

IBM Spectrum Protect for Mail
Version 8.1.2

*Data Protection for Microsoft Exchange
Server
Installation and User's Guide*

IBM

IBM Spectrum Protect for Mail
Version 8.1.2

*Data Protection for Microsoft Exchange
Server
Installation and User's Guide*

IBM

Note:

Before you use this information and the product it supports, read the information in “Notices” on page 237.

This edition applies to version 8, release 1, modification 2 of IBM Spectrum Protect for Mail: Data Protection for Microsoft Exchange Server (product number 5725-X02) and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 1998, 2017.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Tables	vii
-------------------------	------------

About this publication	ix
---	-----------

Who should read this guide	ix
Publications	x
Reading syntax diagrams	x

Updates for V8.1.2	xiii
-------------------------------------	-------------

Chapter 1. Getting started	1
---	----------

Data Protection for Microsoft Exchange Server capabilities	1
Volume Shadow Copy Service framework	3
Data protection in VSS environments	3
Data backup processing.	7
Database backup types	7
Data backup methods	11
Policy management with Data Protection for Microsoft Exchange Server	14
Data restore processing	25
VSS fast restore processing	25
VSS instant restore processing	25
VSS backups that are restored to alternate databases	26
Mailbox restore operations	26
Data Protection for Exchange Server with IBM SAN Volume Controller and IBM Storwize V7000	28
IBM System Storage requirements	29
Automated IBM Spectrum Protect server failover for data recovery.	30

Chapter 2. Installing, upgrading, and migrating.	33
---	-----------

Prerequisites	33
Installing and configuring Data Protection for Microsoft Exchange Server	34
Installing Data Protection for Exchange Server.	35
Completing the installation configuration	35
Verifying the configuration	37
Customizing the configuration	38
Installing Data Protection for Exchange Server on a local system	38
Implementing a silent (unattended) installation of Data Protection for Microsoft Exchange Server.	39
Silently installing Data Protection for Exchange Server with the setup program	40
Silently installing Data Protection for Exchange Server with the Microsoft Installer program	41
Creating and testing a silent installation package on a DVD or a file server.	42
Silent installation error messages	43
Upgrading Data Protection for Microsoft Exchange Server	44
Data Protection for Exchange Server migration	44

Managing migrated backups to a Database Availability Group node	44
---	----

Chapter 3. Configuring	47
---	-----------

Proxy node definitions for VSS backups	47
Required node names for basic VSS operations	48
Required node names for basic VSS offloaded backups	48
Specifying configuration parameters for IBM Spectrum Protect	50
Specifying Data Protection for Exchange Server DAG member name parameters	52
Specifying configuration and options files in non-default locations	53
Setting user preferences	54
Data Protection properties	54
Configuring Data Protection for Microsoft Exchange Server by using the IBM Spectrum Protect Configuration Wizard	64
Verifying the configuration	66
Configuring a Data Protection for Microsoft Exchange Server remote system to integrate with IBM Spectrum Protect	68
Configuring Data Protection for Microsoft Exchange Server by using the Mailbox Restore Only Configuration Wizard	70
Manually configuring Data Protection for Exchange Server for IBM Spectrum Protect Configuration	70
Configuring the computer that runs the Exchange Server	71
Configuring the IBM Spectrum Protect server	72
Configuring the system that runs offloaded backups	74
Configuring your system for mailbox restore operations.	74
Configuring mailbox history handling for improved performance	75
Verifying the configuration of Data Protection for Exchange Server.	76
Transitioning Exchange Server backups from IBM Spectrum Protect Snapshot to IBM Spectrum Protect.	80
Configuring the IBM Spectrum Protect server	81
Configuring the computer that runs the Exchange Server	81
Examples of IBM SAN Volume Controller and IBM Storwize V7000 configuration scenarios	83

Chapter 4. Protecting data	87
---	-----------

Prerequisites	87
Security requirements for backup and restore operations	87
Software requirements for backup and restore operations	88
Software requirements for mailbox restore operations	89

VSS backup methods	89
Database Availability Group backup and restore operations	91
Starting Microsoft Management Console	94
Starting the Data Protection for Exchange Server command-line interface	95
Managing Data Protection for Exchange Server installations remotely	96
Adding remote systems	96
Determining managed storage capacity	98
Backing up Exchange Server data	98
Ensuring successful MAPI connections	98
Backing up Exchange Server data by using VSS	100
Mounting Exchange Server backups	102
Deleting Exchange Server backups	103
Restoring Exchange Server data	104
Setting data restore options in Microsoft Management Console	104
Restoring an Exchange Server database	106
Restoring a Database Availability Group database backup	107
Complete restore or replacement of Exchange Server	108
Restoring mailbox data	108
Individual mailbox recovery	108
Restoring mailbox data	108
Restoring mailbox messages interactively with the Mailbox Restore Browser	114
Restoring mailboxes directly from Exchange database files	119
Restoring a deleted mailbox or items from a deleted mailbox	120
Viewing, printing, and saving reports	120
Chapter 5. Automating	121
Preparing to use Windows PowerShell cmdlets with Data Protection for Exchange Server	121
Cmdlets for Microsoft Management Console	122
Cmdlets for protecting Microsoft Exchange Server data	123
Automating tasks	124
IBM Spectrum Protect task scheduler	125
Chapter 6. Troubleshooting	127
Diagnosing problems	127
Diagnosing VSS issues	127
Determining that the problem is a Data Protection for Exchange issue or a general VSS issue	128
Resolving reproducible problems	131
Troubleshooting VSS backup and restore operations	132
Troubleshooting mailbox restore errors	134
Troubleshooting VSS and SAN Volume Controller, Storwize V7000, or DS8000	138
Resolving problems with IBM Support	139
Gathering trace and log files	140
Gathering installation log files to debug installation problems	141

Gathering traces for the Data Protection client when using VSS technology	141
Gathering information about Exchange with VSS before calling IBM	142
Gathering information about Exchange Server with VSS before you call IBM	143
Viewing and modifying system information	145
Emailing files to IBM Support	146
Online IBM support	147

Chapter 7. Performance tuning 149

LAN-free data movement	150
----------------------------------	-----

Chapter 8. Reference 151

Support for Microsoft Exchange 2016	151
Command-line overview	152
Backup command	153
Backup syntax	154
Backup positional parameters	155
Backup optional parameters	155
Examples: backup command	159
Changetsmppassword command	160
Changetsmppassword syntax	160
Changetsmppassword positional parameters	160
Changetsmppassword optional parameters	161
Example: changetsmppassword command	163
Delete backup command	163
Delete Backup syntax	163
Delete Backup positional parameters	164
Delete Backup optional parameters	165
Help command	167
Help syntax	168
Help optional parameters	168
Mount backup command	168
Mount Backup syntax	169
Mount backup positional parameter	169
Mount Backup optional parameters	170
Query Exchange command	173
Query Exchange syntax	173
Query Exchange optional parameters	173
Query Managedcapacity command	175
Query policy command	176
Query TDP command	176
Query TDP syntax	176
Query TDP optional parameters	177
Examples: query tdp command	178
Query TSM command	178
Query TSM syntax	179
Query TSM positional parameters	180
Query TSM optional parameters	180
Examples: query tsm command	184
Restore command	186
Restore syntax	187
Restore positional parameters	188
Restore optional parameters	189
Restorefiles command	194
Restorefiles syntax	195
Restorefiles positional parameters	196
Restorefiles optional parameters	196
Restoremailbox command	199

Restoremailbox syntax	201
Restoremailbox positional parameters	203
Restoremailbox optional parameters	203
Examples: restoremailbox command	218
Set command	218
Set syntax	219
Set positional parameters	220
Set optional parameters	225
Examples: set command	225
Unmount backup command	226
Unmount Backup syntax	226
Unmount Backup positional parameter	227
Unmount Backup optional parameters	227

Frequently asked questions	229
--------------------------------------	-----

Appendix. Accessibility features for the IBM Spectrum Protect product family.	235
--	------------

Notices	237
--------------------------	------------

Glossary	241
---------------------------	------------

Index	243
------------------------	------------

Tables

1. Data Protection for Microsoft Exchange Server capabilities	1	13. Database restore options	104
2. Type: Full backup	9	14. Database restore options	111
3. Type: Copy backup	9	15. Restore options	111
4. Type: Incremental backup	10	16. Selecting mailboxes to restore	115
5. Type: Differential backup	10	17. Previewing and filtering mailbox items	117
6. Preferred Policy Settings	17	18. Restoring a mailbox to another mailbox or .pst file	117
7. Exchange Server 2013 Recoverable Items folder contents.	27	19. Cmdlets to protect Microsoft Exchange Server data.	123
8. Commands for creating a silent installation package.	43	20. Mailbox restore options	151
9. Required node names for basic VSS operations	48	21. MAPI/CDO changes	152
10. Required node names for basic VSS offloaded backups.	48	22. Backup compressed values	181
11. Diagnostics modes and their usage.	57	23. Backup encryption type values.	181
12. Options for integrity checking	101	24. Backup client-deduplicated values	181
		25. Backup supports instant restore values	181

About this publication

With Data Protection for Microsoft Exchange Server software you can back up online Microsoft Exchange Server databases to IBM Spectrum Protect™ storage.

Data Protection for Microsoft Exchange Server provides a connection between an Exchange Server and a IBM Spectrum Protect server. This connection allows Exchange data to be protected and managed by IBM Spectrum Protect server.

IBM Spectrum Protect is a client/server licensed product that provides storage management services in a multi-platform computer environment.

This publication provides information about installing, configuring, and protecting data with Data Protection for Microsoft Exchange Server.

Who should read this guide

This publication is intended for system installers, system users, IBM Spectrum Protect administrators, and system administrators.

In this publication, it is assumed that you have an understanding of the following applications:

- Microsoft Exchange Server
- IBM Spectrum Protect server
- IBM Spectrum Protect Backup-Archive Client
- IBM Spectrum Protect Application Program Interface
- Microsoft Volume Shadow Copy Service (VSS) technology (knowledge of this application is only assumed if you plan to perform VSS operations)

It is also assumed that if you are using the following operating systems or the directory service, you understand the technology:

- Windows Server 2008
- Windows Server 2008 R2
- Windows Server 2012
- Windows Server 2012 R2
- Active Directory

It is also assumed that you understand one of the following storage systems that is used for the database:

- Any storage device that implements the VSS provider interface as defined in the VSS system provider overview section of this document
- IBM® System Storage® Disk Storage Models DS3000, DS4000®, DS5000
- IBM System Storage SAN Volume Controller (SVC)
- IBM Storwize® V7000 Disk System
- IBM XIV® Storage System Model 2810 (Gen2)
- IBM System Storage DS8000™ series

Publications

The IBM Spectrum Protect product family includes IBM Spectrum Protect Snapshot, IBM Spectrum Protect for Space Management, IBM Spectrum Protect for Databases, and several other storage management products from IBM.

To view IBM product documentation, see IBM Knowledge Center.

Reading syntax diagrams

The section describes how to read the syntax diagrams that are used in this publication. To read a syntax diagram, follow the path of the line. Read from left to right, and top to bottom.


- The ►— symbol indicates the beginning of a syntax diagram.
- The —► symbol at the end of a line indicates the syntax diagram continues on the next line.
- The ►— symbol at the beginning of a line indicates a syntax diagram continues from the previous line.
- The —►◄ symbol indicates the end of a syntax diagram.

