node is an active node or a passive node.

Steps

1. In the Scope pane, click **Backup**.

   SnapManager displays the list of Storage Groups in the **Result** pane and corresponding backup actions in the Actions pane.
2. Click **Backup wizard** in the **Actions** pane.
3. Select the Storage Groups that you want to back up or verify.
4. In the **Select an operation** panel, select the type of operation.

If you want to...	Then do this...
Back up databases and transaction logs	Select the **Back up databases and transaction logs** option.

If you want to...	Then do this...
Create a database backup copy on another CCR node	Select the **Create a copy database backup on another CCR node along with current backup** check box below the **Backup databases and transaction logs** option.
Verify transaction logs and databases	Select the **Verify databases and transaction logs in the most unverified backups** option.

5. Click **Next** to continue and follow the instructions in the Backup wizard to initiate a SnapManager Backup.

Creating a secondary backup on a remote CCR node using the Backup and Verify window

You can use the Backup and Verify window to create secondary backup copies more quickly than using the Backup wizard.

Steps

1. In the SnapManager console, select the Continuous Cluster Replica (CCR) node on which you want to create the backup.
2. Click **Backup** either in the CCR active node or passive node in the Scope pane.
3. Select the Storage Groups you want to back up and verify.
4. Click **Backup and Verify**.
5. Click **Advanced Options**.
6. Select **Create a copy database backup on another CCR node along with current backup**.

   You have created a secondary backup on a remote CCR node.

How SnapManager creates a secondary backup on a remote CCR node

When you create a copy of the secondary backup on a remote CCR node, SnapManager performs a particular sequence of steps.

1. Connects to another CCR node to rename backups and updates the backup timestamp if there is a more recent backup.
2. Creates a copy of database backup on the remote CCR node and deletes the older backups.

   Note: SnapManager does not regroup the backups on the remote CCR node based on the CCR node volume that is used by the databases.

3. Creates a full database backup on the CCR node on which the full backup was initialized and performs verification of the backup sets, if you select the Run Verification option.
4. Connects to the remote CCR node and copies the extra log files, which might be truncated from the primary node to the remote node.

Note: The secondary backup created on the remote is always standard. The secondary backup created on the remote node is not verified even if the primary backup is verified.

Reasons that a SnapManager backup might fail

If your SnapManager backup fails, check the backup report to determine the cause of the failure.

Problem: cluster failover during backup

If a cluster failover or a Windows cluster move group happens during a backup operation, the backup fails, and you need to restart the backup operation.

Problem: Snapshot copy limit reached

You get an error message if you try to back up a LUN that contains more than 255 Snapshot copies. The backup operation fails irrespective of whether SnapManager created the Snapshot copies or not.

Note: Automatic backup deletion is performed only after a successful backup process is complete. Therefore, you must be able to create new Snapshot copies before you begin a new backup procedure.

Problem: SnapInfo directory being accessed

You get an error message if you access the SnapInfo directory while performing a backup operation.

A SnapManager backup operation might include renaming a SnapInfo subdirectory, and Windows does not allow you to change a directory name while it is being accessed. Accessing the SnapInfo directory with Windows Explorer might cause the backup to fail. Ensure that you do not hold exclusive access to the SnapInfo directory on the Exchange host when a backup process is performed.

Problem: SnapInfo directory out of space

You get an error message if your SnapInfo directory runs out of space. Expand the LUN that contains the SnapInfo directory and ensure that enough space remains in the volume for Snapshot copy creation.

Problem: data does not match

You get an error message if you make changes to your Exchange Server storage groups and databases after SnapManager started and you did not refresh your view. You can refresh your view by pressing F5, or you can restart SnapManager.

Problem: busy Snapshot copy

You get an error message if you back up a LUN when a Snapshot copy of the LUN already exists and then you try to delete a Snapshot copy of the LUN. Event 249 is logged by SnapDrive and SnapManager backups fail.

Problem: Snapshot copy already exists

You get an error message either if the system clock on the SnapManager host is not synchronized with the storage system clock or if SnapMirror replication is running when you try to start a backup operation.

Synchronize the system clock on the SnapManager host and the storage system clock so that SnapDrive functions properly. Also, ensure that any SnapMirror replications have enough time to complete before you initiate another SnapManager backup process.

Problem: out of disk space

You get an error message when the database or the transaction logs use all of the available disk space in a volume. Resize the volume or expand the LUN.

Problem: SnapManager server initialization failed

You get an error message either if the SnapManager server account or the server account permissions have changed; or if you exit SnapManager when the smesrvr.exe process is running.

Ensure that you use the correct server account and the correct server account permissions and that you terminate any orphaned SnapManager processes that run when you exit SnapManager.

Verifying database backups

If you back up a database without automatically verifying it, you can verify the database after the backup.

Note that Exchange 2010 database verification is not a support requirement for databases with at least two copies in a database availability group. See *How SnapManager checks database integrity in backup sets* on page 106 for more information.

Decisions to make before database verification

Before you start or schedule verification, you need to gather the information required to complete the Backup Wizard or the Backup and Verify window.

Job-specific parameters

- Which are the databases for which you want to verify unverified Snapshot copies?
- Which are the backup management groups within the databases that you selected?
- What is the number of unverified backup Snapshot copies that you want to verify?

  You can verify more unverified Snapshot copies than you specified earlier in your database and backup management group selections.
- Do you want to run a command after the backup is complete?

  This is usually done to archive backups.
- Do you want to run the verification now or schedule it for later?

  If you want to schedule the verification to run later, you also need to know the job scheduling information.

Verification settings

- Which Exchange server do you want to use to perform database verification?

  Only Exchange Server 2010 can be used to verify Exchange Server 2010 mailbox databases. An Exchange Server 2010 mailbox database cannot be verified on Exchange Server 2003 or Exchange Server 2007, while Exchange Server 2003 or Exchange Server 2007 mailbox databases cannot be verified on Exchange Server 2010.

  You can configure the server by using the Verification Server tab of the Database Verification Settings dialog box. If you specify a remote verification server, ensure that you correctly set up the server.

  Note: Whether you perform a remote or local verification, SnapManager performs the verification on a LUN that is backed up by a Snapshot copy. If you make a Snapshot copy of the same volume on which a LUN that is backed up by a Snapshot copy exists, you create a “busy Snapshot copy.” This might cause problems when you attempt to delete some Snapshot copies. For this reason, you must be careful not to schedule backups while a verification operation is in progress.
- Which method do you want to use to access database Snapshot copies during database integrity verification?
- Do you want to configure throttling of the database checksum verification rate?

  Throttling is available only if SnapManager is installed with Exchange Server 2003 SP2, or later, or Exchange Server 2007.
- When you restore from an unverified backup, do you want to override the verification requirement for restoring databases from backups?

Starting or scheduling database verification

You can run or schedule database verification from the production Exchange server. You cannot use the remote verification server.

Steps

1. On the production Exchange server, click **Backup and Verify** in the **Actions** pane.
2. Select the Storage Group on which you want to run the verification.
3. Select **Verify Databases and Transaction Logs in the Most Recent Unverified Backups**.
4. Select the required, and most recent unverified backups you want to verify.

   Note: Only unverified backups are eligible for verification. For example, if you select to verify the two most recent unverified backups, but the most recent backup's databases have already been verified, then the previous two unverified backups are verified.

5. In the **Backup Management Group** field, select the backup management group of the backups that you want to verify.

   If you want to verify the most recent backups regardless of their backup management group, select **All**.

6. If you want to run a command after the verification is complete, select the **Run Command After Operation** check box.

   The command that you enter in the **Run Command After Operation** option archives backups.

7. If your volume is a SnapMirror source volume and you do not want the destination volume to be updated after this verification is complete, clear the **Update SnapMirror After Operation** check box.
8. If your volume is a SnapMirror destination volume and you want to run integrity verification on the available destination volume, select the **Verify databases and transaction logs on the SnapMirror destination** check box.
9. Click **Verify Now**, and in the **Backup Status** window, click **Start Now**.

   Note: To schedule the verification for later, click **Schedule** instead and follow the prompts in the dialog boxes.

   Verification is run on the databases in the backup that you selected.

After you finish

If any backup database fails in the verification operation, delete that backup and create another backup.

Backup management groups

Creating backup management groups enables you to designate various levels of backup retention, which facilitates your database backup strategy. The backup management group neither depends on nor enforces how often backups are performed. Backup management groups are only a backup labeling convention that you can assign to a backup set.

You can identify and group backups according to the operations to be performed on them: manual deletion, automatic deletion, and database verification.

When you explicitly delete multiple backups, you can specify that only backups belonging to a certain backup management group be deleted.

When you run or schedule a backup, you can specify how many of the most recent backups in a group you want to retain. Only backups of the specified backup management group are deleted.

When you run or schedule a database verification separate from the database backup operation, you can limit the number of backups you want to verify by specifying a particular backup management group.

The database Snapshot copy names and SnapInfo directory Snapshot copy names reflect the management group to which you assigned the backup.

Backup management group assignments

You can assign a backup to a Standard, Daily, or Weekly backup management groups. When you start or schedule a database backup, the Backup wizard and the Backup and Verify window use the Standard management group by default.

The choice of backup management group and the backup Snapshot copy naming convention affects the name that is assigned to the Snapshot copy. The name of each Snapshot copy created during a SnapManager backup operation includes information that identifies the Snapshot copy contents.

Example using backup management groups

You can use backup management groups to perform simultaneous backups, on a schedule that you can customize. You can schedule backups on a Standard, Weekly, or Daily basis using the backup management groups.

Suppose you want to make backups at regular intervals between 7:30 a.m. and 7:30 p.m. You want to keep the last backup of the day and retain it for a few weeks, and you want to keep one backup per week for several months for archiving.

To achieve this using backup management groups, you can use the Standard backup management group for the backups during the day. Use a separate backup job to create one backup in the Daily management group at the end of the day. Then, once a week, you can use another job to create a backup in the Weekly backup management group.

You can then decide how many backup copies to retain independently for each backup management group. For example, you can keep ten standard backups, seven daily backup copies (one week's worth), and four weekly backup copies (one month's worth).

If your daily or weekly backup job fails for any reason, you can replace it with the most recent successful standard backup by changing its backup management group.

Assigning a backup set to a different backup management group

You can use the dialog box Change Backup Management Group to change the backup management group to which a backup set belongs.

About this task

- The backup management group for all these backups is changed if you complete this operation. The change occurs because the backup management group share a common Snapshot copy.
- When you change a backup's backup management group, you also change that backup's name, because the name includes the backup management group.
- The report for the backup management group change is in the Miscellaneous report directory.

Steps

1. Click **Restore** in the Scope pane.

   SnapManager displays the list of Storage Groups in the Result pane.

2. Expand the Storage Group that contains the backup set whose management group you want to change in the Result pane.
3. Double-click the backup whose management group you want to change.
4. Click **Change Management Group** in the **Actions** pane.

   Note: If you have datasets configured in your system, and archiving of the primary Storage Group to the secondary Storage Group is in process, you can change the backup management group only after the database transfer for archiving completes.

5. Review the backups listed in the "Backups sharing this Snapshot" list.

   Note: The backup management group for all of these backups is changed if you complete this operation. The change occurs because the backup management group members share a common Snapshot copy.

6. In the "New Management Group" list, select the backup management group to which you want to move the backup set.
7. Click **OK**.

   The backup management group for this backup and all backups listed in the "Backups sharing this Snapshot" list is changed.

Frequent Recovery Point backup operation

A Frequent Recovery Point backup copy operation backs up new transaction logs and is performed after a full backup operation. You can restore data up to a recovery point that you select. SnapManager combines the restore operation of a full backup copy and the required transaction logs to restore to the selected recovery point.

How the Frequent Recovery Point feature works

Frequent Recovery Point backs up transaction logs at a frequency that you determine to meet the Recovery Point Objective (RPO). Frequent Recovery Point backs up the most recent transaction logs created after the last full backup copy or the previous Frequent Recovery Point backup copy.

The Frequent Recovery Point feature creates backup copies at a specific frequency and names them using a unique naming convention. The lowest value for the interval between any two Frequent Recovery Point backup copies is 10 minutes. The default value is 15 minutes.

You can create a Frequent Recovery Point backup copy after you create a full backup copy. You can trigger a SnapMirror update of the SnapInfo volume after the Frequent Recovery Point backup is complete.

A Frequent Recovery Point backup operation creates a Snapshot copy on the SnapInfo volume and names it using the following convention:

`efrpinfo__<exchserver name>_date_time`

SnapManager retains only one Snapshot copy with this name.

You can use "Run command after operation" with the Frequent Recovery Point feature.

Frequent Recovery Point backup operations

There are various operations that you can perform with Frequent Recovery Point backups.

- Create a Frequent Recovery Point backup
- Schedule a new Frequent Recovery Point backup job with the name FRPBackup
- Modify an existing scheduled Frequent Recovery Point backup job
- Delete an existing Frequent Recovery Point backup job

Frequent Recovery Point backup operation on clustered configurations

The Frequent Recovery Point backup operation that SnapManager performs on clustered configurations depends on the type of clustered configuration you use.

Configuration	Frequent Recovery Point operation
Exchange 2007 LCR	If you select a database for Frequent recovery Point backup operation, SnapManager backs up the transaction logs of the replica database also.

Configuration	Frequent Recovery Point operation
Exchange 2007 CCR	If you select the copy-based backup option when performing a Frequent Recovery Point backup operation, SnapManager backs up the transaction logs on the remote node also.
Exchange 2003 Cluster	SnapManager backs up the transaction logs on the active node.

Frequent Recovery Point backup operation in a DAG

You can perform a Frequent Recovery Point (FRP) backup operation in a Database Availability Group (DAG) by connecting either to the DAG or to a member server.

When you connect to the DAG, the Frequent Recovery Point backup operation creates FRP backups on all member servers in the DAG. You can create FRP backup for either active or passive databases, or for all databases on a specified member server, or on all member servers in the DAG.

Each server has its own FRP backup, independent of other member servers in the DAG.

Verification of Frequent Recovery Point backup copies

SnapManager does not perform transaction log verification when it creates Frequent Recovery Point backup copies. When you verify a full backup copy, SnapManager verifies all its previous Frequent Recovery Point backup copies.

Deletion of Frequent Recovery Point backup copies

When you delete a full backup without enabling up-to-the-minute restore functionality, SnapManager deletes all the subsequent Frequent Recovery Point (FRP) backup copies up to the next full backup. You cannot delete FRP backup copies independent of full backups.

Frequent Recovery Point backup reports

SnapManager creates a separate folder in which to save Frequent Recovery Point backup reports and names it using the naming convention "`FRP Backup [server]`". The naming convention for the file names of reports does not change.

Performing a Frequent Recovery Point backup operation

You can either create a Frequent Recovery Point backup or schedule Frequent Recovery Point backups.

Steps

1. In the Scope pane, select an Exchange server.
2. In the **Actions** pane, click **Frequent Recovery Point Backup**.

If...	Then...
You have not scheduled the backup job of Frequent Recovery Point	SnapManager does not display any job details. You can use the **Create Job** button to create a new job specification.
You have scheduled the backup job of Frequent Recovery Point	SnapManager displays the job details. Use **Update Job** to update the job specifications.

Note: You can disable the current job by selecting the **Disable scheduled job** check box.

3. Under "Selected for backup," click **Select All** to select all the Storage Groups or databases at one time, or click **Unselect All** to clear all your selections.
4. Under "Backup Frequency", specify the maximum frequency at which the recovery point backup needs to be made:
   a) In the "Every" list, select a number of minutes or hours from the list box to specify the interval between the Frequent Recovery Point backups.
   b) In the "Start at" list, select the time to run the Frequent Recovery Point backup.
5. Under "Operation options", select **Run command after operation** if you want to run a command or script of your choice after the backup operation completes.
6. Do one of the following:

If you want to...	Then do this...
Create one Frequent Recovery Point backup	Click **Create Recovery Point**. The Backup Status window is displayed with the list of tasks.
Run Frequent Recovery Point backups at a specific interval	Click **Create Job**. A scheduled job named FRPBackup is created under Windows Scheduled Tasks.

7. When you click **Create Recovery Point** and the backup status is displayed, click **Start Now** to create one new recovery point backup.
8. In the **Backup Status** window, click **Close** when the backup operation completes.

   SnapManager displays a message with the result of the Frequent Recovery Point backup.

Creating multiple FRP backup jobs for different Exchange instances

You can create multiple FRP backup jobs for different Exchange instances in a Microsoft Active/Active cluster configuration using SnapManager. Such a configuration can have two or more active Exchange instances on a single node, so two or more FRP backup jobs are required, one for each instance of Exchange.

Steps

1. Connect to one of the Exchange instances on the cluster.

2. Create the FRP backup job.

   Note: The FRP backup job created by SnapManager is always named FRPBackup.

3. Rename the FRP backup job on all nodes by using **Windows Scheduled Tasks**.
4. Connect to another Exchange instance and repeat Steps 2 and 3.

Advantages of using the remote additional copy backup feature

For database availability groups consisting of three nodes with no more than eight databases per node, you can create a simpler, easier backup strategy than is needed for larger systems by using the SnapManager for Exchange feature remote additional copy backup. The simplified plan still permits up-to-the-minute restore operations, even if a node is unavailable.

With remote additional copy backup enabled, you can do the following:

- Set up gapless DAG backups by requesting a copy of the databases to be backed up using the full backup procedure.

  You should not select duplicated copies of the same database for gapless backups. Although duplicate gapless backups work, typically they take longer than necessary to create several redundant gapless backups. You can select a single copy of each database for the gapless backup using one of the following methods:

  - Using tree view, check a single copy of each database.
  - Name the single database copy using a PowerShell `-database` parameter

  Alternatively, you can optimize the list of copies with the following database filter options:

  - **Only active databases**
  - **Only passive databases with specific activation preference for backup**
- Enable gapless DAG backup by selecting **Create a copy database backup on other DAG nodes along with current backup**.
- Select on which nodes the copy-only backups are to be performed.
- For the copy-only backup on other nodes, select the backup retention policy and request that the policy be executed after the copy backup is created.

Additional considerations for using gapless DAG backup include:

- If you perform a gapless DAG backup, do not also perform a primary backup of each copy of the database; otherwise, redundant and unnecessary backups are created.
- You can perform a copy backup on active as well as passive databases.

When SnapManager performs a remote additional copy backup, it does the following:

1. Creates copy-only backups on those nodes selected for copy backup

2. Creates a full backup on the selected databases
3. Backs up the logs on the selected nodes
4. Makes a request to Microsoft Exchange to truncate the log on the primary backup node
5. Replicates the truncation to the remote copies of the database

Configuring gapless DAG backups using remote additional copy

On Database Availability Groups (DAGs) consisting of three nodes and no more than eight databases per node, you can use a simplified procedure to back up the databases while preserving your ability to perform up-to-the-minute restore operations, even if a node becomes unavailable.

Steps

1. Connect directly to the DAG server.

   You cannot use remote additional backup copy if you directly connect to a member server.

2. Select the databases for which you want to create a backup.

   Do not select multiple copies of the same database for backup.

3. Using the **Advanced Options** window, select the **Remote Additional Copy Backup** tab.
4. Select **Create additional backup (copy based) on CCR (Cluster Continuous Replication) or DAG remote server(s)**.
5. Optionally select either or both of **Delete backups during additional backups on remote servers** and **If the dataset is enabled, archive backups created on remote servers to secondary storage**.
6. Using the **Advanced Backup Options** dialog box, select on which server you want to run copy-based additional remote copy backups.

   For best results, select all DAG nodes (this is the default).

PowerShell new-backup cmdlet parameters for DAG backups

You can use the PowerShell new-backup cmdlet with various parameters to create and manage backups, using the remote additional copy backup feature. For example, you can specify the number of backup copies to keep on remote nodes, how long backup copies are kept, and whether they are archived.

Use the following parameters to the `new-backup` cmdlet to create backups on either DAGs or CCR servers.

If you want to do this...	Then use this parameter...
Specify that any backups also have a remote additional backup copy	`-RemoteAdditionalCopyBackup` `[Boolean]` Short form: `remCopyBk` Default values: false = DAG, true = CCR If this parameter is specified, if databases are replicated to other member servers, a remote additional backup (copy based without log truncation) of the selected databases of the remote nodes is created. This applies only to DAG or CCR server connections.
Specify DAG nodes allowed to have remote additional (copy-based) backups	`-AdditionalCopyBackupDAGNode` `[String]` Short form: `remCopyBkNode` Default value: all DAG member nodes are allowed
Specify the number of remote additional (copy-based) backups to keep on remote nodes.	`-RetainRemoteAdditionalCopyBackup` `Integer` Short form: `rtRemCopyBk` Default: 1 This is used when the backup completes to determine which previous backups to retain.
Specify the number of days to keep backups when doing a remote additional (copy based) backup on a remote node.	`-RetainRemoteAdditionalCopyBackupDays` `Float` Short form: `rtRemCopyBkDays` Default value: no default value Using this parameter indicates that deletion is based on an “older than number of days” policy. If this parameter is not used, the deletion policy is based on the number of retained additional backups.
Create an archive of the remote additional (copy-based) backup created on the remote cluster member nodes. Include this parameter if you have datasets configured on those nodes and want to archive the remote additional copy-based backup to another secondary SnapVault storage system.	`-ArchiveRemoteAdditionalCopyBackup` `[SwitchParameter]` Short form: `archRemCopyBk` Default value: false

Example PowerShell `new-backup`command for remote additional backups

This command creates a backup of the active DAG databases and also performs a remote additional (copy-based) backup of the databases that are replicated to member servers DAG1-NODE1 and DAG1-NODE2. If datasets are configured on DAG1-NODE1 or DAG1-NODE2 member servers, then SnapVault creates additional backups on the secondary storage system.

```
new-backup -Server ‘SM-DAG1’ -ClusterAware -ActiveDatabaseOnly
-RemoteAdditionalCopyBackup $true
-ArchiveRemoteAdditionalCopyBackup
-AdditionalCopyBackupDAGNode DAG1-NODE1, DAG1-NODE2
-RetainRemoteAdditionalCopyBackup 4
-BackupTypeFullBackup
```

control.xml file settings for remote additional copy backup

You can control various aspects of remote additional copy backup by including certain settings in the control.xml file. For example, you can control the number of backup copies kept on a remote node, and how long they are kept.

If you want to do this...	Then use this setting...
Control the deletion of backups on remote nodes	`<DELETE_BACKUPS_OPTION_REMCOPYBACKUP>XXXX</DELETE_BACKUPS_OPTION_REMCOPYBACKUP>` integer 0=no backups are deleted on the remote nodes, 1=backups are deleted on the remote nodes
Control the number of backups to keep on the remote node after a gapless backup is done	`<DELETE_BACKUPS_OPTION_REMCOPYBACKUP> XXXX</DELETE_BACKUPS_OPTION_REMCOPYBACKUP>` integer 0=no backups are deleted on the remote nodes, 1=backups are deleted on the remote nodes
Control how long (days or part days) to keep backups after a gapless backup is done	`<BACKUP_SET_TO_KEEP_IN_DAYS_REMCOPYBACKUP> XXXXX.XXXX </BACKUP_SET_TO_KEEP_IN_DAYS_REMCOPYBACKUP>` real-float Equals days or partial days to keep backups after a gapless backup is done

If you want to do this...	Then use this setting...
Define the member servers on which remote copy backups can be done	`<HOSTS_REMCOPYBACKUP> SERVERNAME1, SERVERNAME2, ...</HOSTS_REMCOPYBACKUP>` string Comma-delimited list of member servers on which to perform remote copy backup; if none is specified, all cluster nodes are processed
Control whether additional remote additional copy backups are done	`<REQUEST_REMCOPYBACKUP>xxxx</REQUEST_REMCOPYBACKUP>` integer 0=no remote additional copy backup is done, 1=remote additional copy backup is done

Example control.xml fragment for remote additional copy backup control

```
<BACKUP>
   <BACKUP_CLIENT_SETTING>
<DELETE_BACKUPS_OPTION_REMCOPYBACKUP>1</
DELETE_BACKUPS_OPTION_REMCOPYBACKUP>
   <BACKUP_SET_TO_KEEP_REMCOPYBACKUP>8</
BACKUP_SET_TO_KEEP_REMCOPYBACKUP>
<BACKUP_SET_TO_KEEP_IN_DAYS_REMCOPYBACKUP>7.0308
   </BACKUP_SET_TO_KEEP_IN_DAYS_REMCOPYBACKUP>
   <HOSTS_REMCOPYBACKUP>SNAPMGR-11-VM30,SNAPMGR-11-VM31</
HOSTS_REMCOPYBACKUP>
   <REQUEST_REMCOPYBACKUP>1</REQUEST_REMCOPYBACKUP></
BACKUP_CLIENT_SETTING>
</BACKUP>
```

Types of database restore operations using SnapManager

Using SnapManager, you can restore a Local Copy Replication (LCR) copy, a Cluster Copy Replication (CCR) copy, a Database Availability Group (DAG) copy, databases created on a different Exchange servers, or an archived backup set. You can also restore a database to a Recovery Storage Group (RSG) or Recovery Database.

You can restore your databases either from a backup copy that you created previously, or from an archive. You can then optionally restore the databases up to their current state by replaying the transaction logs. Replaying the transaction logs is necessary if your Exchange data becomes corrupted or becomes unavailable. With SnapManager 5.0 and later, you can restore the production backup copies, LCR-enabled backup copies, and CCR-enabled backup copies to either of the following locations:

- The same Storage Group
- The Recovery Storage Group

Additionally, with SnapManager 6.0 and later, you can also restore database copies to member servers in the DAG.

SnapManager supports restoring an entire Storage Group or an individual database.

When to choose SnapManager Restore to recover Exchange 2010 mailbox databases

You should know when to use SnapManager to recover an Exchange 2010 mailbox from database failure or data corruption.

The database switchover to a different server in the Database Availability Group (DAG) can be done using the EMC or the Exchange management shell.

Use SnapManager in the following situations:

- When using Exchange Server 2010 and DAG cannot recover the failed database.
- When using Exchange Server 2010 and DAG to recover database will cause data loss.
- When using Exchange Server 2010 and DAG cannot recover the database to a particular point in time: for example, when a virus has spread to all copies of databases.
- When using Exchange Server 2010 and DAG is confusing to the administrator, and using SnapManager Restore is simple, fast and reliable.

How SnapManager Restore works

SnapManager Restore performs a sequence of steps when it restores a backup copy. If you are restoring in a Windows cluster, do not attempt to manage any cluster resources while the restore operation is running.

1. Checks that the backup copy is verified.
   If not, SnapManager checks if you have enabled the Database Verification Override option and continues.

   Attention: You must restore only from verified backup copies to ensure a successful restore operation.

2. Creates a Snapshot copy of the volume to be restored, so that the state before the restore operation is preserved.
3. Restores the data from the Snapshot copy of the LUN to be restored.
4. Restores the transaction logs that you need to replay to restore the data.
5. Uses the Exchange recovery process to play these transaction logs to the restored databases.
   On Windows Server 2003, the soft recovery process is used, which is a transaction log replay process that occurs when a database is re-mounted after an unexpected stop, or when transaction logs are replayed into an offline backup of a database.

How to choose the type of restore operation to perform

You should consider your reasons for performing a restore operation so that you can decide what type of restore operation to perform: point-in-time or up-to-the-minute restore operation.

When to use an up-to-the-minute restore operation

- You want to play forward all the transactions up to the most current time.
- You want to restore individual databases.
- You want to restore backup copies after a point-in-time restore operation of a backup copy that is not the most recent.
- You want to perform an up-to-the-minute restore operation from any backup copy, including the most recent backup copy after a point-in-time restore operation of the most recent backup operation.
  You might lose all transactions that occurred between the time when you created the last backup copy and when you performed the point-in-time restore operation.

Note: For to an up-to-the-minute restore operation to succeed, a contiguous set of all required transaction logs must be in the SnapInfo folder and the transaction log directory of the Storage Group.

When to use a point-in-time restore operation

- You want to recover the databases as they were at a particular point in time: example, when the most recent backup copy was created.
- You want to restore databases to a Recovery Storage Group or to a Recovery Database.
- You want to restore all existing backup copies after a point-in-time restoration of a backup that is not the most recent one.

Types of SnapManager Restore operations

There are two types of SnapManager Restore operations. In a point-in-time restore operation, the transaction logs in the active file system not committed to the databases at the time of backup operation are replayed. In an up-to-the-minute restore operation, all uncommitted transaction logs are replayed and are applied to the databases. This option is selected by default.

Point-in-time restore operation	In this case, only the uncommitted transaction logs that existed in the active file system at the time the backup copy was created are replayed. All the transaction logs beyond the point-in-time (after the backup copy was created) that exist in the transaction log directory and that belong to the restored Storage Group are removed. You can use this method to restore a database or Storage Group back to a time before a corruption occurred.
Up-to-the-minute restore operation	There are some transaction logs that are not committed to the databases. In an up-to-the-minute restore operation, all uncommitted transaction logs, from the time the backup set was created up to the most current time, are played forward and applied to the databases. This includes transaction logs from the backup sets, in addition to the transaction logs in the transaction log directory. A contiguous set of transaction logs is required for an up-to-the-minute restore operation to succeed.

Example: Up-to-the-minute restore operation

Suppose you run SnapManager Backup every day at noon, and on Wednesday at 4 p.m. you decide to restore from a previous backup set. For some reason, the backup set from Wednesday noon failed verification, so you decide to restore from the Tuesday noon backup set. After that backup set is restored, all the transaction logs are played forward and applied to the restored databases. This starts with those that were not committed when you created Tuesday's backup set and continues through the latest transaction log written on Wednesday at 4 p.m.

Snapshot copies created during a restore process

Every time you perform a restore process, a Snapshot copy is created on the volume that contains the transaction logs.

The name of the Snapshot copy that is created during a restore operation contains the prefix `rstrsnap__`. After you verify that a restore process is completed successfully and you are satisfied

with the results, you can delete the restored Snapshot copy using the **Delete** option in the Actions pane.

Methods that can decrease restore process time

You can decrease the time that SnapManager requires to complete a restore operation by making more backup copies and ensuring that any LUN clone split operations in progress are complete.

The more backup operations you perform, the fewer Exchange transaction logs need to be played forward at restore time. At a minimum, you should perform one SnapManager full database backup operation every 24 hours.

Before proceeding with failback to the production site, ensure that any LUN clone split operations that are in progress are complete. The LUN clone split functionality, introduced with Data ONTAP 7.1, supports significantly faster online Snapshot copy restore operation when using SnapManager or SnapDrive to restore databases. By default, this functionality is enabled.

Transaction log sequence verification options

You can make a restore operation faster by omitting the pre-restore transaction log sequence and database metadata checks. However, if you omit these checks, and a problem exists with your transaction log sequencing or your database metadata, the restore operation fails.

To perform an exhaustive check of the metadata, select the **Verify Transaction Log Sequence...** and the **Exhaustive Verification** check boxes in **Restore > Advanced Options**. For backward compatibility among SnapManager versions, the exhaustive check of metadata is selected by default.

To perform a quicker, less comprehensive check of the transaction logs, select **Verify Transaction Log Sequence...** and clear **Exhaustive Verification** in **Restore > Advanced Options**.

To omit metadata verification, clear the **Verify Transaction Log Sequence...** and the **Exhaustive Verification** check boxes in **Restore > Advanced Options**.

You have the following pre-restore transaction log sequence options available:

- Exhaustive transaction log sequence checks
- Non-exhaustive transaction log sequence checks

Exhaustive transaction log sequence

The exhaustive transaction log sequence checks the transaction log headers and mounts the Snapshot copies for further verification.

SnapManager follows this sequence of steps:

1. Analyzes the transaction log headers.
2. Verifies the following conditions:
   - The transaction logs are all present and sequenced correctly.
   - The signatures of all the transaction logs match.

- Each transaction log's create time matches the next log's previous time.

3. Mounts a Snapshot copy of the databases in this backup copy and analyzes the database headers.
4. Checks if the database signature matches the log database signature.

Non-exhaustive transaction log sequence

The non-exhaustive transaction log sequence uses only the file names of the transaction log for verification. SnapManager verifies that the transaction log sequence is correct and that all log files are present.

This option is enabled if you clear the **Exhaustive Verification** check box in **Restore > Advanced Options**.

LUN Clone Split Restore method

SnapDrive uses the LUN Clone Split Restore method for Data ONTAP 7.1 or later. This feature enables you to perform a restore operation quickly, because the only data blocks that are duplicated during the restore operation are those currently in use in the active file system.

Note: The restore destination must exist in the active file system.

SnapManager follows these steps to restore a LUN to its LUN clone that exists in the active file system:

1. SnapManager creates a LUN clone.
   The LUN clone is unavailable for data I/O. The LUN clone inherits the presentation of the LUN in the active file system.
2. SnapManager renames the LUN clone with the name of the original clone.
3. LUN Clone Split (LCS) starts in the background.
   The LUN becomes available for data I/O after LCS completes.

Verifying the LUN Clone Split status

You can determine the LUN Clone Split status of your LUNs during a restore operation using the storage system console or SnapDrive.

Step

1. Use the Data ONTAP `lun clone split status` command on the storage system console.

   Additionally, SnapDrive 6.0 and later displays the LUN Clone Split status in the **Operation Status** column of the SnapDrive user interface. For more information, see the relevant SnapDrive and Data ONTAP documentation.

SnapManager Restore in a Windows cluster

In a Windows cluster, SnapManager does not take the entire Exchange virtual server offline; instead, it takes the online Storage Group that is to be restored offline before performing a restore operation.

SnapManager takes the Storage Group that is to be restored offline before performing the restore process, if that Storage Group is still online.

If a cluster move group operation occurs during the restore process, you must restart SnapManager and the restore process.

SnapManager Restore in a live Exchange virtual server cluster

SnapManager can restore databases in a clustered Exchange virtual server while keeping the cluster physical disk resource or the Exchange virtual server online. Keeping the cluster physical disk resource or the Exchange virtual server online minimizes downtime during the restore operation for databases in Storage Groups that share a cluster group but are not being restored.

This feature applies to LUN restore operations that use the Single-File SnapRestore (SFSR) method as well as to those that use the LUN Clone Split restore method supported with Data ONTAP 7.1.

Guidelines for using SnapManager Restore

You must follow some guidelines for choosing if you want to use an up-to-the-minute restore operation or a point-in-time restore operation; about the restore destination, and the naming of the Storage Groups.

How to choose the type of restore operation to perform

You should consider your reasons for performing a restore operation so that you can decide what type of restore operation to perform: point-in-time or up-to-the-minute restore operation.

When to use an up-to-the-minute restore operation

- You want to play forward all the transactions up to the most current time.
- You want to restore individual databases.
- You want to restore backup copies after a point-in-time restore operation of a backup copy that is not the most recent.
- You want to perform an up-to-the-minute restore operation from any backup copy, including the most recent backup copy after a point-in-time restore operation of the most recent backup operation.

  You might lose all transactions that occurred between the time when you created the last backup copy and when you performed the point-in-time restore operation.

Note: For to an up-to-the-minute restore operation to succeed, a contiguous set of all required transaction logs must be in the SnapInfo folder and the transaction log directory of the Storage Group.

When to use a point-in-time restore operation

- You want to recover the databases as they were at a particular point in time: example, when the most recent backup copy was created.

- You want to restore databases to a Recovery Storage Group or to a Recovery Database.
- You want to restore all existing backup copies after a point-in-time restoration of a backup that is not the most recent one.

Guidelines for restoring from a SnapManager backup copy

You must follow some guidelines about the type of restore operation, the restore destination, and naming of the Storage Groups before you begin to restore from a backup.

- Use the most recent backup copy for an up-to-the-minute restore process; otherwise, more transaction logs need to be replayed, and the restore process takes more time.

  The most recent backup copy's name contains `recent` if you did not use the unique naming convention; otherwise, it uses the most recent timestamp.
- You can restore an entire Storage Group or, if you configured your Exchange Server storage groups and databases appropriately, individual databases.

  If you need to restore more than one Storage Group, run SnapManager Restore more than once.
- If you rename a Storage Group after backing it up, you cannot recover it from that backup copy.

  If you must recover a Storage Group from a backup copy made before the Storage Group was renamed, return the Storage Group to its former name, then recover it from the backup copy by the same name.
- If you rename a Storage Group, make sure that you back it up as soon as possible.

  Note: When you rename a Storage Group, unmount all of the databases in that Storage Group. After they are all unmounted, you can remount the databases. At this point, Exchange recognizes the new Storage Group name, and SnapManager can make new backup copies.
- If you restore the Storage Group to the Recovery Storage Group, you can perform only a point-in-time restore process.
- Store all databases that reside on the same LUN together.