Syntax items, such as a keyword or variable, can be:

- On the line (required element)
- Above the line (default element)
- Below the line (optional element)

Syntax diagram description	Example
Abbreviations:	
Uppercase letters denote the shortest acceptable truncation. If an item is entirely in uppercase letters, it cannot be truncated.	►—KEYWOrd—►
You can type the item in any combination of uppercase or lowercase letters.	
In this example, you can enter KEYWO, KEYWORD, or KEYWOrd.	
Symbols:	
Enter these symbols exactly as they are displayed in the syntax diagram.	* Asterisk
	{ } Braces
	: Colon
	, Comma
	= Equal Sign
	- Hyphen
	() Parentheses
	. Period
	' Single quotation mark
	Space
	" Quotation mark

Syntax diagram description	Example
<p>Variables:</p> <p>Italicized lowercase items (<i>var_name</i>) denote variables.</p> <p>In this example, you can specify a <i>var_name</i> when you enter the KEYWORD command.</p>	
<p>Repetition:</p> <p>An arrow that points to the left means you can repeat the item.</p> <p>A character or space within an arrow means you must separate the repeated items with that character or space.</p>	
<p>Required Choices:</p> <p>When two or more items are in a stack and one of them is on the line, specify one item.</p> <p>In this example, you <i>must</i> choose A, B, or C.</p>	
<p>Optional Choice:</p> <p>When an item is below the line, that item is optional. In the first example, you can choose A or nothing at all.</p> <p>When two or more items are in a stack below the line, all of them are optional. In the second example, you can choose A, B, C, or nothing at all.</p>	
<p>Defaults:</p> <p>Defaults are above the line. The default is selected unless you override it. You can override the default by including an option from the stack below the line.</p> <p>In this example, A is the default. You can override A by choosing B or C. You can also specify the default explicitly.</p>	
<p>Repeatable Choices:</p> <p>A stack of items followed by an arrow pointing to the left means you can select more than one item or, in some cases, repeat a single item.</p> <p>In this example, you can choose any combination of A, B, or C.</p>	

Syntax diagram description	Example
<p>Syntax Fragments:</p> <p>Some diagrams because of their length, must fragment the syntax. The fragment name is displayed between vertical bars in the diagram. The expanded fragment is displayed between vertical bars in the diagram after a heading with the same fragment name.</p>	<p>▶▶ The fragment name ◀◀</p> <p>The fragment name:</p> 
<p>Footnote:</p> <p>A footnote in the diagram references specific details about the syntax that contains the footnote.</p> <p>In this example, the footnote by the arrow references the number of times you can repeat the item.</p>	 <p>Notes:</p> <ol style="list-style-type: none"> 1 Specify <i>repeat</i> as many as 5 times.

Updates for V8.1.2

This document provides information about what's new or what has changed in Data Protection for Microsoft Exchange Server version 8.1.2.

New and changed information is indicated by a vertical bar (|) to the left of the change.

What's new

Automatic loading of Windows PowerShell modules

- Windows PowerShell modules, FmModuleExc and FmModuleMMC, which contain the Data Protection for Exchange Server cmdlets, are imported automatically from the TDPEXchange folder during installation of Data Protection for Exchange Server. For more information, see Preparing to use Windows PowerShell cmdlets.

What has changed

Scheduling wizard updated to support IBM Spectrum Protect client scheduler

- If you use the scheduling wizard on MMC to create a scheduled data protection task, you can now choose a IBM Spectrum Protect client scheduler to manage the scheduled operation. When you select the **IBM Spectrum Protect client scheduler** check box, a macro file is created that automatically contains the relevant IBM Spectrum Protect server commands you require to create a scheduled task. You can specify IBM Spectrum Protect server details, such as administrator user details and the port and address of the IBM Spectrum Protect server. You can also edit the created macro file according to your requirements. You can access the scheduling wizard in MMC by clicking **Manage> Scheduling** in the navigation tree or from the **Actions** pane by clicking the **Run Interactively** button to change the mode to **Run Scheduled**.

Chapter 1. Getting started

With IBM Spectrum Protect for Mail: Data Protection for Microsoft Exchange Server, you can back up and restore Microsoft Exchange Server databases to IBM Spectrum Protect storage or local shadow volumes. A *local shadow volume* contains data that is stored on shadow volumes, which are local to a disk storage system.

Data Protection for Exchange Server provides a connection between an Exchange Server and an IBM Spectrum Protect, which allows Exchange Server data to be protected and managed by IBM Spectrum Protect. Data Protection for Exchange Server protects Exchange Server data and improves the availability of Exchange Server databases.

Data Protection for Exchange Server backs up and restores Microsoft Exchange Server databases to IBM Spectrum Protect storage or local shadow volumes. You can use a command-line interface or graphical user interface (GUI) to back up and restore Exchange Server databases.

Beginning with Exchange Server 2010, Microsoft no longer supports the Microsoft Legacy application programming interface (API) for streaming backup and restore operations. Microsoft supports the use of Volume Shadow Copy Service (VSS) technology for backup and restore operations.

Data Protection for Exchange Server uses the IBM Spectrum Protect API to communicate with the IBM Spectrum Protect, and the Exchange API to communicate with Exchange Server.

In addition to these APIs, Data Protection for Exchange Server VSS operations require the IBM Spectrum Protect backup-archive client (VSS Requestor) and Microsoft VSS to produce an online snapshot (point-in-time consistent copy) of Exchange Server data.

You must install Data Protection for Exchange Server on the same system as the Exchange Server. Data Protection for Exchange Server also supports backup and restore operations in a Database Availability Group (DAG) environment.

Related concepts:

“Prerequisites” on page 33

Data Protection for Microsoft Exchange Server capabilities

Data Protection for Microsoft Exchange Server helps you to protect and manage Exchange Server environments by facilitating the backup, restore, and recovery of Exchange Server data.

The following table lists the tasks that you can perform with Data Protection for Microsoft Exchange Server:

Table 1. Data Protection for Microsoft Exchange Server capabilities

Feature	Referred to as:	More information:
Back up Exchange Server databases by using Microsoft VSS	VSS backup	“VSS data backups” on page 11

Table 1. Data Protection for Microsoft Exchange Server capabilities (continued)

Feature	Referred to as:	More information:
Back up Exchange Server Database Availability Group (DAG) databases to a common node so that you can manage all DAG members with a single policy	Back up to DAG node	"Managing Exchange Database Availability Group members by using a single policy" on page 23
Back up databases to the IBM Spectrum Protect server by using an alternate system to a production system	Offloaded backup	"Offloaded VSS backups" on page 12
Restore database backups that are on IBM Spectrum Protect storage to their original location	VSS restore	"VSS restore characteristics" on page 5
Restore database backups that are on local shadow volumes by using file-level copy mechanisms	VSS Fast Restore	"VSS fast restore processing" on page 25
Restore database backups that are on local shadow volumes by using hardware-assisted volume-level copy mechanisms	VSS Instant Restore	"VSS instant restore processing" on page 25
Restore a database backup to a recovery database, alternate database, or relocated database	Restore into	"VSS backups that are restored to alternate databases" on page 26
Restore individual mailboxes and mailbox item-level data from Data Protection for Microsoft Exchange Server backups	Mailbox restore	"Restoring mailbox data" on page 108
Set up remote computers in the same or a different domain to manage Exchange Server backup and restore operations.	Remote system management	"Managing Data Protection for Exchange Server installations remotely" on page 96
Query the managed capacity for database backups that are on local shadow volumes	query managedcapacity command	"Query Managedcapacity command" on page 175
Delete a backup of an Exchange Server database	delete backup command	"Delete backup command" on page 163
Manage policy for database backups that are on local shadow volumes	policy commands	"Query policy command" on page 176
Integrate with IBM Spectrum Protect Snapshot	Advanced VSS support	"Transitioning Exchange Server backups from IBM Spectrum Protect Snapshot to IBM Spectrum Protect" on page 80
Manage IBM Spectrum Protect database backup policy	Server policy	"How policy affects backup management on Data Protection for Exchange Server" on page 15
Issue the restorefiles command to restore database backups to flat files without involving the Exchange Server	restorefiles command	"Restorefiles command" on page 194

Volume Shadow Copy Service framework

Volume Shadow Copy Service (VSS) provides a common interface model to generate and manage online snapshots of Exchange Server data.

The Microsoft VSS service manages and directs three VSS software components that are used during VSS operations: the VSS writer, the VSS Requestor, and the VSS provider. The VSS writer is the application that stores data on the source volumes. The VSS Requestor is the backup software. The VSS provider is the combined hardware and software that generates the snapshot volume.

The VSS system provider creates and maintains snapshots on local shadow volumes and refers to the default VSS provider that is available with Windows Server. If you use the Windows VSS system provider, no configuration is required. However, you can make changes by using the **VSSADMIN** commands.

With a VSS hardware and software copy provider, you can create shadow copies of running volumes on demand. A hardware provider uses a hardware storage adapter or controller to manage shadow copies at the hardware level. Data Protection for Exchange Server software does not control the VSS hardware provider. The VSS hardware provider is controlled by the hardware vendor. Install and configure the VSS hardware and software provider as required.

Data protection in VSS environments

The characteristics of VSS backups can affect your backup management tasks. As you decide your backup strategies, be aware of the following VSS backup guidelines.

As you decide your backup and restore strategies, be aware of VSS requirements and guidelines.

Related tasks:

“Backing up Exchange Server data” on page 98

VSS backup characteristics

Backups can be stored on local shadow volumes, an IBM Spectrum Protect server, or at both locations. You can define different policy settings for each backup location.

You can offload backups from IBM Spectrum Protect server storage to another system as resource relief for production servers. In addition, you can restore backups to flat files.

Databases must have unique names. If a database has the same name as another database, but the capitalization differs, the software does not differentiate between case.

When you back up Exchange Server data by using Data Protection for Exchange Server, VSS backups have the following characteristics:

- Backups provide an Exchange Server database integrity check function, but do not provide a zeroing function.
- You can run full, copy, differential and incremental backups.
- You can restore a backup to a local disk only on the same system.
- You can back up Exchange Server Database Availability Group (DAG) databases under a common DAG node name, regardless of which DAG member runs the

backup. You can create the backup from an active or passive copy. When you back up data to a common node, the backups are managed by a common policy, and you can restore the backup to any Exchange Server under the same DAG node.

VSS backup requirements

You can plan your VSS backup strategy to optimize the performance of your backup operations and to avoid potential problems. Follow these guidelines when you plan your VSS backups:

- **Planning VSS backups**

- When you perform VSS operations, ensure that at least 200 MB of free disk space is on your Windows System Drive. This space is used to store the metadata files for Data Protection for Exchange Server.
- Use basic disks, which are initialized for basic storage. A basic disk consists of basic volumes, such as primary partitions, extended partitions, and logical drives.
- If you plan to keep VSS snapshot backups only on local shadow volumes, know how to implement the configuration options of your VSS hardware provider.

For example, if your VSS hardware provider supports a full-copy snapshot versus a copy-on-write snapshot mechanism, full-copy type implementations have greater disk storage requirements. However, full-copy type implementations do not rely on the original volume to restore the data and are less risky. Copy-on-write implementations require less disk storage but rely on the original volume to restore the data.

- Do not place multiple volumes on the same LUN. Configure a single volume, single partition, and single LUN as one-to-one.
- Do not set the ASNODENAME option in the dsm.opt file when you use Data Protection for Exchange Server. Setting the ASNODENAME option can cause VSS data backups and VSS restore operations to fail.

- **Running parallel VSS backups**

If you need to run parallel VSS backups, do the following:

- Stagger the start time of the backups by at least 10 minutes. This interval ensures that the snapshot operations do not overlap.

Attention:

If backup operations overlap, a VSS timeout error may occur and the second backup request may fail. Therefore, it is recommended to stagger the start time of the backups.

- Configure the parallel instance backups so that snapshots of the same volumes are not created.
- Ensure that parallel backups do not create a snapshot of the same LUN.

VSS restore characteristics

In a VSS restore operation, VSS backups (Exchange database files and log files) that are on IBM Spectrum Protect server storage are restored to their original location on the Exchange Server.