Restore from a SnapManager backup copy

You can access SnapManager Restore to restore databases or Storage Groups from a backup copy in two ways.

- SnapManager Restore wizard
- Restore window accessed from the Actions pane

Decisions to make before restoring from a SnapManager backup copy

When you perform a restore operation, SnapManager prompts you to provide some information like the restore object and destination, type of restore operation, and verification options that you need to keep ready with you before you make decisions during the operation.

Restore object and destination

- What items do you want to restore?

  You can select the production databases and LCR copies that you want to restore.
- Which database do you want to restore?
- Which Storage Group do you want to restore, or which Storage Group contains the databases that you want to restore?
- To what location do you want to restore the backup sets?

  You can restore the backup sets to either one of the following locations:

  - To the same Storage Group or same server
  - To the Recovery Storage Group or Recovery Database

Type of restore operation

- Do you want to perform an up-to-the-minute restore operation (the default) or a point-in-time restore operation?
- Do you want to perform an actual restore operation or only a test restore operation?

  A test restore operation gives you a preview of the actual restore process to help you decide if you want to go ahead with your current settings.
- If you want to perform an actual restore operation, do you want to mount your databases after the restore operation is complete?

Verification options

- Do you want to verify the transaction log sequence and database metadata before the restore operation?

  For Exchange Server 2007, SnapManager recovers the databases if you select **Recover and mount databases after restore**. For Exchange Server 2003, SnapManager recovers and mounts the databases if you select the **Mount databases after restore** option. If you clear this option, SnapManager recovers only the databases.
- If you want to perform a test up-to-the-minute restore operation, do you want to verify the current transaction logs?

  The option to verify the current transaction logs confirms that all the required transaction logs are present.
- If the backup copy is unverified, do you want to verify it before the restore operation?
- Which verification settings do you want to apply for restoring backup sets?
- Do you want to select SnapMirror destination volumes for integrity verification?

If yes, click the **Verification Server** tab in the Database Verification Settings dialog box to select SnapMirror destination volumes.

Mount options

When you perform a restore operation, you can specify the mount options after restore. You can select the option to enable SnapManager to mount the databases immediately after completing the restore operation or you can mount the databases manually at a later time.

When you are restoring backups that were created on other Exchange servers or restoring backup sets from an archive, you have the additional option to associate all mailboxes in the newly restored databases to the currently connected Exchange server.

Restoring databases using the Restore wizard

You can restore a Storage Group from a backup set by using the SnapManager Restore wizard which guides you through the restore process.

Before you begin

- Close all windows on the Exchange server that runs SnapManager.
- Disable any SnapManager operations that are scheduled to run against the Exchange data that you want to restore, including any jobs scheduled on the remote administration server or the remote verification server.

About this task

You can restore only one database or only one Storage Group or any one of the databases in a Storage Group at a time.

You can optionally perform a SnapManager backup and verification operation after restoration to verify that your restored database is free of physical-level corruption.

Do not perform a restore operation while a backup operation is in progress. If you cancel a current backup operation in progress, SnapManager pauses all the active scheduled backup jobs on the Exchange server, or on all nodes in the cluster environment, and cancels the current backup copy before performing the restore operation. On completing the restore operation, SnapManager reenables the paused scheduled backup jobs. All the other inactive jobs are not changed.

If you want to restore after the current backup operation completes, SnapManager pauses all the active scheduled backup jobs on the current Exchange server (on all nodes in the cluster environment) but waits for the completion of the current backup operation before performing the restore operation. On completing the restore operation, SnapManager reenables the paused scheduled backup jobs. All the other inactive jobs are not changed.

If you select a backup that is or has any archived backups, the Results pane displays the following details about them:

- EDB primary storage
- EDB primary LUN path

- EDB secondary storage
- EDB secondary qtree path
- EDB secondary Snapshot copy name

The EDB secondary storage, EDB secondary qtree path, and EDB secondary Snapshot copy name are used to mount the Snapshot copies using the SnapDrive CLI. To recover a mailbox from the archived backup, use the Single Mailbox Recovery tool.

Steps

1. Select the Exchange server node in the Scope pane, and then click **Restore**.

   SnapManager displays the Exchange Server storage groups and databases and the Storage Groups. When the SnapManager connects to the Database Availability Group (DAG), all backups in the DAG are displayed and when it connects to a member server of the DAG, only the backups on that server are displayed.

2. Double-click the Storage Group or Database that you want to restore.
3. In the **Actions** pane, click **Restore wizard**.
4. Follow the instructions in the Restore wizard and go to the **Restore Status** window.
5. After the restore process is complete, click **OK**.

   Your Exchange server comes back online.

Restoring databases using the Restore window

Unlike using the Restore wizard to restore your databases, using the Restore window enables you to configure the restore process on your own.

Before you begin

- Close all windows on the Exchange server that runs SnapManager.
- Disable any SnapManager operations that are scheduled to run against the Exchange data that you want to restore, including any jobs scheduled on remote management or remote verification servers.

About this task

Do not perform a restore operation while a backup operation is in progress. If you cancel a current backup operation in progress, SnapManager pauses all the active scheduled backup jobs on the Exchange server, or on all nodes in the cluster environment, and cancels the current backup copy before performing the restore operation. On completing the restore operation, SnapManager reenables the paused scheduled backup jobs. All the other inactive jobs are not changed.

If you want to restore after the current backup operation completes, SnapManager pauses all the active scheduled backup jobs on the current Exchange server (on all nodes in the cluster environment) but waits for the completion of the current backup operation before performing the

restore operation. On completing the restore operation, SnapManager reenables the paused scheduled backup jobs. All the other inactive jobs are not changed.

After an actual restore process, you can optionally perform a backup and verification operation to verify that your restored database is free of physical-level corruption.

Steps

1. In the Scope pane, click **Restore**.
2. Select the Storage Group or Database that you want to restore.

   If you want to restore from an archived backup copy and have datasets configured in your system, select the SnapVault and SnapMirror-enabled storage systems from the listed archived backup copies.
3. Double-click the backup copy under the Storage Group or Database that you want to restore.
4. In the **Actions** pane, select **Restore**.
5. In the **Choose databases to restore** pane, select the databases that you want to restore.
6. Under **Type of Restore**, select if you want an up-to-the-minute restore or a point-in-time restore operation.
7. If you do not want to perform exhaustive verification of the transaction log sequence and database metadata before the restore process, click **Advanced Options**.
8. Under **Operation Options**, clear the **Verify Transaction Log Sequence and Database Metadata Before Restore** and **Exhaustive Verification** check boxes.
9. Coordinate the backup and restore processes using the **Job Control Options** pane:

If you want to...	Then do this...
Cancel the current backup operation in progress	Select **Cancel conflicting backup that is in progress**.
Restore after the current backup operation is complete	Select **Wait for running backup to complete**.
End the restore operation	Select **Abort restore if conflicting operation is running**.

10. If the restore server does not have access to the archived backup storage, click **Advanced options**, and then under **Archived Backup Access**, select **Restore server does not have access to the archived backup storage**.
11. If you are performing a test restore operation, click **Test Restore**.
    a) To check the current logs (for up-to-the-minute restore operations only) and run database verification, in the **Select Test Restore Options** window, select **Check current Logs**.
    b) To verify databases and transaction logs on available destination volumes, select **Run Verification on Computer**.

SnapManager displays the Restore Status window, showing the tasks that are performed for the restore process.

12. Click **Start Now**.

Restoring data to a specified Frequent Recovery Point

You can use the Restore window to restore a point-in-time Frequent Recovery Point backup that is created after a full or copy backup operation using the Restore window. SnapManager combines the restore operation of a full backup or copy backup and the required transaction logs to restore to the selected recovery point.

Before you begin

You must know the Frequent Recovery Points of the backup copies that you want to restore.

No backup operations can be running.

About this task

Do not perform a restore operation while a backup operation is in progress. If you cancel a current backup operation in progress, SnapManager pauses all the active scheduled backup jobs on the Exchange server, or on all nodes in the cluster environment, and cancels the current backup copy before performing the restore operation. On completing the restore operation, SnapManager reenables the paused scheduled backup jobs. All the other inactive jobs are not changed.

If you want to restore after the current backup operation completes, SnapManager pauses all the active scheduled backup jobs on the current Exchange server (on all nodes in the cluster environment) but waits for the completion of the current backup operation before performing the restore operation. On completing the restore operation, SnapManager reenables the paused scheduled backup jobs. All the other inactive jobs are not changed.

After an actual restore process, you can optionally perform a backup and verification operation to verify that your restored database is free of physical-level corruption.

Steps

1. In the Scope pane, click **Restore**.
2. Select the storage groups that you want to restore, and then double-click the backup sets that you want to restore.
3. In the **Actions** pane, select **Restore**.
4. In the **Choose databases to restore** pane, select the databases that you want to restore.
5. Under **Type of Restore**, select **Point in time**.
6. Click **Select Recovery Point**.
7. In the Scope for **Recovery Point Selection** list, select the date, start time, and end time.

8. Click **Refresh**.
9. Select a recovery point from the **Available Recovery Points** list.

   SnapManager restores the transaction logs up to the selected recovery point. These logs are later played forward by Exchange during recovery.
10. Click **OK**.
11. To coordinate the backup and restore processes, click **Advanced Options**.

If you want to...	Then do this...
Cancel the current backup operation in progress	Select **Cancel conflicting backup that is in progress**.
Restore after the current backup operation is complete	Select **Wait for running backup to complete**.
End the restore operation	Select **Abort restore if conflicting operation is running**.

12. In the **Operation Options** pane, select one of the following to configure exhaustive verification and mount settings:

If you want...	Then do this...
To perform exhaustive verification of the transaction log sequence and database metadata before the restore process	Select **Verify Transaction Log Sequence and Database Metadata Before Restore and Exhaustive Verification**.
SnapManager to automatically remount your restored databases immediately after the restore operation	Select **Recover and mount database after restore**.

13. If you are performing a test restore operation, click **Test Restore**.
    a) To check the current logs (for up-to-the-minute restore operations only) and run database verification, in the **Select Test Restore Options** window, select **Check current Logs**.
    b) To verify databases and transaction logs on available destination volumes, select **Run Verification on Computer**.

    SnapManager displays the Restore Status window, showing the tasks that are performed for the restore process.
14. Click **Start Now**.

Recovery of a database to a Frequent Recovery Point

When you restore a full or copy backup, you can specify the recovery point from the subsequent Frequent Recovery Point (FRP) backups. SnapManager performs a point-in-time restore operation and recovers the database up to the specified recovery point.

Restoring Exchange 2010 databases in a DAG

You can restore Exchange Server 2010 databases in the Database Availability Group (DAG) using SnapManager Restore operation.

You can restore a database in a DAG from the following restore sources:

- SnapManager backup copies created on the same Exchange server
- Backup sets created on multiple Exchange servers in the DAG
- Unmanaged media

You can restore data to a specified Frequent Recovery Point (FRP).

You can also restore a database to the Recovery Database.

Backups available for restore in a DAG

When the SnapManager MMC snap-in connects to a Database Availability Group (DAG), you get a consolidated view of all backups in the DAG. You can then select a backup copy to restore.

Backups are created on the server where the database is located. When the SnapManager MMC snap-in connects to a member server of the DAG, you can view backups specific to that server. For each database that is backed up, its backups are grouped under two headings: **Local Backups** and **Archived Backups**. Because the backup name includes the name of the Exchange server, it indicates where the backup is located.

Note: The database restore takes place on the server where the backup is located.

Restoring a backup copy from one Exchange server to another in a DAG

You can use the SnapManager Restore wizard to restore to one server an Exchange 2010 mailbox database from a backup created on a different server in the Database Availability Group (DAG).

Before you begin

If a Exchange Server 2010 (mailbox server) was down due to physical server crash, disconnect the LUNs, remap the LUNs to different Exchange Server 2010 in the DAG, then run the restore operation from Restore Wizard to recover the database.

Note: If a member server is removed from a DAG, all its backups will still be valid for restore, even it is added to another DAG.

About this task

If the backup set is not verified, SnapManager gives you the option to verify the backup set before the restore operation.

Note: The option **Restore from archive** is renamed to **Restore from Unmanaged Media** for SnapManager 5.0 and later.

After an actual restore process, you can optionally perform a backup and verification operation to verify that your restored database is free of physical-level corruption.

Steps

1. In the **Scope** pane, select an Exchange server node.
2. In the **Actions** pane, click **Restore wizard**.
3. In the **Which Exchange Server Created the Backups** window, select the **Restore backup created on a different server** option.
4. In the **Exchange Server Where Backups were Created** window, specify the **Exchange server where the databases were backed up** and the **SnapInfo Directory Path**.
5. Select a backup set that you want to restore.
6. In the **Where to restore** panel, select a restore option.
7. Follow the instructions in the Restore wizard, and go to **Mount Options** page.
8. To recover and mount database after the restore process, select **Recover and mount database after restore**.
9. To associate all mailboxes in the newly restored databases to the currently connected Exchange server, select **Update user accounts associated with mailboxes in restored databases to point to mailbox server with the new name**.

   Note: If you restore the backup created on a different server, and select the **Update user accounts associated with mailboxes in restored databases to point to mailbox server with the new name** check box, SnapManager uses the Exchange cmdlet `Move-Mailbox -ConfigureOnly` internally to connect all the mailboxes in the newly restored database to the currently connected Exchange server.

Reseed after a restore operation in a DAG

After a restore operation, reseeding starts automatically.

Reseeding a database using a PowerShell script

The standard Exchange reseed function can be slow with large databases. SnapManager normally uses a more rapid method to reseed restored databases, but if that fails, or if for some other reason you want to reseed a restored database, you can use a PowerShell script. The script requires a healthy database copy, a Snapshot copy, or a SnapManager for Exchange backup of the database.

If you use a script to perform the following steps to restore the database first, when Exchange reseeds the database, much less data is copied and so it completes sooner. A summary of the steps is as follows:

1. Suspend the failed database copy.
2. Remove the failed database copy.
3. Add the database copy with seeding postponed.
4. Restore the database LUN from a recent Snapshot copy, or use SnapManager to restore the database from a SnapManager backup.
5. Resume the database copy.

These steps are most easily performed using a script that leverages Windows OS, Exchange PowerShell, and SnapDrive commands. A template you can customize for your environment is as follows:

```
cls
#This script requires that the SnapManager for Exchange PS.snapin
#and Exchange powershell cmdlets are loaded prior in the session
#prior to use

#Modify your powershell profile to add them with the following:

#Add-PSSnapin NetApp.SnapManager.Exchange.PS.Admin

#$Session = New-PSSession -ConfigurationName Microsoft.Exchange
#-ConnectionUri http://mst-exch-02.mst.local/PowerShell/
#-Authentication Kerberos

#import-pssession $session |out-null


#remove the broken one
Write-Host "Suspend-MailboxDatabaseCopy" -foreground "magenta"
Suspend-mailboxdatabasecopy –Identity “Exch_01\MST-EXCH-03” -Confirm:
$False

Write-Host "Remove-MailboxDatabaseCopy" -foreground "magenta"
Remove-mailboxdatabasecopy –Identity "Exch_01\MST-EXCH-03" -Confirm:$False

Write-Host "The database replica on MST-EXCH-03 is now removed" -
foreground "magenta"
Get-MailboxDatabaseCopyStatus -Identity "Exch_01"

#‘add the new one, but don’t start seeding yet
Write-Host "We now re-add the database on MST-EXCH-03 but with seeding
postponed" -foreground "magenta"
Add-mailboxdatabasecopy –Identity “Exch_01” –Mailboxserver MST-EXCH-03 –
seedingPostponed -Confirm:$False

Write-Host "Get-MailboxDatabaseCopyStatus" -foreground "magenta"
Get-MailboxDatabaseCopyStatus -Identity "Exch_01"

Write-Host ""
Write-Host "It's now Failed and Suspended, as expected. We'll now leverage
NetApp Snapshot Copies to quickly restore the LUNs and then perform
an incremental reseed"

Write-Host "Press a Key to Continue" -foreground "magenta"
```

```
$x = $host.UI.RawUI.ReadKey("NoEcho,IncludeKeyDown")

$backupArray = get-backup -server mst-exch-03 -storagegroup Exch_01 | sort
Backup -descending
$targetSet = $backupArray[0].Backup.ToLower()
Write-Host "Determine lastest backup set with SnapManager PowerShell..."
write-host Most Recent Backup Set : $targetSet

#‘restore the last snapshot of the logs and database with snapdrive
# restore_volume = restore : restore = lun_clone_split
Write-Host "Restore the Database LUN" -foreground "magenta"
sdcli snap restore -m mst-exch-03 -d E -s $targetSet
Write-Host "Restore the Transaction Log LUN" -foreground "magenta"
sdcli snap restore -m mst-exch-03 -d F -s $targetSet

Write-Host "Press a Key to Continue" -foreground "magenta"
$x = $host.UI.RawUI.ReadKey("NoEcho,IncludeKeyDown")

#‘resume replication to play it up to current and get healthy
Write-Host "Resume-MailboxDatabaseCopy" -foreground "magenta"
Resume-mailboxdatabasecopy –Identity “Exch_01\MST-EXCH-03” -Confirm:$False

Write-Host "Get-MailboxDatabaseCopyStatus" -foreground "magenta"
Get-MailboxDatabaseCopyStatus -Identity "Exch_01"

Write-Host "Wait 5 seconds for incremental log reseed, then press any key"
-foreground "magenta"
$x = $host.UI.RawUI.ReadKey("NoEcho,IncludeKeyDown")
Get-MailboxDatabaseCopyStatus -Identity "Exch_01"

Write-Host "Wait 5 seconds replay queue to drain, then press any key" -
foreground "magenta"
$x = $host.UI.RawUI.ReadKey("NoEcho,IncludeKeyDown")
Get-MailboxDatabaseCopyStatus -Identity "Exch_01"

Write-Host "Rapid Reseed complete" -foreground "magenta"

#‘now it’s fixed
#
```

The script runs on the node with the failed database.

In the unlikely event that the reseeding fails, check the Application Event Log on the server where the database copy is being restored. The event log shows which database logs are needed to complete the reseed. You could also run `eseutil /mh` on the restored database to see which log files are needed. Use `xcopy` to copy the missing logs from the Snapinfo directory from a good copy of the database to the transaction log directory for the database copy being reseeded.

Restore from an LCR-enabled Storage Group

You can restore Exchange databases from a backup set created from an LCR-enabled storage group.

Restoring a backup copy from an LCR replica Storage Group

You can use the Restore wizard to restore Exchange databases from a backup set that was created from either an LCR location or in its original location using the Restore wizard.

Before you begin

- Ensure that you close all Windows Explorer windows on the Exchange server running SnapManager.
- Disable any SnapManager operations that are scheduled to run against the Exchange data you are restoring, including any jobs scheduled on remote management or remote verification servers.

About this task

SnapManager uses a special naming convention to indicate the backup sets created in an LCR location. For example, the backup copy created from the LCR copy in Storage Group1 is represented as `Storage Group1(LCR)`.

If you select a backup copy from `Storage Group1(LCR)`, SnapManager uses the backup copy created on the LCR replica. If the backup copy is not verified, SnapManager displays an option in the message box to verify the backup copies before restore operation.

All databases that reside in the same LUN must be restored together.

A database restore operation from an LCR replica Storage Group replaces the database and log file paths but retains the SnapInfo directory path of original Storage Group.

Steps

1. In the Scope pane, click **Restore**.
2. In the **Actions** pane, click **Restore Wizard**.
3. In the **Which Exchange Server Created the Backups** window, select **Restore SnapManager backups that were created on the same Exchange Server**.
4. In the **Choose the backup that you want to restore** window, select the Storage Group tree and the Storage Group that you want to restore.
5. Double-click the backup set under the Storage Group that you want to restore.
6. In the **Choose the target for back up restoration** window, select the location to which you want to restore the SnapManager backup copy.
7. In the **Choose the items to be restored** window, select the check boxes corresponding to the production databases and the LCR replica databases that you want to restore.
8. In the **Type of Restore** window, select the type of restore operation.
9. If you are performing a test restore operation, click **Yes, this is a test restore**.

If you choose to perform an actual restore operation, select **No, this is an actual restore**.

10. Configure the verification settings in the **Verify options** window:
    a) To check the current logs (for up-to-the-minute restore operations only) and run database verification, in the **Select Test Restore Options** window, select **Check current Logs**.
    b) If you do not want to perform exhaustive verification of the transaction log sequence and database metadata before the restore process, clear the **Exhaustive verification** check box.
    c) If the restore server does not have access to the archived backup storage, click **Advanced options**, and then under Archived Backup Access, select **Restore server does not have access to the archived backup storage**.
    d) To coordinate the backup and restore processes, do the following in the Job Control Options pane:

If you want to...	Then do this...
Cancel the current backup operation in progress	Select **Cancel conflicting backup that is in progress**.
Restore after the current backup operation is complete	Select **Wait for running backup to complete**.
End the restore operation	Select **Abort restore if conflicting operation is running**.

11. For an actual restore process, if you want SnapManager to automatically recover and remount your restored databases immediately after the restore process, select **Mount Database Automatically After Restore** in the **Mount Options** window.
12. Verify the restore settings and then click **Finish**.
13. To start the restore process, click **Start Now**.

What to do if corruption occurs in an LCR-enabled Storage Group

When corruption occurs in an LCR enabled Storage Group, see Microsoft Exchange server documentation and use a Microsoft recommended method to recover a Storage Group. If an LCR recovery fails, use a backup set to restore the Storage Group.

SnapManager restores the database from a backup set that was created from another LCR location to its own LCR location. After the restore operation, SnapManager changes the location of the production database to the previous LCR location and disables the LCR process.

Restore from a CCR-enabled Storage Group

You can restore Exchange databases from a backup set created from a CCR location.

CCR backup set behavior

A CCR backup set has a certain pattern of behavior regarding on which node it resides and the port sharing service when you restore it in a CCR configuration.

- A backup set always resides on the node on which it is created.
- When you select the backup set from either an active or a passive node, SnapManager performs the restore operation on the node on which backup set was created.
- Before the restore operation, SnapManager moves the virtual server cluster resource group to the node on which the backup set is located.
- You need the Net.Tcp Port Sharing Service for remote server-to-server communication.

  By default, it is enabled when you install SnapManager. Verify that the Net.Tcp Port Sharing Service is enabled before you create a CCR backup set.

CCR restore rules for reseeding

The processes of up-to-the-minute restore operation and point-in-time restore operation of the primary database and CCR replicas have different reseeding requirements.

Type of restore operation	Type of database	Reseeding requirement
Up-to-the-minute restore operation	Primary database	Not required
Point-in-time restore operation	Primary database	Required
Up-to-the-minute restore operation	CCR replica database	Not required
Point-in-time restore operation	CCR replica database	Required

Point-in-time or up-to-the-minute restore operation of a CCR replica makes the entire cluster group fail over to the passive node.

Restoring databases from a CCR location

You can use the SnapManager Restore wizard to restore all the Exchange databases from a backup set created in either the CCR location or its original location.

Before you begin

- Ensure that all Windows Explorer windows are closed on the Exchange server running SnapManager.
- Disable any SnapManager operations that are scheduled to run against the Exchange data you are restoring, including any jobs scheduled on remote management or remote verification servers.
- Ensure that you install SnapManager on both nodes of a CCR configuration before you start the restore operation.

About this task

Unlike the LCR database replication copies, the SnapManager backup sets created from a CCR location do not follow any special naming convention. All backup sets created on both active and passive nodes appear in the Scope pane, below the Storage Group tree.

All databases that reside in the same LUN must be restored together.

You can choose to restore from the following restore sources:

- SnapManager backup copies that were created on the same Exchange server
- Unmanaged Media
- Backup sets created on a different server (available only with Exchange Server 2007)

If the backup set is not verified, SnapManager gives you an option to verify the backup set before the restore operation.

> **Note:** The option **Restore from archive** is renamed to **Restore from Unmanaged Media** for SnapManager 5.0 and later.

After an actual restore process, you can optionally perform a backup and verification operation to verify that your restored database is free of physical-level corruption.

Steps

1. Select an Exchange server node in the Scope pane.
2. In the **Actions** pane, click **Restore wizard**.
3. In the **Which Exchange Server Created the Backup** window, select the restore source.
4. In the **Choose the backup that you want to restore** window, double-click the backup set present in a Storage Group.
5. In the **Choose the target for backup restoration** window, select **Restore back to the same Storage Group/Database**.
6. In the **Choose the items to be restored** window, select the check boxes corresponding to the production databases and the replica databases that you want to restore.
7. In the **Type of Restore** window, select the type of restore operation.
8. If you are performing a test restore operation, click **Yes, this is a test restore**.

   If you choose to perform an actual restore operation, select **No, this is an actual restore**.
9. Configure the verification settings in the **Verify options** window:
   a) To check the current logs (for up-to-the-minute restore operations only) and run database verification, in the **Select Test Restore Options** window, select **Check current Logs**.
   b) If you do not want to perform exhaustive verification of the transaction log sequence and database metadata before the restore process, clear the **Exhaustive verification** check box.

c) If the restore server does not have access to the archived backup storage, click **Advanced options**, and then under Archived Backup Access, select **Restore server does not have access to the archived backup storage**.

d) To coordinate the backup and restore processes, do the following in the Job Control Options pane:

If you want to...	Then do this...
Cancel the current backup operation in progress	Select **Cancel conflicting backup that is in progress**.
Restore after the current backup operation is complete	Select **Wait for running backup to complete**.
End the restore operation	Select **Abort restore if conflicting operation is running**.

10. For an actual restore process, if you want SnapManager to automatically recover and remount your restored databases immediately after the restore process, select **Mount Database Automatically After Restore** in the **Mount Options** window.

11. Verify the restore settings and then click **Finish**.

12. To start the restore process, click **Start Now**.

What to do if corruption occurs in a CCR-enabled Storage Group

When corruption occurs in a Storage Group that is CCR-enabled, see Microsoft Exchange Server documentation and use a Microsoft recommended method to recover a Storage Group. Use a backup set located on CCR active node or passive node to restore the Storage Group.

If you use the backup set located on a CCR active node, the normal recovery procedure is applicable.

If you use the backup set located on a CCR passive node, SnapManager performs the following tasks:

1. Manually fails over to the CCR passive node on which the backup set is located

The Exchange store writer becomes available on the previous passive node. SnapManager runs the `Restore-StorageGroupCopy` cmdlet to make an extra copy of the log file on the previous passive node before restore.

2. Performs a normal restore operation

When the server comes online, the previous active node becomes a CCR passive node. The previous production database becomes the CCR copy location. To replicate, Reseeding the databases (which were the primary databases before the restore operation) again can enable replication of the data.

Recovery Storage Groups

A Recovery Storage Group allows you to restore databases, to extract mailbox data, and to merge that data to the original Storage Group on the production server. You can mount a second copy of an Exchange mailbox store on a production Exchange server when the original store is still running.

A Recovery Storage Group is useful in the disaster recovery process. SnapManager 5.0 and later for Exchange automates the mailbox store recovery process. In a normal recovery process, you need to manually create the Recovery Storage Group. Then you add the Exchange mailbox store through Exchange System Manager (ESM) on Exchange Server 2003 and management shell cmdlets on Exchange Server 2007. You can then restore the mailbox from the backup copies created and stored in the Recovery Storage Group.

You can restore databases to a Recovery Storage Group by using the `restorebackup` cmdlet with the appropriate parameter in the SnapManager command-line interface.

Limitations of using a Recovery Storage Group

Although you can benefit from using Recovery Storage Groups in SnapManager, there are some limitations with the restore source, the type of restore operation, configuration of the restore operation, and the creation of a Recovery Storage Group.

- You cannot restore a public folder store to the Recovery Storage Group.
- You cannot perform up-to-the-minute database recovery to a Recovery Storage Group.
- You cannot restore databases to a Recovery Storage Group from the backups that are created by other backup utilities.
- You cannot change the configuration settings of a Recovery Storage Group, such as renaming it and changing the database paths.
- You cannot use SnapManager to create and destroy a Recovery Storage Group on Exchange Server 2003.

  You need to manually create and destroy a Recovery Storage Group.
- You cannot use SnapManager to add, mount, and unmount databases in a Recovery Storage Group on Exchange Server 2003.

  You need to add, mount, and unmount databases manually.

Restoring a database to a Recovery Storage Group in Exchange Server 2007

You can use the Restore wizard to restore databases to the Recovery Storage Group. You create a Recovery Storage Group using the Exchange Server Disaster Recovery Analyzer Tool (ExDRA) in

Exchange Server 2007. SnapManager performs only a point-in-time restore process to the Recovery Storage Group during a restore operation.

About this task

By default, the Delete the existing Storage Group check box is selected during the restore process, resulting in the deletion of the group as part of the restore operation. SnapManager does not allow you to restore to an existing Recovery Storage Group.

To restore a backup set to a Recovery Storage Group in an Exchange cluster, connect the mounted LUNs that store the restored databases, transaction logs, and Storage Group system files only to the node that owns the Exchange virtual server.

Add these mounted LUNs to the Exchange cluster group as physical disk resources and keep them online at the end of restore process. These physical disk resources do not have the **Affect the group** check box enabled because these LUNs are mapped to only one of the cluster nodes. If you move the Exchange cluster group to another node, these physical disk resources fail, because they are not connected to the other nodes. After you return the Exchange cluster group to the node on which these resources are connected, you can bring them online.

SnapManager does not display the public folder database in the Restore wizard. You cannot restore public folder databases to the Recovery Storage Group.

As part of the restore to the Recovery Storage Group operation, SnapManager performs the following operations:

1. Destroys the existing Recovery Storage Group
2. Creates a new Recovery Storage Group
3. Adds databases to the new Recovery Storage Group
4. Mounts databases on the new Recovery Storage Group

If the backup set is not verified, SnapManager gives you an option to verify the backup set before the restore operation.

You can restore from the following restore sources:

- Backup sets that were created on the same Exchange server
- Unmanaged Media
- Backup sets created on a different server

After an actual restore process, you can optionally perform a backup and verification operation to verify that your restored database is free of physical-level corruption.

Steps

1. Select an Exchange server node in the Scope pane.
2. In the **Actions** pane, click **Restore Wizard**.

3. In the **Which Exchange Server Created the Backups** window, select the restore source.
4. In the **Choose the backup that you want to restore** window, double-click the backup set present in a Storage Group.
5. In the **Select the target for backup restoration** window, select **Restore to Recovery Storage Group/Database**.
6. In the **Choose the items to be restored** window, select the check boxes corresponding to the production databases and LCR replica databases that you want to restore.
7. In the **Choose Recovery Storage Group/Database** window, select the destination Exchange server to which you want to restore the backup set.
8. Specify the name of the new Recovery Storage Group to which you want to restore the selected backup sets.
9. In the **Type of Restore** window, select a Recovery point for your database.
10. Configure the verification settings in the **Verify options** window:
    a) To check the current logs (for up-to-the-minute restore operations only) and run database verification, in the **Select Test Restore Options** window, select **Check current Logs**.
    b) If you do not want to perform exhaustive verification of the transaction log sequence and database metadata before the restore process, clear the **Exhaustive verification** checkbox.
    c) If the restore server does not have access to the archived backup storage, click **Advanced options**, and then under Archived Backup Access, select **Restore server does not have access to the archived backup storage**.
    d) To coordinate the backup and restore processes, do the following in the Job Control Options pane:

If you want to...	Then do this...
Cancel the current backup operation in progress	Select **Cancel conflicting backup that is in progress**.
Restore after the current backup operation is complete	Select **Wait for running backup to complete**.
End the restore operation	Select **Abort restore if conflicting operation is running**.

11. For an actual restore process, if you want SnapManager to automatically recover and remount your restored databases immediately after the restore process, select **Mount Database Automatically After Restore** in the **Mount Options** window.
12. Verify the restore settings and then click **Finish**.
13. To start the restore process, click **Start Now**.

How Recovery Storage Groups are created in Exchange Server 2007

A Recovery Storage Group is created as part of the SnapManager restore operation. By default, SnapManager provides this group with a default Recovery Storage Group name that has the same name as the original Storage Group name, with `-RSG` appended to it.

For example, if the original Storage Group name is SG2, the Recovery Storage Group name is `SG2-RSG`. You can change the default Recovery Storage Group name either in the user interface or by specifying the name in the command-line interface.

Adding databases to a Recovery Storage Group in Exchange Server 2007

You can add Exchange 2007 databases to the Recovery Storage Group by using the Restore wizard to add databases to the Recovery Storage Group during the SnapManager restore operation.

Steps

1. Select an Exchange server node in the Scope pane.
2. In the **Actions** pane, click **Restore wizard**.
3. In the **Which Exchange Server Created the Backup** window, select any of the following options:
   - **Restore SnapManager backups that were created on the same Exchange server**
   - **Restore from Unmanaged Media**
   - **Restore backup created on a different server**
4. In the **Choose the backup that you want to restore** window, double-click the backup set present in a Storage Group.

   If the backup is not verified, SnapManager gives you an option to verify the backup set before the restore operation.
5. In the **Select the target for backup restoration** window, select **Restore to Recovery Storage Group**.
6. In the **Choose the items to be restored** window, select the check boxes corresponding to the production databases and LCR replica databases that you want to restore.
7. Proceed with the steps as directed by the Restore wizard.

Mounting databases on a Recovery Storage Group in Exchange Server 2007

When the Exchange Information Store service restarts, the databases in the Recovery Storage Group are not mounted automatically. Use the Restore wizard to mount databases on a Recovery Storage Group.

Steps

1. Select an Exchange server node in the Scope pane.
2. In the **Actions** pane, click **Restore wizard**.
3. In the **Which Exchange Server Created the Backup** window, select any of the following options:
   - **Restore SnapManager backups that were created on the same Exchange server**
   - **Restore from Unmanaged Media**
   - **Restore backup created on a different server**
4. In the **Choose the backup that you want to restore** window, double-click the backup set present in a Storage Group.

   If the backup set is not verified, SnapManager gives you an option to verify the backup set before the restore operation.
5. In the **Select the target for backup restoration** window, select **Restore to Recovery Storage Group**.
6. In the **Choose the items to be restored** window, select the check boxes corresponding to the production databases and LCR replica databases that you want to restore.
7. In the **Choose Recovery Storage Group** window, select the destination Exchange server to which you want to restore the backup set.
8. Specify the name of the new Recovery Storage Group to which you want to restore the selected backup sets.
9. In the **Type of Restore** window, select a Recovery point for your database.
10. Configure the verification settings in the **Verify options** window:
    a) To check the current logs (for up-to-the-minute restore operations only) and run database verification, in the **Select Test Restore Options** window, select **Check current Logs**.
    b) If you do not want to perform exhaustive verification of the transaction log sequence and database metadata before the restore process, clear the **Exhaustive verification** check box.
    c) If the restore server does not have access to the archived backup storage, click **Advanced options**, and then under Archived Backup Access, select **Restore server does not have access to the archived backup storage**.
    d) To coordinate the backup and restore processes, do the following in the Job Control Options pane:

If you want to...	Then do this...
Cancel the current backup operation in progress	Select **Cancel conflicting backup that is in progress**.
Restore after the current backup operation is complete	Select **Wait for running backup to complete**.
End the restore operation	Select **Abort restore if conflicting operation is running**.

11. For an actual restore process, if you want SnapManager to automatically recover and remount your restored databases immediately after the restore process, select **Mount Database Automatically After Restore** in the **Mount Options** window.
12. Proceed with the steps as directed by the Restore wizard.

Destroying a Recovery Storage Group in Exchange Server 2007

With SnapManager 5.0 and later, you can destroy the Recovery Storage Group, since you cannot restore to an existing Recovery Storage Group.

About this task

SnapManager performs the following actions when you destroy the Recovery Storage Group.

1. Unmounts the Recovery Storage Group from Exchange Server 2007
2. Disconnects all of the database LUNs in the Snapshot copy
3. Deletes the Recovery Storage Group from the target Exchange Server 2007
4. If your configuration is an MSCS cluster, removes the physical disk resource from the Exchange cluster group

Steps

1. In the Scope pane, click **Backup**.
2. In the **Backup** window, select the Recovery Storage Group that you want to delete.
3. In the **Actions** pane, select **Delete Storage Group**.

Restoring a database to a Recovery Storage Group in Exchange Server 2003

Using Recovery Storage Groups, you can mount a copy of a mailbox store onto a production server and recover data within the restored mailbox store while the current store is running. You must

manually create the Recovery Storage Group and add a mailbox store by using Exchange System Manager (ESM).