The following characteristics are true of a VSS data restore operation:

- If you use a hardware provider, the disks that contain Exchange Server data are configured as basic disks.
- You can restore data by using full, copy, incremental, and differential backup methods.
- You can restore data from a VSS backup to an alternate database.
- Data is restored at the database level.
- You can restore one or more databases from a VSS snapshot backup on IBM Spectrum Protect server storage.
- You can restore a VSS backup directly from IBM Spectrum Protect server storage to an alternate system.
- You can restore data in a Database Availability Group (DAG) environment.
- You can restore data from a DAG replica on Exchange Server 2010 or later backup versions to the production server.
- You cannot use VSS fast restore or VSS instant restore operations with Microsoft Windows Server 2008 or later versions.
- VSS restore operations place data directly into the production database, unless you specify the `/intodb` parameter.

VSS restore requirements

Unless otherwise specified, a *VSS restore* operation refers to all restore types that use VSS, including VSS restore, VSS fast restore, and VSS instant restore operations.

Install any Microsoft VSS-related urgent fixes.

As you decide your restore strategies, consider the following VSS requirements.

- Unless you issue the `/INTODB` parameter on the **restore** command, a VSS restore operation ignores the recovery database and data is stored in the production database.
- If you use a VSS hardware provider, the disks that contain Exchange Server data are configured as basic disks.
- For VSS restore operations, you must dismount the restored database.
- When a VSS restore operation from local shadow volumes is complete, zero bytes are transferred because no data (0) is restored from the IBM Spectrum Protect server.
- When you use Data Protection for Microsoft Exchange Server, do not set the `ASNODENAME` option in the `dsm.opt` file. If you set the `ASNODENAME` option, VSS backups and VSS restore operations might fail.

VSS instant restore

A VSS instant restore operation overwrites the entire contents of the source volumes.

- If you do not want to overwrite the source volumes, ensure that you set the **Instant Restore** option to **No** in Microsoft Management Console (MMC).

- VSS instant restore processing requires that the local disk is not accessed by other applications, for example, Windows Explorer.
- Before you run a VSS instant restore operation in an Exchange Server 2013 environment, stop the Exchange Search Host Controller Service on the active node.
- When you run a VSS instant restore operation, verify that there is no other data on the volumes that are being restored.
- If you perform a VSS restore of a database that was relocated (system file path, log file path, or database file path), you must use the **Restore Into** function and specify the same database name as the one you are restoring. The restore operation fails if you do not specify the same database name.
- When you run the **Restore Into** function, VSS instant restore operations are automatically disabled.
- Before you start a VSS instant restore operation, ensure that any previous background copies that contain the volumes that are being restored are completed. XIV, SAN Volume Controller, or Storwize family with space-efficient target volumes do not need to be completed.

VSS fast restore

In a VSS fast restore operation, if you do not want to overwrite all the files on the original volume, mount the snapshot. Copy only the files that you want to restore.

Database Availability Group (DAG) environment

You can perform data restore operations in a Database Availability Group (DAG) environment including restoring an Exchange Server 2010 or later backup from a DAG replica into the production server.

- Before you run a VSS instant restore operation in a DAG environment, stop the Microsoft Exchange Replication Service on the active node.
- Restore your backups to the active database copy.
- If you back up data to a local system, you can restore the snapshot only to the same system.
- However, backups to a local server can be restored only on the server where the backup is created.
- To restore a backup to a server that is hosting a passive database copy, make the copy active before you restore the backup. When the backup is restored, you can move the active database copy back to the passive state.

VSS operations in IBM N-series and NetApp environments

You must consider storage space limitations when you perform VSS operations in environments that contain IBM N-series and NetApp systems.

Snapshots that are created by using the IBM N-series and NetApp snapshot provider are stored on the same volume where the LUN are located.

Disk space that is used by a local backup consists only of the blocks that changed since the last local backup was created. You can use the following formula to determine how much space is required for each local backup:

Amount of data changed per hour * number of hours before a local backup expires

In addition, Write Anywhere File Layout (WAFL) reserves space, that is, blocks equal to two times the specified size of the LUN to be used. This space reservation

ensures that write operations are allowed for virtual disks. The following example shows how to calculate the size of the volumes:

```
Database size of an Exchange database: 100GB
Number of local backups to be kept: 3
Snapshot for TSM backup: 1
duration for TSM backup: 2hr
Backup frequency: 3hrs
The duration before a local backup is expired: 9 hrs
Amount of data changed/added/deleted per hr: 50MB
Space required for each local backup: 50*9= 450 MB
Space required for 3 local backups + 1 TSM backup: 450*3 + 50*2 = 1450 MB
The volume size required for the database: 100*2 (space reservation) + 1.5 = 201.5 GB
```

Data backup processing

Data Protection for Exchange Server can use the Microsoft Volume Shadow Copy Service (VSS) framework to produce a point-in-time, consistent, online copy of Exchange Server data.

Database backup types

With IBM Spectrum Protect for Mail: Data Protection for Microsoft Exchange Server, you can use the common interface in the Volume Shadow Copy Service (VSS) framework to create database backups.

To back up Exchange Server data, you can use the following backup types and backup expiration policies on the IBM Spectrum Protect server.

Backup types	<ul style="list-style-type: none">• Full (for more information, see Table 2 on page 9)• Copy (for more information, see Table 3 on page 9)• Incremental (for more information, see Table 4 on page 10)• Differential (for more information, see Table 5 on page 10)
--------------	--

<p>Expiration policies</p>	<ul style="list-style-type: none"> • Version-based expiration Version-based expiration uses the VERExists and the VERDeleted copy group parameters. <ul style="list-style-type: none"> – VERExists The maximum number of Exchange Server database backup versions to retain for the databases that exist on the protected Exchange Server system. – VERDeleted The maximum number of Exchange Server database backup versions to retain for the databases that were deleted from the protected Exchange Server system after they were backed up by IBM Spectrum Protect. When you deactivate database backups, any existing backups on IBM Spectrum Protect server are subject to deletion, as specified by the VERDeleted setting. <p>Attention:</p> <ul style="list-style-type: none"> – The IBM Spectrum Protect server considers backups as corresponding to a deleted database if there are no active backups of this database. – A new full database backup object deactivates all prior legacy active backup objects for the same Exchange Server database. This deactivation includes any active full backup object and any active file, group, set, differential, and log backup objects. • Retention-based expiration Retention-based expiration uses the RETExtra and the RETOnly copy group parameters. <ul style="list-style-type: none"> – RETExtra The number of days to retain an Exchange Server database backup version after that version becomes inactive. This parameter applies to backup types where it is possible to have more than one version, for example, full, copy, or differential backup types. – RETOnly The number of days to retain the last Exchange Server database backup version of a database that was deleted from the protected Exchange Server system. The RETOnly parameter applies to all backup types, including incremental backup objects that can never have more than one version. <p>For more information, see Recommended IBM Spectrum Protect policy settings.</p>
----------------------------	--

Policy settings on the IBM Spectrum Protect server

The following tables summarize the policy settings for each VSS backup type.

Tip: When you enable circular logging, you cannot use differential or incremental backups.

Table 2. Type: Full backup

Full backup	
Description	<ul style="list-style-type: none"> • Data Protection for Exchange Server backs up the specified database and associated transaction logs. • Each version of a full backup has the same name because the server recognizes each backup as a new version of the same backup object.
Expiration policies available for selection	Both retention-based and version-based policies
Recommended usage	If you use only a full backup type, you can use either a retention-based or version-based expiration policy to retain the version of the backup in the IBM Spectrum Protect server.

Table 3. Type: Copy backup

Copy backup	
Description	<ul style="list-style-type: none"> • Data Protection for Exchange Server backs up the transaction logs and does not delete the log files after the backup. Otherwise, this backup type is similar to a full backup. • You can create a full backup of the Exchange Server database without disrupting any backup processes that use an incremental or differential backup. • Each version of a copy backup has the same name because the server recognizes each backup as a new version of the same backup object.
Expiration policies available for selection	Only retention-based policies
Recommended usage	<p>You can keep copy backups of the Exchange Server database for retention periods that are different from the periods that you set for the full backup operations</p> <p>For example, legal regulations might require that you keep a monthly backup for several years. To meet this requirement, you can set those monthly backup processes as copy backups. You must define a management class and set the retention parameters to use for copy backups. These definitions must be different than the parameters set for full backups.</p>

Table 4. Type: Incremental backup

Incremental backup	
Description	<ul style="list-style-type: none"> Transaction log files are not deleted if the backup fails. Only one version of an incremental backup object exists at a time because each incremental backup is named with a unique time stamp. Data Protection for Exchange Server software deactivates all incremental backups (and the active differential backup, if one exists) that are associated with a full backup operation, whenever a new full backup operation is done. When you restore an Exchange Server database from an incremental backup, you must complete the following tasks: <ul style="list-style-type: none"> Restore the last full backup. Restore any other incremental backups that occur between the full backup and the incremental backup. Restore the incremental backup.
Expiration policies available for selection	Only retention-based policies
Recommended usage	<p>Because each incremental backup has a unique name, you cannot use a version-based expiration policy. For incremental backups, you must use retention-based policies.</p> <p>To ensure that incremental backups do not expire before the full backup on which they depend, you must specify the following parameters:</p> <ul style="list-style-type: none"> In the management class that you use for incremental backups, set a value for the RETOnly parameter. In the management class you use for full backups, set the value of the RETExtra copy group parameter to be equal to the value you set for the RETOnly parameter.

Table 5. Type: Differential backup

Differential backup	
Description	<ul style="list-style-type: none"> Data Protection for Exchange Server backs up transaction logs but does not delete the log files after the backup. For a full backup with only differential backups, the last full backup and the last differential backup contain all the data that is required to restore the database to its most recent state. Subsequent backups create a new version of the differential backup object on the Exchange Server. When you restore an Exchange Server database from a differential backup, you must complete the following tasks: <ul style="list-style-type: none"> Restore the last full backup. Restore this differential backup, but no other differential backups. <p>Tip:</p> <p>You can use both version-based and retention-based policies to control the expiration of differential backups.</p>
Expiration policies available for selection	Retention-based and version-based policies

Table 5. Type: Differential backup (continued)

Differential backup	
Recommended usage	<p>To ensure that differential backups expire at the same time as the full backup on which they depend, it is recommended (for differential backups) that you use a retention-based policy similar to the one used for incremental backups.</p> <p>To use a retention-based expiration policy, in the management class that you use for differential backups, you must specify the following parameters:</p> <ul style="list-style-type: none"> • Set the VERDeleted and VERExists copy group parameters to no limit. • Set the RETExtra and RETOnly parameter values to match the parameter values in the management class that you use for full backups. <p>To limit the number of differential backups for an Exchange Server database, you can use a version-based policy. If you choose to use a version-based expiration policy, do so in combination with retention-based policies to ensure that old, inactive backups expire, even if you no longer use differential backups.</p> <p>If you want to use both a version-based and a retention-based policy, you must specify the following parameters and version control settings:</p> <ul style="list-style-type: none"> • Set the RETExtra and RETOnly parameter values to match the parameter values in the management class that you use for full backups. • Set the VERExist parameter to the value you want. • Set the VERDeleted parameter to be equal to VERExist.

Data backup methods

You can use Volume Shadow Copy Service (VSS) to back up Data Protection for Exchange Server data.

You can run Exchange Server backup operations in a Database Availability Group (DAG) environment.

VSS data backups

You can store VSS backups on local VSS shadow volumes, or, when integrated with IBM Spectrum Protect, in IBM Spectrum Protect server storage.

VSS backups eliminate the need for the server or file system to be in backup mode for an extended time. The length of time to complete the snapshot is measured in seconds, not hours. In addition, a VSS backup allows a snapshot of large amounts of data at one time because the snapshot works at the volume level.

You must ensure that sufficient space is available for the snapshot at the storage destination. Both storage destinations require space to store the snapshot until the data transfer to the IBM Spectrum Protect server is complete. After the data transfer to the server is complete, VSS backups that are stored locally on VSS shadow volumes are directly accessible by the system. The snapshot volume is released and the space can be reused.

- For data that is backed up to local VSS shadow volumes, the snapshot backup is on the shadow copy volume.

- For data that is backed up only to IBM Spectrum Protect server storage, a local snapshot backup is run and the data on the local snapshot volume is sent to the IBM Spectrum Protect server.
- For data that is backed up to VSS shadow volumes and IBM Spectrum Protect server, the local snapshot volume is retained as a local backup after the transfer to the IBM Spectrum Protect server is complete.

If you store VSS backups both locally and to IBM Spectrum Protect server, and the maximum number of local backup versions to be maintained is reached, the oldest local backup version expires to create the new snapshot for the backup to IBM Spectrum Protect server storage. The maximum number of local backup versions that are maintained is set in the IBM Spectrum Protect policy.

Offloaded VSS backups

By running an offloaded backup, you can move the backup load from the production system to another system. You can reduce the load on network, I/O, and processor resources during backup processing.

Use the **RemoteDSMAGENTNode** parameter to run an offloaded system. Ensure that you install a VSS hardware provider, which supports transportable shadow copy volumes, on the production and secondary systems.

Database Availability Group backups

You can use the high-availability feature of Database Availability Group (DAG) backups for enhanced data and service availability, and automatic recovery from failures. You can use Exchange Server 2010 or later versions with DAG backups to improve Exchange Server data backups and data recovery.

Beginning with Exchange 2013 SP1, you can also back up Exchange Server databases in a Database Availability Group (DAG) environment without a Cluster Administrative Access Point (CAAP).

A DAG environment includes the following functions:

- A group of up to 16 mailbox servers that can host up to 100 mailbox databases
- Up to 16 online copies of a database (1 active database and up to 15 passive databases)
- Synchronous or lagged replication. With lagged replication, you can delay the replaying of logs on target databases if, for example, there are time differences between source and target databases.
- Automatic migration and failover of active database copies

The following figure illustrates a DAG environment:

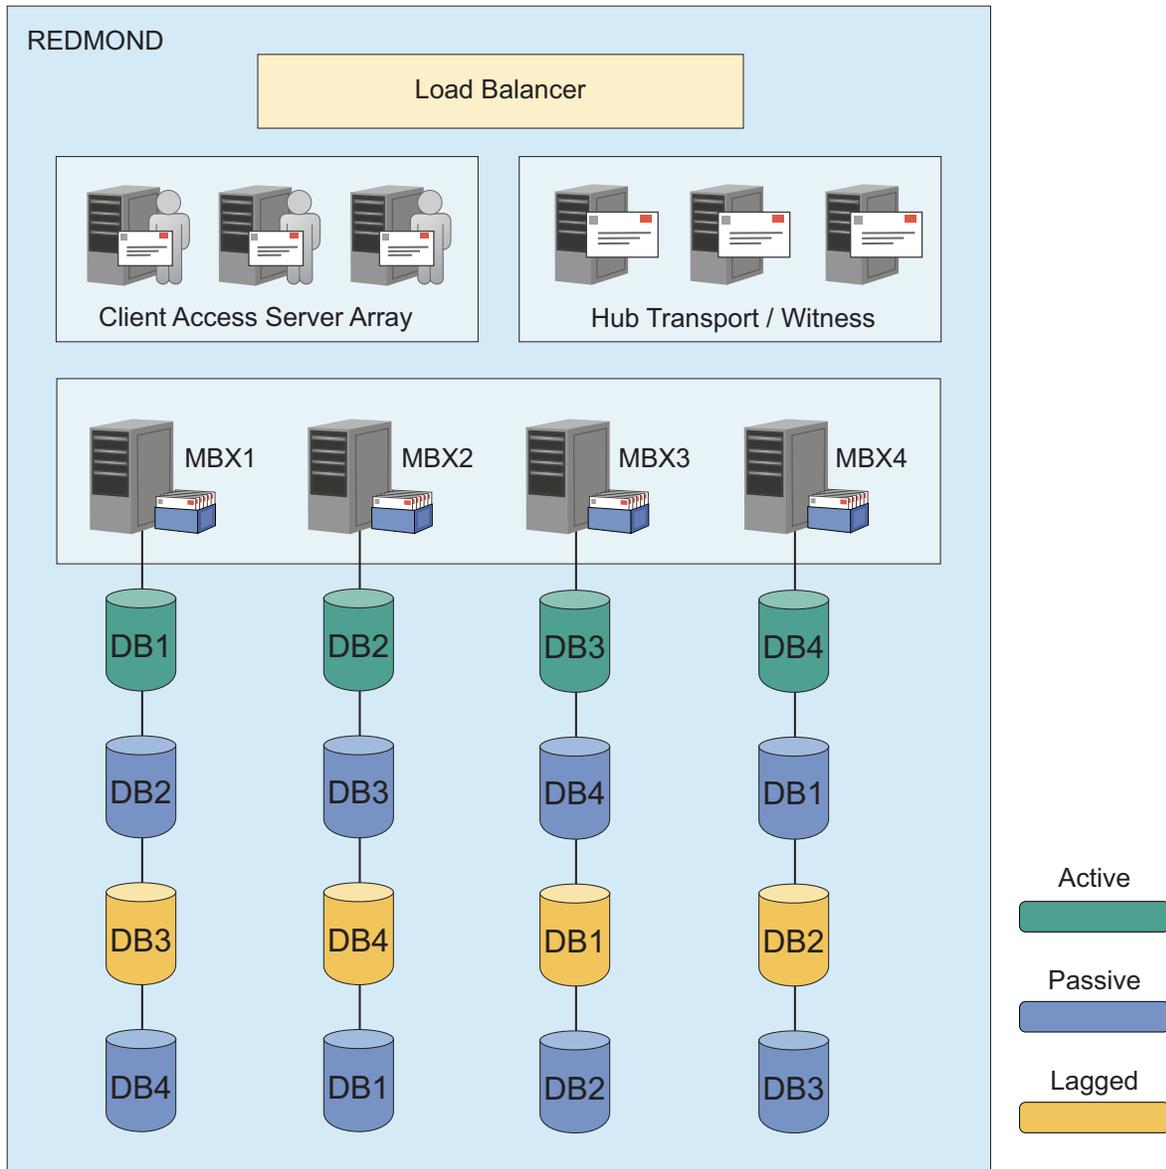


Figure 1. Sample DAG environment

Database copies are mirrored on any node within the DAG. You can complete the following tasks:

- Query DAG database copies, including status.
- Manage full, copy, incremental, and differential backups of active and passive databases within a DAG. You can create a backup from any active database copy, any passive synchronous copy, or any lagged copy within the DAG. If you back up a lagged database copy, it might take more time to restore the backup because the lagged copy can have more transaction logs to restore and replay. As a best practice, create your backup from a passive synchronous copy and not a lagged copy.
- Move an active database copy to other nodes.
- Query all DAG database copy backups.
- Restore all DAG database copy backups.

- Restore data into an active database, from either active or passive database copy backups.
- Restore data into a recovery or alternate database.
- Process Individual Mailbox Restore (IMR) operations from a DAG database copy backup.
- Delete DAG database copy backups.

Policy management with Data Protection for Microsoft Exchange Server

With Data Protection for Microsoft Exchange Server, you can manage and configure storage management policies for backups. A backup policy determines how backups on local shadow volumes are managed and retained.

Although IBM Spectrum Protect policy determines how Data Protection for Microsoft Exchange Server backups are managed on IBM Spectrum Protect storage, backup retention on local shadow volumes is determined by version and time-based policies. Ensure that sufficient local storage space is available on local shadow volumes for a VSS backup. In addition, verify that enough available storage space is assigned to the volumes to accommodate your backup operations. The shadow copy volume that is the storage destination of a snapshot must have sufficient space for the snapshot.

Environment and storage resources also affect how many backup versions are maintained on local shadow volumes. The amount of space that is required depends on the VSS provider that you use.

How backups expire based on policy

Backups expire based on Data Protection for Exchange Server policy.

Expiration is the process by which Exchange Server backup objects are identified for deletion when the expiration date is past or the maximum number of backup versions that must be retained is reached.

The date on which data expires depends on the business needs that are identified by the recovery point objective (RPO) and the recovery time objective (RTO) of your enterprise. For example, legal, operational, and application requirements affect how data must be protected to meet these RPO and RTO demands. With Data Protection for Exchange Server, you can specify the number of snapshot backups to retain and the length of time to retain them.

Backups can expire during a query, backup, or restore operation of a Data Protection for Exchange Server session.

For Exchange Database Availability Group (DAG) backups that use the DAG node, only the system on which the backup is created can cause a local backup to expire. If a new backup is created on a different system, and it exceeds the number of backups to be retained, the oldest backup expires from the Data Protection for Exchange Server. An expired backup can no longer be restored. However, the physical storage for that backup version is not released until the next time the original system runs a backup, query, or delete operation.

You specify the number of backup copies that are retained. When the maximum number of backup copies is reached, the oldest backup expires and is deleted. You can specify the maximum number of backup copies in a Data Protection for Exchange Server policy.

A backup copy is retained for a maximum number of days. The maximum number of days that a backup can be retained is specified in the Data Protection for Exchange Server policy.

How policy affects backup management on Data Protection for Exchange Server

IBM Spectrum Protect policies determine how Data Protection for Exchange Server backups are managed on IBM Spectrum Protect storage and on local shadow volumes when the environment is configured for VSS operations.

The IBM Spectrum Protect server recognizes Data Protection for Exchange Server as a *node*.

Data that is backed up to IBM Spectrum Protect storage from the Data Protection for Exchange Server node is stored and managed according to settings that you specify in the IBM Spectrum Protect server policy.

IBM Spectrum Protect policies manage the VSS backups that are placed in IBM Spectrum Protect server storage pools. The server manages VSS backups.

If you use IBM Spectrum Protect for Copy Services and upgrade to Data Protection for Exchange Server, with the license for IBM Spectrum Protect for Copy Services, you can store VSS backups to local shadow volumes.

IBM Spectrum Protect requires that sufficient storage space is available to create shadow volumes for VSS backup processing. Even when the VSS backup destination is the IBM Spectrum Protect server, storage space to create a shadow volume is still required temporarily.

The number of local backup versions that are maintained by the IBM Spectrum Protect server is determined by the value that is specified by the IBM Spectrum Protect server **verexists** parameter, which is defined in the copy group of the management class to which the local backup belongs. It is not necessary to allocate target sets when you use the VSS system provider. When you do not use the VSS system provider, the number of target volume sets that are allocated for local backups must be equal to the value of the **verexists** parameter. Target volume sets are not applicable to IBM XIV Storage Systems.

For example, if **verexists**=3, then at least three sets of target volumes must be allocated for the backup to complete successfully. If only two sets of target volumes are allocated, the third and subsequent backup attempts fail. If more sets of target volumes exist than the number specified by the **verexists** parameter, these sets are ignored by the IBM Spectrum Protect server. A high number of local backup versions cannot be stored. If you want to have n number of local backup versions, set the **verexists** parameter to $n + 1$.

If you keep only one backup, the same disk is reused. The process initially removes the existing backup and attempts the new backup. If the new backup fails, no backups exist.

If you retain multiple backups, the oldest backup is removed before another backup is created. If the new backup fails, you might have one less backup than specified by the policy. For example, if you specify that you want to retain five backups, but the last backup fails, you might have only four backup versions.

Ensure that you specify a **verexists** value that meets your VSS backup goals. If you have limited storage space for VSS operations and are restricted to a `verexists=1` setting, set the backup destination to `BOTH`. This option stores the backup on local shadow volumes and sends a copy to IBM Spectrum Protect server storage.