About this task

When you run the SnapManager restore operation, SnapManager does the following:

1. SnapManager mounts the databases and transaction log LUNs onto the Snapshot copy that is created by the selected backup.

   Do not move the Exchange group. If you move the Exchange group to another node, the newly added physical disk resources fail.

2. SnapManager performs a soft recovery on the databases during the restore operation.

   SnapManager tracks the LUN information for unmounting. If your configuration is a cluster configuration, SnapManager adds mounted LUNs in a Snapshot copy to the Exchange cluster group as a physical disk resource.

   After the soft recovery is complete, the databases are in a Clean Shutdown state.

Steps

1. Start Exchange System Manager (ESM).
2. Choose the option to create a recovery storage group.
3. Enter the transaction log path by clicking **Browse** to locate and select the LUN that SnapManager connected to during the restore operation.
4. Copy the transaction log path to the System path field.
5. Click **OK** to save your changes.
6. Select the database you want to recover and then click **OK**.
7. Click the **Database** tab.
8. Enter the path to the database located on the LUN that SnapManager connected to during the restore operation.
9. Add a unique database file name to the path.

   The database file name can be any name except the existing database file name. The Exchange database file name must end in .edb: for example, `SG2-DB1-RSG.edb`.

10. Copy the Exchange database path and file name to the Exchange streaming database field and replace the .edb suffix of the database file name with .stm.
11. Click **OK**.
12. Start Windows Explorer and navigate to the location of the database LUN that SnapManager connected to during the restore.
13. Rename the existing database files to the unique database file names you just created.

14. Mount the databases to the Recovery Storage Group.
15. To restore more than one database to the same Storage Group, repeat Steps 1 through 15.
16. To restore the database from a different Storage Group, delete all of the databases in the Recovery Storage Group and remove the Recovery Storage Group through ESM.

Creating a Recovery Storage Group in Exchange Server 2003

You can restore databases to the Recovery Storage Group using the Exchange System Manager in Exchange Server 2003. Before that, you must manually create the Recovery Storage Group.

Before you begin

Delete the existing Recovery Storage Group and disconnect all the database and log LUNs of the existing Recovery Storage Group before you restore a backup set to a new Recovery Storage Group. Run the **Restore to Recovery Storage Group** option in the SnapManager Configuration wizard before adding the databases to the Recovery Storage Group.

Steps

1. In Exchange System Manager, right-click the server on which you intend to place the Recovery Storage Group, click **New**, and then click **Recovery Storage Group**.
2. Name the Recovery Storage Group.
3. Specify the transaction log location and the system path location.
4. Click **OK**.

Adding databases to a Recovery Storage Group in Exchange Server 2003

You can use the Restore wizard to add database to a Recovery Storage Group to restore the database. Exchange automatically determines which databases can be added to the recovery storage group and presents you with a list from which to choose.

About this task

If the Recovery Storage Group already contains a database, Exchange limits the list of databases to those in the same Storage Group as the database that you already added to the Recovery Storage Group.

Steps

1. In Exchange System Manager, find the server on which you created the Recovery Storage Group, right-click the Recovery Storage Group, and then click **Add Database to Recover**.
2. In the **Select database to recover** dialog box, click the database that you want to recover and then click **OK**.

3. Type the path to the database located at the LUN mounted on the Snapshot copy created by the selected backup set.

   The database file name can be any name (for example, `SG2-DB1-RSG.edb`) except the database file name (for example, `SG2-DB1.edb`) that exists on that location.

4. Name the database, and define the paths for the database and streaming database files.
5. Click **OK**.
6. Rename the existing database to the specified name (for example, `SG2-DB1-RSG.edb`).

Disconnecting LUNs of a Recovery Storage Group in Exchange Server 2003

You can temporarily disconnect the Recovery Storage Group LUNs using SnapManager. In Exchange Server 2003, SnapManager can only disconnect the LUNs; you need to use the Exchange System Manager (ESM) to delete the Recovery Storage Group.

Steps

1. In the Scope pane, click **Backup**.
2. Select the Recovery Storage Group LUN that you want to disconnect.
3. In the **Actions** pane, select **Disconnect LUNs**.

   SnapManager disconnects all the LUNs in the Snapshot copy that belong to that Recovery Storage Group.

Restoring multiple databases to a Recovery Storage Group

You can use SnapManager to restore multiple databases to a Recovery Storage Group. This option saves you from performing repeated restore operations.

About this task

You cannot add more than one database from the same Storage Group to the existing Recovery Storage Group. If you attempt to perform this operation in an Exchange Server 2007 environment, after you select all of the databases that you want to recover, SnapManager deletes the existing Recovery Storage Group and creates a new Recovery Storage Group, into which it then adds all the databases that you chose to recover. In an Exchange Server 2003 environment, SnapManager does not allow you to proceed if the Recovery Storage Group exists. You need to manually remove the Recovery Storage Group through ESM.

Steps

1. In the **Actions** pane, click **Restore wizard**.
2. In the **Choose the target for back up restoration** window, select the **Restore to the Recovery Storage Group** option.

3. In the **Choose the items to be restored**, select the check boxes corresponding to the production databases and the LCR replica databases that you want to restore.
4. Follow the instructions in the Restore wizard to complete the restoration of multiple databases to a Recovery Storage Group.

Restoring an unverified backup copy to a Recovery Storage Group

As a part of a normal restore process, SnapManager verifies the backup Snapshot copies before the actual restore operation. If you need to perform a quick restore operation of a backup copy, you can override the database verification requirement. In such cases, SnapManager restores the unverified backup copy directly to the Recovery Storage Group.

About this task

You should restore only from verified backup copies to ensure a successful restore operation.

Steps

1. In the **Actions** pane, select **Backup Verification Settings**.
2. In the **Database Verification Settings** dialog box, click the **Override Verification** tab.
3. Select the **Override database verification requirement for restore** check box.
4. Click **OK**.

Restore of backups created at different Exchange server locations

The SnapManager Restore wizard enables you to select the Exchange 2007 locations on which the databases were backed up, and then restore the databases to the current Exchange server. You can restore those backups to either the same Storage Group or to a Recovery Storage Group. To restore backups that were created on different servers, specify the details in the Exchange Servers where backups were created window during the restore operation.

Restoring backup copies that were created on other Exchange servers in Exchange 2007

You can restore SnapManager backup copies that were created on other Exchange servers to either the same Storage Group or the Recovery Storage Group. By default, SnapManager restores the backup sets to the Exchange server on which the restore operation is initiated.

Before you begin

To restore the backup copies from a different Exchange server to the Exchange server where the restore operation is initiated, you must first remap the source LUNs to the Exchange server by using

either the same drive letters, or volume mountpoints (both drive letters and volume mountpoints) that were assigned for the original Exchange server.

To use the restored backup set created on another Exchange server for a particular Storage Group, all database LUNs, transaction log LUNs, and SnapInfo LUNs must be remapped to the destination.

Steps

1. In the Scope pane, click **Restore**.
2. In the **Actions** pane, click **Restore wizard**.
3. Go through wizard to the **Which Exchange Server Created the Backups** page and select the **Restore backup created on a different server** option.
4. From the **Exchange Server where the databases were backed up** list, select the Exchange server to which the databases were backed up.
5. Under **Enter the SnapInfo directory path of the backups**, type the SnapInfo directory path.
6. Select a backup set that you want to restore.
7. In the **Where to restore** panel, select a restore option.
8. Follow the instructions in the Restore wizard, and go to **Mount Options** page.
9. To recover and mount database after the restore process, select **Recover and mount database after restore**.
10. To associate all mailboxes in the newly restored databases to the currently connected Exchange server, Select **Update user accounts associated with mailboxes in restored databases to point to mailbox server with the new name**.

    Note: If you restore the backup created on a different server, and select the **Update user accounts associated with mailboxes in restored databases to point to mailbox server with the new name** check box, SnapManager uses the Exchange cmdlet `Move-Mailbox -ConfigureOnly` internally to connect all the mailboxes in the newly restored database to the currently connected Exchange server.

Restoring backup sets from unmanaged media

You can restore backup sets from an archive created by third-party software to the same Storage Group or server or to the Recovery Storage Group or Recovery Database. Using the Recovery Storage Group or Recovery Database enables you to mount a mailbox store copy on a production server and recover data within the restored mailbox while the current mailbox store is still running.

Before you begin

While performing a restore operation from, restore the archived SnapInfo directory to a temporary location. This SnapInfo directory must be part of the same backup set as the restored LUN that

contains the Exchange databases. The restored SnapInfo directory must have the SnapInfo file and the `.xml` files.

Ensure that you have ready the SnapInfo directory path for the backup set that you want to restore.

About this task

When you are using the Restore wizard to configure a restore operation, you can decide whether to restore back to the same storage group/database, or to restore to a recovery storage group/database. When you are restoring to a recovery storage group or database, you must do the following before launching the Restore wizard:

1. Put the database and logs files restored from tape on NetApp LUNs. Use SnapDrive to manually create the LUN.
2. The number of temporary LUNs must be equal to the number of LUNs used by the original database and log LUNs.
3. The database and log file path on the temporary LUNs must be the same as the original database and log path.

Steps

1. In the SnapManager for Exchange console root, start the Restore wizard:
   a. In the Scope pane, click **Restore**.
   b. In the **Actions** pane, click **Restore wizard**.
2. In the Restore wizard **Start** screens:
   a. The **Welcome** screen displays. Click **Next**.
   b. Click **Restore from Unmanaged Media** and click **Next**.
   c. If **Use/verify transaction log files** is checked, uncheck it. This is a point-in-time restore, and transaction logs do not need to be verified for point in time restores.
   d. Use the **Backup Exchange Server** drop-down list to specify the name of the server on which the archived backup was originally created.
   e. Select or browse to the **SnapInfo Directory Path** and click **Next**.
3. In the Restore wizard **Choose Backup** screen, double-click the backup from which you want to restore.
4. In the Restore wizard **Restore Target** screen:
   a. Verify the server name.
   b. Verify the name of the backup from which you want to restore.
   c. Choose one of the following restore options:

   a. Restore back to the same Storage Group/Database
   b. Restore to the Recovery Storage Group/Database

5. In the Restore Wizard **Choose Items** screen click the check-boxes corresponding to the Mailbox databases or public folders you want to restore.

   Note: If you are restoring to a Recovery Storage Group or database, change the drive letter or mount point of **New EDB Path** and the **New Log Path** to correspond with the temporary LUN you have already created.

6. If you are restoring to the same Storage Group/Database:
   a. Specify whether you want to perform an up-to-the-minute restore operation or a point-in-time restore operation.
   b. Specify if the restore operation is a test restore operation.
7. If you are restoring to a Recovery Storage Group or database:
   a. Specify the Destination Server and the new Recovery Storage Group/Database name. Click **Next**.
   b. Set the recovery point time range, and select from the available recovery points.
8. In the **Verify Options** screen:
   a. In the **Operation Options** field, select **Verify transaction log sequence and database metadata before restore** to verify transaction log sequence and database metadata before the restore operation. You can also select **Exhaustive verification** to verify all signatures and metadata before the restore operation.
   b. In the **Job Control Options** field, choose the conflict management option you want.
   c. In the **Archived Backup Access** field, specify whether the restore server has access to the remote archived storage. Click **Next**.
9. In the **Mount Options** screen:
   a. To recover and mount the database after the restore process, select **Recover and mount database after restore**.
   b. To associate all mailboxes in the newly restored databases to the currently connected Exchange server, select **Update user accounts associated with mailboxes in restored databases to point to mailbox server with the new name**. Click **Next**.

      Note: If you restore the backup created on a different server, and select the **Update user accounts associated with mailboxes in restored databases to point to mailbox server with the new name** check box, SnapManager uses the Exchange cmdlet `Move-Mailbox -ConfigureOnly` internally to connect all the mailboxes in the newly restored database to the currently connected Exchange server.

10. In the **Completion** screen, review the information you have selected. When you have finished, click **Finish**.

Up-to-the-minute restore from an archive backup

As newer backups are made, older backups and logs for older backups get deleted. Because an up-to-the-minute restore operation is possible only if the logs are present, you can perform an up-to-the-minute restore operation from archive backup only if local log backup exists. Therefore, restoring from an archive backup is usually a point-in-time restore operation, and not an up-to-the-minute restore operation.

Recovery Database

A Recovery Database is a special type of mailbox database that allows you to restore, recover, and mount your database to an Exchange server, while the original database is still serving the user. It is a feature available with Exchange Server 2010. It is not bound to any server or database and each Exchange Server 2010 can have only one such Recovery Database.

A Recovery Database is useful in the disaster recovery process. SnapManager automates the mailbox store recovery process. You can add the Exchange mailbox store through management shell cmdlets on Exchange Server 2010, then you can restore the mailbox from the backup copy created and stored in the Recovery Database.

You can restore a database to a Recovery Database by using the restore-backup cmdlet with the appropriate parameter in the SnapManager command-line interface.

Limitations of using a Recovery Database

Although a Recovery Database allows you to restore, recover, and mount your database to an Exchange server to retrieve mail while the original database is still serving the user, it has some limitations in terms of the number of databases that you can restore, the type of restore operation that you can perform, and the configuration of the restore operation.

- You can use the Recovery Database for Exchange 2010 mailbox databases only.
- You can use the Recovery Database as a target for restore operations but not for backup operations.
- You can mount only one Recovery Database at any time on an Exchange Server 2010.
- You cannot use the Recovery Database to recover mailbox databases of Exchange Server 2003 or Exchange Server 2007.
- You cannot restore public folder data to the Recovery Database.

Restoring a mailbox database to the Recovery Database in Exchange Server 2010

You can use the Restore wizard to restore an Exchange 2010 database to the Recovery Database.

About this task

SnapManager does not display the public folder database in the Restore wizard. You cannot restore public folder database to a Recovery Database.

If the backup set is not verified, SnapManager gives you the option to verify the backup set before the restore operation.

You can restore from the following restore sources:

- Backup that was created on the same Exchange server
- Unmanaged media
- Backup that was created on a different Exchange server

After an actual restore process, you can optionally perform a verification operation to verify that your restored database is free of physical-level corruption.

Steps

1. Either select an Exchange server node in the **Scope** pane or connect to the DAG.
2. In the **Actions** pane, click **Restore Wizard**.
3. In the **Which Exchange Server Created the Backups** window, select the restore source.
4. In the **Choose the backup that you want to restore** window, double-click the backup set listed under the database.
5. In the **Select the target for backup restoration** window, select **Restore to the Recovery Storage Group/Database**.
6. In the **Choose the items to be restored** window, select the check boxes corresponding to the production databases that you want to restore.
7. In the **Choose Recovery Storage Group/Database** window, select the destination Exchange server to which you want to restore the backup set.

   Note: You cannot specify a Database Availability Group (DAG) name as a destination Exchange server.

8. Specify the name of the new Recovery Database to which you want to restore the selected backup set.
9. In the **Type of Restore** window, select a recovery point for your database.
10. Configure the verification settings in the **Verify options** window:

a) To check the current logs (for up-to-the-minute restore operations only) and run database verification, in the **Select Test Restore Options** window, select **Check current Logs**.
b) If you do not want to perform exhaustive verification of the transaction log sequence and database metadata before the restore process, clear the **Exhaustive verification** check box.
c) If the restore server does not have access to the archived backup storage, click **Advanced options**, and then under Archived Backup Access, select **Restore server does not have access to the archived backup storage**.
d) To coordinate the backup and restore processes, make the following selections in the Job Control Options pane:

If you want to...	Then do this...
Cancel the current backup operation in progress	Select **Cancel conflicting backup that is in progress**.
Restore after the current backup operation is complete	Select **Wait for running backup to complete**.
End the restore operation	Select **Abort restore if conflicting operation is running**.

11. For an actual restore process, if you want SnapManager to automatically recover and remount your restored databases immediately after the restore process, select **Mount Database Automatically After Restore** in the **Mount Options** window.
12. Verify the restore settings and then click **Finish**.
13. To start the restore process, click **Start Now**.

When to delete a Recovery Database

Because each Exchange Server 2010 can have only one Recovery Database mounted at any given time, after you recover the mailbox from the Recovery Database, you must delete the Recovery Database. You can then recover other mailbox databases.

You can delete the Recovery Database using the SnapManager MMC.

Troubleshooting restore operations

If a restore operation fails with an error message similar to `Error: 0xC00413CA, SnapManager detected the following Exchange writer error. Please retry SnapManager operation. VSS_E_WRITERERROR_RETRYABLE:` (and additional information), you might be able to clear the problem and try to restore again.

About this task

A restore.env log file is created during a restore operation. If you later restore the same data and the log file is still present, the operation fails.

Steps

1. Check the transaction log LUN for the database for a restore.env file: for example, `E01restore.env`.
2. If such a file is present, delete it and retry the restore operation.

Mailbox restore using Single Mailbox Recovery

You can use SnapManager for Microsoft Exchange to launch the Single Mailbox Recovery (SMBR) application to restore your mailbox data.

Single Mailbox Recovery works in conjunction with SnapManager for Microsoft Exchange. You can use SnapManager for Exchange to perform online backups of Exchange databases. When you need to restore previously deleted, individual Exchange mailbox items, you can launch Single Mailbox Recovery through SnapManager for Exchange to locate and then restore items at any level of granularity, directly to an existing mailbox on your Exchange server. For more information, see Single Mailbox Recovery User Guide.

Recovering mailbox data

When you need to restore mailbox items to your Exchange server, you can launch the Single Mailbox Recovery application via SnapManager for Microsoft Exchange to locate and restore these items.

Before you begin

Ensure that Single Mailbox Recovery version 5.0 or later is installed on the server that is running the SnapManager for Microsoft Exchange user interface, MMC Snap-in. Also, if you are using Protection Manager, ensure that both the source and destination Single Mailbox Recovery servers are configured for Protection Manager access. See the `sdcli dfm_config` command for help on how to configure SnapDrive to access a Protection Manager server.

Steps

1. Select the Exchange server node in the **Scope** pane, and then click **Restore**.

   SnapManager for Microsoft Exchange displays the Exchange server storage groups and databases.

2. Click the name of the backup that you want to restore.

   You can select both local backups and archive backups.

3. In the **Actions** pane, click **Run SMBR**.

   SnapManager for Microsoft Exchange displays the Start SMBR dialog box.

4. Provide the information required by the Single Mailbox Recovery software, such as the database name and optionally usage of the transaction logs. If **Usage of transaction logs** is checked, uncheck it. (The transaction logs require the software mount the log LUN, but because this is a point-in-time restore, the log LUN has already been deleted, and so the mount fails.)

For more information about using Single Mailbox Recovery software, see the Single Mailbox Recovery User Guide.

5. Click **Start SMBR** to launch the SMBR software and start the mailbox recovery.

After you finish

After the Single Mailbox Recovery application restores your mailbox items, you must unmount the snapshots you restored.

Cleaning up after Single Mailbox Recovery

In order to recover mailbox data, Single Mailbox Recovery needed to mount the mailbox backup copy that contained the appropriate data. After your mailbox data is recovered, you need to unmount the backup copy. Mounted backup copies use resources and cannot be deleted.

About this task

For Microsoft Windows 2007, unmounting the backup copies actually renames the backup copy of the resource group by appending the back up copy name with a timestamp.

Steps

1. Select the Exchange server node in the **Scope** pane, and then click **Restore**.

   SnapManager displays the Exchange server data store and the Storage Groups.

2. In the **Actions** pane, click **SMBR Cleanup**.
3. Select the LUNs that you want to unmount, and then click **Dismount**.

   To select all of the LUNs that were mounted during the restore process, click **Select All**.

4. Close the window when the progress bar indicates that the cleanup process is complete.

Deletion of Snapshot copies

You can use the Delete Backup dialog box to explicitly delete multiple Snapshot copies from one or more Storage Groups or databases. You can delete either entire backup sets, or only the SnapInfo Snapshot copies related to the selected backup sets, and you can delete Snapshot copies created during previous restore operations.

You can use the Delete Backup dialog box to delete backups when the SnapManager MMC snap-in connects to a member server of the Database Availability Group (DAG), but this box is not available when the snap-in connects to the DAG itself.

You should not use SnapDrive or storage system administration tools to delete Snapshot copies created by SnapManager. Doing so leaves behind unwanted data that cannot be removed.

Criteria for deleting backups

To avoid reaching the limit of 255 Snapshot copies per volume, you should delete the backups that you no longer need. You can delete backups by using either the age-based criteria or the quantity-based criteria.

Age-based Snapshot copy deletion	You can delete backups that are assigned to a particular backup management group and that are older than a specified number of days. You can specify the deletion criteria by selecting the “Delete backups older than” option. Enter the number of days for which you want to keep the most recent backups for the backup management group.
Quantity-based Snapshot copy deletion	You can delete backups that are assigned to a particular backup management group so that only a specified number of backups remain for that backup management group. Specify the deletion criteria by selecting the “Delete (oldest) backups in excess of” option. Enter the number of the most recent backups that you want to keep for the backup management group.

Note: Automatic deletion of older backups in a management group is the recommended method for managing the number of Snapshot copies stored on your system.

When SnapManager counts the number of stored backups, it also counts the backups that are shared by multiple Storage Groups. Therefore, more backups than you specify in the “Delete backups older than” or “Delete backups in excess of” box are retained.

For example, assume that you back up Storage Groups A and B. Both reside on the same volume and contain the following backup sets.

Backup set	Description
Storage Group A	
`exchsnap__orbit3_01-23-2006_16.21.07`	Old backup—good
`exchsnap__orbit3__recent`	Recent backup—good
Storage Group B	
`exchsnap__orbit3_01-23-2006_16.21.07`	Old backup—good
`exchsnap__orbit3__recent`	Recent backup—inconsistent

You set the “Delete backups in excess of” box to 2 to preserve the most recent two backup sets.

To preserve two good backups for Storage Group B, SnapManager does not delete the Snapshot copy `exchsnap__orbit3_01-23-2006_16.21.07`, which is old and good. Therefore, because both Storage Groups reside on the same volume, three backups for Storage Group A remain instead of two.

Automatic deletion of Snapshot copies

You can manage the number of Snapshot copies you store by configuring SnapManager to delete backups automatically, based on how old the backups are or based on how many of them are stored.

Automatic deletion deletes a backup only if the backup has the following characteristics:

- The backup is in the same management group as the management group of the backups that you just created.
- The backup is the oldest backup of the Storage Group
- The number of backups exceeds the backup retention level that you specified in the “Delete backups older than” option or the “Delete backups in excess of” option.

Note: If a storage group is no longer available in the Exchange server, then backups associated with it cannot be removed by the delete backup module.

If you do not select automatic backup deletion, backups that are created after the current backup are retained. This would require manual removal of backups, or enough storage capacity for all backups and transaction logs. You can delete the retained backups by selecting automatic backup deletion in the next backup that you take.

Explicit deletion of Snapshot copies

You can explicitly delete any backups or Snapshot copies of LUNs created during a restore operation by selecting them. You can delete either an individual backup or multiple backups.

When you explicitly delete backups, you can also ensure that the deletion is limited so as not to create a break in the continuity of the transaction logs between the previous backup and the present time. Doing so allows you to retain up-to-the-minute restore ability for older backups from other management groups.

> **Attention:** Do not use SnapDrive or the storage system administration tools to delete Snapshot copies created by SnapManager. Doing so leaves behind unwanted data that you cannot remove.

You can delete individual selected Snapshot copies for either full database backups or transaction logs. You can also select a database to be deleted, the types of backup set components to be deleted (full database backups or SnapInfo directory backups), and the type of backup management group to be deleted for multiple backups.

When you explicitly delete backup sets (or related SnapInfo Snapshot copies) that contain selected Storage Groups, you can expand the scope of the deletion to include backups that contain only some of the selected Storage Groups (in addition to backups that contain all the selected Storage Groups).

Option to retain up-to-the-minute restore ability

If you delete backups and transaction logs that are not the oldest backups in your backup list, the oldest backups are no longer available for up-to-the-minute restore. Ensure that you do not break the continuity of transaction logs between the previous backup and the present time.

When you delete backups of a particular backup management group, limiting the backup deletion allows you retain the ability to use the older backups in an up-to-the-minute restore operation.

To limit the backup deletion, use the "Retain up-to-the-minute restore ability for older backups in other backup management groups" option in the Advanced Options dialog box. Selecting this option consumes more space on your storage system for the transaction logs that are not deleted, and it is selected by default.

If all backups have the same management group designation, there is no effect if you clear the option. SnapManager deletes the SnapInfo directory transaction logs according to the Delete backups setting, and ignores the Retain up-to-the-minute restore ability option.

By default, SnapManager backs up selected database files and all associated transaction logs, so that up-to-the-minute restores are possible from all backups. If you do not need to perform an up-to-the-minute restore operation from the older backups, delete the transaction logs to free space on the storage system that contains the backups. The "Retain up-to-the-minute restore ability for older backups in other backup management groups" option is automatically selected the next time you start SnapManager.

As an example, assume that you have the following backups stored:

- exchsnap_WINSRV3_01_05_2004_19.05.37_Daily (Oldest)
- exchsnap_WINSRV3_01_06_2004_08.23.14
- exchsnap_WINSRV3_01_06_2004_10.22.19
- exchsnap_WINSRV3_01_06_2004_12.18.58
- exchsnap_WINSRV3_01_06_2004_14.28.03 (Newest)

If you perform a Standard backup, and specify that you want to retain only the four most recent Standard backups the oldest Standard backup is deleted, and the daily backup becomes point-in-time only:

- exchsnap_WINSRV3_01_05_2004_19.05.37_Daily (Oldest) (Daily backup becomes point-in-time only)
- exchsnap_WINSRV3_01_06_2004_08.23.14 (Oldest Standard backup deleted)
- exchsnap_WINSRV3_01_06_2004_10.22.19
- exchsnap_WINSRV3_01_06_2004_12.18.58
- exchsnap_WINSRV3_01_06_2004_14.28.03
- exchsnap_WINSRV3_01_06_2004_16.23.48 (Newest) (New Standard backup created)

To avoid breaking the continuity of transaction logs, you can enable the "Retain up-to-the-minute restore ability for older backups in other backup management groups" option.

Explicitly deleting individual backup copies

You can delete an individual backup copy.

About this task

If you select the **Retain up-to-the-minute restore ability for older backups in other backup management groups** option, transaction logs are deleted only from this backup management group; transaction logs are not deleted from other backup management groups. If you clear this option, transaction logs are deleted from other backup management groups as well.

Steps

1. In the Scope pane, Select **Restore**.
2. Select the Storage Group in which the backup copy resides.
3. Click **Delete** in the **Actions** pane.
4. In the **Delete Single Backup** window, select the backup copies that you want to delete.

5. Review the list of backup copies that share the Snapshot copy you want to delete, because all of these backup copies are deleted simultaneously.
6. If you want to retain up-to-the-minute restore ability for older backups in other backup management groups, select **Retain up-to-the-minute restore ability for older backups in other backup management groups**.
7. Click **OK**.

Explicitly deleting backup sets or SnapInfo Snapshot copies

You can delete backup sets (or only the SnapInfo Snapshot copies related to those backup sets) that contain one or more selected Storage Groups or databases. If you need to free some space or reduce your Snapshot copy count without deleting a SnapManager Backup, you can delete the SnapInfo Snapshot copies without affecting the associated backup.

About this task

The Backup Component selection is set to Backup Data Sets by default, so that you delete the entire backup set, including related transaction logs and SnapInfo directories. You can narrow this selection further by selecting **SnapInfo Snapshot Copies Only**.

The Management Group selection further narrows the scope of the deletion by specifying the backup management group of the backups that you want to delete.

For easiest backup administration, it is best if you always back up an entire Storage Group set.

If you select the **Retain up-to-the-minute restore ability for older backups in other backup management groups** option, transaction logs are deleted only from the specified backup management group; transaction logs are not deleted from other backup management groups. On the other hand, if you clear this option, transaction logs are deleted from other backup management groups as well.

Steps

1. In the **Actions** pane, click **Delete Backup**.

   Note: You can access **Delete Backup** when you connect to a member server, but not when you connect through the Database Availability Group (DAG).

2. If you want to display a list of the backup components on your SnapManager system, select **Backup Component**.

   Note: You cannot select an individual backup component for deletion by using this list.

3. Select one or more Storage Group sets or databases whose backups are to be deleted.

4. To delete backups that contain some of the selected Storage Groups (in addition to the backups that contain all the selected Storage Groups), click **Advanced > Delete backups that contain databases belonging to one or more storage groups**.

   Note: This option is useful only if you created backups that did not contain all Storage Groups in a Storage Group set.

   For this backup deletion operation only, multiple backup deletions delete backups that contain any one or more of the selected databases.

5. Select the Backup Component.

6. Select the Management Group.

   It is set to Standard by default. You can change Management Group to Daily, Weekly, or All.

7. If you want to retain up-to-the-minute restore ability for older backups in other backup management groups, click **Advanced > Retain up-to-the-minute restore ability for older backups in other backup management groups**.

   Note: If you selected All for the backup management group, this option has no effect and is disabled.

8. In the specified Storage Group and backup management groups, specify those that you want to delete.

If you want to delete...	Then...
The oldest backups	In the **Delete oldest backups in excess of** box, specify how many of the newest backups in the specified backup management group you want to preserve.
All the backups in the specified backup management group	Select the **Delete all backups in the specified management group** option.
Only the backups that are older than a specified number of days	In the **Delete backups older than** box, specify the number of days for which you want to keep the most recent backups.

9. To see what backup components would be deleted with the parameters you have entered, without deleting them, click **Delete Preview**.

   The Delete Backup Data Set Preview window is displayed. After a moment, a count and list of the backup components identified for deletion are displayed.

10. If, after previewing the deletion, you want to delete the backup components listed in the preview, click **Delete** in the **Delete Backup Data Set Preview** window.

    The backup components listed in the preview window are deleted.

11. If you chose not to preview and delete the backup components, as described in Steps 9 and 10, click **Delete** now to delete them.

    When the deletion is complete, a status popup is displayed.

Explicitly deleting Snapshot copies created by SnapManager Restore

You can delete Snapshot copies created during previous restore operations.

Steps

1. In the **Actions** pane, click **Delete Backup**.
2. Select **Delete snapshot of LUNs created during restore**.
3. If you want to display a list of the backup components on your SnapManager system, select **Backup Component**.

   Note: You cannot select an individual backup component for deletion by using this list.

4. Specify the restore Snapshot copies that you want to delete.

If you want to delete...	Then...
The oldest restore Snapshot copies	In the **Delete oldest snapshots in excess of** box, specify how many of the newest restore Snapshot copies you want to preserve.
All the restore Snapshot copies	Select the **Delete all snapshots created during restore** option.
Only the restore Snapshot copies that are older than a specified number of days	In the **Delete backups older than** box, specify the number of days for which you want to keep the most recent restore Snapshot copies.

5. To see what restore Snapshot copies would be deleted without deleting them, click **Delete Preview**.

   The **Delete Backup Data Set Preview** dialog box appears. After a moment, a count and list of the restore Snapshot copies identified for deletion are displayed.

6. If, after previewing the deletion, you want to delete the restore Snapshot copies listed in the preview, click **Delete** on the **Delete Backup Data Set Preview** dialog box.

   The restore Snapshot copies listed in the preview window are deleted.

7. If you chose not to preview and delete the restore Snapshot copies, as described in Steps 5 and 6, click **Delete** now to delete them.

   When the deletion is complete, a status popup is displayed.

Problem deleting backups due to busy Snapshot copy error

If you delete a backup copy of a LUN that was already backed up by another Snapshot copy, you get an error message saying that the Snapshot copy is busy and cannot be deleted. In this case, you must delete the most recent backup copy before the older backup can be deleted.

You can view Snapshot copy status in OnCommand System Manager or use the storage system `snap list` command. For more information about deleting busy Snapshot copies, see *Data ONTAP Block Access Management Guide for iSCSI and FC* for your version of Data ONTAP.

Note: To avoid this situation, ensure that you do not make backup copies of LUNs that are already backed up by Snapshot copies (for example, during a verification or while archiving from a LUN backed by a Snapshot copy).

How SnapManager uses SnapMirror

SnapManager 5.0 and later enables you to verify databases that are stored on the LUNs of the destination SnapMirror volumes.

Volume replication using SnapMirror

SnapMirror creates replicas of volumes. SnapMirror mirrors a Snapshot copy of data on a source volume to one or more destination volumes. SnapMirror automatically reflects incremental changes of the source volume on the destination volume.

SnapMirror can replicate a source volume to a destination volume on the same storage system or on a different storage system. The destination storage system can be in a different geographical location. This ability to duplicate data in different locations is a key component of a sound disaster recovery plan.

Data stored in a destination volume can be accessed through SnapDrive. Because the duplication is volume wide, Snapshot copies of other datasets on the source volume are also mirrored. And because SnapMirror updates the destination volumes to reflect incremental changes on the source volume, the result is an online, read-only destination volume that contains the same data as the source at the time of the most recent update.

Attention: SnapManager uses SnapMirror in asynchronous mode. Any disk writes on the source volume after the most recent SnapMirror replication update are not available if a catastrophic failure occurs before the next SnapMirror update.

Where to find more information about configuring and using SnapMirror

You can obtain more information about configuring and using SnapMirror in Data ONTAP documentation.

- See the *Data ONTAP Data Protection Online Backup and Recovery Guide* for your version of Data ONTAP for the following SnapMirror information:
  - What SnapMirror volumes are
  - How SnapMirror volumes work
  - When to use or access data stored on a SnapMirror destination volume
  - How to set up and configure SnapMirror on a storage system

Requirements for using SnapMirror with SnapManager

To use SnapMirror with SnapManager, you should ensure that your configuration satisfies some requirements regarding the source and destination volumes, licenses, and configuration, as well as the replication schedule.

- Ensure that there is at least one or more SnapMirror source volume.
- Ensure that there is at least one or more SnapMirror destination volume for each source volume.
- Ensure that the size of the destination volumes is equal to or greater than the size of source volumes.
- Configure SnapMirror first on the source volume and then its destination volume.
  See the relevant SnapDrive documentation for more information.
- Enable SnapMirror licenses on both the source and destination storage systems.
- Manually configure and initialize SnapMirror replication between source and destination volumes.
- Configure SnapMirror replication to be asynchronous.
- Disable the SnapMirror replication schedule on your storage system.

How SnapManager uses SnapMirror and SnapDrive

SnapManager uses SnapMirror to enable you to replicate backups to mirrored volumes. You can also perform concurrent backup verification. SnapManager coordinates with SnapDrive to perform asynchronous replication using SnapMirror. The changes reflected in a restore operation depend upon your backup schedule.

If a backup resides in a volume that is configured as a SnapMirror source volume, SnapDrive requests a SnapMirror update for that volume. SnapMirror replication is asynchronous, and SnapManager does not support SnapMirror qtree replication.

The schedule you define for backups sets the schedule for mirror replication. The changes made between consecutive backups are not reflected in the SnapMirror destination volume. Any restore operation from the destination volume restores the databases to their state at the time of the last backup. You can use SnapDrive to minimize the time between two SnapMirror replications.

How SnapMirror replication works

You can configure a SnapMirror relationship and request a mirror update. SnapMirror then coordinates with SnapDrive to reflect the incremental changes of the source volume on the destination volume.

With SnapManager 3.0 for Microsoft Exchange and later, after the SnapManager backup is complete, the mirror update request is delayed until SnapManager explicitly requests it. If you do not want to update the SnapMirror destination volume after a backup, you can override the update by using the Backup wizard or the Backup and Verify window.

The following process describes SnapMirror destination replication:

1. You configure a SnapMirror relationship between source volumes and destination volumes. Create a Custom schedule with schedule parameters ("- - - - -"). Initialize the mirror to perform initial transfer of contents from a source volume to a destination volume.
2. You select the "Update SnapMirror after operation" option for the operation.
3. If any volume whose data is captured in the backup is a SnapMirror source volume, SnapDrive requests information about all SnapMirror destination volumes for the source volume.
4. SnapDrive sends a SnapMirror destination update request to all the related destination volumes.
5. SnapMirror updates the destination volumes to reflect incremental changes on the source volume.

Integrity verification on SnapMirror destination volumes

SnapManager 5.0 and later enables you to verify the Exchange databases that are stored on the LUNs of the destination SnapMirror volumes.

When you verify the integrity of a destination volume, SnapManager automatically detects the SnapMirror relationships with the appropriate source volume in the volumes and selects the available SnapMirror relationship for the selected destination volume.

If you have changed the volume SnapMirror relationship by mirroring data to a new destination storage system, you can set the new path as the volume SnapMirror destination for verification purposes. To change the volume SnapMirror destination, perform the following steps:

- Click the SnapManager for Exchange Backup verification settings.
- Click the Verification Server tab.
- Click the Verification on destination volumes button.