You can change and delete VSS backups that Data Protection for Exchange Server creates and stores on local shadow volumes. From the command-line interface, for example, issue the Microsoft **VSSADMIN DELETE SHADOWS** command to remove a VSS backup that is managed by IBM Spectrum Protect. IBM Spectrum Protect is not able to prevent the backup from being removed, and, in this instance, it detects that the backup is removed and reconciles its index of available backups with what is on local shadow volumes. Because backups can be removed, establish a strategy that protects VSS backup data that is stored on local shadow volumes from being compromised.

When you use the configuration wizard in the GUI, the **VSSPOLICY** parameter is set in the `tdpexc.cfg` file.

Depending on the policy management settings, you can reuse a logical unit number (LUN) for a new backup. When a backup is requested and the maximum number of versions is reached, the software deletes the oldest snapshot (backup) to make space for the snapshot. If the new request fails after the oldest snapshot is deleted, you have one less backup version than expected.

You must manage the policy for local backups to reconcile the local backup repository with the information that is stored on the IBM Spectrum Protect server. For example, if target volume LUNs that are used for a local backup are removed from the storage system, the information that represents the backup on the IBM Spectrum Protect server must be reconciled. Similarly, if IBM Spectrum Protect server policy determines that a local backup copy is no longer needed, the local backup manager must free the target volume LUNs to the storage system. The local backup manager is released so that these LUNs can be used for future backup operations. IBM Spectrum Protect automatically detects when these situations occur and completes the reconciliation.

Consider the scenario where you use a two-member Database Availability Group (DAG), named `MEMBER1` and `MEMBER2`. When you complete a backup to `LOCAL` on `MEMBER1` and complete more backups on `MEMBER2`, the backups to `LOCAL` on `MEMBER1` do not expire until the next time you back up, query, or delete data on `MEMBER1`. In this scenario, you might use more storage than specified by **verexists**.

Preferred settings for IBM Spectrum Protect policies

Within an IBM Spectrum Protect storage environment, you can define policies to help ensure that the storage environment meets your organization's requirements for data protection and retention. Before you start using Data Protection for Microsoft Exchange Server, review the preferred settings for IBM Spectrum Protect policies.

Policy settings

Table 6. Preferred Policy Settings

Setting	Definition	Guidelines	Additional information
Domain	A policy domain contains policy sets, management classes, and copy groups.	Create a policy domain on the IBM Spectrum Protect server to be used exclusively for Exchange Server backups.	
Policy sets	Policy sets contain management classes (which contain copy groups) that determine the rules by which protected Exchange Server backups are run and managed.	Define the policy set to the policy domain to which protected Exchange Server backups belong. The policy set must be activated and only one policy set can be active in the policy domain.	

Table 6. Preferred Policy Settings (continued)

Setting	Definition	Guidelines	Additional information
Management class	A management class is a policy object that users can bind to each file to specify how the file is managed.	<p>Ensure that you plan your backup strategy before you define management classes.</p> <p>Define a management class for backups on local shadow volumes, and a management class for backups on IBM Spectrum Protect server storage.</p> <p>Important: Because VSS backup processing requires sufficient storage space to create shadow volumes, ensure that you specify <i>verexists=N+1</i> to keep <i>n</i> backups on local shadow volumes.</p> <p>You can have multiple active backups of the same database because legacy backups on IBM Spectrum Protect server storage and VSS backups on IBM Spectrum Protect server storage (Copy and Full) all have different IBM Spectrum Protect server naming. Therefore, each can have their own management class.</p>	<p>Different management classes provide the opportunity for specialized policies for each storage destination. For example, you can maintain six versions of local VSS backups of a particular database (<i>verexists=6</i>) while you maintain only two versions of the same database on IBM Spectrum Protect server storage (<i>verexists=2</i>).</p> <p>In addition, you can create a separate management class for copy backup types for use in long-term storage. Such policies can maximize storage resources and provide more control over your storage strategy.</p>
Storage pool	A storage pool is a named set of storage volumes and the destination that is used by the IBM Spectrum Protect server to store data.	<p>Use collocation if backups are stored on removable media. Specify collocation by file space (define stgpool COLlocate=FILEspace) if you plan to restore multiple Exchange databases in parallel.</p> <p>Tip: As a best practice, use collocation because data for any one database is stored within one IBM Spectrum Protect server file space.</p>	A single restore operation can require a full backup, a differential backup, and multiple incremental backups.

Copy Group Parameters

A copy group controls how backup versions are generated, located, and expired. Define the copy group as a backup copy group and not as an archive copy group. Because Data Protection for Exchange Server stores all objects as backup objects on IBM Spectrum Protect in backup storage pools, an archive copy group is not required, although an archive copy group can exist.

The following backup copy group parameters significantly influence your backup policy:

VERExists

Determines the maximum number of Exchange Server database backup versions to retain for the databases that exist on the protected Exchange Server system.

VERDeleted

Determines the maximum number of Exchange Server database backup versions to retain for the databases that were deleted from the protected Exchange Server system after they were backed up by IBM Spectrum Protect.

The **VERDeleted** copy group parameter can control both differential and incremental backups. In the case of differential backups, the **VERDeleted** parameter applies when a new full backup is done (which deactivates an active differential backup) until the next differential backup is done. To maintain consistent version-expiration behavior, you can set the **VERDeleted** and **VERExists** parameters to the same value in the management class used for differential backups.

Note: If you set the **VERDeleted** parameter to **nolimit**, the expiration of backup versions is controlled instead by the time-based retention parameters **RETEExtra** and **RETOnly**.

RETEExtra

Determines the number of days to retain an Exchange Server database backup version after that version becomes inactive.

RETOnly

Determines the number of days to retain the last Exchange Server database backup version of a database that was deleted from the protected Exchange Server system.

The **RETOnly** parameter applies to all backup types, including incremental backup objects that can never have more than one version because they are always uniquely named with a unique time stamp. However, all legacy backup objects for an Exchange Server database are deactivated when a new full backup of that Exchange Server database is completed. VSS backup objects remain active.

The retention period that is set in the **RETOnly** parameter controls the expiration of incremental backup objects. When you set the value of the **RETOnly** parameter for incremental backups, the value must be, at a minimum, as long as the value that is set for the full backup on which the incremental backup depends. You can use the same management class for incremental backups and the full backup objects (that are retained the longest) to ensure that an adequate value is used. However, when a new legacy full backup of that Exchange Server database is completed, all

legacy backup objects for an Exchange Server database are deactivated. In this scenario, VSS backup objects remain active.

MODE, SERIALIZATION, FREQUENCY

These parameter settings do not apply to Data Protection for Exchange Server. Therefore, you can accept the default values.

When you plan a backup strategy, as a best practice, consult with the IBM Spectrum Protect server administrator about preferred parameter settings.

Creating a local backup policy

A local backup policy determines how different backup versions are retained on local shadow volumes.

Before you begin

Backup retention on local shadow volumes is determined by your overall backup strategy, the type and number of VSS backup version on IBM Spectrum Protect and on the local shadow volumes, and time-based policies. Ensure that there is sufficient local storage space on local shadow volumes. The amount of space that is required depends on the VSS provider that you use.

Procedure

1. Start Microsoft Management Console (MMC).
2. In the navigation tree, click **IBM Spectrum Protect**.
3. Select an **Exchange Server**, **SQL Server**, or **File System** instance.
4. In the Actions pane, click **Properties**.
5. From the list of available property pages, select **Policy Management**.
6. Add, delete, or update local policies for data retention. When you add a policy, specify a unique policy name. Double-click the policy to edit a policy field. To retain an unlimited number of snapshots, or to retain snapshots for an unlimited number of days, specify NL.
7. Click **Save**.

What to do next

After you add a policy, you can bind a backup to that policy. Updates to existing, bound policies do not take effect until the next backup is run.

Specifying policy binding statements

Bind policy statements to associate Microsoft Exchange Server backups to a management policy.

About this task

A default policy binds any backups that are not explicitly bound to a named policy. Policy binding is available in environments with or without an IBM Spectrum Protect server.

For Exchange Database Availability Groups (DAG), all the DAG members that share the DAG node must use the same VSS policy.

Tip: Use the same policy binding method for Exchange Server backups. Define a policy statement in the configuration file.

Procedure

Specify the policy-binding statements to use to bind snapshots to a policy. Manually add the binding statements in the respective configuration file that defines the policy statements.

Policy-binding statements in the Data Protection for Exchange Server configuration files might look similar to the information in the following table.

	<i>server name</i>	<i>object name</i>	<i>backup type</i>	<i>backup dest</i>	<i>mgmt class</i>
VSSPOLICY	*	"Accounting"	FULL	LOCAL	MC_1
VSSPOLICY	SERVER_3	"Human Resources"	INCR	LOCAL	MC_6

Binding backups to a policy

You can add, update, delete, or change the processing order of binding statements.

About this task

A backup policy determines how backups on local shadow volumes are managed and retained.

Procedure

1. Start Microsoft Management Console (MMC).
2. In the navigation tree, click **IBM Spectrum Protect**.
3. Select an **Exchange Server** instance.
4. In the Actions pane, click **Properties**.
5. From the list of available property pages, select **VSS Policy Binding**.
6. Add, update, delete, or change the processing order of existing binding statements.

Tip: You can use an asterisk (*) as a wildcard character to represent all characters.

For example, in the **Server** field, enter the asterisk to bind the policy to all Exchange Servers.

7. Optional: To change the processing order, use **Move Up** and **Move Down**. Policies are processed from the bottom to the top of the file, and processing stops at the first match.

Tip: To ensure that more specific statements are processed before general statements, list the more general specification before the more specific statement.

8. Save the binding statement.
9. Optional: Verify new or updated policies and bindings.
 - a. Run one or more test backup operations.
 - b. On the **Recover** tab, verify the management classes that are bound to the test backups.

VSSPOLICY statements for backup types

For VSS backups, VSSPOLICY statements are used to associate VSS backups with management classes. When you change from legacy backups to VSS backups, consider the VSSPOLICY statements that you set for the backup.

The VSSPOLICY statements are in a configuration file, for example, `tdpexc.cfg`. A configuration file can include multiple VSSPOLICY statements. The configuration file is read from the bottom to the top of the file. VSSPOLICY statements in the `tdpexc.cfg` file are similar to the INCLUDE statements that are specified in the IBM Spectrum Protect backup-archive client in the `dsm.opt` file.

If no VSSPOLICY statements are included in the configuration file, or if the VSSPOLICY statements do not match the type of backup that is created, the default management class for the policy domain is used. Backup expiration parameters for the default management class might differ from the settings that are used for preexisting legacy backups. For example, the backup expiration period might be set to 30 days. This setting means that after 30 days, the backup is deleted. Verify that the backups expire according to the business needs of your environment.

If you change the `tdpexc.cfg` file, you must restart the IBM Spectrum Protect client acceptor daemon (CAD), IBM Spectrum Protect remote client agent (DSMAgent), and the IBM Spectrum Protect Scheduler Service for Exchange Server. If the DSMAgent service state is set to **Manual (Started)**, stop the service. The DSMAgent service starts when a VSS backup is initiated, but if the service is started and you change the policy settings, the policy settings do not take effect until you restart the service.

Sample VSSPOLICY statements

The following example shows the syntax of a VSSPOLICY statement:

```
VSSPOLICY srv_name db-name backup-type backup-dest mgmtclass
```

where:

- *srv_name* defines the Exchange Server name. You can enter an asterisk (*) as a wildcard character to match all Microsoft Exchange Server.
- *db-name* defines the database name. You can enter an asterisk (*) as a wildcard character to match all Microsoft Exchange Server groups. Because the name can include a space, use the quotation marks to encapsulate the database name.
- *backup-type* defines the backup type for example, FULL or COPY. You can enter an asterisk (*) as a wildcard character to match all backup types.
- *backup-dest* defines the backup destination. Use the TSM option to back up data to IBM Spectrum Protect, the LOCAL option to back up data to a local disk, enter an asterisk (*) as a wildcard character to match both backup destinations.
- *mgmtclass* defines the IBM Spectrum Protect management class that is used to bind the types of specified backups.

In the following example, the VSSPOLICY statement is commented out. Before you can use this policy statement, you must remove the asterisk character (*) from the first column of each line.

```

-----
* Sample VSSPOLICY Statements
* -----
* These statements are used to bind VSS backup to specific TSM
* Server management classes. Adjust the statements to meet your
* needs and remove the leading asterisks to make them operational.
* Note: Matching of these policy bindings are from the bottom up.
*****
* Server      Database      Name      BU Type   BU Dest.  Mgmt Class
* -----
VSSPOLICY *          *          FULL      TSM       IUG_TSM
VSSPOLICY *          *          COPY      TSM       IUG_TSM_COPY
VSSPOLICY *          *          COPY      LOCAL     IUG_COPY
VSSPOLICY *          *          FULL      LOCAL     IUG_LOCAL
VSSPOLICY *          "HR"      FULL      LOCAL     MCLASS3
VSSPOLICY SERVER1 "ACT" *      LOCAL     MCLASS2
VSSPOLICY SERVER1 "SG 1" *      TSM       IUG1
*****

```

In the preceding example, the following policy rules are specified:

- Any VSS backups of the *SG 1* database on the Exchange Server *SERVER1* to IBM Spectrum Protect are bound to the management class *IUG1*.
- Any VSS backups of the *ACT* database on the Exchange Server *SERVER1* to *LOCAL* are bound to the management class *MCLASS2*.
- Full VSS backups of the *HR* database on any Exchange Server to *LOCAL* are bound to the management class *MCLASS3*.
- Full VSS backups of any other database on any other Exchange Server to *LOCAL* are bound to the management class *IUG_LOCAL*.
- Copy VSS backups of any other database on any other Exchange Server to *LOCAL* are bound to the management class *IUG_COPY*.
- Copy VSS backups of any other database on any other Exchange Server to IBM Spectrum Protect are bound to the management class *IUG_TSM_COPY*.
- Full VSS backups of any other database on any other Exchange Server to IBM Spectrum Protect are bound to the management class *IUG_TSM*.
- Any type of backup matches a rule because of the wildcard *VSSPOLICY* statements at the top of the file. Use these types of statements so that you explicitly state what management class is used.

Managing Exchange Database Availability Group members by using a single policy

For Microsoft Exchange Server databases in a Database Availability Group (DAG) environment, several online copies of a database are maintained for high availability. To reduce the number of database backups that are created, set up Data Protection for Exchange Server to back up database copies from different DAG members under a single DAG node.

About this task

You can prevent Data Protection for Exchange Server from backing up each database copy separately by backing up the database copies under a single Database Availability Group (DAG) node.

All database copies can be managed as a single entity regardless of where the database copies are backed up from, and whether the backup copies are active or passive at the time of the backup. You can set up a minimum interval between

database backups, which ensures that the database copies are not backed up at the same time or backed up too frequently.

Procedure

1. Use the IBM Spectrum Protect configuration wizard to configure the DAG node. Ensure that all the DAG members are configured with the same DAG node name.

- For VSS backups to IBM Spectrum Protect, specify a node name in the **DAG Node** field on the IBM Spectrum Protect Node Names page in the wizard. This node is used to back up all the DAG.

2. Grant permission to the DAG member server to act as a proxy for the DAG node. Issue the **proxynode** command for each member server in the DAG. For example, issue the following commands:

```
register node backup_archive_client_node password
userID=backup_archive_client_node
register node data_protection_node password userID=data_protection_node
grant proxynode target=data_protection_node agent=backup_archive_client_node
grant proxynode target=data_protection_node agent=backup_archive_client_node
register node DAG_node password userID=DAG_node
grant proxynode target=DAG_node agent=backup_archive_client_node
grant proxynode target=DAG_node agent=data_protection_node
```

Tip: If you do not use the configuration wizard to configure the IBM Spectrum Protect server, you must define the proxies and assign proxynode authority to the backup-archive client node and the Data Protection node.

3. For a stand-alone configuration, ensure that the DAG node and the node are in the same policy domain.
4. Create a backup schedule and specify the **/MINIMUMBACKUPINTERVAL** parameter with the **backup** command. For example, to back up one copy of a database that contains multiple copies, complete the following steps:
 - a. Create a command script that is named C:\BACKUP.CMD by issuing this command:

```
TDPEXCC BACKUP DB1 FULL /MINIMUMBACKUPINTERVAL=60
```
 - b. Copy the BACKUP.CMD file to all the DAG members.
 - c. Create one schedule and associate all the nodes with this schedule.
5. Run the schedule by using the IBM Spectrum Protect scheduler. When the backup schedule runs, the minimum backup interval is observed and only one backup is created.

What to do next

To decrease the load on the production Exchange Server, you can specify that the backups are taken from a valid passive database copy. If a valid passive copy is not available, the backup copy is created from the active copy of the database. To add this specification, specify the **/PREFERDAGPASSIVE** on the **backup** command, for example:

```
TDPEXCC BACKUP DB1 FULL /MINIMUMBACKUPINTERVAL=60 /PREFERDAGPASSIVE
```

Data restore processing

Data Protection for Exchange Server can use the Microsoft Volume Shadow Copy Service (VSS) framework to complete fast and instant restores of database backups. You also restore VSS backups to an alternate database and complete Exchange mailbox restore operations.

In a VSS restore operation, you restore one or more databases from a VSS backup on IBM Spectrum Protect server storage to the original location on the Exchange Server.

VSS fast restore processing

A VSS fast restore operation restores data from a local snapshot. A VSS fast restore operation overwrites any files that exist at the time of the snapshot on the original source location. The file is overwritten with the version stored on the snapshot. Data is overwritten even if a file is marked read-only.

You can use VSS fast restore operations for the following tasks:

- Restore Exchange Server VSS backups, full, copy, incremental, and differential backup types.
- Restore data at the database level.
- Restore a VSS backup to an alternate database.
- Restore a backup from a DAG replica to the production server.
- Restore a backup from a relocated database to the production server.
- Restore local database backups to only the system that created the backup.

VSS instant restore processing

A VSS instant restore operation restores data by using a hardware-assisted restore method. A FlashCopy® operation is an example of a hardware-assisted restore method. Instant restore processing is done at the volume level.

Even though Exchange Server data is restored relatively quickly, the transaction logs must be replayed after a restore operation. The time of recovery for the database increases as the number of logs to be replayed increases.

VSS instant restore is typically used to restore local VSS backups of the SAN-attached volumes from IBM System Storage SAN Volume Controller, IBM System Storage DS8000® series, IBM Storwize, or XIV. The VSS instant restore process requires a VSS hardware provider. For Exchange Server data, you cannot use instant restore processing for differential and incremental backups.

The data that is to be restored must be on a storage system that is valid for VSS instant restore operations. If data is not on an XIV, SAN Volume Controller, or Storwize family systems with space-efficient target volumes, you must ensure that background copies that use the volumes are restored.

You can manually disable VSS instant restore processing so that IBM Spectrum Protect Snapshot uses VSS fast restore. Instant restore processing is automatically disabled for these VSS restore operations:

- Restore data to an alternate location
- Restore data to an recovery database
- Restore files by issuing the **restorefiles** command

When you plan to implement VSS instant restore operations, consider the following guidelines:

- IBM System Storage DS8000 series requires IBM System Storage Support for Microsoft Volume Shadow Copy Service software.
- SAN Volume Controller requires IBM System Storage Support for Microsoft Volume Shadow Copy Service software.
- Storwize family requires IBM System Storage Support for Microsoft Volume Shadow Copy Service software.
- XIV has separate VSS Provider software.
- Backups can be restored only to the same storage system from which they are backed up.

For instant restore operations, use the devices that are listed here: Storage Architecture Support for Tivoli® Storage FlashCopy Manager and IBM Spectrum Protect Snapshot (<http://www.ibm.com/support/docview.wss?uid=swg21455924>).

VSS backups that are restored to alternate databases

Data Protection for Exchange Server can restore an Exchange Server database backup or DAG active or passive database copy backup, to a recovery database or to an alternate (or relocated) database.

This restore method is called *restore into*. If you are restoring a relocated database, use the *restore into* function. You must specify the same database name as the one you are restoring.

Note: If you use the *restore into* function, VSS instant restore capability is automatically disabled.

Backups to local shadow volumes can be restored only to the system where the backups are created.

Mailbox restore operations

By using Data Protection for Microsoft Exchange Server, you can restore an entire mailbox or selected mailbox items to the original mailbox, another online mailbox on the same server, or to a Unicode or non-Unicode .pst file.

Restriction:

In the Mailbox Restore view of Microsoft Management Console (MMC), you can restore mail items to Unicode and non-Unicode .pst files. In the Mailbox Restore Browser view, however, you can restore mail items to only non-Unicode .pst files.

The Recoverable Items folder is a storage area in a mailbox that contains operational data about the mailbox. Depending on the Exchange Server functions that you enabled for the mailbox, you can recover and restore the following types of mail items in the Recoverable Items folder:

- Deleted item retention
- Single item recovery
- In-place hold
- Litigation hold
- Mailbox audit-logging
- Calendar logging

Typically, a mailbox is set up to protect mail items from being accidentally or maliciously deleted, or to retrieve mail items during litigation or investigations.

Mailbox-enabled functions

You can verify whether a mailbox is enabled for mailbox restore operations by running the following Exchange Powershell cmdlets. In the examples, the mailbox is for George Clark:

Deleted item retention

```
Get-Mailbox "george clark" | FL RetentionHoldEnabled,
    RetainDeletedItemsFor, RetainDeletedItemsUntilBackup
```

Single item recovery

```
Get-Mailbox "george clark" | FL SingleItemRecoveryEnabled
```

In-place hold

```
Get-Mailbox "george clark" | FL InPlaceHolds
```

Litigation hold

```
Get-Mailbox "george clark" | FL LitigationHoldEnabled
```

Mailbox audit-logging

```
Get-Mailbox "george clark" | FL AuditEnabled,
    AuditLogAgeLimit
```

Calendar logging

```
Get-Mailbox "george clark" | FL CalendarVersionStoreDisabled
```

Mail items in the Recoverable Items folder

In the mailbox restore views in Microsoft Management Console (MMC), you can recover and restore mail items from the subfolders within the Recoverable Items folder. You can also complete this task by issuing the **restoremailbox** command. The following table lists the subfolders that are included in the Recoverable Items folder.

Table 7. Exchange Server 2013 Recoverable Items folder contents

Recoverable Items subfolder	Mailbox-enabled functions	Subfolder contents
Deletions	Deleted item retention	Contains mail items that a user deleted from the Deleted Items folder in their mailbox
Versions	<ul style="list-style-type: none"> In-place hold Litigation hold Single item recovery 	Contains the original and modified copies of the deleted mail items
Purges	<ul style="list-style-type: none"> Litigation hold Single item recovery 	Contains all mail items that a user <i>hard deleted</i> , that is, purged from their mailbox
Audits	Mailbox audit-logging	Contains audit log entries
Discovery Holds	In-place hold	Contains mail items that are to be protected from deletion and match <i>hold</i> query parameters
Calendar Logging	Calendar logging	Contains calendar changes that occur within a mailbox

Restriction:

- You cannot restore the Recoverable Items folder and subfolder hierarchy to a mailbox restore destination. You can restore only the contents of the email folders.
- You cannot create a subfolder in the Recoverable Items folder in a mailbox.
- You can restore the Recoverable Items content for a public folder mailbox but not for each public folder in the public folder mailbox.