Selecting the SnapMirror destination volumes for verification

You must select a destination volume when you establish a SnapMirror relationship. SnapManager automatically detects and selects the available SnapMirror relationships between the SnapMirror source and SnapMirror destination volumes.

Steps

1. Click **Backup Verification Settings** in the **Actions** pane.
2. Click **Verification Server**.
3. In the Verification Server tab, click **Verification on destination volumes**.

   Note: If the SnapMirrored volume is not available, SnapManager displays an appropriate error message.

You cannot configure SnapMirror for deferred verification through the Database Availability Group (DAG). You can configure SnapMirror for deferred verification only when you connect to a member server of the DAG.

SnapManager displays the Choose SnapMirror Destination Volumes for Integrity Verification window.

4. Select the destination volume for each SnapMirrored volume.

   Note: By default, SnapManager displays the Number of Relationships field. You cannot edit this value. If the destination volume is not in the SnapMirrored state or does not have FlexClone license installed, SnapManager displays an error message when you click **Apply**, and the integrity verification fails for that volume.

5. Click **Apply**.
6. Click **OK**.

Requirements to run destination volume integrity verification

You must ensure that your system meets the system requirements for the SnapMirror and FlexClone licenses before you start integrity verification.

Before beginning integrity verification, ensure that a SnapMirror license is enabled on the source volume and a FlexClone license is enabled on the destination volume. SnapManager uses SnapDrive to verify that the required licenses are enabled on the source and destination storage system.

If SnapDrive is using RPC to communicate with the destination storage system, the destination storage system must have CIFS service started and configured correctly, to be accessible by SnapDrive (note that only a limited CIFS license is needed for this; it is sufficient that the storage system can be a member of the Active Directory domain). SnapDrive provides access to the Exchange databases that are stored on the destination volume, and SnapManager performs the integrity verification on the backups of those databases.

Troubleshooting integrity verification failure on SnapMirror destination volumes

If you have changed the volume SnapMirror relationship by mirroring data to a new destination storage system, SnapManager is not able to find the destination storage system and issues an error. You can set the new path as the volume SnapMirror destination for verification purposes.

About this task

SnapManager for Exchange integrity verification on volume SnapMirror destination fails when the destination storage system is no longer available or it has changed. If the volume SnapMirror destination storage system is no longer available, SnapManager for Exchange provides the following error message:

The destination volume filer2:vol1 of the source Logical Disk e does not have valid FlexClone state or SnapMirrored state. FlexClone State is Unknown and SnapMirror State is Unknown. Please check

the SnapMirror relationship filer1:vol1 filer2:vol1_dest. Unable to mount snapshot, abort verification.

If you have changed the volume SnapMirror relationship by mirroring data to a new destination storage system (for example filer3), SnapManager for Exchange might still try to contact filer2 which might not necessarily exist. This is very likely to happen when you perform a storage system hardware upgrade and replace an old storage system head with a new one. When you run the SnapMirror update command on new destination storage system it will work, but SnapManager for Exchange might not be up to date with the new destination storage system.

To resolve this problem, ensure that the following requirements are met:

- SnapMirror status on both the volume SnapMirror source and the volume SnapMirror destination storage system is updated with the new destination and that you have run at least one successful manual update (from storage system console) after the volume SnapMirror destination has changed.
- The new destination storage system holds the SnapManager for Exchange user in its built in Administrators group.

Complete the following steps to set the new path as a volume SnapMirror destination for verification purposes.

Steps

1. Click SnapManager for Exchange Backup verification settings.
2. Click the **Verification server** tab.
3. Click the **Verification on destination volumes** button.

Types of destination volume integrity verification

You can run different types of SnapMirror destination volume integrity verifications for different SnapManager operations.

- Backup with verification
- Deferred integrity verification
- Test restore operation
- Restore operation
- Remote verification

Backup with verification

When you run integrity verification on the SnapMirror destination volume, SnapManager requests SnapMirror updates using SnapDrive, verifies the backup, and monitors the SnapMirror update activity through SnapDrive.

When you run integrity verification on the SnapMirror destination volume, SnapManager performs the following operations:

1. Backs up the databases
2. Requests a SnapMirror update through SnapDrive to replicate the data across destination volumes
3. Verifies the integrity of databases and transaction logs from LUNs that are located in the selected destination volumes
4. Updates verification results to the SnapInfo directory
5. Updates the SnapMirror instance after the operation, to replicate verification results to the SnapInfo volume
6. When the SnapMirror update replicates the backup to the selected destination, SnapManager continuously monitors the SnapMirror update activity through SnapDrive, as follows:
   - SnapDrive provides the SnapMirror update progress information continuously to SnapManager during the update.
   - SnapManager logs the SnapMirror update activity to the Windows Application event log and to the backup report at every defined interval.

   Note: If you have passthrough LUNs on a hyperV client, you must either run the verification operation on a separate host that has direct connection to the destination storage system, or you must create a direct connection on the local machine for verification.

   Note: If the SnapMirror update operation does not have any progress within a defined interval, SnapManager aborts monitoring it and leaves the backup unverified.

Integrity verification for test restore operations

When you run an integrity verification for the test restore operation, SnapManager performs integrity verification and updates the verification results to the source SnapInfo volumes.

SnapManager performs the following tasks for integrity verification for a test restore operation:

1. Runs the metadata verification and integrity verification on the backups that are stored on the destination volumes
2. Verifies the transaction log signature and sequence in the source SnapInfo directory
3. Verifies the database metadata
4. Updates the verification results to the source SnapInfo volumes

Note: SnapManager does not verify the backup on the source volume, if an unverified SnapManager backup is not available on the destination volume. If you try to run integrity verification in such a case, SnapManager displays an error message.

Integrity verification for a restore process

When you run an integrity verification for a restore process, SnapManager performs integrity verification, updates source SnapInfo volumes, and logs additional steps to event logs and restore reports.

1. SnapManager performs integrity verification on the destination SnapMirror volumes, if a backup is available on the destination volume.
2. SnapManager displays an appropriate error message, if a backup is not available on the destination volume.
3. SnapManager updates the verification results to the source SnapInfo volumes.
4. SnapManager logs the additional steps to the Windows Application event log and to the SnapManager restore report.

Remote destination volume integrity verification

When you run integrity verification on a remote destination volume, SnapManager performs integrity verification and updates source SnapInfo volumes.

SnapManager performs the following actions:

1. Runs database metadata verification on the backups that are stored on the source volumes
2. Verifies the transaction log signature and sequence in the source SnapInfo directory
3. Runs integrity verification on the backups that are stored on the destination volumes
4. Updates the verification results to the source SnapInfo volumes

Deferred integrity verification

In deferred integrity verification, SnapManager checks if a backup copy exists and then updates SnapMirror if you select to update SnapMirror after the backup operation. If a backup Snapshot copy does not exist, SnapManager displays an appropriate error message.

If you do not select the “SnapMirror update after operation” option, SnapManager updates the verification results only on the source SnapInfo volumes.

If you select the **SnapMirror update after operation** option, the following actions occur:

1. SnapManager verifies the backup copy on the selected destination.
2. SnapManager updates the verification results on the source SnapInfo volumes.
3. SnapMirror replicates the verification results on the source SnapInfo volumes and sends them to all the destination SnapInfo volumes.

If you select the **SnapMirror update after operation** option but the backup Snapshot copy is not available on the destination volume, the mount operation fails for integrity verification and leaves the backup copy unverified. In this case, SnapManager does not request for the SnapMirror update.

Concurrent backup verification

You can run multiple backup verification jobs concurrently on the verification server, when verification requests originate on a different server than the server where the backup is created.

Note: SnapManager executes multiple verification requests from the same host serially.

SnapManager workflow for concurrent backup verification

The workflow for concurrent backup verification is as follows:

1. After you create a new backup copy at the source backup server, SnapManager immediately sends it for verification to the remote verification server.

   The remote server uses the same Job ID as the source server.

2. The backup job sent to the remote server is displayed as running on the source server.

   The remote server shows the active job as running or queued, depending on its position in the queue.

3. After SnapManager verifies the first backup set at the remote server, the job runs until all the backup sets are created and verified.

   Note: There is a maximum of four backup jobs from different servers that SnapManager can verify simultaneously. SnapManager places the subsequent jobs in the queue.

   The job can be a full backup job that includes backup verification, a deferred backup verification job, or a backup verification job that is initiated as part of a restore job.

You can create a new full backup when the deferred integrity verification job is running if all production volumes are FlexClone enabled.

Concurrent backup verification during a restore operation

You can restore a backup when the verification job is running. When you submit a restore job with verification, SnapManager performs one of the following tasks based on the scenario:

1. Aborts the restore job if any backup or deferred integrity verification jobs are running
2. Cancels all the running jobs and then starts the restore operation

   Note: The status of the Storage Group of all the canceled verification jobs remains unverified.

3. Waits for all the currently running jobs to finish and then starts the restore operation

When you select to have the restore operation cancel all the currently running jobs, the restore operation performs the following tasks:

1. Stops processing new jobs that are in the queue
2. Disables all SnapManager scheduled tasks in the Windows scheduler
3. If SnapManager is creating a full backup, cancels the full backup operation
4. If SnapManager is running a Frequent Recovery Point operation, cancels the Frequent Recovery Point operation
5. If one or more verification jobs are running, cancels all running jobs
6. Waits for the full backup operation to stop
7. Runs the restore operation

Concurrent backup verification and Frequent Recovery Point backup

Concurrent backup verification involves Frequent Recovery Point backup running in parallel with full backup.

- If SnapManager is waiting for VSS Snapshot copy to start, then Frequent Recovery Point backup is running.
- If SnapManager is Running backup verification for the new VSS Snapshot copy backup, then Frequent Recovery Point backup starts.
- If SnapManager is Creating a new VSS Snapshot copy backup, then Frequent Recovery Point backup fails to start.

You can select restore options for such backup from Job Control Options in the Restore window.

Managing integrity verification jobs

You can view, move, and cancel queued and running integrity verification jobs. You can perform all the tasks from the Current Job Status pane. By default, job management is enabled in SnapManager.

Steps

1. In the Scope pane, select an Exchange server.

   SnapManager displays the queued and running jobs for that server in the Current Job Status pane.
2. Select the job you want to manage.
3. Position your mouse over the job.

   A floating menu appears.
4. Manage the jobs as described in the following table.

If you want to...	Then do this...
Retrieve job information	Select the job. SnapManager displays the job information in the Results pane.

If you want to...	Then do this...
Move a job up within a queue	Select **Move Up** in the floating menu.
Move a job down within a queue	Select **Move Down** in the floating menu.
Cancel a queued or running job	Select **Cancel Job** in the floating menu.

Disaster recovery using SnapManager

You can use SnapManager to prepare for disasters, to perform failover and failback operations, and to restore Exchange databases that have been destroyed or compromised.

For the most current information about planning for disaster recovery, see *Technical Report 4066: SnapManager for Exchange Business Continuance Module for Microsoft Exchange Server 2003/2007 Disaster Recovery*.

SnapManager backup archiving

You can use SnapManager to create offline archives of Snapshot copies containing SnapManager backup sets. Archiving data enables you to create a complete, self-consistent replica of your data, should you have to recover it. You can use several methods to archive your data.

Why organizations archive data

The main reason for archiving data is disaster recovery. Archiving helps you to recover data damaged or accidentally deleted due to human error, due to accidental deletion, hardware failure, or natural calamity. Space constraints, historical analysis, and litigation often require the older data to be archived.

Guidelines for archiving SnapManager backups

Before you start archiving your data, consider the guidelines for choosing the type of data to be archived, naming, and using the appropriate protocols.

- Archive only verified backups.

  If you are not sure whether a backup copy is verified, use the SnapManager Restore window to check.
- Archive a complete backup set.
- Archive the most recent backup copy.
- Archive all Storage Groups together.

  Do not archive individual Storage Groups or databases unless you know which Snapshot copies contain the appropriate Storage Groups and transaction logs for a specific time.
- Keep in mind that single backup version is used during backup archiving when multiple storage groups or databases are selected to create a backup. This means that a backup version in Protection Manager includes backups from multiple storage groups or databases. When deleting a remote backup, all backups in that backup version will be deleted.
- If you use the unique naming convention, look for the Snapshot copy with the most recent date and time.
- If you did not use the unique name option when you created the backup, look for the most recent Snapshot copy in the storage system's LUN drive volume named `/exchsnap_`*`servername`*`_recent` or `/exchsnap_`*`servername`*`_recent_`*`backupmgmtgroup`*.This is for backward compatibility with earlier versions of SnapManager, which did not include the **Run Command After Operation** feature.
- Do not use the CIFS or NFS protocols to archive LUNs.

  Use the storage system's `dump` command or an NDMP backup application to archive LUNs.

Note: If the system is busy, the network is slow, or the load is primarily on the DataFabric Manager server or the storage system, there is a time lag between the creation of a backup and the appearance of the archive in the Restore view.

- Consider the following factors:
  - The archive method you use
  - Service Level Agreements for disaster recovery
  - The number of SnapManager backups performed per day
  - Exchange client activity schedules
  - Backup verification time

Methods of archiving SnapManager backups

You can use different methods for selecting the components of your backup to archive.

- Use NDMP or the storage system's `dump` command to archive LUNs directly from the storage system to the archive medium.
- Mount the LUNs in a SnapManager backup Snapshot copy and share it, then use NTBackup or another Windows backup utility to copy the LUNs' contents to the archive medium.
- Create an Exchange copy-type backup directly on the archive medium using an Exchange-aware backup application.

Archives created with NDMP or the dump command

You can use Network Data Management Protocol (NDMP) or the Data ONTAP `dump` command to archive each of the LUNs that contain data for the backup set that you want to archive directly from the storage system to the archive medium, without involving Exchange or the Exchange server at all.

NDMP and the `dump` command are the most efficient ways to archive LUN drive files. LUN Snapshot copies are made, copied to the archive medium, and deleted.

For more information about backing up storage system data to tape, see the *Data ONTAP System Administration Guide* for your version of Data ONTAP.

When using NDMP or the `dump` command to archive your SnapManager backups, specify the database LUN by using its absolute path name `/vol/`*`volume_name`*`/.snapshot/`*`snapshot_name`*`/`*`LUN_name`* name, in which the three variables represent the following text strings:

volume_name	Name of the volume containing the data to be archived
snapshot_name	Name of the Snapshot copy containing the LUNs to be archived
LUN_name	Name of the LUN containing the data to be archived

For example, the absolute path name of the LUN `exch1db.lun` in the Snapshot copy `exchsnap__SRVR3_01-20-2006_12.05.58_Daily` on the volume ExchVoln is represented as follows:

```
/vol/ExchVol/.snapshot/exchsnap__SRVR3_01-20- 2006_12.05.58__Daily/
exch1db.lun
```

Evaluation of the NDMP and dump command method of archiving

The NDMP and the `dump` command are the most efficient methods of creating archives of the LUN drive files, you archive more data than you need.

Advantages

- Because the NDMP and the `dump` command methods do not rely on mounting a Snapshot copy, you do not risk creating of busy Snapshot copies.
- Because the NDMP and the `dump` command methods archive the entire raw LUN, restoring involves replacing the LUNs.
- If your archive procedure does not send the data over the network, the NDMP and the `dump` command methods can be significantly faster than other methods.

Disadvantages

- Because you are archiving raw LUNs, the entire LUN containing the Exchange data is archived, so you archive more data than you need.

  If archiving extra data is undesirable, you can use NTBackup or another Windows backup utility to back up the corresponding SnapInfo directory. Coordinate this so that the two pieces of the archive are kept together for later retrieval.
- If you archive the SnapInfo directory separately, you must ensure that you get the SnapInfo directory backed up directly from the Exchange server and the Exchange data extracted from the LUN backed by Snapshot copy from different locations into the same archive.

Example using NDMP or dump to archive SnapManager backups

You can run a script to archive your SnapManager backups to tape either by using the Run Command After Operation feature or by hard coding the Snapshot copy names into the script.

Assume that you want to run a script in the following environment:

- Run the script on the computer running Exchange and SnapManager.
- The name of the Exchange server is `SRVR3`.
- The name of the storage system is `storagesystem1`.
- The name of the LUN containing the Exchange databases is `exch1db.lun`.
- The name of the LUN containing the Exchange transaction logs and the SnapInfo directory is `exch1logs.lun`.
- The name of the volume that contains the Exchange database LUN is `ExchDBvol` .

- The name of the volume that contains the Exchange transaction log LUN is `ExchLogvol`.
- The NetApp Data ONTAP PowerShell Toolkit is installed on the Exchange server.

To run a script using the SnapManager Run Command After Operation feature, use the following value with the `-RunCommand` parameter:

`C:\scripts\archive.bat $ExchSnapshot $InfoSnapShot`

The values passed into the script provide values similar to the following, respectively:

`exchsnap__SRVR3_01-20-2012_12.05.58__Daily`

`exchloginfo__SRVR3_01-20-2012_12.05.58__Daily`

The archive.bat script calls a Windows PowerShell script that passes the values into the PowerShell script:

`powershell.exe C:\Scripts\archive.ps1 %1 %2`

If you do not use the Run Command After Operation feature and you do not use the unique naming convention for your backups, you can hard code the Snapshot copy names into the script.

For example, the `dump` command for the database might look like this PowerShell command, for an Exchange server named SRVR3:

`invoke -nassh -command "dump 0f nrst0a" /vol/ExchDBVol/.snapshot/exchsnap__SRVR3__recent__Daily/exch1db.lun`

In the following script the `archive.ps1` script performs the archive of the SnapManager backup to tape via ndmp:

```
#Backup up the LUN in the snapshot containing the Exchange database.
#For this example, a level 0 backup is performed
#to the tape storage device named nrst0a.
import-module dataontap

$ExchSnapshot = $args[0]
$InfoSnapshot = $args[1]

$ExchDBPath = “/vol/ExchDBVol/.snapshot/” + $ExchSnapShot + “/
exch1db.lun”
$ExchLogPath = “/vol/ExchLogVol/.snapshot/” + $ExchSnapShot + “/
exch1db.lun”
$ndmpDBcmd = “dump 0f nrst0a” + $ExchDBPath
$ndmpLogcmd = “dump 0f nrst0a” + $ExchLogPath

Connect-NaController -Name storagesystem1 -Transient
```

```
invoke-nassh –command $ndmpDBcmd
invoke-nassh –command $ndmpLogcmd
```

Archives created using a Windows backup utility

You can use a Windows backup utility to archive your SnapManager backups. To do so, you mount the LUNs backed up by the Snapshot copy that you want to archive, and then use Windows NTBackup or another Exchange-aware Windows backup utility to copy the archive data to your archive medium.

In this case, the NTFS data is backed up, rather than the raw LUNs, as shown in the following diagram:

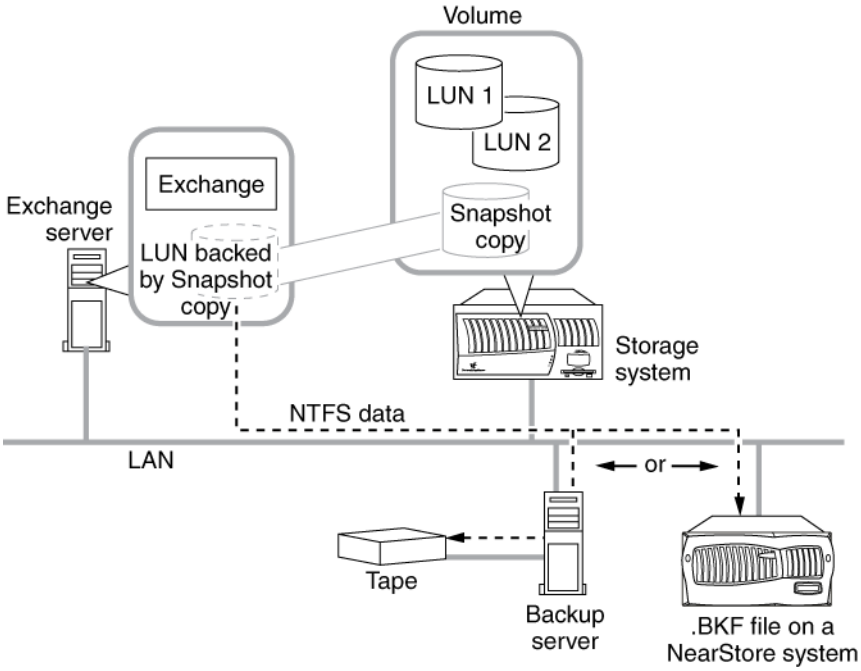

Note: You do not need to mount the LUN on the Exchange server; you can use another computer for archiving.

Your archive must include the following two components:

- The SnapInfo directory backed up directly from the Exchange server
- The Exchange data extracted from the LUN backed up by a Snapshot copy

Evaluation of the Windows backup utility method of archiving

Using the Windows backup utility to archive your SnapManager backups enables you to select exactly which data you archive; however, you must be careful to avoid scheduling any backups while the archiving is performed.

Advantages

- Because you are archiving NTFS data rather than raw LUNs, you can archive exactly the data that you need, and no more.
- The procedures and tools used for this method are familiar and available to you.

Disadvantages

- Because this method relies on mounting a Snapshot copy, you must be careful to avoid scheduling any backups during the archive process, because creating a Snapshot copy of a mounted Snapshot copy results in a Snapshot copy that cannot be deleted.
- You must make sure that you get the SnapInfo directory backed up directly from the Exchange server and the Exchange data extracted from the LUN backed by a Snapshot copy from different locations into the same archive.

Example: Using a Windows backup utility to archive SnapManager backups

You can use a script to mount a Snapshot copy of the LUNs that contain the Exchange databases, then back up the databases using Windows NTBackup, unmount the LUNs, and back up the SnapInfo directory.

Assume that you use a script in the following environment:

- The script is run on the computer running Exchange and SnapManager.
- The drive letter for the Exchange database is S:
- The Snapshot copy is mounted as drive V:

  The drive letter used for the LUN mount must be available when you run the script.

The following script mounts a Snapshot copy of the LUNs that contain the Exchange databases, and it backs them up using Windows NTBackup. It then unmounts the LUNs and backs up the SnapInfo directory. The `/N` and `/D` options are used to show the name of the tape and the description options.

```
REM Mount a LUN backed by the Snapshot copy of the Exchange
REM database as drive letter V:\.
sdcli snap mount -k s -s %1 -d v
REM Use NTBackup to back up the database files. The path
REM to the databases in your environment might be different.
ntbackup backup "V:\Program Files\Exchsrvr\mdbdata" /N %1 /D %2
REM Dismount the Snapshot copy mounted as drive letter V:\.
REM The below example is shown with an optional parameter (-f)
REM to forcefully disconnect the drive letter.
sdcli snap unmount -d v -f
REM Use NTBackup to back up the snapinfo directory.
REM This backup appends the media so as not to
```

```
REM overwrite the database backup. The path to the
REM snapinfo directory is passed as the third
REM parameter on the command line used to launch
REM this script.
ntbackup backup %3 /T%1 /A
```

To run this script using the SnapManager Run Command After Operation feature, use the following command:

`C:\SnapManager Scripts\scriptname.txt $ExchSnapshot $InfoSnapshot`

The command-line parameters %1 and %2 might provide values similar to the following, respectively:

- `exchsnap__SRVR3_01-20-2006_12.05.58__Daily`
- `eloginfo__SRVR3_01-20-2006_12.05.58__Daily`

If you prefer not to use the Run Command After Operation feature and you are not using the unique naming convention for your backups, you can hard code the Snapshot copy names into the script.

For example, the `dump` command for the database might look like this, for an Exchange server named SRVR3:

`rsh storagesystem1 dump 0f nrst0a`

`/vol/Exch/.snapshot/exchsnap__SRVR3__recent__Daily/exch1db.lun`

How to use the Exchange Backup Agent to archive Exchange backups

You can use Exchange backup copies to archive your SnapManager backup copies. You use NTBackup (which, in turn, uses the Exchange Backup Agent) to create a backup copies of your Exchange databases. In this case, the Exchange data itself is captured and archived.

Because you do not use SnapManager to create the backup copy, you also do not use SnapManager to perform the restore process, so the SnapInfo directory does not need to be archived.

> **Note:** Use a copy or differential backup for this method. Performing any other type of backup operation causes your existing SnapManager backup copies to become unusable for an up-to-the-minute restore operation.

Evaluation of the Exchange Backup Agent method of archiving

When you use Exchange Backup Agent to archive Exchange backup copies, you do not need to archive the SnapInfo directory using the Exchange Backup Agent method, and Exchange data itself is captured and archived.

Advantage

The advantage of this approach is that the procedures and tools used for this method are familiar and available to you.

Disadvantage

The disadvantage of this approach is that no SnapManager backup operation can be initiated while this method is in use.

Example: Using Exchange Backup Agent to archive Exchange backup copies

You can run a script to archive your SnapManager backup copies to tape using Exchange Backup Agent. The script launches NTBackup and tells it to perform a copy-type backup operation through the Exchange Backup Agent.

Example script

Before you can use this script, you need to use the Windows NTBackup GUI to create a `.BKS` file that contains the Storage Groups you want to back up. This script assumes that it is run on the computer running Exchange and SnapManager.

```
REM Launch NTBackup
ntbackup backup "@c:\MyConfig.BKS" /M copy
```

Note: Although this method does not use SnapManager backup copies, it can be run using the Run Command After Operation feature.

If you use a centralized backup model

When you use a centralized backup model, you can initiate a backup job by running a command on a centralized backup server, or by running a command on the Exchange server through the third-party agent that came with your backup software.

Use your centralized backup software documentation to determine which command to run and which server to run it on. Then, you can use the Run Command After Operation feature to trigger the backup.

If you plan to run the command on the Exchange server, then you need take no extra steps.

If you plan to run the command on a centralized backup server, install SnapManager on it. Configure the SnapManager application on the Exchange server to run the command on a remote machine (the centralized backup server).

> **Note:** To run the command remotely, you need to install SnapManager on the remote server even if it does not have Exchange server installed. Although you need to install SnapManager, you do not need to completely configure it.

You can run a command remotely in many ways, but this document only covers doing it through SnapManager.

Automatic backup archiving using the Run Command After Operation feature

SnapManager allows you to automatically run your own program or script after a backup or database verification operation, but only after the operation finishes successfully. This feature, called Run Command After Operation, is typically used to automatically archive a backup copy.

Make sure that you archive all components of a backup set together. If any one of the Storage Groups is offline and you a initiate a backup process using the Run Command After Operation, the backup operation fails for that particular job.

Command arguments supported by the Run Command After Operation feature

The Run Command After Operation feature supports some pre-defined variables that can pass operation-specific information to your program or script.

`$ExchSnapshot` Expands to the name of an Exchange database Snapshot copy, as shown in the following examples:

`exchsnap__winsrvr3__01-31-2006_15.03.09`

`exchsnap__winsrvr3__recent`

The number of database Snapshot copies in a SnapManager backup set depends on the number of volumes used to store the Storage Groups included in the backup operation.

`$InfoSnapshot` Expands to the name of a SnapInfo directory Snapshot copy, as shown in the following examples:

`eloginfo__winsrvr2__01-31-2006_15.03.09`

`eloginfo__winsrvr2__recent`

`$SnapInfoName` Expands to the name of the SnapInfo directory, as shown in the following examples:

`WINSRVR3__01-23-2006_16.21.07__Daily`

`WINSRVR3__recent`

If you use this variable, you must also provide the correct path to the directory.

`$SnapInfoPath` Expands to the name of the SnapInfo subdirectory, as shown in the following examples:

```
I:\SME_SnapInfo\EXCH__WINSRVR3\SG__First Storage Group
\WINSRVR5_01-06-2006_11.52.36
```

The `$SnapInfoPath` variable is automatically enclosed within double quotes so that the actual path name can contain spaces without affecting the script invocation on the Windows command line.

If you do not want the double quotes to appear in your command line, remove them from the Command Arguments field in the Run Command After Operation dialog box.

Specifying the command to be run by the Run Command After Operation feature

While specifying a backup or verification operation, you can use Run Command After Operation to specify the details of the command that you want to run after the operation finishes. Using default values for the command arguments frees you from entering the command information each time you initiate a backup or verification operation.

Before you begin

- You must install SnapManager on a computer that you specify, but you do not need to configure it.
- If your script is stored on a network share, use the UNC path rather than the network drive letter to specify the script's location.

  Add the network location to Internet Explorer's list of trusted sites; otherwise, your command might fail.

About this task

If more than one Snapshot copy of the same type exists for your command, include that Snapshot copy's variable name once for every Snapshot copy. If you do not, you cannot create a complete archive. For example, if you enter `myscript $exchSnap $exchSnap`, the command that is generated is `myscript snap1 snap2`.

Steps

1. Use the Backup wizard or the Backup and Verify window to specify a full database backup, a transaction log only backup, or a database verification.
2. Select **Yes, Run a command after this operation** in the Backup wizard or **Run Command After Operation** in the Backup and Verify window.

3. Under **Specify a computer where...**, select the host on which your program or script resides.
4. Under **Specify the full path...**, select your program or script.
5. Type the command input string in **Command Arguments**.

   You can directly enter text or select the variable you want to enter from the SnapManager Variables list.
6. Repeat Steps 4 and 5 as needed until the Command Arguments box contains all of the arguments that you want to pass to your program or script.
7. Click **OK** and return to the backup or verification setup process.

   After the backup or verification finishes successfully, the command that you specified, with the Snapshot copy and SnapInfo information for this backup operation inserted, runs automatically.

Example: script and the required parameters

Suppose that you want to run the following script:

```
C:\SnapManager Scripts\scriptname.bat
```

You need the following parameters:

- Name of the database Snapshot copy
- Name of the specific SnapInfo directory
- NTFS path to the SnapInfo directory

Enabling the launch of SnapManager scripts from a UNC path

If SnapManager is installed on a Windows Server system and you want to launch a script from a UNC (Uniform Naming Convention) path, you need to add the host where the script is located to the Internet Explorer list of trusted sites.

About this task

- If SnapManager is installed on a Windows Server system, and you attempt to launch a script from a UNC (Uniform Naming Convention) path, SnapManager might hang. When this happens, the log file shows no indication of the failure and the last line of the log text “The specified command has been launched successfully.” does not appear.
- Windows Server 2003 or Windows Server 2008 ships with Internet Explorer's Enhanced Internet Explorer Security Configuration enabled. This setting is restrictive and prevents batch files located on a network share from running. The network location needs to be added to Internet Explorer's list of trusted sites.

- The system you log on to can be the same system that is running SnapManager, or it can be a different server. Log on using the same user account that SnapManager is configured to use.
- The security settings are specific to each user account on each machine. You need to repeat this procedure for on all machines that are used to run scripts and for all users who might launch the script on each machine. You will also need to repeat these steps if you start using a different user account to run SnapManager.

Steps

1. Log in to the Windows Server 2003 or Windows Server 2008 system that will be running the script.
2. Launch Internet Explorer.
3. In the menu bar, select **Tools > Internet Options**.
4. In the **Internet Options** dialog box, select **Security > Local intranet > Sites**.
5. In “Add this Web site to the zone,” enter the host name of the machine from which the script will be launched and click **Add**.
6. Click **OK**.
7. To verify your changes, browse to the network on which the script resides and launch the script.

   The script should now run normally without security prompts. The script might fail to complete properly because the SnapManager variables are not passed to the script in this test.

SnapManager reports and the report directory

SnapManager reports list step-by-step details of every SnapManager operation that you perform, their final statuses, and any error messages that you encounter during the operation.

The SnapManager Report Directory provides folders that group the reports for each of these operation types:

- Backup
- Config
- Debug
- Frequent Recovery Point
- Delete Snapshot
- Restore
- Miscellaneous

Note: There might be other folders in the directory if you are running SnapManager in a Windows cluster or you have a previous version of SnapManager.

SnapManager reports in a Database Availability Group (DAG)

When the SnapManager MMC snap-in is connecting to the Database Availability Group (DAG), you can view all the reports from all the member servers of the DAG.

The report folders reside on each member server of the DAG. The report folders for the DAG reside on the "owner node" of the DAG. Whenever configuration, backup, or restore operations are performed, the reports generated for the DAG are stored in the report folders of the server which is the owner node of the DAG at the time these operations were performed.

Reasons to change the report directory location

You might need to change the report directory location, either because of limited space or because you want to share the directory between SnapManager and the Exchange server running in an Microsoft Cluster Services (MSCS) cluster.

If you find you have limited space in the current report directory, you can move it to a different location that has more available disk space.

If you are running Exchange and SnapManager in an MSCS cluster, storing the SnapManager reports in the default location (in a directory named Report under the SnapManager installation directory) does not allow the report directory to be shared between the nodes in the cluster. Furthermore, you

would not see the same reports from different nodes. To avoid these problems, you can move the report directory to a disk that belongs to the same group as your Exchange virtual server.

Changing the SnapManager report directory

If you have limited space in your original report directory, you can change the report directory to store your reports.

About this task

If you change the directory, you can no longer see the reports that were created before you changed the path. To view them, change the report directory path back to the original path and refresh the view.

Steps

1. In the **Actions** pane, click **Report Directory Settings**.
2. Enter or browse to the path name for your new Report Directory.
3. In the **Actions** pane, click **Refresh**.

Locating the report directory in a Windows cluster

By default, the SnapManager Report Directory is on the disk on which SnapManager is installed. In a Windows cluster, the report directory is not shared among the nodes in the cluster, and you cannot see the same reports from different nodes.

About this task

- Change the directory from every SnapManager node.
- Use a disk that does not contain Exchange or SnapManager data for the report directory.
  The report directory is restored from its Snapshot copy when you perform a restore operation.

Steps

1. In the **Actions** pane, click **Report Directory Settings**.
2. Enter or browse to the new report directory path to a disk that belongs to the same group as your Exchange virtual server.
3. In the **Actions** pane, click **Refresh**.

Viewing SnapManager reports

You can view SnapManager reports from the SnapManager GUI.

Steps

1. In the Scope pane, click **Reports**.

   SnapManager displays the report folders in the Results pane.

2. In the report folders, select the database for which you want to view a report.

   SnapManager displays the report in the Results pane.

Printing SnapManager reports

You can print SnapManager reports either from within the displayed report, or you can open the report in Notepad and print it.

Steps

1. In the Scope pane, click **Reports**.
2. Click the directory that contains the report you want to print.

If you want to...	Then...
Print the displayed report directly	**a.** Select the report that you want to print. **b.** Right-click anywhere within the displayed report in the Result pane, and select **Print**.
Print from Notepad	**a.** Right-click the report name in the Result pane and select **Open with Notepad**. **b.** From Notepad, select **Print** in the File menu.

Deleting SnapManager reports

You can delete reports from the Report Directory to increase the space available in the directory.

Steps

1. Click the reports in the Scope pane.

2. Select the report you want in the SnapManager Report Directory.
3. Right-click the report directory or the individual report that you want to delete and select **Delete All** or **Delete.**
4. Click **Yes** in the dialog box that appears.

   The report is deleted.

Dataset and SnapVault integration

Dataset and SnapVault integration with SnapManager provides an integrated, rapid way to create, restore, and manage remote backup sets and archives. SnapManager coordinates with Protection Manager for the integration.

The functionality for dataset and SnapVault integration is available only if you are using Data ONTAP 7.3 or later, along with Protection Manager 3.7 or later. SnapManager uses Data ONTAP Snapshot technology to create and restore local backup copies.

Note: Protection Manager integration can be enabled without reinstalling SnapManager for Exchange anytime by rerunning the Configuration wizard. You must first ensure that you have configured SnapDrive for Protection Manager integration. If you have not already done so, you can set up SnapDrive Protection Manager integration via the `sdcli` command, and then restart the SnapDrive service, or you can rerun the SnapDrive setup. (See the SnapDrive help for more information on Protection Manager integration.)

Dataset concepts

Associating data protection, disaster recovery, a provisioning policy, or a storage service with a dataset enables storage administrators to automate tasks, such as assigning consistent policies to primary data, propagating policy changes, and provisioning new volumes, qtrees, or LUNs on primary and secondary dataset nodes.

Configuring a dataset combines the following objects:

Dataset of physical storage objects	For protection purposes, a collection of physical resources on a primary node, such as volumes, flexible volumes, and qtrees, and the physical resources for copies of backed-up data on secondary and tertiary nodes. For provisioning purposes, a collection of physical resources, such as volumes, flexible volumes, qtrees, and LUNs, that are assigned to a dataset node. If the protection policy establishes a primary and one or more nonprimary nodes, each node of the dataset is a collection of physical resources that might or might not be provisioned from the same resource pool.
Dataset of virtual objects	A collection of supported VMware virtual objects that reside on storage systems. These virtual objects can also be backed up locally and backed up or mirrored on secondary and tertiary nodes.
Resource pool	A collection of physical resources from which storage is provisioned. Resource pools can be used to group storage systems and aggregates by attributes, such as performance, cost, physical location, or availability. Resource pools can be assigned directly to the primary, secondary, or tertiary nodes of datasets of physical storage objects.

	They can be assigned indirectly both to datasets of virtual objects and to datasets of physical storage objects through a storage service.
Data protection policy	A set of rules that define how to protect primary data on primary, secondary or tertiary storage, as well as when to create copies of data and how many copies to keep. Protection policies can be assigned directly to datasets of physical storage objects. They can be assigned indirectly to both datasets of virtual objects and to datasets of physical storage objects through a storage service.
Provisioning policy	A set of rules that define how to provision storage for the primary or secondary dataset nodes, and provides rules for monitoring and managing storage space and for allocating storage space from available resource pools. Provisioning policies can be assigned directly to the primary, secondary, or tertiary nodes of datasets of physical storage objects. They can be assigned indirectly to both datasets of virtual objects and datasets of physical storage objects through a storage service.
Storage service	A single dataset configuration package that consists of a protection policy, provisioning policies, resource pools, and an optional vFiler template (for vFiler unit creation). You can assign a single uniform storage service to datasets with common configuration requirements as an alternative to separately assigning the same protection policy, provisioning policies, and resource pools, and to setting up similar vFiler unit attachments to each of them. The only way to configure a dataset of virtual objects with secondary or tertiary backup and mirror protection and provisioning is by assignment of a storage service. You cannot configure secondary storage vFiler attachments for datasets of virtual objects.
Protection or provisioning related objects	Snapshot copies, primary volumes, secondary volumes, or secondary qtrees that are generated as a result of protection jobs or provisioning jobs.
Naming settings	Character strings and naming formats that are applied when naming related objects that are generated as a result of protection jobs or provisioning jobs.

Available functionalities of dataset and SnapVault integration with SnapManager

By using dataset and SnapVault integration with SnapManager, you can create, manage, verify and restore remote backup copies. You can select dataset policies, manage remote backup copies, perform temporary restore operations and remote backup integrity verification operations.

The following functionalities of dataset and SnapVault integration with SnapManager are available:

- You can create and restore remote backup copies.
- You can select policies related to the dataset created by Protection Manager.
- You can protect created datasets, by doing the following:
  - Creating remote backup copy on the SnapVault secondary.
  - Using topologies supported by SnapManager and Protection Manager.
- You can delete individual remote backup copies based on the backup version.
- You can display remote backup copies that are available for restore.
- You can perform a restore operation to a Recovery Storage Group using SnapVault remote backup technology.
- You can perform remote backup integrity verification.

Dataset and SnapVault integration with SnapManager

A dataset is a collection of storage sets. SnapManager integrates with datasets and SnapVault to archive these storage sets to secondary storage. By replicating Snapshot copies of the storage sets to secondary storage, SnapVault provides you with a centralized disk-based backup solution.

The dataset being a data management concept introduced in Protection Manager includes the backup copies and replica of the primary data and configuration information, along with data protection policies that determine how the data is protected. Datasets enable you to keep backup online for faster restore.

The following capabilities of Protection Manager make it a good option for integration with SnapManager:

- Automatic setting up of SnapVault relationships and complex replication topologies with resource pools
- Monitoring of data transfer
- Management of remote backup retentions

If Protection Manager is available and SnapDrive is configured for DataFabric Manager, SnapManager automatically becomes aware of the dataset. If Protection Manager is not available, SnapDrive informs SnapManager of its unavailability. SnapManager continues in normal mode, and does not support remote backup operations.

Protection Manager integration can be enabled without reinstalling SnapManager for Exchange anytime by rerunning the SnapManager Configuration wizard. However, you must first ensure that you have configured Snapdrive for Protection Manager integration. If you have not already done so, you can set up Snapdrive Protection Manager integration with the sdcli command, and then restart the SnapDrive service, or by rerunning the SnapDrive setup. (Refer to the SnapDrive Help for more information on Protection Manager integration.)

Dataset policies

Dataset policies control the protection of data in datasets. A policy defines characteristics such as the replication topology, backup retention, replication lag, and throttle.

Remote backup retention policies control the backup copies that are created at the remote site. The remote backup retention policies are controlled by Protection Manager, not SnapManager.

Datasets associate the LUNs that are used by an Exchange server to the related set of protection policies. The LUN association with the policies enables the administrator to protect the data through remote backup copy and to relate to the corresponding resource pool.

Prerequisites for dataset and SnapVault integration with SnapManager

Before you integrate datasets and SnapVault with SnapManager, ensure that you have all required software installed. Also ensure that you meet the storage system, license, SnapMirror, and LUN placement requirements.

Ensure that you have the following softwares installed:

- Protection Manager 3.7 or later
- NetApp Management Console 3.7 or later
- SnapDrive for Windows 6.0 or later
- Data ONTAP 7.3 or later

Note: You can upgrade SnapManager from an earlier version that did not support datasets to a later version that supports datasets. You can also revert to the older version without any adverse effects on the system.

Ensure that you have the following configuration requirements met:

- Two storage systems
  The primary storage system is the archiving source; the secondary storage system is the archiving destination. Both should have Data ONTAP 7.3 or later installed. One system should have the SnapVault primary license, and the other should have the SnapVault secondary license.
- All LUNS on qtrees
- Each LUN on its own qtree and should contain only a single LUN
- Protection Manager and NetApp Management Console software installed on a dedicated server other than the Exchange server that you are working with SnapDrive for Windows installed
- If you set up SnapMirror relationships using Protection Manager, the `snapmirror.conf` and `.allow` files have been updated manually, so that Business Continuance works.

Limitations of dataset and SnapVault integration with SnapManager

When you integrate SnapManager with datasets and SnapVault, you should understand the limitations regarding dataset configuration, retention policies, Business Continuance, and verification.

- You cannot perform remote backup and archiving operations without dataset configuration.
- You cannot control the archived backup retention policy by using SnapManager.
  Archived backup retention policies are controlled by Protection Manager.
- You cannot use SnapManager to create and manage datasets for disaster recovery or business continuance.

- SnapManager does not support LUNs residing on one qtree.
- SnapManager does not support LUNs that do not reside on a qtree.
- After you apply a dataset policy, you cannot change it using SnapManager.
  If you want to change a dataset policy, then change it using Protection Manager, which automatically updates the policy in SnapManager as well.

Dataset configuration

A storage set grouped with its configuration information makes a data set. Data sets associate the LUNs used by an Exchange server to the related set of protection policies. This enables the administrator to protect the data through remote backup and relate to the corresponding resource pool. One data set is created for each Exchange server on the server host.

You can create and configure datasets when you run the SnapManager Configuration wizard for the first time on a system with Protection Manager installed. If you upgrade SnapManager from an earlier version, re-run the Configuration wizard to configure a dataset.

You cannot configure datasets through the Database Availability Group (DAG). You can configure datasets on all member servers of the DAG. You must enable the dataset option on a member server during SnapDrive for Windows installation to create a dataset on that server.

You cannot change the names of the datasets. The following is the example for the naming convention for a dataset:

`SnapMgr_Exchange_exchange1`

For a server running on Microsoft Clustered Server, a virtual server is used to name the dataset. In the case of Continuous Cluster Replica (CCR) configuration, the Exchange server name is used along with the name of the node.

Dataset member information is a list of drive letters and mountpoints related to SnapManager. The information is stored and tracked by Protection Manager, and its information is retained even after SnapDrive is uninstalled. The member information is retained on all cluster nodes.

If you are using Microsoft Exchange 2003 or 2007, you can view the current dataset status for your server by selecting the server and then selecting **Backup** in the Scope pane. SnapManager gives you the dataset description and status, the protection policy description, the protection status, the conformance status, and lists the dataset members in the Results pane.

You can also view the dataset status from the Microsoft Management Console. When you click **SnapManager for Exchange**, you can see a list of servers. If datasets are configured for a server, MMC displays the dataset status. If datasets are not configured, the message
`Dataset not configured`
is displayed.

Note: If you are using Microsoft Exchange 2010, the dataset status will not be accurate.

Creating a dataset using SnapManager

When you run the Configuration wizard for the first time with Protection Manager configured, SnapManager creates a dataset for each server internally that you configure with SnapManager using the Configuration wizard.

Before you begin

Ensure that you are assigned an administrator role that enables you to create a dataset.

Ensure that you go through documentation on automatic e-mail notification settings before you start the configuration.

About this task

For configuring archived database by using Protection Manager with SnapManager, SnapManager currently allows only the Backup and Remote backups only policies to be implemented. Back up is selected as the default option. You can change the policy by using the NetApp Management Console.

For a CCR configuration you can create two datasets, one for each CCR node.

Steps

1. If you have datasets configured in your system, in the **Configure Dataset for Backup Archival** page of the Configuration wizard, select the Storage Groups that you want to add to your dataset.
2. Choose the default protection policy as **Back up** or **Remote backups only** as per your requirement.
3. In the **Add Microsoft iSCSI Service Dependency** window, select **Yes, add the..** if you want to add iSCSI service as a dependency for the MSExchangeSA service.
4. Use the **Configure Automatic Event Notification** page to configure the automatic event notification options for SnapManager.
5. Review the configuration summary in the **Completing the Configuration Wizard** page, and click **Finish**.
6. Click **Start Now** to migrate your databases and their transaction logs and SnapInfo files to the LUNs you specified.

   Note: If you are moving the location of Exchange databases in this step, it could take some time to complete.

7. Click **OK**.

After you finish

Edit the dataset that you have created.

Editing a dataset using Protection Manager

It is mandatory that you use Protection Manager to add individual physical resources to the dataset (that you created with Protection Manager) after you run the SnapManager Configuration wizard. Only after adding resources you can start archiving the database. SnapManager cannot associate a resource pool to the dataset.

Before you begin

After the dataset is created, use Protection Manager to check the Protection status and the Conformance status of the dataset.

Steps

1. In the Scope pane, click **Dataset**.
2. Select the dataset.
3. In the **Dataset** window, click **Edit**.
4. In the **Edit Dataset** window, click **Physical Resources** under **Backup**.
5. From the list of available resources, add the resources you want in the dataset.
6. Click **Next**.
7. Click **Finish**.
8. In the **Edit Dataset** window, click **Provisioning/Resource Pools** under **Backup**.
9. From the list of available resource pools, assign the resource pools for the archive backup.
10. Click **Next**.
11. Click **Finish**.

SnapVault relationships

After you create the dataset, determine its policies, and add secondary resource pools to the dataset, Protection Manager creates SnapVault relationships for archiving. You cannot restore a remote backup if you change or modify the SnapVault relationship.

If you upgrade from earlier versions of SnapManager to SnapManager 5.0 and later, and archive the backup database by using SnapVault, import the SnapVault relationships that exist to the dataset. If you do not import the SnapVault relationship, a new one is created. For more information, see the relevant Protection Manager documentation.

Related information

Using SnapVault to Archive SnapManager for Exchange Backups Sets

Local backup protection using dataset and SnapVault integration

SnapManager uses datasets to create remote backup copies of local data residing in the primary Storage Group. You must schedule your backup operations such that their transfer to SnapVault secondary storage does not get affected.

Ensure that you meet the following conditions before SnapManager starts creating remote backup copies:

1. A dataset exists.
2. The **Archive local backup using SnapVault** option is enabled.
3. The dataset protection status is Protected and conformance status as Conformant.
   If the configuration contains non-Exchange LUNs, the qtrees containing the non-Exchange database are not updated during archiving. This changes the dataset protection status to Lag Warning or Lag Error.

For more information, see the relevant Protection Manager documentation.

Scheduling backup operations according to SnapVault transfers

You must schedule your backup operations so that they cannot be deleted before their SnapVault transfer to the secondary storage is complete.

You configure a backup operation so that only two backup copies can be retained on the primary storage. You schedule the backup operations at 12 PM, 1 PM, 2 PM, and 3 PM. The backup operations take three hours before their SnapVault transfer to the secondary storage is complete. When the 3 PM backup operation is running, the 12 PM backup operation is still transferring. Hence SnapManager deletes the 1 PM backup, and the Protection Manager job fails as it fails to find the deleted 1 PM backup to transfer. You can avoid this by using number of days for local backup retention instead of using number of backup copies for local backup retention.

Information used to create remote backups

When configured for archiving, SnapManager creates remote backups after it creates local backups. SnapManager sends certain information like the local backup version, the local backup management group, and the list of backup LUNs to SnapDrive before the remote backup operation can take place.

- The version number of the backup
  The version number acts as both the timestamp and locator for the backup copy during a restore operation.
- The type of backup management group

  - Local management group

    The local management groups can be standard, daily, and weekly.
  - Remote management group

    The remote management groups can be hourly, daily, weekly, monthly, and unlimited. The default management group is daily.

    If you select the hourly management group for remote backup, SnapManager shows a message that conveys that hourly archived backups are deleted the next time the Protection Manager monitoring service runs.
- A list of LUNs with their corresponding Snapshot names

You can defer remote backup to a time after the local backup is created.

SnapManager uses the unique backup naming convention for archives created on secondary storage. In the Backup wizard, if you configure a dataset and the archival process starts, SnapManager does not change the generic backup naming convention to the unique backup naming convention. If you select to keep generic naming, no archives are created.

If the backup naming convention is generic and you archive the backup, SnapManager provides options to either continue to archive to secondary storage with the unique naming convention, or to disable archiving and continue the local backup operations with generic naming convention.

Remote backup retention

Remote backup retention refers to the duration of time that the remote backups that are retained at a secondary Storage Group. SnapManager and Protection Manager work together to retain or delete your backups, based on the protection policy that you select for a dataset.

You can specify the number of local backups to retain by using the backup management groups. Protection Manager controls remote backup retention. When SnapManager deletes a backup, it deletes the metadata only, after confirming with Protection Manager that the archive backup has also been deleted.

New backups are continuously created. If the number of backups or the duration of retention exceeds the management group setup, the policy deletes the oldest backup on the secondary storage.

Deferred database integrity verification with SnapVault

You can defer integrity verification of either the local or remote backup copies.

Deferred integrity verification runs on the management group that you selected when you created the backup. You cannot change the remote backup management group after you create the remote backup copy. You can perform deferred verification on the SnapVault secondary Storage Group from either the local application server or the remote verification server.

Note: For deferred verification of the remote backup copies, verification is based only on the Remote retention type. The number of backups that you select with the specified retention type is verified.

Restoring from a remote backup

The process for restoring from a remote backup copy is almost the same as that for restoring a local backup copy, except that the remote backup copy needs to be restored from the archived backup.

Before you begin

Prior to restoring from a remote backup copy, you must enter the following information on the remote storage system:

```
options snapvault.access <filename>
```

For more information on the `options snapvault.access` command, see the Data ONTAP documentation.

Steps

1. Select the Exchange server from the Scope pane.
2. If you don't have direct access to the SnapVault secondary, you must cancel the backup verification so that the remote backup will not mount on your local machine before the restore. To cancel the backup verification, complete the following steps:
   a) Click **Restore**.
   b) Click the **Advanced options** tab.
   c) Uncheck the **Verify log sequence and database metadata before restore** checkbox, then click **Okay**.
   d) Navigate to the Backup Verification Setting tab.
   e) Uncheck the Override database verification requirements for restore checkbox, then click **Okay**.
3. In the **Actions** pane, select **Restore wizard**.
4. In the **Which Exchange Server Created the Backups** page, select **Restore from unmanaged media**.
5. Continue with the instructions provided in the Restore wizard.
6. Click **Finish**.

SnapManager application settings configuration

You can configure or change SnapManager application settings at any time after you install SnapManager, as long as you run SnapManager from the system console and not from a Terminal Services client.

Where to access SnapManager application settings

Using the SnapManager user interface, you can easily configure and change SnapManager application settings at any time after you install SnapManager. There are multiple ways with which you can access your application settings.

From the Configuration wizard, you can access settings for only the verification server. To access other verification settings (Override Verification, Throttling, and Access LUN in Snapshot copy), you must open the **Verification Settings** dialog box.

When you open the **Run Command After Operation** dialog box from the Actions pane, you can view or configure only the default settings. However, from within the context of a specific operation, the default settings are presented and then can be modified for this operation only. As an option, the default settings can be updated.

The following table shows the GUI components that you can use to configure SnapManager application settings.

Application settings	Where the setting can be accessed
Add servers to be managed	Actions pane Configuration wizard
Exchange user account	Actions pane
Migrate databases to local disk	Actions pane
Database verification settings • Verification server • Override for restore operations • Verification throttling • Mountpoint	Actions pane Backup wizard Restore wizard Configuration wizard

Application settings	Where the setting can be accessed
Backup settings	Actions pane Backup wizard Backup and Verification
Run Command After Operation default settings	Actions pane Within the context of a backup or database verification operation: Backup wizard, or Backup and Verify window
Fractional space reservation policy settings • Current status • Policy settings	Actions pane
Report directory settings	Actions pane
Event notification settings • e-mail notification • Logging	Configuration wizard Actions pane

Adding Exchange servers to be managed

You can manage one Exchange server or multiple Exchange servers with SnapManager using the **Add Servers to be Managed** option. You cannot run SnapManager until you successfully add an Exchange server. You can also manage multiple servers.

Steps

1. In the **Actions** pane, select **Add Servers to be Managed**.
2. Type the name of or browse to the Exchange server that you want to manage.

   This setting remains in effect, specifying the default Exchange server, until, or unless you change it.

Result

Whenever the SnapManager program starts, SnapManager automatically connects to the default Exchange server by using the default security authentication method.

Note: If you want to manage a different Exchange server later, use the **Add Servers to be Managed** option to connect and manage an Exchange server.

Enabling database migration back to local disks

Once you have migrated your database for SnapManager configuration, you can enable migration back to local disks.

About this task

If you enable the **Enable databases and transaction logs to be migrated back to local disk** option, SnapManager disables it the next time you start SnapManager.

Steps

1. In the **Actions** pane, click **Configuration Wizard**.
2. Select the **Enable databases and transaction logs to be migrated back to local disk** option to enable the migrate-back-to-local-disk feature.
3. Click **OK**.

Disabling database migration back to local disks

You can disable the migration of your database back to local disks after you have configured SnapManager.

Steps

1. In the **Actions** pane, click **Configuration Wizard**.
2. Select the **Enable databases and transaction logs to be migrated back to local disk** option to disable the migrate-back-to-local-disk feature.
3. Click **OK**.

Considerations for selecting the database verification server

You can manage verification your backup sets on either your production Exchange server or on a remote verification server, and you can use any of deferred verification, remote verification, and verification throttling to do so.

Verification of databases in a backup set can be done on the production Exchange server (the Windows host system running the Exchange server used to service the users) or on a remote verification system (a different Exchange server with Exchange management tools installed). Running database verification on a production Exchange server is CPU-intensive for the Windows host and also involves a substantial amount of activity on the storage system.

Note: When you change the database verification server, this change does not affect any database backup (with verification) or database verification-only jobs that are already scheduled. You must re-create the scheduled backup jobs for the change to apply.

SnapManager offers three methods for managing database verification load that you can use separately or in any combination:

- Deferred verification
- Remote verification
- Verification throttling

Both the deferred verification and remote verification manage database verification load by separating database verification from the backup operation. Verification throttling manages database verification load by slowing down the verification throughput rate.

Configuring the verification server

After you have configured SnapManager, you can configure your verification server to be the same as your host server or a remote server. Configuring a remote server as the verification server reduces the load on the host server. You can optionally configure the verification server after you have completed SnapManager configuration.

Before you begin

You must have SnapDrive, SnapManager, and Exchange installed on your verification server. If you configure a remote server as your verification server, the versions of SnapManager and SnapDrive must be the same on both the host and the remote server.

You can also use the following path to configure the verification server: **Actions > Backup Verification Settings > Verification Server**.

About this task

You can omit the configuration of the verification server during configuration by selecting the **Select the verification server later...** check box in the Configuration wizard.

If Exchange is not installed on the computer that you select, specify the `eseutil.exe` filepath and copy the necessary files to the computer before proceeding.

Steps

1. In the **Actions** pane, click **Configuration wizard**.
2. In the **Database Verification Server** window, enter or browse to the name of the server that you want to use as the verification server.

Remote verification prerequisites

Before you use a remote verification server, ensure that you establish the Windows host requirements, LUN requirements, and verification server designation.

Remote verification uses the same mechanisms as local verification, except that the verification occurs on a different host than the one that initiated the backup operation. This is why you need SnapDrive and SnapManager installed on your remote verification server, in addition to FC or iSCSI connectivity to the storage system.

If you use a remote verification server to verify the databases in multiple databases in a single job, you need an additional LUN.

You can designate your verification server from the production Exchange server.

How remote verification works

SnapManager initiates a backup operation at the primary host, which then contacts the remote verification server. The remote verification server then uses SnapDrive for verification and sends the results back to the primary host.

The basic steps of this process are as follows:

1. The SnapManager backup with verification (or verification only) is initiated on the primary SnapManager host, which is configured to run verifications on the remote verification server.
2. The primary SnapManager host contacts the remote verification server and initiates the verification job.
3. The remote verification server uses SnapDrive to connect to a LUN backed by a Snapshot copy containing the databases to be verified.
4. The remote verification server performs the database verification on the LUN backed by Snapshot copy.
5. When the remote verification server completes the verification, it sends the results back to the primary SnapManager host.

Note: You must be careful not to schedule backup operations while a verification operation is in progress. Verification is always performed on a LUN that is backed by a Snapshot copy. If you make a Snapshot copy of the same volume while a LUN backed by Snapshot copy exists, you create a “busy Snapshot copy,” which might cause problems when you attempt to delete some Snapshot copies.

Viewing or changing the verification server

You can view or change your verification server for verification load management.

Before you begin

You need to connect to the production Exchange server to view or to change the verification server.

About this task

Until you specify verification settings, database verification is run from the Exchange server you selected. Verification does not necessarily run on the system from which you opened the **Database Verification Settings** dialog box. A change in the verification server does not affect any database verification jobs that are already scheduled.

Steps

1. Click **Backup Verification Settings** in the **Actions** pane.

   You can also use these paths:

   - **Backup wizard > Verification Settings**
   - **Restore wizard > Verification Settings**

   The Verification Server tab is active by default and displays the host name of the current verification server.

2. In the Verification Server box, type or browse to the Exchange server you want to use as the database verification server.

   Note: If you plan to specify a remote verification server, ensure that the server is set up correctly.

3. To make the database verification server verify backups from both Windows Exchange Server 2003 and Windows Exchange Server 2007, select the **Verification server will verify both Exchange Server 2003 and 2007 backups** check box.

4. Click **OK**.

Selecting the Snapshot copy access method for database verification

Use the **Access LUN in Snapshot** tab to specify how SnapManager should access database backup Snapshot copies during database integrity verification. Assign either a drive letter or directory path to access the backup Snapshot copy as a mounted LUN.

Steps

1. Click the **Access LUN in Snapshot** tab.

   You can access the **Access LUN in Snapshot** tab from the Verification Settings window of the Configuration wizard, the Backup Verification Settings window, the Backup wizard or the Restore wizard.

2. Assign either a drive letter or directory path to access the backup Snapshot copy as a mounted LUN.

If you want to...	Then do this...
Mount the Snapshot copy on the next available drive letter	Select **Automatically assign available drive letter**.
Mount the Snapshot copy on a specific NTFS mountpoint	Do the following: **a.** Select the **Mount in an empty NTFS directory** option. **b.** Enter or browse to the directory path of an NTFS mountpoint. **Note:** This mountpoint is used if SnapManager is configured to use drive letters but runs out of available drive letters.

3. Click **OK**.

Database verification throttling

SnapManager allows you to throttle the database checksum verification rate with Exchange Server 2003 SP2 or later, or Exchange Server 2007. Database verification throttling enables you to manage your verification load.

How database verification throttling works

`Eseutil.exe`, the Microsoft Exchange consistency checker utility, inserts a one-second pause after a given number of input output (I/O) operations during the database physical consistency verification. `ChkSgFiles`, the Microsoft Exchange integrity verification library, inserts a one-second pause after a given number of I/O operations during the database physical consistency verification.

`Eseutil.exe` and `ChkSgFiles` read 512 KB for each database checksum verification I/O operation. Therefore, when configuring the throttling value, the maximum throughput rate for database checksum verification is decreased to the following:

512 KB per I/O × *x* I/Os per second = 512 × *x* KBps

You can decrease the maximum throughput rate by decreasing the number of input output operations (*x*) that elapse between one-second pauses.

Number of I/O operations (*x*) between one-second pauses	Maximum possible database verification speed	
	Calculation	**Maximum speed**
100	512 KB/IO × 100 IO/sec = 51,200 KBps	50 MBps
150	512 KB/IO × 150 IO/sec = 76,800 KBps	75 MBps
200	512 KB/IO × 200 IO/sec = 102,400 KBps	100 MBps

250	512 KB/IO × 250 IO/sec = 128,000 KBps	125 MBps

Note: Decreasing the `Eseutil.exe` and `ChkSgFiles` database verification throughput causes the database checksum verification to take longer to complete. This also means that a backup job that is configured with verification also takes longer to complete. However, decreasing verification throughput does not cause a backup job without verification configured to take any longer to complete.

Database verification throttling options

SnapManager 5.0 and later for Microsoft Exchange supports two throttling options: Eseutil with Microsoft Exchange Server 2003, and ChkSgFiles with Microsoft Exchange Server 2007.

Attention: Attempting to use the Eseutil throttling feature with earlier versions of SnapManager 3.2 or Exchange has no effect.

Calculating the verification throttling sleep interval

You must calculate the appropriate sleep interval value to use for each server. Monitor the read and write performance of the LUNs on which the databases reside, and then calculate the verification throttle setting so that the average and peak physical disk performance is below the maximum values determined by Microsoft.

About this task

- To monitor the SnapManager progress, watch the Backup Status window that is displayed during a manually launched backup job.
- Monitor the counters during the busiest time of day that a verification or backup with verification job might occur.

Steps

1. Launch the Windows Performance Monitor utility (also known as “Perfmon”) on the machine that is the Exchange server.
2. Add the Perfmon counters PhysicalDisk\Average Disk sec/Read, and PhysicalDisk\Average Disk sec/Write for each LUN that contains a database.
3. Monitor the pairs of values, Physical Disk\Average Disk sec/Read and PhysicalDisk\Average Disk sec/Write, while a SnapManager verification job or a backup job with verification is in progress.

   If monitoring a backup with verification job, be sure to monitor the counters during the verification portion of the job. Multiple databases on the same Exchange server are verified sequentially, so be sure to monitor the counters for all of the database verifications.

4. The throttle setting is global, so select the reading from the database that exhibits the worst performance (highest readings).
5. Compare the measured values to the maximum read and write values recommended by Microsoft.

   Physical Disk\Average Disk sec/Read: Average read time should be below 20 ms, Maximum peak read time should be below 50 ms.

   Physical Disk\Average Disk sec/Write: The maximum physical disk performance times determined by Microsoft are as follows: Average write time should be below 20 ms, Maximum peak write time should be below 50 ms.

   Compare the measured values to the maximum read and write values recommended by Microsoft: less than 20 ms. average time and less than 50 ms. peak time.
6. If the measured values exceed the maximum values recommended by Microsoft, adjust the throttle setting to decrease the verification throughput rate.

Configuring database verification throttling

You have to configure database verification throttling on all of the Windows host systems that are to perform database verification and is running Exchange Server.

Before you begin

Throttling is accessible only if the verification server is installed with Microsoft Exchange Server 2003 SP2 or later, or Microsoft Exchange Server 2007.

Steps

1. From the SnapManager console, click **Backup Verification Settings** in the **Actions** pane.

   You can also use these paths:

   - **Backup Wizard > Backup or Verify Databases and Transaction Logs > Verify databases and transaction Logs > Backup management group > Database Verification Server > Verification Settings**
   - **Restore wizard > Mount options > Verification Settings**
2. Click the **Verification Throttling** tab.
3. Select or clear the **Throttle database checksum verification** check box to enable or disable the throttling feature.
4. If you enabled the throttling feature, enter a positive integer value in the **Pause for 1 second after x I/O operations** box.

   The default value is 150. Enter the number Input Output (I/O) operations to complete before pausing for one second during checking of the database physical consistency.

   Using a value in the range of 100 to 250 suits most environments.
5. Click **OK**.

Throttling entries in the SnapManager backup and verification report

If you enable either Eseutil throttling or ChkSgFiles throttling, the SnapManager Backup and verification report logs messages indicating that the throttling feature is enabled before each database verification operation.

If you use Eseutil throttling, before each database verification, the SnapManager backup and verification report includes the following entry:

```
ESEUTIL throttling feature is enabled.
```

If you use ChkSgFiles throttling, before each database verification, the SnapManager and verification report includes the following entry:

```
Running Integrity Verification using ChkSgFiles API Throttle (Pause): 1000
ms per X I/O's
```

Verification override entry in the SnapManager restore report

You must restore only from verified database to ensure a successful restore operation. Before each database verification, the SnapManager restore report indicates if verification override is enabled. If the feature is enabled, the report includes the warning "Database verification before restore was overridden."

Impact of database verification on performance

Database verification can impact the performance of both the Exchange server and the storage system. To overcome the performance impact, you can separate the database verification process and database backup operations.

Verification can degrade Exchange Server response, particularly during peak work hours. With SnapManager 3.1 and earlier releases, there are two options for distributing this load by separating database verification and database backup operations: deferred database verification and remote database verification.

SnapManager 3.2 introduced database verification throttling, which you can combine with either or both of the existing options to reduce the load even further. When installed with Microsoft Exchange Server 2003 SP2 or later, SnapManager 3.2 and later supports the ability to throttle Eseutil (the Microsoft Exchange consistency checker utility) in database checksum verification mode.

SnapManager 5.0 and later for Microsoft Exchange supports a new option for managing database verification that uses the same database checksum algorithm as Eseutil throttling. When installed with Microsoft Exchange Server 2007, SnapManager 5.0 and later supports the ability to throttle ChkSgFiles (the Microsoft Exchange integrity verification library) in database checksum verification mode. In Exchange Server 2003, ChkSgFiles is installed only if the Exchange 2007 management tool is installed.

Database verification override during restore operation

You must not select an unverified backup copy as the source of the restore operation. If a verified backup copy is not available when you need to restore and you cannot wait for a verification to complete before restoring, you can override verification and restore from an unverified backup copy.

Configuring the database verification override option

You can configure SnapManager to override the verification requirement and restore directly from an unverified backup copy.

Steps

1. In the **Actions** pane, click **Backup Verification Settings > Verification Settings**.

   Other ways to open Verification Settings are as follows:

   - From the Backup Wizard, go to the **Verify the Databases and Transaction Logs in this Backup** window and click **Verification Settings**.
   - From the Restore Wizard, go to the **Verify the Database Integrity in this Backup** window (displayed only if you select to restore from an unverified backup copy) and click **Verification Settings**.

2. Select the **Override Verification** tab.
3. Configure the **Override Database Verification Requirement for Restore** option.

   Note: If you enable this option, it is reset the next time you start the SnapManager application.

4. A message appears that says "An unverified backup may contain an image of an Exchange database that is physically corrupt. If the backup copy contains physical database corruption, this corruption will persist in the restored database." If you want to enable the **Override Database Verification Requirement for Restore** option, click **OK**; otherwise, click **No**.
5. Click **OK** to close the dialog box and apply your change.

   Note: If you are using the Restore wizard, by selecting the Override Database Verification Requirement for Restore option causes the Verify the database integrity of this backup prior to restoring it option to become accessible so that you can disable it.

If...	Then...
The override option is disabled (the default and recommended setting)	You cannot select an unverified backup copy as the source of the restore operation.

If...	Then...
The override option is enabled (not recommended)	If you select an unverified backup copy as the source of the restore operation, you are asked to proceed with the restore operation. **Note:** This option is reset when you exit SnapManager.

Verification override entry in the SnapManager restore report

You must restore only from verified database to ensure a successful restore operation. Before each database verification, the SnapManager restore report indicates if verification override is enabled. If the feature is enabled, the report includes the warning "Database verification before restore was overridden."

Verification override entry in the SnapManager restore report

You must restore only from verified database to ensure a successful restore operation. Before each database verification, the SnapManager restore report indicates if verification override is enabled. If the feature is enabled, the report includes the warning "Database verification before restore was overridden."

Configuring default settings for the Run Command After Operation option

You can configure default values to populate the **Run Command After Operation** dialog box when you open it from the Actions pane, or the Backup and Verify window, or the Backup wizard, to run a command or a script after a backup or verification operation.

Steps

1. In the **Actions** pane, click **Run Command After Operation**.
2. In the **Specify a computer where...** box, enter or browse to the name of the host on which your program or script resides.
3. In the **Specify the full path...** box, browse to your program or script.
4. Enter the command input string in the **Command Arguments** box.

   You can do this using any combination of the following methods:

   - To enter text directly into the **Command Arguments** box, click the box and type the desired text.
   - To enter a SnapManager variable into the **Command Arguments** box, do the following:

a. Click the **Command Arguments** box to position the cursor.

b. In the **SnapManager Variables** list, select the variable you want to enter.

c. Click **Select**.

Note: The $*SnapInfoPath* variable is enclosed within double quotes so that the path name can contain spaces without affecting the script invocation on the Windows command line. If you do not want the double quotes to appear in your command line, remove them from the Command Arguments box.

5. Repeat steps 1 to 4 as required, until the **Command Arguments** box contains the arguments you want to pass to your program or script.

6. Click **OK**.

   Note: SnapManager verifies that the specified program exists on your system. SnapManager does not run the command until the backup or verification operation is complete.

Whenever the **Run Command After Operation** dialog box is opened from either the Backup and Verify window or the Backup wizard, the boxes are populated with the values you specified as default settings.

Fractional space reservation

When you create a LUN, Data ONTAP reserves space in the volume containing that LUN so that write operations to that LUN do not fail due to lack of disk space. With fractional reserve, this space is set to less than 100 percent of the total size of the LUNs.

SnapDrive creates and manages LUNs with space reservation enabled. Operations such as creating a Snapshot copy or creating new LUNs can occur only if there is enough available unreserved space. These operations are restricted from using reserved space.

While space reservation is enabled at the LUN level, fractional overwrite reserve amounts are configured at the volume level; that is, fractional space reservation does not control how the total amount of space reserved for overwrites in a volume is applied to individual LUNs in that volume.

The volume has the guarantee option set to `volume` rather than `file`. Fractional reserve is supported by Data ONTAP 7.1 or later. For more detailed information, see the *Data ONTAP Block Access Management Guide for iSCSI and FC* for Data ONTAP 7.1 or later.

Additional space that is not space-reservation-enabled on the volume is automatically reserved for overwriting blocks that belong to a LUN. By default this additional space is equal to 100 percent of the total size of all space-reserved LUNs in the volume. If space reservation is disabled, write operations to a LUN might fail due to insufficient disk space in the volume and the host application might terminate, report I/O errors, or experience unexpected behavior.

With fractional reserve, the space reserved for overwrites is set to less than 100 percent and the space that is preallocated for space reservation is reduced to that percentage. Fractional reserve is generally used for volumes with LUNs that store data with a low rate of change.

What can happen with a fractional-space-reserved volume

If a fractional-space-reserved volume runs out of overwrite reserve space, write operations to a LUN fail and the host application might terminate, report I/O errors, or exhibit unexpected behavior. Data ONTAP uses automatic expansion of flexible volumes and automatic deletion of Snapshot copies from flexible volumes to avoid this situation.

When a LUN is fully space reserved, write operations cannot fail due to an out-of-space condition. When the overwrite reserve for a volume is set to less than 100 percent, write operations to the LUNs on that volume might fail when the volume runs low in free disk space.

The automatic expansion of flexible volumes and the automatic deletion of Snapshot copies from flexible volumes monitor the reserved space and take action if the free space becomes scarce. For more detailed information, see the *Data ONTAP Block Access Management Guide for iSCSI and FC* for Data ONTAP 7.1 or later.

Automatic expansion of flexible volumes

Data ONTAP automatically expands a nearly full volume into the space preallocated for it in the aggregate. The volume must be a flexible volume with the guarantee option set to **Volume**. You can enable automatic deletion of Snapshot copies and FlexVol expansion features separately or together, with one policy to be applied before the other. When fractional-space-reserved volumes hold LUNs that store Exchange database files, however, you can only use only the automatic FlexVol expansion feature.