Related reference:

“Restorefiles command” on page 194

Related information:

 <https://technet.microsoft.com/en-us/library/ee364755%28v=exchg.150%29.aspx>

Data Protection for Exchange Server with IBM SAN Volume Controller and IBM Storwize V7000

The way in which you configure the VSS provider for IBM SAN Volume Controller and IBM Storwize V7000 controls the type of FlashCopy operation that runs when you create a VSS snapshot.

The VSS provider that you use with IBM SAN Volume Controller and IBM Storwize V7000 must have the following characteristics:

- If the VSS provider is configured to use incremental FlashCopy, you can take only one backup version. Each VSS snapshot request for a source volume causes an incremental refresh of the same target volume.

When you delete the VSS snapshot, it is removed from the VSS inventory. If you create another VSS snapshot of the same source volume, the process results in an incremental refresh of the target volume.

The following guidelines apply when you use Data Protection for Exchange Server with SAN Volume Controller- based storage:

- Do not use a combination of space-efficient and fully allocated target volumes. Choose to use either space-efficient or fully allocated volumes for FlashCopy targets. Provision enough target volumes in the SAN Volume Controller VSS_FREE volume group for the backup versions you require. If you use fully allocated target volumes, the capacity size of those volumes must match the size of the source volumes.
- If space-efficient virtual disks (VDisks) are used for backup targets, set the IBM VSS provider background copy value to zero by entering the `ibmvcfg set backgroundCopy 0` command. To activate the changes, restart the IBM VSS system service after you enter the command.

You can transition your data from fully allocated targets to space-efficient targets by using fully allocated targets as if those targets are space-efficient when the background copy rate is set to 0.

- To determine how much storage space is required for each local backup, the backup LUNs require the same amount of storage space as the original LUNs. For example, if you have a 100 GB database on a 200 GB LUN, you need a 200 GB LUN for each backup version.
- Do not use a combination of persistent and nonpersistent VSS snapshots.
- Do not mix COPY and NOCOPY FlashCopy relationships from the same source volume or volumes.

- Enable the autoexpand option for the space-efficient target volumes to avoid out-of-space conditions.
- Allocate enough space for space-efficient target volumes to hold 120 % of the data that is expected to change on the source volume in the time between snapshots. For example, if a database changes at a rate of 20 % per day, VSS backups complete every six hours, and a steady rate of change throughout the day is assumed. The expected change rate between snapshots is 5 % of the source volume (20/4). Therefore, the allocated space for the space-efficient target volumes is to be 1.2 times 5 % equal to 6 % of the source volume size. If the rate of change is not consistent throughout the day, allocate enough space to the target volumes to accommodate the highest expected change rate for the period between snapshots.

You can use VSS instant restore operations with Data Protection for Exchange Server when multiple backup versions exist on IBM SAN Volume Controller and IBM Storwize V7000 space-efficient target volumes.

- Do not delete snapshots manually. Allow Data Protection for Exchange Server when multiple backup versions exist on IBM SAN Volume Controller and IBM Storwize V7000 space-efficient target volume to delete backup versions that are based on the defined policy to ensure that deletion occurs in the correct order.

IBM System Storage requirements

If you use IBM System Storage DS8000 series, SAN Volume Controller, or Storwize family storage systems, be aware of database, log, file, and LUN settings.

Follow these guidelines when you plan for IBM System Storage:

- Place database files on a separate and dedicated logical volume.
- Place logs on a separate logical volume.
- Do not place non-Exchange data on storage volumes that are dedicated to Exchange.
- When you use hardware snapshot providers, ensure that the database LUNs are dedicated to only one database or application.
- If you delete a local snapshot that is stored on an IBM SAN Volume Controller or IBM Storwize V7000 space-efficient volume (SEV) that has multiple dependent targets, delete the snapshots in the same order in which you created the snapshots. You must delete the oldest one first, followed by the second oldest.
- In an IBM SAN Volume Controller or IBM Storwize V7000 environment, if you use multiple target FlashCopy mappings, a mapping might stay in the copying state after all the source data is copied to the target. This situation can occur if mappings that started earlier and use the same source disk are not yet fully copied. In this instance, schedule local backups for IBM SAN Volume Controller and IBM Storwize V7000 storage systems at intervals that are greater than the time required for the background copy process to complete.

Automated IBM Spectrum Protect server failover for data recovery

If you use Data Protection for Exchange Server with the IBM Spectrum Protect configuration, Data Protection for Exchange Server can automatically fail over to the secondary server for data recovery when there is an outage on the IBM Spectrum Protect server.

The IBM Spectrum Protect server that Data Protection for Exchange Server connects to for backup services is called the *primary server*. If the primary server is set up for node replication, the client node data on the primary server can be replicated to another IBM Spectrum Protect server, which is the *secondary server*.

Depending on your configuration, the following nodes must be set up for replication on the primary server:

- Data Protection node
- VSS requestor node (also called the DSM agent node)
- Remote DSM agent node (for offloaded backups to the primary server)
- Exchange Server Database Availability Group (DAG) node for backups of databases in a DAG

During normal operations, connection information for the secondary server is automatically sent to Data Protection for Exchange Server from the primary server. The secondary server information is saved to the client options file (`dsm.opt`). No manual intervention is required by you to add the information for the secondary server.

Each time the backup-archive client logs on to the server for backup services, it attempts to contact the primary server. If the primary server is unavailable, the backup-archive client automatically fails over to the secondary server. In failover mode, you can restore data that is replicated to the secondary server. When the primary server is online again, the backup-archive client automatically fails back to the primary server the next time the backup-archive client connects to the server.

Requirements: To ensure that automated client failover can occur, Data Protection for Exchange Server must meet the following requirements:

- Data Protection for Exchange Server must be at least at V7.1 level or later.
- The primary server, secondary server, and backup-archive client must be at least at V7.1 level or later.
- The primary and secondary servers must be set up for node replication.
- The following nodes must be configured for replication with the `replstate=enabled` option in each node definition on the server:
 - Data Protection node
 - VSS requestor node
 - Remote DSM agent node for offloaded backups
 - DAG node, if applicable
- Before the connection information for the secondary server can be sent to Data Protection for Exchange Server, the following processes must occur:
 - You must back up data at least one time to the primary server.
 - The following nodes must be replicated at least one time to the secondary server:
 - Data Protection node
 - DAG node, if applicable

Restriction: The following restrictions apply to Data Protection for Exchange Server during failover:

- Any operation that requires data to be stored on the IBM Spectrum Protect server, such as backup operations, are not available. You can use only data recovery functions, such as restore or query operations.
- Schedules are not replicated to the secondary server. Therefore, schedules are not run while the primary server is unavailable.
- If the primary server goes down before or during node replication, the most recent backup data is not successfully replicated to the secondary server. The replication status of the file space is not current. If you restore data in failover mode and the replication status is not current, the recovered data might not be usable. You must wait until the primary server comes back online before you can restore the data.

Chapter 2. Installing, upgrading, and migrating

Before you start the installation process, review the appropriate prerequisite information, including hardware and software requirements.

Prerequisites

Before you install Data Protection for Microsoft Exchange Server, ensure that your system meets the minimum hardware, software, and operating system requirements.

Hardware and software requirements change over time due to maintenance updates and the addition of operating system, application, and other software currency support.

For the latest requirements, review the Hardware and Software Requirements technote that is associated with the level of your Data Protection for Exchange Server program. This technote is available at this web page: All Requirements(<http://www.ibm.com/support/docview.wss?uid=swg21219345>)

Follow the link to the requirements technote for your specific release or update level.

Installation process might require a reboot

If you do not install all of the prerequisites before starting the installation process, the installation process might require a reboot. As part of the installation process, one or more Microsoft C++ redistributable packages are installed, if they are not already installed on the Windows workstation. These packages can also be automatically updated by the Windows Update service. If the packages are updated, the update can cause the system to reboot when you start the installation program.

Additionally, because the Microsoft Visual Studio C++ redistributable package is a shared Windows component, other applications that have dependencies on the package might be stopped or restarted by Windows as part of the installation or upgrade of the C++ redistributable package. Schedule installations and upgrades during a maintenance window when other applications are not be adversely affected if they are stopped or restarted when the C++ redistributable package is installed. Monitor other applications after the installation is complete. If applications stopped and did not restart, restart the applications.

Minimum hardware requirements

The following hardware is required to install Data Protection for Exchange Server:

Hardware for an x64 system

Compatible hardware that is supported by the Windows operating system and Exchange Server.

Virtualization environment resources

If you operate in a virtualization environment with Data Protection for Exchange Server, review these resources.

For more information about virtualization environments that can be used with Data Protection for Exchange Server, see this web page: IBM Tivoli Storage Manager (TSM) and IBM Spectrum Protect™ guest support for Virtual Machines and Virtualization (<http://www.ibm.com/support/docview.wss?uid=swg21239546>)

Installing and configuring Data Protection for Microsoft Exchange Server

You can quickly install and configure Data Protection for Exchange Server to start protecting your Exchange Server data.

Before you begin

Before you install and configure Data Protection for Exchange Server, verify that you satisfy the hardware and software requirements.

You can obtain the installation package from the product DVD or from an IBM download site.

- If you obtain the package from the product DVD, ensure that the installation files are visible on the DVD drive.
- If you obtain the package from an IBM download site, you must extract the installation files.

About this task

Data Protection for Exchange Server is available in both licensed and maintenance packages. The installation process differs based on the package type.

Licensed package

Includes a license enablement file that is only available from your software distribution channel, such as Passport Advantage®, and includes the initial General Availability release of a product or component.

Maintenance update (fix pack or interim fix package)

Available from the maintenance delivery channel, and can sometimes be used to refresh the software distribution channel. Maintenance packages do not contain license enablement files and must be installed after a licensed package.

For information about how to install a fix pack or interim fix package, see the README.FTP file. The README.FTP file is available in the same directory where the maintenance package is downloaded.

Related concepts:

“Security requirements for backup and restore operations” on page 87

Installing Data Protection for Exchange Server Procedure

1. Log on to the system as an administrator.
2. Download the appropriate package file from one of the following websites, or you can access the files from the product DVD:
 - For a first time installation or a new release go to Passport Advantage at IBM Passport Advantage. Passport Advantage is the only website from which you can download a licensed package file.
 - For a maintenance fix, go to this FTP site and to the directory that contains the maintenance fix version that you require, Index of Data Protection for Microsoft Exchange Server patch files (<ftp://public.dhe.ibm.com/storage/tivoli-storage-management/patches/tivoli-data-protection/ntexch/>).
3. If you download the package from one of the download sites, complete the following steps:
 - a. Verify that you have enough space to store the installation files when they are extracted from the product package.
 - b. Change to the directory where you placed the executable file.

Tip: In the next step, the files are extracted to the current directory. The path must contain no more than 128 characters. Extract the installation files to an empty directory. Do not extract the files to a directory that contains previously extracted files, or any other files.

- c. Either double-click the executable file, or enter the following command on the command line to extract the installation files. The files are extracted to the current directory.

```
package_name.exe
```

where package_name is like this example:

```
8.1.1.2.0-TIV-TSMEXC-Win.exe
```

4. Follow the installation instructions that are displayed.
5. Click **Finish**.
6. If you plan to use VSS operations, you must install the most recent version of the IBM Spectrum Protect backup-archive client. The backup-archive client is also the VSS Requestor and is available separately.

Completing the installation configuration Procedure

1. To start the MMC, click **Start > All Programs > IBM Spectrum Protect > Data Protection for Microsoft Exchange Server > DP for Exchange Management Console**. If you did not previously configure Data Protection for Exchange Server, the IBM Spectrum Protect configuration wizard starts automatically.
2. If the IBM Spectrum Protect configuration wizard does not start automatically, click **Manage > Configuration > Wizards** in the navigation tree, select the wizard, and click **Start** in the Actions pane.
3. Complete the following pages of the wizard:

Data Protection Selection

Select **Exchange Server** as the application to protect.

Requirements Check

Click any **Failed** or **Warnings** links to resolve errors.