Automatic deletion of Snapshot copies from flexible volumes

Data ONTAP automatically deletes one or more Snapshot copies on a nearly full volume, when you enable the Snapshot copy automatic deletion policy. If the trigger condition is detected, the oldest or newest Snapshot copies are deleted until a configured percentage of the volume is free space. If you do not want to automatically delete Snapshot copies on the volume, you can set the overwrite reserve to 100 percent, by setting the fractional space reserve to 100 percent on the storage system.

This Data ONTAP feature is not designed specifically to support backup and restore operations on Exchange databases. The options for selecting Snapshot copies to be deleted do not have visibility to the automatic backup Snapshot copy deletion criteria configured in SnapManager. You must always retain at least one online backup copy for each database.

Fractional space reservation policies

Fractional space reservation policies include specific thresholds that determine when SnapManager must delete Exchange backup sets or unmount Exchange databases (or both), because the overwrite reserve utilization for the volume is running low.

If overwrite reserve space runs low for a fractional space-reserved volume, SnapManager prevents the overwrite reserve from becoming fully depleted.

The default fractional space reservation policy

SnapManager sets a default policy for fractional space reservation. You can use the default values, or change the values that would apply to all storage system volumes.

The default fractional space reservation policy is automatically enabled for any traditional or flexible volume that has an overwrite reserve that is set to less than 100 percent. The volume must also contain LUNs that store Exchange database files, Exchange transaction log files, or SnapManager SnapInfo directories.

Default policy	You can use the default policy as-is, allowing the factory default values to be applied to all volumes that contain fractional space-reserved LUNs.
Default policy with customized settings	Optionally, you can customize the default policy that is applied to all storage system volume that contains fractional space-reserved LUNs.
Volume-specific policies	Optionally, you can override the default policy for any particular volume that contains fractional-space-reserved LUNs, by applying a custom policy.

Fractional space reservation policy settings

The fractional space reservation policy settings enable you to indicate when SnapManager should begin automatically deleting Snapshot copies and unmounting Exchange databases due to overwrite reserve utilization. You can also specify how many Snapshot copies to retain.

Enabling automatic deletion of Exchange backup Snapshot copies does not necessarily prevent an out-of-space condition on the volume. Therefore, database unmounting is always enabled. If Snapshot copy deletion is enabled, you must configure it to trigger before unmounting the database.

Deletion of Exchange backup Snapshot copies

SnapManager fractional space reservation policy setting	Factory default value	Configurable values
	Status: enabled	Status: enabled or disabled
Trigger on overwrite reserve utilization	70%	1% through 99%
Number of Snapshot copies to retain	5	1 through 256

Unmounting of Exchange databases

SnapManager fractional space reservation policy setting	Factory default value	Configurable values
	Status: enabled	Status: enabled
Trigger on overwrite reserve utilization	90%	1% through 99%

Configuring fractional space reservation policies

You can enable fractional space reservation and also set the value of the fractional space reservation. With fractional reserve, the space reserved for overwrites is set to less than 100 percent of the total size of the space-reserved LUNs in a traditional volume or a flexible volume.

About this task

- Although automatic deletion of Exchange backup Snapshot copies does not necessarily prevent an out-of-space condition on the volume, it is recommended that the automatic deletion of backups be enabled for every volume that contains fractional-space-reserved LUNs that store Exchange data.
- Data ONTAP includes a separate Snapshot copy autodelete feature. The SnapManager autodelete feature can be used in place of or along with the Data ONTAP autodelete feature.

Note: You cannot configure fractional space reservation policies through the Database Availability Group (DAG). You can configure them only when you connect to a member server of the DAG.

Steps

1. In the **Actions** pane, select **Fractional Space Reservation Settings**.
2. Click the **Policy Settings** tab.
3. Specify which policy you want to view or change.

If you want to access this...	Then do this...
The default policy	In the navigation tree, select **Default Policy**.
A volume-specific policy	In the navigation tree, select the storage system and then the volume.

4. To enable fractional space reservation monitoring, select the **Enable Fractional Space Reservation Monitoring** check box.
5. Use **Automatically delete backup sets** to enable or disable automatic deletion of Exchange backup Snapshot copies in fractional-space-reserved LUNs on the volume.

If you want to...	Then...
Enable automatic deletion of Exchange backup Snapshot copies	Select **Delete backups that include LUNs which have less than 100% overwrite reservation**, and then skip ahead to Step 8.
Disable automatic deletion of Exchange backup Snapshot copies	Clear **Delete backups that include LUNs which have less than 100% overwrite reservation**, and then proceed to Step 6.

6. In the **Trigger point for overwrite reserve utilization** field, type the level of overwrite reserve utilization (in percentage of total reserve) that is to trigger deletion of Exchange backup Snapshot copies.

The value must not be a negative integer that is less than the **Trigger point for overwrite reserve utilization** value in the **Automatic dismount of databases** panel.

7. In the **Number of most recent backup sets to retain** field, type the number of backup sets to be retained if automatic backup set deletion is triggered.

   The value must be an integer from 1 through 256 and should be based on the backup creation and verification schedule.

8. Use the **Automatically dismount databases** panel to configure automatic unmounting of Exchange databases in fraction-space-reserved LUNs on the volume.

   Note: Because automatic deletion of Exchange backup Snapshot copies does not necessarily prevent an out-of-space condition on the volume, SnapManager does not allow you to disable unmounting of databases for any fractional space reservation policy.

9. In the **Trigger point for overwrite reserve utilization** field, type the level of overwrite reserve utilization (in percentage of total reserve) that is to trigger unmounting of Exchange databases.

   The value must be an integer from 0 through 99.

   Note: If Snapshot copy autodelete is enabled, SnapManager requires that this threshold be set to a higher level than the threshold that triggers automatic Snapshot copy deletion. Setting the threshold at a higher value ensures that Snapshot copy autodelete is triggered first.

10. Click **OK**.

Fractional space reservation policies to manage Exchange data

Fractional space reservation policies enable you to monitor overwrite reserve utilization on fractional space-reserved volumes that contain your Exchange data.

If you store Exchange data on LUNs in a fractional space-reserved volume in a SnapManager environment, you need to avoid an out-of-space condition on the volume such that you have explicit or implicit Exchange-aware control over the deletion of Exchange backup set components.

To address this need, SnapManager provides its own space management tool for monitoring overwrite reserve utilization on the volumes. If overwrite reserve space runs low for a fractional space-reserved volume, SnapManager can take action to prevent the overwrite reserve from becoming fully depleted.

Fractional space reservation policies include specific thresholds that act as trigger points. SnapManager can delete Exchange backup sets or unmount Exchange databases (or both) when the overwrite reserve utilization for the volume reaches the trigger point.

Note: If you enable SnapManager e-mail notification, SnapManager sends SMTP e-mail message after an event of SnapManager fractional space reservation policy is complete.

Automatic unmounting of Exchange databases

Automatic unmounting is triggered if overwrite reserve utilization on the volume reaches the threshold specified by the fractional space reservation policy. SnapManager automatically unmounts the databases and stops the write operations to LUNs in that volume.

The threshold for overwrite reserve is specified by the policy for fractional space reservation. Another component of the fractional space reservation policy is the last-resort action that prevents further consumption of overwrite reserve; hence, automatic unmounting is always enabled.

SnapManager first uses backup set deletion to free some overwrite reserve. If this is not enough, unmounting the affected database prevents further consumption of overwrite reserve. This happens when both components of a fractional space reservation policy are enabled; and the unmounting of databases is triggered at a later level of overwrite reserve utilization than the level that is used to trigger the deletion of Exchange backup Snapshot copies.

Attention: If another host or client continues to write data to the affected volume, the overwrite reserve space might still run out and the volume goes offline. For this reason, you must use dedicated volumes for Exchange data.

Automatic deletion of Exchange backup copies

You can enable automatic deletion of LUN backup copies that store Exchange data. SnapManager checks for the level of overwrite utilization on the volume and triggers automatic deletion of backups if the level of overwrite utilization reaches the threshold.

Automatic deletion of LUN backup copies serves as the Exchange-aware replacement for or adjunct to the feature of Data ONTAP Snapshot copy deletion. SnapManager follows the following sequence of steps:

1. Deletes the oldest Snapshot copies.
2. Retains the specified number of total Snapshot copies on the volume.
3. Retains the most recent backup of any database (if it resides on the volume).
4. Retains any backup copies of databases no longer in existence.

You must select the backup retention level based on your backup copy creation and verification schedule. It is important that at least one verified backup copy remains on the volume if Snapshot copy deletion is triggered. Due to its Exchange-aware features, the automatic deletion of Snapshot copies does not necessarily prevent an out-of-space condition on the volume.

You must set the same number of backup sets to be deleted on database LUNs and transaction log LUNs. If there is a mismatch in this number, SnapManager attempts to delete backup sets based on the fractional reserve policy settings.

Example: automatic deletion of backup sets that span multiple volumes

If you have a backup copy that spans multiple volumes, with a different automatic deletion threshold configured on each volume, then for a specific volume, SnapManager deletes Snapshot copies based on the policy for that volume.

In this example, the automatic deletion settings for each volume are configured to take the following actions:

- Volume 1: Delete all but 2 Snapshot copies if 20 percent overwrite reserve utilization is exceeded.
- Volume 2: Delete all but 5 Snapshot copies if 20 percent overwrite reserve utilization is exceeded.
- Volume 3: Delete all but 10 Snapshot copies if 20 percent overwrite reserve utilization is exceeded.

If the 20 percent overwrite reserve utilization threshold for Volume 1 is exceeded, SnapManager deletes all but two Snapshot copies, regardless of the policies for Volumes 2 and 3. If the 20 percent overwrite reserve utilization threshold for Volume 2 is exceeded, SnapManager deletes all but five Snapshot copies, regardless of the policies for Volumes 1 and 3.

Viewing current fractional space reservation data for a LUN

You can view current fractional space reservation settings to ensure that the policy in force for each LUN is configured appropriately.

About this task

Only the **Drive Letter** or **Mount Point** column displays LUN-specific information. All other columns in the **Current Settings** tab display information that applies across the volume that contains the LUN.

The SnapManager fractional space reservation policy includes a separate, Exchange-aware automatic deletion feature. The SnapManager automatic deletion feature can be used in place of or along with the Data ONTAP automatic deletion feature; you can also select to disable the SnapManager automatic deletion feature.

Steps

1. In the **Actions** pane, click **Fractional Space Reservation Settings**.
2. In the **Current Settings** tab, note the space consumption status for each LUN that stores database or SnapInfo directories.

   The information displayed in this tab automatically refreshes every 60 seconds.
3. If the **Snapshot Autodelete** column is enabled, investigate the cause of it being enabled and take one of the following preventive actions:

- Disable the Data ONTAP Automatic deletion of Snapshot copies feature.
- Ensure that the Data ONTAP Automatic deletion of Snapshot copies feature is configured in such a way that it does not delete Exchange backup set components.

For details about the `snap automatic delete storage system` command, see the *Data ONTAP Block Access Management Guide for iSCSI and FC* for Data ONTAP 7.1 or later.

4. To close the dialog box, click **OK**.

Fractional space reservation status data

You can view data about your current space-reservation status in the **Current Settings** tab of the **Fractional Space Reservation Settings** dialog box

If the Storage Snapshot Autodelete option is enabled, the LUN is contained in a FlexVol volume that has overwrite reserve set to less than 100 percent and that also has the Data ONTAP automatic Snapshot copy deletion feature enabled and configured to trigger when the overwrite reserve is nearly full. If Exchange data or SnapManager SnapInfo directories are stored on LUNs contained in a volume with these characteristics, the Data ONTAP Snapshot copy automatic deletion policy might delete Exchange backup set components.

The following columns display SnapManager configuration information:

Drive Letter or Mountpoint	A SnapManager configuration setting. The drive letter or NTFS mountpoint on which the LUN is mounted.
Backup Autodelete Trigger (percentage)	A SnapManager fractional space reservation policy setting. The percentage of overwrite reserve utilization that triggers automatic deletion of Exchange backup sets for the volume that contains the LUN.
Disable Database Trigger (percentage)	A SnapManager fractional space reservation policy setting. The percentage of overwrite reserve utilization that triggers automatic disabling of Exchange databases for the volume that contains the LUN.

The following columns display the fractional overwrite reserve settings and status:

Fractional Reserve (percentage)	The amount of space reserved for overwrites on the volume that contains this LUN, expressed as a percentage of the total size of all space-reserved LUNs in the volume
Used Reserve (percentage)	For the volume that contains this LUN, the amount of overwrite reserve in use, expressed in two ways: as a percentage of the total size of all space-reserved LUNs in the volume, and in megabytes
Available Reserve (MB)	For the volume that contains this LUN, the amount of overwrite reserve available
Snapshot Autodelete	For the volume that contains this LUN, the state of the Data ONTAP Snapshot copy automatic deletion feature: enabled or disabled

If this LUN stores Exchange data files and is contained in a volume for which the Data ONTAP Snapshot copy automatic deletion feature is enabled, disable this feature on that volume or ensure that it is configured so that it does not delete SnapManager backup set components.

Note: The SnapManager fractional space reservation policy triggers are not applicable to fully space-reserved LUNs. If Fractional Overwrite Reserve (percentage) is 100, the LUN is contained in a fully space-reserved volume rather than a fractionally space-reserved volume.

Event notification options

You can use either the Configuration wizard or the **Auto Notification Settings** dialog box to configure e-mail notifications, syslog event logging, and AutoSupport notifications.

Operations with e-mail notification support

SnapManager can notify you through e-mail messages (using SMTP) about the success or failure of the following types of events:

- SnapManager backup operation
- Database integrity verification
- SnapManager restore operation
- SnapManager configuration
- SnapManager fractional space reservation policy event execution

You can select one of these body messages to include in the body of the e-mail: Send operation results summary, or Send verbose operation results.

You must enable the e-mail notifications option, which is disabled by default.

SnapManager syslog event logging

SnapManager events are posted to the storage system syslog by default. You can disable this option to reduce the load on the network or when you are troubleshooting your system.

AutoSupport notification

If AutoSupport is enabled on both the storage system and SnapManager, technical support receives automatic e-mail notification about any SnapManager events or storage system problems that might occur. This option is enabled by default. You can disable this option to reduce the load on the network or when you are troubleshooting your system.

The AutoSupport daemon monitors the storage system’s operations and sends automatic messages to technical support to alert them to potential storage system problems. If necessary, technical support contacts you by e-mail to help resolve a potential system problem. The AutoSupport daemon is

allowed by default on the storage system. For more information, see the *Data ONTAP System Administration Guide for 7-Mode*.

Limitation of AutoSupport notification to failure events only

If AutoSupport is enabled, you can limit the SnapManager events that are posted to the storage system syslog and AutoSupport (if allowed for SnapManager) to failure events only. The option is enabled by default.

Configuring automatic event notification settings

You can configure automatic event notification settings for SnapManager. You can enable and configure e-mail notification, advanced event notification settings, advanced e-mail notification settings, storage system syslog settings, and AutoSupport notification settings.

Before you begin

You must have the following ready before you configure automatic event notification settings:

- IP address of the SMTP e-mail server or gateway
- E-mail address of each recipient to whom the notification is to be sent.
- E-mail address of the sender of the notification that you want to use

  By default, SMEAutoSender is the name of the notification sender. To specify a sender other than the default, use one of the following formats:

  - *`SenderAlias<SenderName@SenderDomain>`*
  - *`SenderAlias`*
  - *`SenderName@SenderDomain`*
- The text to be appended to the standard subject line, which is included in all notification messages:

  `Backup status at mm_dd_yyyy-hh.mm.ss from MachineName`

  By default, the string `SnapManager for Exchange` is appended.

If you select to send the operational results in summary format rather than in verbose format, you can also select the **Include SnapManager Operation Report as an Attachment** option.

About this task

By default, the automatic e-mail notification feature is disabled.

SnapManager relies on and requires an external e-mail host at your site to send e-mail. The e-mail host is a host that runs a mail server that listens on the SMTP port (25).

You can configure the SMTP mail notification settings by selecting either or both of the following two options:

Only send notification when operation fails	Specifies that you want e-mail notification sent only when a backup process or a verification process fails (cleared by default).

Include SnapManager operation report as attachment	Specifies that you want the status report to be attached to the e-mail notification (cleared by default).

Steps

1. In the **Actions** pane (or the **Configure Automatic Event Notification** screen of the Configuration wizard), select **Notification Settings**.

   The **Auto Notification Settings** dialog box opens.

2. To enable e-mail notification, select the **Send Email Notification** option.

3. In the relevant text boxes, type the following information.
   a) In the **SMTP Server** text box, type the host name or the IP address of the SMTP e-mail server or gateway to be used.
   b) In the **From** text box, type the e-mail address of the sender of the notification.
   c) In the **To** text box, type the e-mail address of each recipient.

      To send to more than one recipient, use a semicolon (;) to separate the addresses.
   d) In the **Subject** text box, type the text to be appended to the standard subject line.

4. Click **Advanced**.

   The **Advanced Event Notification Settings** dialog box opens.

5. In the **E-mail Message Content** pane, select the types of body messages to include in the e-mail.

6. If you choose the summary format rather than the verbose format, you can also select the **Include SnapManager Operation Report** as an Attachment option.

7. Click **Apply** to commit your settings.

8. To configure the SMTP mail notification settings, select either or both of the following two options:

   - **Only send notification when operation fails** check box
   - **Include SnapManager operation report as attachment** check box

9. Click **OK** to apply your settings and close the **Advanced E-mail Notifications Settings** dialog box.

10. Click **Send a Test Email**.

    SnapManager sends a test e-mail notification that uses the settings you specified and displays what that e-mail message looks like.

11. If you want to post SnapManager events to the storage system syslog, select **Log SnapManager Events to Storage System Syslog**.

12. If you want to enable automatic notification of syslog entries to technical support, and if SnapManager is configured to log events to the storage system syslog, select **Send AutoSupport Notification**.
13. If you want to limit SnapManager event logging to failure events, select **On failure only**.
14. Click **OK** (or **Next**, if you are using the Configuration wizard).

SnapManager control file XML schema

You can set and edit storage layout, notification, verification, report directory, backup, and SnapMirror settings through the control file using an XML schema.

Storage layout settings XML schema

Use the storage layout settings XML schema to set and edit storage layout settings.

```
<?xml version="1.0" encoding="utf-8"?>
<xsd:schema xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns:xs="http://www.w3.org/2001/XMLSchema"
xmlns:xsd="http://www.w3.org/2001/XMLSchema"
attributeFormDefault="unqualified"
elementFormDefault="qualified">
    <xs:element name="SMECONFIG">
        <xs:complexType>
            <xs:sequence>
                <xs:element name="HOST_NAME" type="xs:string" />
                <xs:element name="SERVER_NAME" type="xs:string" />
                <xs:element name="STORAGE_LAYOUT">
                    <xs:complexType>
                        <xs:sequence>
                            <xs:element name="MTA_SYSTEM_FILES">
                                <xs:complexType>
                                    <xs:sequence>
                                        <xs:element
name="MTA_DB_DIR" type="xs:string" />
                                        <xs:element
name="MTA_RUN_DIR" type="xs:string" />
                                    </xs:sequence>
                                </xs:complexType>
                            </xs:element>
                            <xs:element
name="SMTP_SYSTEM_FILES_VS">
                                <xs:complexType>
                                    <xs:sequence>
                                        <xs:element
name="SMTP_SVC_PATH">
                                            <xs:complexType>
                                                <xs:sequence>
                                                    <xs:element
name="STMP_DISK" type="xs:string" />
                                                </xs:sequence>
                                            </xs:complexType>
```

```
                                        </xs:element>
                                    </xs:sequence>
                                </xs:complexType>
                            </xs:element>
                            <xs:element
name="ADD_MSISCSI_DEPENDENCY" type="xs:boolean" />
                            <xs:element name="STORAGE_GROUPS">
                                <xs:complexType>
                                    <xs:sequence>
                                        <xs:element
maxOccurs="unbounded" name="STORAGE_GROUP">
                                            <xs:complexType>
                                                <xs:sequence>
                                                    <xs:element
name="SG_NAME" type="xs:string" />
                                                    <xs:element
name="SG_SYS_PATH" type="xs:string" />
                                                    <xs:element
name="SG_LOG_PATH" type="xs:string" />
                                                    <xs:element
name="SG_SNAPINFO" type="xs:string" />
                                                    <xs:element
name="DATABASES">
<xs:complexType>
<xs:sequence>
<xs:element maxOccurs="unbounded" name="DATABASE">
 <xs:complexType>
     <xs:sequence>
         <xs:element name="DB_NAME" type="xs:string" />
             <xs:element name="EDB_PATH" type="xs:string" />
             <xs:element name="STM_PATH" type="xs:string" />
         </xs:sequence>
     </xs:complexType>
 </xs:element>
                                                                    </
xs:sequence>
                                                                  </
xs:complexType>
                                                              </
xs:element>
```

```
                                                                    </xs:sequence>
                                                                </xs:complexType>
                                                            </xs:element>
                                                        </xs:sequence>
                                                    </xs:complexType>
                                                </xs:element>
                                            </xs:sequence>
                                        </xs:complexType>
                                    </xs:element>
```

When you have datasets configured in the host, and the Storage Group is configured with the dataset policy, the exported control file contains the following information in the storage layout settings:

```
<STORAGE_GROUPS>
    <STORAGE GROUP>
        <SG_NAME>First Storage Group</SG_NAME>
        <SG_SYS_PATH>Q:\Program Files\Exchsrvr\mdbdata\</
SG_SYS_PATH>
        <SG_LOG_PATH>Q:\Program Files\Exchsrvr\mdbdata\</
SG_LOG_PATH>
        <SG_SNAPINFO_PATH>Q:\SIF\SME_SnapInfo\</SG_SNAPINFO_PATH>
        <SG_DATASET_POLICY>Backup up</SG_DATASET_POLICY>
        ......
    </STORAGE_GROUP>
</STORAGE_GROUPS>
```

Notification settings XML schema

Use the notification settings XML schema to specify notification settings.

```
<xs:element name="COMMON_SETTINGS">
<xs:complexType>
    <xs:sequence>
        <xs:element name="NOTIFICATION">
            <xs:complexType>
                <xs:sequence>
                    <xs:element name="SEND_EMAIL_NOTIFICATION">
                        <xs:complexType>
                            <xs:sequence>
                                <xs:element name="SMTPSERVER"
type="xs:string" />
                                    <xs:element name="FROM"
type="xs:string" />
                                        <xs:element name="TO"
type="xs:string" />
                                            <xs:element
name="SUBJECT" type="xs:string" />
```

```
                                                <xs:element
name="NOTIFY_AUTO" type="xs:boolean" />
                                                <xs:element
name="LONG_MSG" type="xs:boolean" />
                                                <xs:element
name="AS_ATTACHMENT" type="xs:boolean" />
                                                <xs:element
name="SEND_ON_FAILURE" type="xs:boolean" />
                                            </xs:sequence>
                                        </xs:complexType>
</xs:element>
<xs:element name="EMS_ENABLED" type="xs:boolean" />
                            <xs:element name="ASUP_ENABLED"
type="xs:boolean" />
                            <xs:element name="ASUP_ON_FAIL"
type="xs:boolean" />
                        </xs:sequence>
                    </xs:complexType>
                </xs:element>
```

Verification settings XML schema

Use the verification settings XML schema to specify the verification settings.

```
<xs:element name="VERIFICATION">
    <xs:complexType>
        <xs:sequence>
            <xs:element name="VERIFICATION_CLIENT_SETTING">
                <xs:complexType>
                    <xs:sequence>
                        <xs:element name="VER_SERVER"
type="xs:string" />
                            </xs:sequence>
                                </xs:complexType>
                                    </xs:element>
                                        <xs:element
name="VERIFICATION_SERVER_SETTING">
                                            <xs:complexType>
                                                <xs:sequence>
                                                    <xs:element
name="ESEUTIL_PATH" type="xs:string" />

<xs:element name="AUTO_DRIVELETTER" type="xs:boolean" />

<xs:element name="MP_DIR" type="xs:string" />

<xs:element name="THROTTLE" type="xs:boolean" />

<xs:element name="IO_PAUSE" type="xs:unsignedByte" />
```

```
                                                                    </
xs:sequence>
                                                                </
xs:complexType>
                                                            </xs:element>
                                                        </xs:sequence>
                                                    </xs:complexType>
                                                </xs:element>
```

Report directory settings XML schema

Use the report directory settings XML schema to specify report directory settings.

```
<xs:element name="REPORT_DIRECTORY" type="xs:string" />
    <xs:element name="BACKUP">
        <xs:complexType>
            <xs:sequence>
                <xs:element name="BACKUP_CLIENT_SETTING">
                    <xs:complexType>
                        <xs:sequence>
                            <xs:element name="NAMING_CONVENTION">
                                <xs:complexType>
                                    <xs:attribute name="GENERIC"
type="xs:boolean" use="required" />
                                        <xs:attribute
name="UNIQUE" type="xs:boolean" use="required" />
                                            </xs:complexType>
                                                </xs:element>
                                                    <xs:element
name="BACKUP_SET_TO_KEEP" type="xs:unsignedByte" />

<xs:element name="BACKUP_SET_TO_KEEP_IN_DAYS"
type="xs:unsignedByte" />

<xs:element name="BACKUP_SET_TO_VERIFY" type="xs:unsignedByte" />
                                                                    </
xs:sequence>
                                                                    </
xs:complexType>
                                                                        </
xs:element>
```

Backup settings XML schema

Use the backup settings XML schema to specify backup settings.

```
<xs:element name="BACKUP_SERVER_SETTING">
    <xs:complexType>
        <xs:sequence>
            <xs:element name="RUN_CMD_PATH" type="xs:string" />
                <xs:element name="RUN_CMD_ARGUMENT"
type="xs:string" />
                    </xs:sequence>
                        </xs:complexType>
                            </xs:element>
                                </xs:sequence>
                                    </xs:complexType>
                                        </xs:element>
```

SnapMirror relationship settings XML schema

Use the SnapMirror relationship settings XML schema to specify SnapMirror relationship settings.

```
<xs:element name="VERIFICATION_ON_DESTINATION">
    <xs:complexType>
        <xs:element name="SELECTED_DESTINATIONS">
            <xs:complexType>
                <xs:sequence>
                    <xs:element maxOccurs="unbounded"
name="SELECTED_DESTINATION">
                        <xs:complexType>
                            <xs:sequence>
                                <xs:element name="SOURCE_FILER"
type="xs:string" />
                                    <xs:element
name="SOURCE_VOLUME" type="xs:string" />
                                        <xs:element
name="DESTINATION_FILER" type="xs:string" />
                                            <xs:element
name="DESTINATION_VOLUME" type="xs:string" />
                                                </xs:sequence>
                                                    </
xs:complexType>
                                                        </
xs:element>
                                                    </xs:sequence>
```

```
                                                    </xs:complexType>
                                                </xs:element>
                                                <xs:element
name="OTHER_AVAILABLE_DESTINATIONS">
                                                    <xs:complexType>
                                                        <xs:sequence>

<xs:element maxOccurs="unbounded"
name="AVAILABLE_DESTINATION">

<xs:complexType>

<xs:sequence>

 <xs:element name="SOURCE_FILER" type="xs:string"/>

 <xs:element name="SOURCE_VOLUME" type="xs:string"/>

<xs:element name="DESTINATION_FILER" type="xs:string"/>

<xs:element name="DESTINATION_VOLUME" type="xs:string"/>
                                                                    </
xs:sequence>
                                                                 </
xs:complexType>
                                                        </xs:element>
                                                    </xs:sequence>
                                                </xs:complexType>
                                            </xs:element>
                                        </xs:sequence>
                                    </xs:complexType>
                                </xs:element>
                            </xs:sequence>
                        </xs:complexType>
                    </xs:element>
                </xs:sequence>
            </xs:complexType>
        </xs:element>
</xsd:schema>
```

SnapManager command-line reference

SnapManager 6.0 for Microsoft Exchange provides SnapManager command-line functionality, allowing you to create scripts to run SnapManager without using the graphical user interface (GUI).

Guidelines for using the SnapManager for Exchange PowerShell command-line tool

You need to keep in mind some points before you start using the PowerShell command-line tool.

To run the Microsoft Exchange cmdlets, use Microsoft Exchange Management Shell.

Observe the following guidelines when using the SnapManager command-line tool:

- All parameters and options are case-insensitive. So using the option “-Daily” gives the same results as using “-daily”.
- Some of the options must be invoked in a particular order. For best results, use the order given in the syntax for all options.
- When a parameter value string contains spaces, ensure to enclose it in single quote. For example, use “First Storage Group” rather than First Storage Group.

Launching SnapManager for Exchange PowerShell

You can launch the SnapManager for Exchange PowerShell from the Windows Start menu to use the command-line interface to perform various SnapManager operations.

Step

1. Go to **Start > Programs > NetApp > SnapManager for Exchange PowerShell**.

   The SnapManager for Exchange PowerShell command-line interface is displayed.

After you finish

On a Windows Server 2008 system with **User Access Control** enabled, use the menu option **Run As Administrator** to prevent any access errors related to the Windows Registry, or any other Windows server resources.

new-backup

The `new-backup` command enables you to back up the Exchange databases.

```
                new-backup                -Clusteraware <True|
False>                -lcr <True|False>                -VerifyOnDestVolumes
```

```
<src_storage_system_list:src_vol:dest_storage_system:dest_vol>
   -Verify <True|False>                -Server
<server_name>                -StorageGroup <storage_grp1,
storage_grp2, ...>                -ManagementGroup <Standard|Weekly|
Daily>                -ActiveDatabaseOnly <True|False>                -
PassiveDatabaseOnly <True|False>                -BackupTargetServer
<server name>                -ActivationPreference
<ActivationPreferenceNum>                -UpdateMirror <True|
False>                -VerDestVolume <True|False>                -
NoUTMRestore <True|False>                -NoTruncateLogs
<False>                -Throttle <throttle_val>                -
VerificationServer <server_name>                -UseMountPoint <True|
False>                -CCRActiveNode Boolean <True|False>                -
MountPointDir <mountpoint_dir>                -RetainBackups
<no_of_days_to_retain_backup>                -RetainDays
<no_of_days_delete_backup>                -Command <True|
False>                -RunCommand
<win_path_and_script_name>                -GenericNaming <True|
False>                -BackupCopyRemoteCCRNode Boolean <True|
False>                -RecoveryPoint
<win_path_and_script_name>                -ReportProgress <True|
False>                -ArchiveBackup <True|False>                -
ArchiveBackupCopyRemoteCCRNode <True|False>                -
ArchivedBackupRetention <Hourly|Monthly|Daily|Weekly|Unlimited>-
RetainUtmBackups
 <no_of_log_backups_to_retain>                common parameters
```

Description

The `new-backup` command enables you to initiate a backup or verification job, with all of the options available through the SnapManager GUI.

This command also supports the following common parameters:

- `-Debug` (-db)
- `-ErrorAction` (-ea)
- `-ErrorVariable` (-ev)
- `-OutBuffer` (-ob)
- `-OutVariable` (-ov)
- `-Verbose` (-vb)
- `-Confirm`

To learn more about common parameters, see Help about_ubiquitous_parameters.

Parameters

`-Clusteraware <True|False>`

Short name: `cl`

Assumes significance only when scheduling jobs in cluster configurations, by facilitating scheduling the same job in multiple cluster nodes to improve fault tolerance.

In the case of Single Copy Cluster (SCC), if a job is scheduled with `-Clusteraware`, the job runs only if the host in which it is scheduled is the active node of the cluster.

In the case of Cluster Continuous Replica (CCR), the processing of `-Clusteraware` depends on whether the host in which it is scheduled is an active or passive node, and on the value of `-CCRActiveNode`. The job runs if the host is a CCR active node and the job targets it with `-CCRActiveNode` as true. If the host is a CCR passive node, the job runs if the job targets the host with `-CCRActiveNode` as false.

In the case of DAG, if a job is scheduled with `-Clusteraware`, the job runs only if the host in which it is scheduled is the active node of the DAG.

`-lcr <True|False>`

Processed only if the storage groups are not specified explicitly using `-StorageGroup`. In Exchange Server 2007, `-lcr` is false by default and in such cases only the production Storage Groups are backed up. Note the following points:

- In Exchange Server 2007, set `-lcr` as true, if you want to back up only the replica Storage Groups.
- If you specify the list of Storage Groups explicitly, the `-lcr` Storage Group is ignored.

Note: In Exchange Server 2003, `-lcr` is not supported.

`-VerifyOnDestVolumes` *<src_storage_system_list:src_vol:dest_storage_system:dest_vol>*

Short name: `vermirror`

Overrides the existing SnapMirror relationships.

`-Verify <True|False>`

Short name: `ver`

Verifies the backed up SnapManager databases and transaction logs.

`-Server` *<server_name>*

Short name: `svr`

Specifies the target Exchange server name. You can also specify the Database Availability Group (DAG) name.

In a cluster configuration, you must specify `-Server` explicitly in all of the cmdlets to perform all operations. If `-Server` is not specified explicitly in a stand-alone server, SnapManager uses the local machine as the default to run the following cmdlets: `new-backup`, `verify-backup`, `restore-backup`, `get-backup`, and `delete-backup`.

`-StorageGroup` *<storage_grp1, storage_grp2, ...>*

Short name: `sg`

Specifies the list of Storage Groups to be verified during the backup operation.

If invalid Storage Groups are specified in `new-backup` during backup operation, PowerShell cancels the `new-backup` operation. SnapManager displays an appropriate error message and does not selectively back up valid Storage Groups. Scheduled jobs and backup scripts also fail when there is an invalid Storage Group configuration.

`-ActiveDatabaseOnly <True|False>`

Short name: `activedb`

Backs up the active databases. This is a switch parameter.

> **Note:** If neither the `ActiveDatabaseOnly` nor the `PassiveDatabaseOnly` option is specified, all passive and active databases in the DAG are included in the backup.

`-PassiveDatabaseOnly <True|False>`

Short name: `passivedb`

Backs up the passive databases. This is a switch parameter.

> **Note:** If neither the `ActiveDatabaseOnly` nor the `PassiveDatabaseOnly` option is specified, all passive and active databases in the DAG are included in the backup.

`-BackupTargetServer` *<server name>*

Short name: `bkupsvr`

Backs up the databases on the specified server. If you do not specify `BackupTargetServer`, mailbox databases on all member servers in the DAG are backed up. For example, the command `new-backup -server DAG1` backs up all databases on all member servers in the DAG.

`-ActivationPreference` *<ActivationPreferenceNum >*

Short name: `actpref`

Back up the databases with the specified ActivationPreference number.

> **Note:** You can use the following cmdlet from the Microsoft Exchange 2010 PowerShell to get the list of `ActivationPreference` numbers on the member servers for a database: `Get-MailboxDatabase -Identity databasename | fl .`

`-ManagementGroup <Standard|Weekly|Daily>`

Short name: `mgmt`

Specifies the frequency of a backup or verify operation that is scheduled to be performed on a daily, weekly, or standard basis.

`-UpdateMirror <True|False>`

Short name: `updmir`

Starts a SnapMirror synchronization after the backup operation.

SnapMirror updates the specified volume to reflect incremental updates to a source volume. If a SnapManager volume is enabled for SnapMirror use, the SnapMirror destination is updated from the source volume.

`-VerDestVolume <True|False>`

Short name: `verdest`

Verifies the SnapMirror destination volume.

`-NoUTMRestore <True|False>`

Short name: `noutm`

Denies the retention of up-to-the-minute restore ability.

The logs are also deleted for any backup copies that you delete as part of this backup operation. This does not retain the up-to-the-minute restore ability for older backup copies that remain after the delete phase of a backup operation.

`-NoTruncateLogs <`*`False`*`>`

Short name: `notrunc`

Denies the backup of truncated transaction logs. This option can be used to conserve space on the LUN containing the backed up Exchange transaction logs.

`-Throttle <`*`throttle_val`*`>`

Short name: `throt`

`throttle_var` is an integer value that defines the throttling value to be used during the verification operation.

`-VerificationServer <`*`server_name`*`>`

Short name: `versvr`

Overrides the preconfigured SnapManager verification settings. It denotes the host to be used as the verification server for the verification phase of a backup operation. This is a switch parameter.

`-UseMountPoint <True|False>`

Short name: `mp`

Mounts the Snapshot copy to a NTFS directory. During a SnapManager verification operation, the Snapshot copies are mounted in a default NTFS directory for database verification. This option is effective when there are no drives available to mount the Snapshot copies during database verification. The value of this parameter overrides the preconfigured SnapManager verification settings.

`-CCRActiveNode <Boolean True|False>`

Short name: `ccrnode`

Specifies whether the connection is to an active or a passive node, respectively. This is a Boolean parameter.

-MountPointDir ***<mountpoint_dir>***

Short name: `mpdir`

Specifies which mount-point directory a Snapshot copy is to be mounted to, during database verification.

-RetainBackups ***<no_of_days_to_retain_backup>***

Short name: `rtbackups`

Specifies the number of backup copies to retain after the delete phase of a SnapManager backup operation.

-RetainDays ***<no_of_days_to_delete_backup>***

Short name: `rtdays`

Specifies the number of days after which a backup copy is to be deleted. Specifies that the deletion of backup copies is to be based on an "older than number of days" policy. Use this option in conjunction with `RetainBackups`. Absence of this option denotes a deletion policy based on retained backup count.

-Command <True|False>

Short name: `cmd`

Indicates that `RunCommand` is to be used after the current operation. This is a switch parameter.

-RunCommand***<win_path_and_script_name>***

Short name: `runcmd`

Specifies the complete path name of and runs the specified command after the SnapManager backup or verification operation is complete.

> **Note:** You must specify this command explicitly; the preconfigured command does not run after the backup or verification operation.

-GenericNaming <True|False>

Short name: `gen`

Specifies the generic naming convention to be used for the SnapManager backup sets.

-BackupCopyRemoteCCRNode Boolean <True|False>

Creates a secondary backup on a remote node. This is a Boolean parameter.

-RecoveryPoint ***<win_path_and_script_name>***

Indicates a Frequent Recovery Point backup.

Note: If you specify the parameter `new-backup`, all the other parameters except `-Server`, `-StorageGroup`, `-UpdateMirror`, `-Clusteraware`, `-lcr`, `-BackupCopyRemoteCCRNode` are ignored.

`-ReportProgress <True|False>`

Short name: `repprog`

Displays the operation status and progress information in the PowerShell output. If you do not use this switch parameter, the progress information is logged only to the report file, and not to the PowerShell output.

Note: Avoid using this switch for scheduled backup and verification jobs.

`-ArchiveBackup <True|False>`

Short name: `arch`

Creates an archive of the backup copy created on the primary node. The primary node is the node on which the backup operation is initiated. Include this parameter if you have datasets configured on your primary node and want to archive the backup copy to the SnapVault secondary storage system.

`-ArchiveBackupCopyRemoteCCRNode <True|False>`

Short name: `arch`

Creates an archive of the copy backup created on the secondary node. Secondary node is the remote node of the primary node. Include this parameter if you have datasets configured on the secondary node and want to archive the copy backup to the SnapVault secondary storage system.

`-ArchivedBackupRetention <Hourly|Monthly|Daily|Weekly|Unlimited>`

Short name: `archret`

Determines the retention time for the archives that were created using the `ArchiveBackup` parameter. The retention can be hourly, monthly, daily, weekly, or unlimited.

Note: Do not use `ArchivedBackupRetention` without using the parameter `ArchiveBackup`. If you use the parameter `ArchiveBackup` only, the daily retention type is used by default.

`ArchiveBackupCopyRemoteCCRNode <True|False>`

Short name: `archbkupcopy`

Archives the backup copies at the remote node specifically for a CCR configuration, because the remote node could be an active node if you run the backup copy from passive node.

`-RetainUtmBackups` *<no_of_log_backups_to_retain>*

Defines the number of log backups to retain.

Example: Creating a new backup copy of two storage groups in standard backup management group, and deleting older backup copies

This command creates a new backup copy of two Storage Groups on EXCHSRVR in the Standard backup management group (the default), using the unique naming convention (the backup copy is named using the date-time stamp). Run Command After Operation is performed, and all older backup copies are deleted except for the eight most recent.

```
new-backup -Server 'EXCHSRVR' -ManagementGroup 'Standard' -
NoTruncateLogs $False -RetainBackups 8 -Command - RunCommand 'C:
\WINDOWS\system32\svripti.bat'-CommandServer 'SNAPMGR-48' -
StorageGroup 'alaska','3' -Verify -VerificationServer 'Snapmgr-48' -
Throttle 150 -UseMountPoint -MountPointDir 'C:\ProgramFiles\NetApp
\SnapManager for Exchange\SnapMgrMountPoint'
```

Example: Creating a new backup of two storage groups in standard backup management group, and retaining older backup copies

This command creates a new backup of two Storage Groups on EXCHSRVR3 in the Standard backup management group (the default), using the unique naming convention (the backup are named using the date-time stamp). No verification or other command is performed after this operation, and no older backup copies are deleted.

```
new-backup -Server 'EXCHSRVR3' -StorageGroup SG1, SG2
```

Example: Creating a new backup in Daily backup management group, verifying backup copies, and deleting older backup copies

This command creates a new backup in the Daily backup management group. The backup is verified on SRVR7, and all older Daily backup copies are deleted except for the three most recent. No SnapMirror replication is initiated after the backup, even if the volume is a SnapMirror source volume.

```
new-backup -Server 'EXCHSRVR3' -StorageGroup SG1, SG2 -
VerificationServer -SRVR7 -retainbackups 3 -ManagementGroup daily
```

Example: Creating a new backup and verifying destination volume

This command creates a backup and performs verification on the destination volumes:

```
new-backup -Server SNAPMGR-55 -ManagementGroup Standard -
NoTruncateLogs $False -StorageGroup 'First Storage Group' -Verify -
VerificationServer SNAPMGR-55 -VerDestVolume -UpdateMirror -
BackupCopyRemoteCCRNode $False
```

Example: Creating a backup of an LCR database

This command creates a backup of the LCR database.

```
new-backup -Server SNAPMGR-55 -ManagementGroup Standard -
NoTruncateLogs $False -StorageGroup First Storage Group\LCR -
BackupCopyRemoteCCRNode $False
```

Example: Creating a frequent recovery backup

This command creates a frequent recovery backup from the specified Storage Group in "SNAPMGR-50":

```
new-backup -Server SNAPMGR-50 -StorageGroup FirstStorageGroup, SG2,
SG3 -UpdateMirror -RecoveryPoint
```

Example: Creating archive of primary backup at the secondary Storage Group

This command creates an archive of the primary backup at the secondary Storage Group location:

```
new-backup -Server exchange1 -CCRActiveNode $False -GenericNaming -
ManagementGroup Standard -NoTruncateLogs $False -RetainBackups 8 -
StorageGroup sg1,sg2 -Verify -VerificationServer exchange2-Throttle
200 -UseMountPoint - MountPointDir 'C:\Program Files\NetApp
\SnapManager for Exchange\SnapMgrMountPoint'-BackupCopyRemoteCCRNode
$False -ArchiveBackup -ArchivedBackupRetention Monthly
```

Example: Backing up databases on the specified server in the DAG

This command creates a backup of the specified databases on the server "SNAPMGR06":

```
new-backup -Server 'SNAPMGR06-DAG1' -ClusterAware -ManagementGroup
'Standard' -NoTruncateLogs $False -RetainBackups 3 -dbs 'Mailbox
Database 0294565900','Mailbox Database 0793619176','DB1' -Verify -
UseMountPoint -BackupTargetServer snapmgr-06
```

Example: Backing up the databases with the specified ActivationPreference number on the specified server in the DAG

This command creates a backup of the databases with the ActivationPreference number 2 on the server "SNAPMGR06":

```
new-backup -Server 'SNAPMGR06-DAG1' -ClusterAware -ManagementGroup
'Standard' -NoTruncateLogs $False -RetainBackups 3 -Verify -
UseMountPoint -ActivationPreference 2 -BackupTargetServer snapmgr-06
```

verify-backup

The `verify-backup` command enables you to verify the backup sets using the SnapManager for Exchange PowerShell command-line interface.

```
                verify-backup                -Clusteraware <True|
False>                -lcr Boolean <True|False>                -
VerifyOnDestVolumes
<src_storage_system_list:src_vol:dest_storage_system:dest_vol>
   -Server <server_name>                -ManagementGroup <Standard|Weekly|
Daily>                -StorageGroup <storage_grp1,
storage_grp2, ...>                -ActiveDatabaseOnly <True|
False>                -PassiveDatabaseOnly <True|False>                -
BackupTargetServer <server name>                -ActivationPreference
<ActivationPreferenceNum>                -UpdateMirror <True|
False>                -VerDestVolume <True|False>                -Throttle
<throttle_val>                -VerificationServer
<server_name>                -UseMountPoint <True|False>                -
MountPointDir <mountpoint_dir>                -CCRActiveNode <True|
False>                -VerifyBackups
<no_of_backups_to_verify>                -ReportProgress <True|
False>                -ArchiveBackup <True|False>                -
ArchivedBackupRetention <Hourly|Monthly|Daily|Weekly|
Unlimited>                -VerifyArchiveBackup <True|False>
common parameters
```

Description

This command enables you to verify the SnapManager backup sets with all of the options available through the SnapManager GUI.

This command also supports the following common parameters:

- `-Debug` (`-db`)
- `-ErrorAction` (`-ea`)
- `-ErrorVariable` (`-ev`)
- `-OutBuffer` (`-ob`)
- `-OutVariable` (`-ov`)
- `-Verbose` (`-vb`)
- `-Confirm`

To learn more about common parameters, see help about_ubiquitous_parameters.

Parameters

`-Clusteraware <True|False>`

Short name: `cl`

Assumes significance only when scheduling jobs in cluster configurations, by facilitating scheduling the same job in multiple cluster nodes to improve fault tolerance.

In the case of Single Copy Cluster (SCC), if a job is scheduled with `-Clusteraware`, the job runs only if the host in which it is scheduled is the active node of the cluster.

In the case of Cluster Continuous Replica (CCR), the processing of `-Clusteraware` depends on whether the host in which it is scheduled is an active or passive node and the value of `-CCRActiveNode`. The job runs if the host is a CCR active node and the job targets it with `-CCRActiveNode` as true. If the host is a CCR passive node, the job runs if the job targets the host with `-CCRActiveNode` as false.

In the case of DAG, if a job is scheduled with `-Clusteraware`, the job runs only if the host in which it is scheduled is the active node of the DAG.

`-lcr <True|False>`

Processed only if the Storage Groups are not specified explicitly using `-StorageGroup`. In Exchange Server 2007, `-lcr` is false by default and in such cases only the production Storage Groups are backed up. Note the following points:

- In Exchange Server 2007, set `-lcr` as true, if you want to back up only the replica Storage Groups.
- If you specify the list of Storage Groups explicitly, the `-lcr` Storage Group is ignored.

Note: In Exchange Server 2003, `-lcr` is not supported.

`-VerifyOnDestVolumes` *<src_storage_system_list:src_vol:dest_storage_system:dest_vol>*

Short name: `vermirror`

Overrides the existing SnapMirror relationships.

`-Server` *<server_name>*

Short name: `svr`

Specifies the target Exchange server name.

In a cluster configuration, you need to specify `-Server` explicitly in all of the cmdlets to perform all operations. If `-Server` is not specified explicitly in a stand-alone server, SnapManager uses the local machine as the default to run the following cmdlets: `new-backup`, `verify-backup`, `restore-backup`, `get-backup`, and `delete-backup`.

`-StorageGroup` *<storage_grp1, storage_grp2, ...>*

Short name: `sg`

Specifies the list of Storage Groups to be verified during the backup operation.

If invalid Storage Groups are specified in `new-backup` during backup operation, PowerShell cancels the `new-backup` operation. SnapManager displays an appropriate error message and does not selectively backup any valid Storage Groups. Scheduled jobs and backup scripts also fail when there is an invalid Storage Group configuration.

`-ActiveDatabaseOnly <True|False>`

Short name: `activedb`

Verifies the active databases. This is a switch parameter.

> **Note:** If neither the `ActiveDatabaseOnly` nor the `PassiveDatabaseOnly` option is specified, backups of all active and passive databases in the Database Availability Group (DAG) are included in the verify-backup operation.

`-PassiveDatabaseOnly <True|False>`

Short name: `passivedb`

Verifies the passive databases. This is a switch parameter.

> **Note:** If neither the `ActiveDatabaseOnly` nor the `PassiveDatabaseOnly` option is specified, backups of all active and passive databases in the Database Availability Group (DAG) are included in the verify-backup operation.

`-BackupTargetServer` *<server name>*

Short name: `bkupsvr`

Verifies the databases on the specified server.

If you do not specify `BackupTargetServer`, backups of mailbox databases on all member servers in the DAG are verified. For example, the command `verify-backup -server DAG1` will verify the backups of databases on all member servers in the DAG.

`-ActivationPreference` *<ActivationPreferenceNum >*

Short name: `actpref`

Verifies the backups of databases with the specified ActivationPreference number on each member server.

> **Note:** You can use the following cmdlet from the Microsoft Exchange 2010 PowerShell to get the list of `ActivationPreference` numbers on the member servers for a database: `Get-MailboxDatabase -Identity databasename | fl`

`-ManagementGroup <Standard|Weekly|Daily>`

Short name: `mgmt`

Specifies the backup or verify operation that is scheduled to be performed on a daily, weekly, or standard basis.

-UpdateMirror <True|False>

Short name: `updmir`

Starts a SnapMirror synchronization after the backup operation.

SnapMirror updates the specified volume to reflect incremental updates to a source volume. If a SnapManager volume is enabled for SnapMirror use, the SnapMirror destination is updated from the source volume.

-VerDestVolume <True|False>

Short name: `verdest`

Verifies the SnapMirror destination volume.

-Throttle *<throttle_val>*

Short name: `throt`

Defines the throttling value to be used during the verification operation.

-VerificationServer *<server_name>*

Short name: `versvr`

Overrides the preconfigured SnapManager verification settings. It denotes the host to be used as the verification server for the verification phase of a backup operation. This is a switch parameter.

-UseMountPoint <True|False>

Short name: `mp`

Mounts the Snapshot copy to a NTFS directory. During a SnapManager verification operation the Snapshot copies are mounted in a default NTFS directory for database verification. This option is effective when there are no drives available to mount the Snapshot copies during database verification. The value of this parameter overrides the preconfigured SnapManager verification settings.

-MountPointDir *<mountpoint_dir>*

Short name: `mpdir`

Specifies which mountpoint directory a Snapshot copy is to be mounted to, during database verification.

-CCRActiveNode Boolean <True|False>

Short name: `ccrnode`

Specifies whether the connection is to an active or a passive node, respectively. This is a Boolean parameter.

-VerifyBackups *<no_of_backups_to_verify>*

Short form: `verbkups`

Verifies the unverified SnapManager backup sets. The default value is one, but you can specify the number of backup copies that you want to verify.

-ReportProgress <True|False>

Short form: `reppro`

Gets the operation status and progress information in the PowerShell output. If this switch is not used, the progress information is logged on only to the report file and not to the PowerShell output.

> **Note:** It is recommended that you do not use this switch for scheduled backup and verification jobs.

-ArchiveBackup <True|False>

Short name: `arch`

Creates an archive of the backup copy created on the primary node. Primary node is the node where backup operation is initiated. Include this parameter if you have datasets configured on your primary node and want to archive the backup copy to the SnapVault secondary storage system.

-ArchivedBackupRetention <Hourly|Monthly|Daily|Weekly|Unlimited>

Short name: `archret`

Determines the retention time for the archives that were created using the `ArchiveBackup` parameter. The retention can be hourly, monthly, daily, weekly, or unlimited.

> **Note:** Do not use `ArchivedBackupRetention` without using `ArchiveBackup`. If you use the parameter `ArchiveBackup` only, the daily retention type is used by default.

-VerifyArchiveBackup <True|False>

Short name: `verarch`

Specifies that the backup copy that has to be verified is an archived backup copy. If you do not specify `VerifyArchiveBackup`, and one local and one archived backup copy exist with the same name, SnapManager verifies the local backup copy.

Example: Verifying a backup copy

This command verifies the backup copies on a Storage Group alaska.

```
verify-backup -Server SNAPMGR-48 -ManagementGroup Standard -
StorageGroup alaska,3 -VerifyBackups 1 - VerificationServer Snapmgr-48
-Throttle 150 -UseMountPoint - MountPointDir 'C:\Program Files\NetApp
\SnapManager for Exchange\SnapMgrMountPoint'
```

Example: Verifying backup copies using another verification server

This command verifies the backup copies using SNAPMGR-48 as the verification server.

```
verify-backup -Server SNAPMGR-55 -ManagementGroup Standard -
StorageGroup 'Second Storage Group','First Storage Group' -
VerifyBackups 5 -VerificationServer SNAPMGR-48 - UseMountPoint -
MountPointDir 'C:\Program Files\NetApp\SnapManager for Exchange
\SnapMgrMountPoint'
```

Example: Verifying the backups of active databases

This command verifies the backups of active databases.

```
verify-backup -Server 'SNAPMGR06-DAG1' -ManagementGroup 'Standard' -
dbs 'Mailbox Database 0294565900','Mailbox Database 0793619176','DB1'
-VerifyBackups 1 -VerificationServer 'SNAPMGR-06' -UseMountPoint -
MountPointDir 'C:\Program Files\NetApp\SnapManager for Exchange
\SnapMgrMountPoint' -ActiveDatabaseOnly
```

delete-backup

The `delete-backup` command enables you to delete backup sets.

```
                delete-backup                -Server
<server_name>                -backup <backup_name>                -
NoUTMRestore <True|False>                -StorageGroup
<storage_grp_name>                -CCRActiveNode <True|
False>                -RemoteBackup <True|False>                -
ArchiveBackup <True|False>                common parameters
```

Description

Use this command to delete backup sets.

This command also supports the following common parameters:

- `-Debug` (`-db`)
- `-ErrorAction` (`-ea`)
- `-ErrorVariable` (`-ev`)
- `-OutBuffer` (`-ob`)
- `-OutVariable` (`-ov`)
- `-Verbose` (`-vb`)
- `-Confirm`

To learn more about common parameters, see help about_ubiquitous_parameters.

Parameter

`-Server` *<server_name>*

Short form: `svr`

Use this parameter to specify the name of the backup server.

Note: In a clustered configuration, specify `Server` explicitly in all of the cmdlets to perform an operation. If you do not specify `Server` explicitly in a stand-alone server, SnapManager uses the local machine, as default, to run the following cmdlets: `new-backup`, `verify-backup`, `restore-backup`, `get-backup`, and `delete-backup`.

Note: For Exchange Server 2010, you cannot specify the Database Availability Group (DAG) name in this parameter as it is not supported.

-backup <*backup_name*>

Specifies the name of the backup set.

-NoUTMRestore <True|False>

Short name: `noutm`

Denies the retention of up-to-the-minute restore ability.

The logs are also deleted for any backups that you delete as part of this backup operation. This does not retain the up-to-the-minute restore ability for older backups that remain after the delete phase of a backup operation.

-StorageGroup <*storage_grp1*, *storage_grp2*, ...>

Short name: `sg`

Specifies the list of Storage Groups to be verified during the backup operation.

If invalid Storage Groups are specified in `new-backup` during backup operation, PowerShell cancels the `new-backup` operation. SnapManager displays an appropriate error message and does not selectively backup any valid Storage Groups. Scheduled jobs and backup scripts also fail when there is an invalid Storage Group configuration.

-CCRActiveNode Boolean <True|False>

Short name: `ccrnode`

Specifies whether the connection is to an active or a passive node, respectively. This is a Boolean parameter.

-RemoteBackup <True|False>

Deletes the archived backup copies.

If `RemoteBackup` is not specified, the local backup copies get deleted.

Example: Deleting a specific backup set of a specified Storage Group

This command deletes the backup set exchsnap__SNAPMGR-55_11-10- 2006_13.39.16.

```
delete-backup -backup exchsnap__SNAPMGR-55_11-10- 2006_13.39.16
```

Example: Deleting the Snapshot copy of a specified LCR backup set

This command deletes the Snapshot copy of the backup set exchsnap__KRISHNA-SVR18 _11-10-2006_13.39.16 that belongs to the LCR-enabled Storage Group sg4 (LCR).

```
delete-backup -storagegroup sg4 (LCR) -backup exchsnap__KRISHNA-SVR18
_11-10-2006_13.39.16 -verbose -confirm
```

Example: Deleting a specific backup set that belongs to the Storage Group 'sg1' without up-to-the-minute restore ability

This command deletes the backup set exchsnap__KRISHNA-SVR18__11-01-2006_18.12.22 without up-to-the-minute restore ability.

```
delete-backup -storagegroup sg2 -backup exchsnap__KRISHNA-
SVR18__11-01-2006_18.12.22 -NoutmRestore -verbose -confirm
```

get-backup

The `get-backup` command enables you to get the backup sets for the specified criteria.

```
get-backup                -Server <server_name>                -
ManagementGroup <Standard|Weekly|Daily>               -StorageGroup
<storage_grp1, storage_grp2, ...>               -CCRActiveNode Boolean
<True|False>               -Backup <name_of_the_backup>               -
RecoveryPoint <True|False>               -Details <True|
False>               common parameters
```

Description

This cmdlet enables you to retrieve backup sets, depending on the input criteria specified in the PowerShell command-line interface.

This command also supports the following common parameters:

- `-Debug` (-db)
- `-ErrorAction` (-ea)
- `-ErrorVariable` (-ev)
- `-OutBuffer` (-ob)
- `-OutVariable` (-ov)
- `-Verbose` (-vb)
- `-Confirm`

To learn more about common parameters, see help about_ubiquitous_parameters.

Parameter

`-Server` *<server_name>*

Short name: `svr`

Specifies the target Exchange server name.

In a cluster configuration, you need to specify `-Server` explicitly in all of the cmdlets to perform all operations. If `-Server` is not specified explicitly in a stand-alone server, SnapManager uses the local machine as the default to run the following cmdlets: `new-backup`, `verify-backup`, `restore-backup`, `get-backup`, and `delete-backup`.

`-ManagementGroup <Standard|Weekly|Daily>`

Short name: `mgmt`

Specifies the backup copy or verify operation that is scheduled to be performed on a daily, weekly, or standard basis.

`-StorageGroup <`*storage_grp1*, *storage_grp2*, ...`>`

Short name: `sg`

Specifies the list of Storage Groups to be retrieved during the get-backup operation.

If invalid Storage Groups are specified in `new-backup` during backup operation, PowerShell cancels the `new-backup` operation. SnapManager displays an appropriate error message and does not selectively back up valid Storage Groups. Scheduled jobs and backup scripts also fail when there is an invalid Storage Group configuration.

`-CCRActiveNode<Boolean True|False>`

Short name: `ccrnode`

Specifies whether the connection is to an active or a passive node, respectively. This is a Boolean parameter.

`-Backup <`*name_of_the_backup*`>`

Shows the specified full backup details.

If you do not specify `Backup`, `get-backup` displays all of the full backup copies.

`-RecoveryPoint <True|False>`

Shows the recovery point.

`-Details <True|False>`

Shows the details of the full backup copy and the recovery point.

Example: Showing the backup copies of a management group

This command shows the backup copies of the standard management group in the Storage Group alaska.

```
get-backup -Server SNAPMGR-48 -ManagementGroup Standard -StorageGroup
alaska,3
```

Example: Showing all full backup copies

This command shows all of the backup copies of the Storage Group First Storage Group.

```
get-backup -Server SNAPMGR-48 -StorageGroup 'First Storage Group'
```

Example: Showing all of the full backup copies with recovery points

This command shows all of the backup copies of the Storage Group First Storage Group with recovery points.

```
get-backup -Server SNAPMGR-48 -StorageGroup 'First Storage Group' -
RecoveryPoint
```

Example: Showing a specific full backup copy with recovery points

This command shows a specific backup copy of the backup set exchsnap_snapmgr-50_03-01-2007_08.00.00 in the Storage Group First Storage Group with recovery points.

```
get-backup -Server SNAPMGR-48 -StorageGroup 'First Storage Group' -
Backup exchsnap_snapmgr-50_03-01-2007_08.00.00 -RecoveryPoint
```

Example: Showing a specific full backup copy with complete details

This command shows the backup set exchsnap_snapmgr-50_03-01-2007_08.00.00 with complete details.

```
get-backup -Server SNAPMGR-48 -StorageGroup 'First Storage Group' -
Backup exchsnap_snapmgr-50_03-01-2007_08.00.00 - Details
```

Example: Showing both full backup copy and recovery points in detail

This command shows the backup set exchsnap_snapmgr-50_03-01-2007_08.00.00 with the recovery points in detail.

```
get-backup -Server SNAPMGR-48 -StorageGroup 'First Storage Group' -
Backup exchsnap_snapmgr-50_03-01-2007_08.00.00 - RecoveryPoint -
Details
```

get-mirrors

The `get-mirrors` command enables you to retrieve the SnapMirror status for the set of volumes that you configure as mirrors. You can also retrieve all SnapMirror relationships in the specified Business Continuance plan.

```
        get-mirrors-ExchServer <server_name>-BcServer <bc_server_name>-
BcPlan <bc_plan_name>-FailoverTo <bc_failover_destination>-Mirrors
```

```
<[storage system1:mirror1:storage system2:mirror1], [storage
system1:mirror2:storage system2:mirror2], ...>common parameters
```

Description

The `get-mirrors` command enables you to retrieve the SnapMirror status for the set of volumes, with all of the options available through the SnapManager GUI. This command also supports the following common parameters:

- `-Debug` (`-db`)
- `-ErrorAction` (`-ea`)
- `-ErrorVariable` (`-ev`)
- `-OutBuffer` (`-ob`)
- `-OutVariable` (`-ov`)
- `-Verbose` (`-vb`)
- `-Confirm`

Parameters

`-ExchServer` *<server_name>*

Short name: `exsvr`

Specifies the Exchange server name.

This is an optional parameter.

`-BcServer` *<bc_server_name>*

Short name: `drsvr`

Specifies the name of the Business Continuance server that you want to connect to.

This is not a mandatory parameter.

`-BcPlan` *<bc_plan_name>*

Short name: `plan`

Specifies the name of the Business Continuance plan that you want to implement.

SnapManager takes all of the SnapMirror relationships constituted by the plan as input. This is not a mandatory parameter.

`-FailoverTo` *<bc_failover_destination>*

Short name: `to`

Specifies the name of the Business Continuance failover destination site.

`-Mirrors` < *[storage system1:mirror1:storage system2:mirror1], [storage system1:mirror2:storage system2:mirror2], ...>*

Short name: `ms`

Specifies the list of the mirrors for the Business Continuance plan as an array-separated list.

This parameter is mutually exclusive of the parameter sets `BcPlan`, `FailoverTo`, and `ExchServer`. If you specify `BcPlan`, SnapManager takes it as input, otherwise SnapManager uses the list of mirrors that you explicitly specify as input. Do not use them in tandem.

Example: Retrieving the SnapMirror status

This command retrieves the SnapMirror status for the "boston_nyc_drplan" Business Continuance plan.

```
-exsvr cms_Ex -bcserver drcluster -bcplan boston_nyc_drplan -to nyc -
verbose -confirm
```

resync-mirrors

The `resync-mirrors` command enables you to resynchronize SnapMirror relationships for the set of volumes that you configure as mirrors. You can also resynchronize all SnapMirror relationships in the specified Business Continuance plan.

```
        resync-mirrors-ExchServer <server_name>-BcServer <bc_server_name>-
BcPlan <bc_plan_name>-FailoverTo <bc_failover_destination>-Mirrors <
[storage system1:mirror1:storage system2:mirror1], [storage
system1:mirror2:storage system2:mirror2], ...>common parameters
```

Description

The `resync-mirrors` command enables you to resynchronize SnapMirror relationships with all of the options available through the SnapManager GUI. This command also supports the following common parameters:

- `-Debug` (`-db`)
- `-ErrorAction` (`-ea`)
- `-ErrorVariable` (`-ev`)
- `-OutBuffer` (`-ob`)
- `-OutVariable` (`-ov`)
- `-Verbose` (`-vb`)
- `-Confirm`

Parameters

`-ExchServer` *<server_name>*

Short name: `exsvr`

Specifies the Exchange server name.

This is not a mandatory parameter.

-BcServer ***<bc_server_name>***

Short name: `drsvr`

Specifies the name of the Business Continuance server that you want to connect to.

This is not a mandatory parameter.

-BcPlan ***<bc_plan_name>***

Short name: `plan`

Specifies the name of the Business Continuance plan that you want to implement.

SnapManager takes all of the SnapMirror relationships constituted by the plan as input. This is not a mandatory parameter.

-FailoverTo ***<bc_failover_destination>***

Short name: `to`

Specifies the name of the Business Continuance failover destination site.

-Mirrors < *[storage system1:mirror1:storage system2:mirror1]*, *[storage system1:mirror2:storage system2:mirror2]*, ...>

Short name: `ms`

Specifies the list of the mirrors for the Business Continuance plan as an array separated list.

This parameter is mutually exclusive with the parameter set `BcPlan`, `FailoverTo`, and `ExchServer`. If you specify `BcPlan`, SnapManager takes it as input, else SnapManager uses the list of mirrors that you explicitly specify as input. Do not use them in tandom.

Example: Resynchronizing SnapMirror relationships

This command resynchronizes SnapMirror relationships in "boston_nyc_drplan" Business Continuance plan in the direction of storage resources in "nyc".

```
 resync-mirrors -exsvr cms_Ex -bcserver drcluster -bcplan
boston_nyc_drplan -to nyc -verbose -confirm
```

release-mirrors

The `release-mirrors` command releases the SnapMirror relationships. You can also release all SnapMirror relationships in the specified Business Continuance plan.

```
        release-mirrors-ExchServer <server_name>-BcServer <bc_server_name>-
BcPlan <bc_plan_name>-FailoverTo <bc_failover_destination>-Mirrors
<[storage system1:mirror1:storage system2:mirror1], [storage
system1:mirror2:storage system2:mirror2], ...>common parameters
```

Description

`release-mirrors` is typically used after you fail back to the production site. This command also supports the following common parameters:

- `-Debug` (`-db`)
- `-ErrorAction` (`-ea`)
- `-ErrorVariable` (`-ev`)
- `-OutBuffer` (`-ob`)
- `-OutVariable` (`-ov`)
- `-Verbose` (`-vb`)
- `-Confirm`

Parameters

`-ExchServer` *<server_name>*

Short name: `exsvr`

Specifies the Exchange server name.

This is not a mandatory parameter.

`-BcServer` *<bc_server_name>*

Short name: `drsvr`

Specifies the name of the Business Continuance server that you want to connect to.

This is not a mandatory parameter.

`-BcPlan` *<bc_plan_name>*

Short name: `plan`

Specifies the name of the Business Continuance plan that you want to implement.

SnapManager takes all of the SnapMirror relationships constituted by the plan as input. This is not a mandatory parameter.

`-FailoverTo` *<bc_failover_destination>*

Short name: `to`

Specifies the name of the Business Continuance failover destination site.

`-Mirrors` < *[storage system1:mirror1:storage system2:mirror1]*, *[storage system1:mirror2:storage system2:mirror2]*, ...>

Short name: `ms`

Specifies the list of the mirrors for the Business Continuance plan as an array separated list.

This parameter is mutually exclusive with the parameter set `BcPlan`, `FailoverTo`, and `ExchServer`. If you specify `BcPlan`, SnapManager takes it as input, else

SnapManager uses the list of mirrors that you explicitly specify as input. Do not use them in tandom.

Releasing SnapMirror relationships

This command releases SnapMirror relationships for the "boston_nyc_drplan" Business Continuance plan.

```
 release-mirrors -exsvr cms_Ex -bcserver drcluster -bcplan
boston_nyc_drplan -to nyc -verbose -confirm
```

break-mirrors

The `break-mirrors` command enables you to break SnapMirror relationships for the set of volumes that you configure as mirrors. You can also break all SnapMirror relationships in the specified Business Continuance plan.

```
        break-mirrors-ExchServer <server_name>-BcServer <bc_server_name>-
BcPlan <bc_plan_name>-FailoverTo <bc_failover_destination>-Mirrors <
[storage system1:mirror1:storage system2:mirror1], [storage
system1:mirror2:storage system2:mirror2], ...>common parameters
```

Description

The `break-mirrors` command enables you to break SnapMirror relationships with all of the options available through the SnapManager GUI. This command also supports the following common parameters:

- `-Debug` (`-db`)
- `-ErrorAction` (`-ea`)
- `-ErrorVariable` (`-ev`)
- `-OutBuffer` (`-ob`)
- `-OutVariable` (`-ov`)
- `-Verbose` (`-vb`)
- `-Confirm`

Parameters

`-ExchServer` *<server_name>*

Short name: `exsvr`

Specifies the Exchange server name.

This is not a mandatory parameter.

`-BcServer` *<bc_server_name>*

Short name: `drsvr`

Specifies the name of the Business Continuance server that you want to connect to.

This is not a mandatory parameter.

-BcPlan ***<bc_plan_name>***

Short name: `plan`

Specifies the name of the Business Continuance plan that you want to implement.

SnapManager takes all of the SnapMirror relationships constituted by the plan as input. This is not a mandatory parameter.

-FailoverTo ***<bc_failover_destination>***

Short name: `to`

Specifies the name of the Business Continuance failover destination site.

-Mirrors ***< [storage system1:mirror1:storage system2:mirror1], [storage system1:mirror2:storage system2:mirror2], ...>***

Short name: `ms`

Specifies the list of the mirrors for the Business Continuance plan as an array separated list.

This parameter is mutually exclusive with the parameter set `BcPlan`, `FailoverTo`, and `ExchServer`. If you specify `BcPlan`, SnapManager takes it as input, else SnapManager uses the list of mirrors that you explicitly specify as input. Do not use them in tandem.

Example: Breaking SnapMirror relationships

This command breaks SnapMirror relationships in the "boston_nyc_drplan" Business Continuance plan:

```
-exsvr cms_Ex -bcserver drcluster -bcplan boston_nyc_drplan -to nyc -
verbose -confirm
```

restore-backup

The `restore-backup` command enables you to restore the Storage Group and databases.

```
                restore-backup                -Backup
<name_of_the_backup>                -RestoreLastBackup
<restore_last_backup>                -VerifyOnDestVolumes
<src_storage_system_list:src_vol:dest_storage_system:dest_vol>
   -Verify <True|False>                -VerifyMetadata <True|
False>                -ExhaustiveVerification <True|False>                -
Server <server_name>                -DestinationServer
<dest_server_name>                -AutoMount <True|False>                -
TestRestore <True|False>                -storagegroup
<storage_grp>                -targetstoragegroup
<dest_storage_grp>                -Database <database1,
```

```
database2, ...>                -RestoreFromUnmanagedMedia <True|
False>                -Rehomemailbox <True|False>                -
BkUpServer <backup_server_name>                -SnapInfoDirectory
<snapinfo_dir_path>                -PointInTime <True|
False>>                -VerDestVolume <True|False>                -
VerificationServer <verf_server_name>                -CCRActiveNode
Boolean <True|False>                -OverrideVer <True|
False>                -CheckLog <True|False>                -Destination
<storage_grp_name>                -DestinationServer
<dest_server_name>                -RecoveryPointTime
<recvry_pt_time_stamp>                -CancelBackup <True|
False>                -WaitForBackupComplete <True|False>                -
RecoveryPoint <True|False>                -RestoreArchivedBackup <True|
False>-NoAccessToRemoteBackup <True|False>                -
VerifyArchiveBackup <True|False>                common parameters
```

Description

This command enables you to restore backup sets with all of the options available through the GUI.

This command also supports the following common parameters:

- `-Debug` (`-db`)
- `-ErrorAction` (`-ea`)
- `-ErrorVariable` (`-ev`)
- `-OutBuffer` (`-ob`)
- `-OutVariable` (`-ov`)
- `-Verbose` (`-vb`)
- `-Confirm`

To learn more about common parameters, see help about_ubiquitous_parameters.

Parameters

`-Backup` *<name_of_the_backup>*

Short form: `bkup`

The name of the backup set that you want to restore.

`-RestoreLastBackup` *<restore_last_backup>*

Short form: `rstlast`

Restores backup copies without specifying the name.

If you try to use `Backup` and `RestoreLastBackup` together, SnapManager ignores `RestoreLastBackup` and uses `Backup` during the restore operation.

A typical usage example of the `RestoreLastBackup` parameter is as follows:

```
restore-backup -restorelastbackup = 1 -backup = “backup name”
```

If the value of `RestoreLastBackup` is 1, SnapManager ignores this parameter and uses the `Backup` during the restore operation.

Note: The default value of this parameter is 0, which means that SnapManager restores the latest backup. If the value is 1, SnapManager restores the second-to-latest backup.

-VerifyOnDestVolumes ***<src_storage_system_list:src_vol:dest_storage_system:dest_vol>***

Short form: `vermirror`

Overrides the existing SnapMirror relationships.

-Verify <True|False>

Short name: `ver`

Verifies the backed up SnapManager databases and transaction logs. `-Verify` is a switch parameter.

-VerifyMetadata <True|False>

Short form: `vermetadata`

Verifies the metadata and transaction logs.

-ExhaustiveVerification <True|False>

Short form: `exhver`

Performs exhaustive database verification.

-Server ***<server_name>***

Short name: `svr`

Specifies the target Exchange server name.

In a cluster configuration, you need to specify `-Server` explicitly in all of the cmdlets to perform all operations. If `-Server` is not specified explicitly in a stand-alone server, SnapManager uses the local machine as the default to run the following cmdlets: `new-backup`, `verify-backup`, `restore-backup`, `get-backup`, and `delete-backup`.

-DestinationServer ***<dest_server_name>***

The name of the target server where the Recovery Storage Group is to be created.

`DestinationServer` is specified to restore to a Recovery Storage Group during the restore operation.

-AutoMount <True|False>

Short from: `mt`

Mounts the databases automatically after the restore operation.

-TestRestore <True|False>

Short form: `test`

Performs test restore operation.

The default value for `TestRestore` is False.

-StorageGroup ***<storage_grp1, storage_grp2, ...>***

Short name: `sg`

Specifies the list of Storage Groups to be verified during the backup operation.

If invalid Storage Groups are specified in `new-backup` during backup operation, PowerShell cancels the `new-backup` operation. SnapManager displays an appropriate error message and does not selectively backup any valid Storage Groups. Scheduled jobs and backup scripts also fail when there is an invalid Storage Group configuration.

-targetstoragegroup ***<dest_storage_grp>***

Short form: `tarsg`

Specifies the destination Storage Group.

-Database ***<database1, database2, ...>***

Short form: `dbs`

Specifies the list of databases separated by commas.

If you do not specify `Database`, SnapManager restores all of the databases.

-RestoreFromUnmanagedMedia <True|False>

Short form: `rstumm`

Restores backup sets that are archived on a server, other than the server they were created on. This is a switch parameter.

> **Note:** You can restore the backup copies created on a different Exchange server to the server on which the restore operation was initiated. Before you can restore databases to server other than your current Exchange server, you must remap the source LUNs to the current Exchange server, using the same drive letter that was assigned to the original Exchange server.

-Rehomemailbox <True|False>

Updates the user accounts associated with mailboxes in restored databases to point to the mailbox server with the new name.

This is an optional parameter with the `restore-backup` cmdlet of another server, or `RestoreFromUnmanagedMedia` in Exchange Server 2007. You cannot use it to restore to a Recovery Storage Group.

-BkUpServer ***<backup_server_name>***

Short form: `bksvr`

Specifies the name of the server on which the backup copy was created. Use this parameter only with `RestoreFromUnmanagedMedia` and `RestoreFromServer` where the backup copy was originally created.

-SnapInfoDirectory ***<snapinfo_dir_path>***

Short form: `sifdir`

Specifies the SnapInfo directory path for the archived backup set during the restore operation. Use this parameter only with `RestoreFromUnmanagedMedia` parameter and the `RestoreFromServer` parameters.

`-PointInTime <True|False>>`

Short form: `pit`

Performs a point-in-time restore operation.

`-VerDestVolume <True|False>`

Short form: `verdest`

Verifies the SnapMirror destination volume.

`-VerificationServer` *<verf_server_name>*

Short form: `versvr`

Overrides the preconfigured SnapManager verification settings. It specifies the host to be used as the verification server for the verification phase of a backup operation.

`-CCRActiveNode` Boolean `<True|False>`

Short form: `ccrnode`

Specifies whether the connection is to an active, or a passive node, respectively. This is a Boolean parameter.

`-OverrideVer <True|False>`

Short form: `ovr`

Overrides the verification of the databases.

`-CheckLog <True|False>`

Short from: `chklog`

Specifies the transaction logs to be restored.

`-Destination` *<storage_grp_name>*

Specifies where the backup copy need to be restored.

You can restore to the same Storage Group or to the Recovery Storage Group. The default value is `tosamesg`. To restore the backup copy to the Recovery Storage Group, enter `torsg`.

`-DestinationServer` *<dest_server_name>*

Short form: `dstsvr`

Specifies the name of the destination Exchange server.

`-RecoveryPointTime` <recvry_pt_time_stamp>

Specifies the recovery point timestamp.

The timestamp for each recovery point can be seen from the output of the `get-backup` cmdlet. If the specified timestamp does not match any of the recovery points shown in the backup copies, the `restore-backup` cmdlet returns an error message showing the available recovery points before and after the timestamp.

Note: `-RecoveryPointTime` option overrides `PointInTime` if you specify both.

-CancelBackup <True|False>

Pauses all the active scheduled backup jobs on the current Exchange Server, or on all nodes in the cluster environment, and cancels the current backup copy before performing the restore operation.

When the restore operation completes, SnapManager enables the paused scheduled backup jobs only. All the other inactive jobs do not change.

-WaitForBackupComplete <True|False>

Pauses all the active scheduled backup jobs on the current Exchange Server, or on all nodes in the cluster environment, and waits for the current backup operation to complete before performing the restore operation. When the restore operation completes, SnapManager enables only the paused scheduled backup jobs. All the other inactive jobs do not change.

-RecoveryPoint <True|False>

Specifies if the backup set is a Frequent Recovery Point backup.

-RestoreArchivedBackup <True|False>

Short form: `rstarchbkup`

Restores database from an archived backup.

-NoAccessToRemoteBackup <True|False>

Short form: `noaccessarchivebkup`

Specifies that there is no direct access to the secondary storage system.

SnapManager uses the proxy server to access the secondary storage system.

-VerifyArchiveBackup <True|False>

Short name: `verarch`

Specifies that the backup copy that has to be verified is an archived backup copy. If you do not specify `VerifyArchiveBackup`, and one local and one archived backup copy exist with the same name, SnapManager verifies the local backup copy.

Example: Restoring database

This command restores exchsnap__SNAPMGR-55_11-10- 2006_13.36.24 to the Storage Group First Storage Group.

```
restore-backup -server SNAPMGR-48 -storagegroup 'First Storage Group'
-backup exchsnap__SNAPMGR-55_11-10- 2006_13.36.24
```

Example: Restoring from an archive

This command restores exchsnap__SNAPMGR-54_11-10-2006_14.47.18 that was created on the archived server SNAPMGR-54.

```
restore-backup -server SNAPMGR-48 -StorageGroup 'First Storage Group'
-RestoreFromUnmanagedMedia -BkUpServer SNAPMGR-54 -backup
exchsnap__SNAPMGR-54_11-10-2006_14.47.18 - SnapInfoDirectory 'K:
\SME_Snap\InfoEXCH__SNAPMGR-48A\SG__WZ00\12- 04-2006_14.47.18'
```

Example: Restoring backup sets created on different Exchange Server

This command restores exchsnap__SNAPMGR-54_11-10-2006_14.47.18 that was created on the server SNAPMGR-54.

```
restore-backup -server 'SNAPMGR-48' -StorageGroup 'First Storage
Group' -RestoreFromUnmanagedMedia -BkUpServer SNAPMGR-54 -backup
exchsnap__SNAPMGR-54_11-10-2006_14.47.18 - SnapInfoDirectory 'K:
\SME_SnapInfo'
```

Example: Restoring backup sets to a Recovery Storage Group

This command restores exchsnap__SNAPMGR-55_11-10- 2006_13.36.24 to the Recovery Storage Group Test_rsg.

```
restore-backup -server SNAPMGR-48 -storagegroup 'First Storage Group'
-backup exchsnap__SNAPMGR-55_11-10- 2006_13.36.24 -destination torsg -
TargetStorageGroup Test_rsg
```

Example: Restoring a specified recovery point time backup

This command restores exchsnap__snapmgr-50_03-01-2007_08.00.00 at the recovery point time 03-01-2007_08:55:00.

```
Restore-backup -Server snapmgr-50 -Storagegroup 'First Storage Group'
-backup exchsnap__snapmgr-50_03-01-2007_08.00.00 - RecoveryPointTime :
03-01-2007_08:55:00
```

Get-JobStatus

The `Get-JobStatus` command enables you to view the status of the queued, running, and finished jobs.

```
        Get-JobStatus-Server <Exchange_server_name>-ShowChildJobs <True|
False>common parameters
```

Description

Specify the server name to view a particular job status. This command also supports the following common parameters:

- `-Debug` (`-db`)
- `-ErrorAction` (`-ea`)
- `-ErrorVariable` (`-ev`)
- `-OutBuffer` (`-ob`)
- `-OutVariable` (`-ov`)
- `-Verbose` (`-vb`)
- `-Confirm`

To learn more about common parameters, see help about_ubiquitous_parameters.

Parameters

`-Server` *<Exchange_server_name>*

Short name: `svr`

Specifies the name of the Exchange server for which you monitor the job status.

If you do not specify this parameter, the name of the server that runs this cmdlet becomes the default host name. `-Server` is an optional parameter

`-ShowChildJobs <True|False>`

Short name: `cj`

Displays all the child jobs of the running and the finished jobs.

`-ShowChildJobs` is an optional parameter.

Example: Displaying all jobs

This command displays all the jobs that are handled by the Exchange server Exchange1.

```
Get-JobStatus -Server Exchange1
```

Example: Displaying child jobs

This command displays all the child jobs of the running jobs, the finished jobs, and the parent-level jobs that are managed by the Exchange server Exchange1.

```
Get-JobStatus -Server Exchange1 -ShowChildJobs
```

Change-JobPriority

If a job is queued, `Change-JobPriority` enables you to move a SnapManager job to a different priority in the queue.

```
        Change-JobPriority-Server<Exchange_server_name>-
JobID<numeric_job_id>-Priority<position_of_the_job>-
SourceBackupServer<name_of_server_that_creates_the_backup>
```

Description

`Change-JobPriority` enables you to move a job into a different priority in the queue. You can view the current queue with the `Get-JobStatus` command.

This command also supports the following common parameters:

- `-Debug` (`-db`)
- `-ErrorAction` (`-ea`)
- `-ErrorVariable` (`-ev`)
- `-OutBuffer` (`-ob`)
- `-OutVariable` (`-ov`)
- `-Verbose` (`-vb`)
- `-Confirm`

To learn more about common parameters, see help about_ubiquitous_parameters.

Parameters

`-Server` *<server_name>*

Short name: `svr`

Specifies the target Exchange server name.

In a cluster configuration, you need to specify `-Server` explicitly in all of the cmdlets to perform all operations. If `-Server` is not specified explicitly in a stand-alone server, SnapManager uses the local machine as the default to run the following cmdlets: `new-backup`, `verify-backup`, `restore-backup`, `get-backup`, and `delete-backup`.

`-JobID`*<numeric_job_id>*

Short name: `id`

Used to identify a particular job that is being handled by the SnapManager server. `-JobID` is a required parameter if you do not specify `-AllJobs`.

`-Priority`*<position_of_the_job>*

Short name: `p`

Specifies the position to which you want to move it. It is a required parameter.

`-SourceBackupServer`*<name_of_server_that_creates_the_backup>*

Short name: `bksvr`

Specifies the name of the Exchange server that creates the backup. This is an optional parameter. If you do not specify this parameter, the name of the source backup server specified by the parameter `-Server` becomes the default name for the server.

Example: Changing job priority

This command changes the priority of the deferred integrity verification job that is queued in the remote verification server VerificationServer1 (with Job ID 123) to priority 1.

```
Change-JobPriority -Server VerificationServer1 -SourceBackupServer
Exchange1 -JobID 123 Priority 1
```

Cancel-Job

The `Cancel-Job` command enables you to cancel the jobs that are in queued or in running state.

```
Cancel-Job-Server <Exchange_server_name>-JobID <numerical_job_id>-AllJobs
<True|False>common parameters
```

Description

If the job is in the queue, SnapManager removes the job from the queue. If the job is running, the cmdlet cancels the running job. This command also supports the following common parameters:

- `-Debug` (`-db`)
- `-ErrorAction` (`-ea`)
- `-ErrorVariable` (`-ev`)
- `-OutBuffer` (`-ob`)
- `-OutVariable` (`-ov`)
- `-Verbose` (`-vb`)
- `-Confirm`

To learn more about common parameters, see help about_ubiquitous_parameters.

Parameters

`-Server` ***<Exchange_server_name>***

Short name: `svr`

Specifies the name of the Exchange server.

If you do not specify this parameter, the default host is the name of the server that runs this cmdlet. `-Server` is an optional parameter.

`-JobID` ***<numerical_job_id>***

Short name: `id`

Identifies a particular job that is handled by the SnapManager server.

`-JobID` is a required parameter if you did not specify `-AllJobs`.

`-AllJobs <True|False`

Short name: `all`

If this parameter is set to true, all the jobs including those in running and in queued states are cancelled.

`-AllJobs` is a required parameter if you did not specify `-JobID`.

Example: Cancelling a job

This command cancels the job running in the Exchange server Exchange1 with Job ID 123.

```
Cancel-Job -Server Exchange1 -JobID 123
```

Example: Cancelling a job managed by a remote integrity verification server

This command cancels all the queued and running jobs managed by the remote integrity verification server.

```
Cancel-Job -Server VerificationServer1 -AllJobs
```

exec-bc

The `exec-bc` command enables you to execute the SnapManager Business Continuance plan.

```
        exec-bc-ExchServer <Exchange_instance>-BcServer
<destination_Business_Continuance_server>-BcPlan
<Business_Continuance_plan>-FailoverTo <destination_storage_site>-Tasks
<name_of_task>-cleanup_site <True|False>common parameters
```

Description

The `exec-bc` cmdlet enables you to execute the disaster recovery failover of an Exchange instance to a destination Business Continuance server.

This command also supports the following common parameters:

- `-Debug` (`-db`)
- `-ErrorAction` (`-ea`)
- `-ErrorVariable` (`-ev`)
- `-OutBuffer` (`-ob`)
- `-OutVariable` (`-ov`)
- `-Verbose` (`-vb`)
- `-Confirm`

To learn more about common parameters, see help about_ubiquitous_parameters.

Parameters

`-ExchServer` *<Exchange_instance>*

Short name: `exsvr`

Specifies the name of the Exchange instance.

By default, it assumes the name of the local machine.

`-BcServer` *<destination_Business_Continuance_server>*

Short name: `drsvr`

Specifies the name of the destination Business Continuance server.

By default, it assumes the name of the local machine.

`-BcPlan` *<Business_Continuance_plan>*

Short name: `plan`

Specifies the name of the Business Continuance plan.

The Business Continuance plan must exist in the disaster recovery server and must be valid for the specified Exchange Server.

`-FailoverTo` *<destination_storage_site>*

Short name: `to`

Specifies the name of the destination storage server for Business Continuance failover.

`-FailoverTo` is an optional parameter. I

`-Tasks` *<name_of_task>*

The following tasks are executed using this parameter:

- validate_site
- break_mirrors
- connect_luns
- recreate_ex

- restore_bkup
- offline_exchange

If you do not specify a particular task, this cmdlet executes all the tasks.

`-Tasks` is an optional parameter.

`-cleanup_site <True|False>`

Performs the cleanup tasks of the destination host server.

Example: Executing a Business Continuance failover

This command executes a Business Continuance failover of an Exchange instance exchtstx to the destination Business Continuance server using the Business Continuance plan plantest.

```
exec-bc -exsvr exchtstx -bcserver DR54 -plan plantest -confirm -
verbose
```

Example: Validating a Business Continuance plan

This command validates a Business Continuance plan for failover of an Exchange instance exchtstx to the destination Business Continuance server using the Business Continuance plan plantest.

```
exec-bc -exsvr exchtstx -bcserver DR54 -plan plantest -confirm -
verbose -tasks validate_site
```

Example: Connecting LUNs and re-creating Exchange instance tasks in a Business Continuance server

This command connects LUNs and re-creates Exchange instance tasks in the Business Continuance Server.

```
exec-bc -exsvr exchtstx -bcserver DR54 -plan plantest -confirm -
verbose -tasks connect_luns, recreate_ex
```

Example: Skipping validation

This command skips the validation operation and executes the tasks that you explicitly specify.

```
exec-bc -exsvr exchtstx -bcserver DR54 -plan plantest -confirm -
verbose -tasks break_mirrors,connect_luns,recreate_ex,restore_bkup
```

Export-config

This cmdlet enables you to export the SnapManager configuration control file.

```
Export-config                -Server <Exchange_server_name>
-ControlFilePath <name_of_control_file_and_path>                -Section
<comma_separated_list_of_section_names>                <common_parameters>
```

Description

The `Export-config` cmdlet enables you to export the SnapManager control file that contains the configuration information which you can later use to configure SnapManager on other systems by using the `import-config` command.

Parameters

`-Server`<***Exchange_server_name***>

Short name: svr

Specifies the name of the Exchange server whose configuration you want to export as an XML control file.

`-ControlFilePath`*<name_of_control_file_and_path>*

Short name: config

Specifies the output XML file name and path.

`-Section`*<comma_separated_list_of_section_names>*

Short name: sect

Specifies the list of section names separated by comma to export.

You can export the following sections of the control file:

storage, notification, verification, report, backup, scheduledjob, snapmirrorvolume.

If you do not specify `-Section`, SnapManager assumes that all sections should be exported.

Example: Exporting a control file

This command exports the specified configuration to the `SMEConfig_12_18_2007_01.12.57.xml` control file.

```
export-config -Server Exchange1 -ControlFilePath "C:\Program Files
\NetApp\SnapManager for Exchange\SMEConfig_12_18_2007_01.12.57.xml"
```

Import-config

The `Import-config` cmdlet enables you to import the SnapManager configuration control file.

```
Import-config                -Server <Exchange_server_name>
-ControlFilePath <name_of_control-file_and_path>                -Section
<comma_separated_list_of_section_names_to_import                -AllowLocal
<true | false>                -ValidateAndApply <true |
false>                -Username <username>                -Password
<password>                -ClusterAware <true | false>
```

Description

This command enables you to import the SnapManager control file that contains the server configuration information. You can import either a section of the control file, or the complete control file.

This command also supports the following common parameters:

- `-Debug` (`-db`)
- `-ErrorAction` (`-ea`)
- `-ErrorVariable` (`-ev`)
- `-OutBuffer` (`-ob`)
- `-OutVariable` (`-ov`)
- `-Verbose` (`-vb`)
- `-Confirm`

To learn more about common parameters, see help about_ubiquitous_parameters.

Parameters

`-Server` <***Exchange_server_name***>

Short name: svr

Specifies the name of the Exchange server to which you want to import the control file.

`-ControlFilePath` <***name_of_control-file_and_path***>

Short name: config

Specifies the location of the control file to import.

Specify `-ControlFilePath` with the control file name. If the control file is not in the current directory, the full file name path must be given.

`-Section` ***<comma_separated_list_of_section_names_to_import>***

Short name: sect

Specifies the comma-separated list of names of sections to import.

You can import the following sections from the control file:

storage, notification, verification, report, backup, scheduledjob, snapmirrorvolume.

If you do not specify particular sections, SnapManager imports all sections in the control file.

`-AllowLocal <true | false>`

Short name: tolocal

Migrates the databases to the local disk.

`-ValidateAndApply <true | false>`

Short name: apply

This is an optional parameter. By default, the value is false, that indicates to perform only validation. If the value is set to true, it indicates to perform validation and apply it if the validation is successful.

`-Username <`*`username`*`>`

Short name: usr

Verifies the user name before creating a scheduled job.

`-Password<`*`password`*`>`

Short name: pwd

Verifies the user credentials before creating a scheduled job.

`-Clusteraware <True|False>`

Short form: `cl`

Assumes significance only when scheduling jobs in cluster configurations, by facilitating scheduling the same job in multiple cluster nodes to improve fault tolerance.

In the case of Single Copy Cluster (SCC), if a job is scheduled with `-Clusteraware`, the job runs only if the host in which it is scheduled is the active node of the cluster.

In the case of Cluster Continuous Replica (CCR), the processing of `-Clusteraware` depends on whether the host in which it is scheduled is an active or passive node and the value of `-CCRActiveNode`. The job runs if the host is a CCR active node and the job targets it with `-CCRActiveNode` as true. If the host is a CCR passive node, the job runs if the job targets the host with `-CCRActiveNode` false.

Example: Importing sections of the control file

This command imports the specified sections from the control file `sme_config.xml` to Exchange server Exchange1.

```
import-config -Server Exchange1 -ControlFilePath "C:\Program Files
\NetApp\SnapManager for Exchange\sme_config.xml" -
Section
"storage,notification,verification,report,backup,scheduledjob,snapmirr
orvolume" -ValidateAndApply
```

Copyright information

Copyright © 1994–2012 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP 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.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, FAServer, FastStak, FilerView, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the web at *www.ibm.com/legal/copytrade.shtml*.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by email to *doccomments@netapp.com*. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

A

Actions pane 19
Activation Preference Number 126
Active databases 124
age-based Snapshot copy deletion 186
application
    settings 231
application settings
    accessing 231
archives 228
archiving
    Exchange Backup Agent 212
    automatic 213
    Exchange Backup Agent
        example 212
    guidelines 205
    methods 206
    NDMP and dump
        evaluation 207
    Windows backup utility
        evaluation 210
        example 210
automatic deletion of backups 111, 187
automatic event notification
    configuring 252
AutoSupport
    configuring 252

B

backup archives
    creating with Exchange Backup Agent 211
backup copies
    archiving 211
    problem deleting 193
    restoring from different servers 176
backup copy
    unmounting 185
backup jobs
    planning for databases in DAGs 120
    using the new-backup command 262, 263
backup management group
    changing 136
backup model
    centralized 212
backup sets
    deleting 186
    names 100
backup settings
    XML schema 260
backup tasks 102
backups
    typical process using VSS 102
break-mirrors 285
Business Continuance
    limitations in SnapManager 224
    SnapManager control file
        importing or exporting 86
busy Snapshot copy 109, 193

C

Cancel-Job command 295, 296
CCR
    restoring 164
CCR backup set
    behavior 164
CCR replica database 127
CCR server
    connecting 55
CCR-enabled databases backups 126
CCR-enabled Storage Group
    corruption 166
centralized backup model 212
Change-JobPriority command 294
checksum verification 106
ChkSgFiles integrity verification library 240
ChkSgFiles.dll 106
cleaning up after
    mailbox data 185
cluster continuous replication databases
    migration of 91
cmdlets
    managing remote additional backups 141
command
    specifying 214
commands
    new-backup 262, 263
concurrent verification 201, 202
configuration
    migrating 96
    migrating from drive letters to mount points in Exchange 2007 96
    of Exchange data 72

of gapless DAG backups 141
configuration and version check
SnapManager 53
Configuration wizard
migrating Exchange data
local disk to LUN 73, 74
LUN to local disk 73, 74
LUN to LUN 73, 74
migrating Exchange databases 73
settings 75
when to use 74, 75
Configuration Wizard 60
conformance status 228
Continuous Cluster Replica (CCR) 38
control file
exporting 86
XML schema 255
conventional backup processes 19
creating LUNs on qtrees
requirements 71

D

DAG
backups 140
See also Database Availability Group
DAG backup jobs
planning 120
DAG backups
configuring using remote additional copy 141
dashboard
view 55, 56
Dashboard view 19
data store configurations 61
data store rules 60
database
automatic unmounting 248
changing location 81
migration 233
migration of 78
restore 145
restoring
Restore window 154
Restore Wizard 153
restoring Frequent Recovery Point 156
restoring LCR replica Storage Group 162
verification
server 233
throttling 237
Database Availability Group
SnapManager requirements in 23
See also DAG
Database Availability Group (DAG) 57, 83, 86, 94, 117, 124, 138, 158, 159, 181, 196
Database Availability Groups
backups 140
planning backups for more than three nodes 120
planning backups for three nodes 120
database backup jobs
planning for in DAGs 120
Database Filter 124
database migration 58
database path
viewing 82
database restoration
individual within multiple storage groups 69
database seeding
LCR (Local Continuous Replication) 90
Local Continuous Replication (LCR) 90
database verification
concurrent 201, 202
deferred 107
load management 107
scheduling 134
settings 133
databases
deleting 186
reseeding large 159
reseeding using PowerShell 159
restoring multiple 175
dataset
configuring 225
creating 226
functionality 222
naming 226
dataset and SnapVault integration
prerequisites 224
dataset configuration
limitations in SnapManager 224
dataset integration with SnapVault 221
dataset protection policy 83
datasets
editing 227
general concepts 221
datasets concepts 223
default server
connecting 53
deferred database integrity verification 229
Deferred database verification 108
delete backup dialog box 186

delete-backup command 276
deleted item retention 115
disaster recovery 204
Disaster Recovery
    SnapManager control file
        importing or exporting 86
Drive letter limitations 63
drive letters
    migrating to mount points for LUN mapping in Exchange 2007 96
dump command
    using to archive LUNs 206

E

e-mail notification
    configuring 252
errors
    busy Snapshot copy 193
    during restore operations 182
    OxC00413A 182
    VSS_E_WRITERERROR_RETRYABLE 182
Eseutil consistency checker 240
ESM
    See Exchange Server Manager
events
    configuring notification 252
examples
    unattended upgrade 44
Exchange
    adding servers 53
    permission level 25
    Service identity account 26
Exchange 2010 94
Exchange and SnapManager components
    placement of 76
    viewing placement of 77
Exchange Backup Agent
    using to archive Exchange backup copies 211
Exchange backup copies
    archiving 211
Exchange backups
    automatic deletion 248
Exchange configuration requirements 58, 59
Exchange configurations 68
Exchange data
    configuration 72, 78
    migration 72, 78
    moving back to local disk 46
Exchange database migration
    Using the Configuration wizard to perform 73
Exchange databases
    backing up using new-backup 262, 263
Exchange page zeroing 115
Exchange server
    configuring the control file 87, 88
    default server 232
    LCR (Local Continuous Replication)
        configuration prerequisites 89
    Local Continuous Replication (LCR)
        configuration prerequisites 89
Exchange server 2010 83
Exchange Server 2010 86
Exchange Server Manager
    using to restore a database to a Recovery Storage Group 172
Exchange Server requirements 27
Exchange Service
    identity account 26
    permissions 25
Exchange Service identity account
    requirements 26
Exchange System Manager 74, 75, 114, 115
Exchange virtual server 149
Exchange Writer 126
Exchanger Server 2010 94
exec-bc command 296, 297
Exhaustive transaction log sequence checks 148
explicit deletion 188
external backups 19

F

firewall configuration 22
firewall rules 22
FlexClone license
    requirements for integrity verification 197
flexible volumes
    automatic expansion 244
FlexVol volumes
    automatic expansion 244
fractional space reservation
    configuring policy 245
    monitoring 250
    policies
        configuring 246
    status data 250
    viewing current data 249
fractional space reservation policies
    managing Exchange data 247

fractional space reservation policy
default 245
settings 245
Fractional space reservation policy settings 245
fractional space reserve 244
Frequent Recovery Point
backup 138
clustered configurations 137
operations 137
reports 138
restoration 138
verification 138
working 137
Frequent Recovery Point backups 156
FRP
See Frequent Point backups

G

gapless DAG backups
configuring using remote additional copy 141
get-backup command 278
Get-JobStatus command 293
get-mirrors 280, 281

H

hardware requirements 27

I

individual database
move to LUN 81
individual database restoration
within multiple storage groups 69
installation
interactive 33, 34
unattended 33, 35
using a script 35
installing SnapManager
preparation for 21
Windows cluster 38
integration limitations
dataset and SnapVault 224
with SnapManager 224
integrity verification
destination SnapMirror volume 107, 108
remote 200
requirements 197
troubleshooting 197
interactive SnapManager installation 34
interactive upgrade
compared with unattended upgrade 42
IP addresses
multiple for storage system 28

J

job-specific parameters 133
JobID parameter 294

L

LCR (Local Continuous Replication) enabled database
migration 89
LCR replicas
restore operation 162
LCR-enabled database backups 126
LCR-enabled Storage Group
corruption 163
legacy scheduled jobs 45
licenses
applying 25
requirements 25
licensing
Per Server 23
Per Storage system 23
licensing requirements 27
limitations
Recovery Storage Group 167
local backup protection 228
local continuous replication
database seeding 90
Local Continuous Replication (LCR) enabled database
migration 89
log database signature 148
LUN Clone Split Restore method 149
LUN clone split status
verifying 149
LUN creation
on qtrees with UNIX secure volumes 71
LUN mapping in Exchange 2007
migrating drive letters to mount points 96
LUNs
disconnecting 175
migrating to mountpoints 96
write operations stopped 244

M

mailbox data
    cleaning up after 185
    restoring 184
mailbox database 94
mailbox databases 94
mailboxes
    restoring from archive 153
MCSC clusters
    restore restriction 217
message tracking 62
migrating 94
migration
    of Exchange data 72
Migration 94
mount options 153
mount points
    migration to from drive letters for LUN mapping in Exchange 2007 96
mounted volume 64, 65
mounted volume restrictions 64
mountpoints
    migrating LUNs 96
move group operation 149
MSCS 62
MSCS clusters
    report directory location 217
multiple FRPs 139
multiple SnapInfo directories 67
multiple Snapshot copies
    deleting 186
multiple storage groups
    restoring individual databases within 69

N

naming settings
    definition of 221
NDMP
    using to archive LUNs 206
NDMP and dump command
    example 207
Net.Tcp port sharing service 164
network port configuration 22
new-backup command 262, 263
Non-exhaustive transaction log sequence checks 149
notification
    XML schema 257
NTFS hard links 66
NTFS volume mount points 62
NTFS volume mountpoints 62

O

operating system requirements 27

P

page zeroing 115
Passive database copies 124
per-storage-system license 25
planning backups 105
point-in-time restore 147
policies
    fractional space reservation 244
Priority parameter 294
problem launching scripts 215
production database 127
production Exchange server 109
production site
    maintenance 204
Protection Manager 223, 226
protection or provisioning related objects
    definition of 221
protection policies
    overview 221
protocol requirements 27
provisioning policies
    overview 221

Q

qtrees
    requirement for creating LUNs on 71
quantity-based Snapshot copy deletion 186
queue location
    SMTP and MTA 78

R

Recover Storage Group
    disconnecting LUNs
        Exchange Server 2003 175
recovering
    mailbox data 184
Recovery Database
    delete 182
Recovery Storage Group
    adding database

Exchange Server 2003 174
adding databases
Exchange Server 2007 170
creating
Exchange Server 2003 174
creating using Exchange Server 2003 172
destroying
Exchange Server 2007 172
Exchange Server 2007 170
limitations 167
mounting databases
Exchange Server 2007 171
overview 167
restoring
Exchange 2007 167
multiple databases 175
unverified backup copy 176
restoring a database to 172
reinstallation
interactive 50
unattended 50
reinstalling SnapManager
in unattended mode 51
remote additional backup
control.xml example 143
remote additional backups
controlling retention 143
defining servers 143
deleting 143
managing 141
number retained 143
remote additional copy
using to configure gapless DAG backups 141
remote administration
server requirements 29
remote backup
restoring 230
remote backup retention 229
remote backups 228
remote database verification 107, 108
remote verification
prerequisites 235
server requirements 30
working 235
remote verification server 109
report directories
Windows cluster 78
report directory
changing 217, 218
clustered configuration 218
sharing
account permissions 26
XML schema 259
reports
DAG 217
deleting 219
printing 219
viewing 219
requirements
Exchange Service identity account 26
integrity verification 197
remote administration server 29
remote verification server 30
reseeding 159
restore
backup copies
different server 176
backup sets from unmanaged media 177
decisions 151, 152
destination 152
live Exchange virtual server cluster 150
mailbox data 184
point-in-time 147
Snapshot copies 147
type 152
up-to-the-minute 147
verification options 152
Windows cluster 149
working 146
restore error 0xC00413CA 182
restore error VSS_E_WRITERERROR_RETRYABLE 182
restore from different server 158
restore operation
Frequent Recovery Point 156
guidelines 151
LCR replicas 162
multiple databases 175
Restore window method 154
Restore wizard method 153
unverified copies 176
restore operation failures 182
restore process
CCR replicas
reseeding requirements 164
primary database
reseeding requirements 164
restore Snapshot copies
explicit deletion 192
restore time

decreasing 148
restore troubleshooting 182
restore-backup command 286, 287
restoring 184
restrictions
mounted volume 64
Results pane 19
retention policies
limitations in SnapManager 224
rstrsnap__ files
deleting 147
Run Command After Operation option
command arguments 213
configuring 242

S

scenarios
unattended upgrade 44
scheduled jobs 57
Scope pane 19
scripts
launching from UNC paths 215
server cluster mountpoints 62
Server parameter 294
settings
application 231
SFSR (Single-File SnapRestore) method 150
single copy cluster 38
Single Mailbox Recovery (SMBR) 184, 185
Single-File SnapRestore (SFSR) method 150
SMTP and MTA
queue location 78
SnapDrive
version verification 53
SnapDrive for Windows
requirement in DAG environment 23
SnapInfo directory 103
SnapInfo directory Snapshot copy names
location 99
name 99
subdirectory 99
SnapInfo files
migration of 78
SnapInfo Snapshot copies
explicit deletion 190
SnapInfo Snapshot copies 190
SnapManager
builds 33
command-line reference 262
coordination with Protection Manager 226
installation 33
starting 53, 55
upgrade 33
SnapManager Backup 99
SnapManager control file
import or export
Business continuance 86
Disaster Recovery 86
SnapManager for Exchange PowerShell
launching 262
SnapManager installation
interactive 34
SnapManager reinstallation
considerations 49, 50
interactive 50
SnapManager snap-in 19, 86, 217
SnapManager software license agreement
displaying 37
SnapManager uninstallation
interactive 47
prerequisites 45
unattended 47
SnapManager upgrade
fractional space reserve monitoring 42
interactive mode 43
preparation 41
unattended mode 43
SnapManager upgrade path 42
SnapMirror
destination volumes 196
documentation 194
SnapMirror license
requirements for integrity verification 197
SnapMirror relationship
XML schema 260
Snapshot copies
automatic deletion 244
Snapshot copies per volume 110
Snapshot copy
access method 236
Snapshot copy busy 193
Snapshot copy naming conventions 101, 104
SnapVault 223
SnapVault integration 221
SnapVault relationships 227
SourceBackupServer parameter 294
Storage Group
CCR (Continuous Cluster Replica)
configuration prerequisites 91

changing location 81
cluster continuous replication
configure 92
considerations 92
Continuous Cluster Replica (CCR)
configuration prerequisites 91
move to LUN 80
Storage Group sets 101
storage groups
deleting 186
restoration within multiple 69
Storage Groups
renaming limitation 151
restore operation guidelines 151
storage layout
XML schema 255
storage services
overview 221
storage system
multiple IP addresses 28
requirements 29
syslog event logging
configuring 252
system resources
back up 23

T

test restore 199
throttling report
entries 240
transaction log archiving 66
Transaction log sequence verification 148
transaction logs
change configuration 84
changing location 84
deleting 111–113
migration of 78
moving to LUN 84
truncated 114
viewing full path 85
transport database 83
troubleshooting
busy Snapshot copy 193
integrity verification 197
problem deleting backup copies 193
restore operations 182
type of restore
guidelines 146, 147, 150

U

unattended installation
examples 38, 52
unattended reinstallation 51
unattended uninstallation
examples 49
unattended upgrade
compared with interactive upgrade 42
examples 44
UNC path 215
uninstallation
interactive 46
unattended 46
unmounting
back up copy 185
up-to-the-minute restore 147
Up-to-the-minute restore 180
up-to-the-minute restore ability 188
upgrade modes
interactive 42
unattended 42
upgrades
examples of unattended 44
upgrading SnapManager
preparation for 21

V

verification
accessing Snapshot copies 236
deferred 200
limitations in SnapManager 224
override 241
performance
impact 240
selecting the server 235
sleep interval
calculating 238
throttling
configuring 239
options 238
working 237
XML schema 258
verification override
report entry 240, 242
verification server
configuring 234
verification status reporting 109
verification throttling 107
Verification throttling 108
verify-backup command 271
verifying multiple backup sets 106
version verification

SnapDrive 53
volume mountpoints 63
volumes
fractional reserve 243
VSS
hardware provider 18
hardware requirement 18
typical backup process 102
VSS Snapshot copy 103

W

Windows
host system requirements 27
Windows backup utility 209
Windows cluster
impact on SnapManager reports 218
report directories 78
restore operation 149
Windows host system requirements 27