If you do not have all the user roles that are required for individual mailbox restore operations, click the **Warnings** link and follow the wizard prompts to add the missing Exchange Server roles. If you are a member of the Exchange Organization Management group and have sufficient role-based access control (RBAC) permissions, you can automatically add the missing roles. If you are not a member of the Exchange Organization Management group and have insufficient RBAC permissions, you must manually add the missing roles.

TSM Node Names

Specify the IBM Spectrum Protect node names to use for the applications that you want to protect.

- In the **VSS Requestor** field, enter the node name that communicates with the VSS Service to access the Exchange data. This node name is the IBM Spectrum Protect client node name, also known as the DSM agent node name.
- In the **Data Protection for Exchange** field, enter the node name where the Data Protection for Exchange Server application is installed. This node stores the Data Protection for Exchange Server backups. If you configure the **DAG Node**, the DAG database backups are not stored under the Data Protection node. The backups are stored under the DAG node. Regardless, the Data Protection node must be defined.
- In the **DAG Node** field, enter the node name that you want to use to back up databases in an Exchange Server Database Availability Group. With this setting, all active and passive copies of the databases are backed up to the same file space on the IBM Spectrum Protect server. The database copies are managed as a single entity, regardless of which Database Availability Group member they were backed up from. This setting can prevent Data Protection for Exchange Server from making too many backups of the same database.

Important: On the IBM Spectrum Protect server, ensure that you register the DAG node. All DAG members need proxy authority to run backups on behalf of the DAG node.

TSM Server Settings

Specify the IBM Spectrum Protect server address, and choose whether to have the wizard configure the IBM Spectrum Protect server. Alternatively, you can view and change the commands that the configuration wizard uses to configure the IBM Spectrum Protect server, or run manually run the commands.

Custom Configuration

Click **Default** in most situations, or click **Custom** to enter all service-related information.

TSM Configuration

Wait for all components to be provisioned and configured. Click **Re-run** if there are any problems. Click the **Failed** or **Warnings** link for more information if any problems remain.

Completion

The configuration status is displayed. Select the **VSS Diagnostics** check box to begin VSS verification.

If you do not use the wizard to configure the IBM Spectrum Protect server, the IBM Spectrum Protect administrator must configure the server before verification can be completed. If the wizard does not configure the server, it provides a link to a macro that can be provided to the IBM Spectrum Protect administrator as an example of one way to configure the server.

Verifying the configuration

Procedure

1. Verify that VSS is working correctly.

If the **VSS Diagnostics** check box is selected at the completion of the configuration wizard, the VSS Diagnostics wizard is displayed. You can also start this wizard by clicking **Manage > Diagnostics**, and clicking **VSS Diagnostics** in the Actions pane.

Do not run these tests if you are already using SAN Volume Controller or Storwize V7000 space-efficient snapshots on your computer. Doing so can result in the removal of previously existing snapshots.

2. Complete the following pages in the VSS Diagnostics wizard:

Snapshot Volume Selection

Select the volumes that you want to test and review the VSS provider and writer information.

VSS Snapshot Tests

Review event log entries that are logged as the persistent and non-persistent snapshots are taken, and resolve any errors.

Completion

Review the test status and click **Finish**.

3. Verify that Data Protection for Exchange Server is configured properly:
 - a. Click the **Automate** tab. An integrated command line is available in the task window. You can use the interface to enter PowerShell cmdlets or command-line interface commands. The output is displayed in the main window.
 - b. Change **PowerShell** to **Command Line**.
 - c. Click the folder icon, and select the `verify_exc.txt` file. Then, click **Open**. These commands are displayed in the command-line pane:

```
query tdp
query tsm
query exchange
```
 - d. With the cursor in the command-line pane, press **Enter** to run the commands to verify your configuration. The configuration is verified when these commands run without warnings or errors.
 - e. When verification is complete, you can use Data Protection for Exchange Server to back up and restore Exchange Server data.
 - f. Back up and restore a set of test data.

Customizing the configuration Procedure

After you successfully configure Data Protection for Exchange Server, define your policy settings and scheduled operations to meet your business requirements.

Related concepts:

“Security requirements for backup and restore operations” on page 87

Installing Data Protection for Exchange Server on a local system

You can install Data Protection for Exchange Server from a DVD. The setup wizard guides you through the process of installing Data Protection for Exchange Server.

Before you begin

Before you install and configure Data Protection for Exchange Server, verify that you satisfy the hardware and software requirements.

You can obtain the installation package from the product DVD or from an IBM download site.

- If you obtain the package from the product DVD, ensure that the installation files are visible on the DVD drive.
- If you obtain the package from an IBM download site, you must extract the installation files.

About this task

Data Protection for Exchange Server is available in both licensed and maintenance packages. The installation process differs between these two package types.

Licensed package

Includes a license enablement file that is only available from your software distribution channel, such as Passport Advantage, and includes the initial General Availability release of a product or component.

Maintenance update (fix pack or interim fix package)

Available from the maintenance delivery channel, and can sometimes be used to refresh the software distribution channel. Maintenance packages do not contain license enablement files and must be installed after a licensed package.

For information about how to install a fix pack or interim fix package, see the README.FTP file. The README.FTP file is available in the same directory where the maintenance package is downloaded.

Procedure

Install Data Protection for Exchange Server by using the setup wizard. The wizard installs the product and any prerequisites such as the .NET Framework and Report Viewer.

1. Log on as an administrator.
2. Download the appropriate package file from one of the following websites, or you can access the files from the product DVD:

- For a first time installation or a new release go to Passport Advantage at IBM Passport Advantage. Passport Advantage is the only website from which you can download a licensed package file.
 - For a maintenance fix, go to this FTP site and to the directory that contains the maintenance fix version that you require, Index of Data Protection for Microsoft Exchange Server patch files (<ftp://public.dhe.ibm.com/storage/tivoli-storage-management/patches/tivoli-data-protection/ntexch/>).
3. If you download the package from one of the download sites, complete the following steps:
 - Verify that you have enough space to store the installation files when they are extracted from the product package.
 - Change to the directory where you placed the executable file.

Tip: In the next step, the files are extracted to the current directory. The path must contain no more than 128 characters. Extract the installation files to an empty directory. Do not extract the files to a directory that contains previously extracted files, or any other files.

- Either double-click the executable file, or enter the following command on the command line to extract the installation files. The files are extracted to the current directory.

```
package_name.exe
```

where package_name is like this example:

```
8.1.2.0-TIV-TSMEXC-Win.exe
```

4. Follow the installation instructions that are displayed.
5. If prompted, restart your system before the installation is completed.
6. Click **Finish**. Microsoft Management Console (MMC) is shared among Data Protection for Exchange Server, Data Protection for SQL Server, and IBM Spectrum Protect Snapshot. If one of these products is installed in a location other than the default location, the setup wizard defaults to the existing installation directory. Use the same directory when you install any of these products on the same computer. The default base directory is `c:\program files\tivoli`.

Related concepts:

“Prerequisites” on page 33

Implementing a silent (unattended) installation of Data Protection for Microsoft Exchange Server

To implement a silent installation of Data Protection for Exchange Server, you can use the setup program or the Microsoft Installer program. The installation package can be made available on a DVD or in a shared directory on a file server.

About this task

A silent installation is useful when Data Protection for Exchange Server must be installed on a number of different computers with identical hardware. For example, a company might distribute 25 Exchange Server installations across 25 different sites.

To ensure a consistent configuration and to avoid having 25 different people enter Data Protection for Exchange Server parameters, an administrator can choose to produce an unattended installation package and make it available to the 25 sites.

The installation package can be placed on a DVD and sent to each of the remote sites, or the package can be placed in a shared directory on a file server for distribution across the different sites.

To implement a silent installation of Data Protection for Exchange Server, you can use the setup program or the Microsoft Installer program.

Silently installing Data Protection for Exchange Server with the setup program

You can use the setup program to silently install Data Protection for Exchange Server.

Before you begin

You must install two components: Data Protection for Exchange Server Management Console and Data Protection for Exchange Server. The setup programs for these components are on the installation media (where x: is your DVD drive):

Data Protection for Exchange Server Management Console setup program

(64-bit) x:\fcm\x64\mmc\8120\enu\spinstall.exe

Data Protection for Exchange Server setup program

(64-bit) x:\fcm\x64\exc\8120\enu\spinstall.exe

The Data Protection for Exchange Server Management Console and Data Protection for Exchange Server must be installed from an account that is a member of the local Administrators group for the system on which the Exchange Server is running.

Procedure

1. Issue the following commands to silently install both components to the default installation directories:

```
x:\fcm\x64\mmc\8120\enu\spinstall.exe /s /v/qn
x:\fcm\x64\exc\8120\enu\spinstall.exe /s /v/qn
```

where x: is your DVD drive.

2. The following examples are commands that specify the target directory, the features, start suppression, and logging. Specify each command on a single line.

```
x:\fcm\x64\mmc\8120\enu\spinstall.exe /s /v"INSTALLDIR="C:\Program Files\Tivoli\
ADDLOCAL="Client" TRANSFORM=1033.mst REBOOT=ReallySuppress /qn /!*v
"C:\Temp\DpExcMmcspinstallLog.txt\"
x:\fcm\x64\exc\8120\enu\spinstall.exe /s /v"INSTALLDIR="C:\Program Files\Tivoli\tsm\
ADDLOCAL="Client" TRANSFORM=1033.mst REBOOT=ReallySuppress /qn /!*v
"C:\Temp\DpExcspinstallLog.txt\"
```

3. Review these guidelines as you complete the installation process:
 - You must place a backslash (\) character before each quotation mark that is within an outer set of quotation marks (").
 - For a single-line command, press **Enter** only when all the parameters are entered.
 - You must place quotation marks (") around the following text:
 - A directory path that contains spaces.
 - An argument that specifies multiple features. Although you must use quotation marks around the complete argument, you must still place a backslash (\) character before each internal quotation mark.

- All features that are listed in a custom installation must be listed after the **addlocal** option.
- Setting the **rebootyesno** option to *No* applies only to the installation of the Data Protection for Exchange Server software. The installation package includes a number of prerequisites that is installed by Data Protection for Exchange Server. Ensure that all the prerequisites are installed before starting the silent installation, then set the **rebootyesno** option to *No* so that no system restart is required after the silent installation process finishes.

Batch files usage in silent installations

You can create a batch file to begin the silent installation with the parameters that you want to use.

The following script is a sample script (c:\spinstall.bat) of an unattended installation:

```
@echo off
rem =====
rem sample silent install script
rem
call x:\fcm\x64\mmc\8120\enu\spinstall.exe /s
/v"INSTALLDIR=\"C:\Program Files\Tivoli\"
ADDLOCAL=\"Client\" TRANSFORM=1033.mst
REBOOT=ReallySuppress /qn /! *v \"C:\Temp\DpExcMmcspinstallLog.txt\"
rem
call x:\fcm\x64\exc\8120\enu\spinstall.exe /s
/v"INSTALLDIR=\"C:\Program Files\Tivoli\tsm\"
ADDLOCAL=\"Client\" TRANSFORM=1033.mst
REBOOT=ReallySuppress /qn /! *v \"C:\Temp\DpExcspinstallLog.txt\"
rem =====
rem code could be added after the
rem installation completes to
rem customize the dsm.opt files
rem if desired
rem =====
```

Silently installing Data Protection for Exchange Server with the Microsoft Installer program

You can use the Microsoft Installer (MSI) program, `msiexec.exe`, to implement a silent installation of Data Protection for Exchange Server.

Before you begin

Data Protection for Exchange Server must be installed from an account that is a member of the local Administrators group for the system on which the Exchange Server is running.

Important: Unlike the `spinstall.exe` and `setupfcm.exe` programs, the `msiexec.exe` program does not include a number of prerequisites that are required by Data Protection for Exchange Server. When you use `msiexec.exe`, you must install all prerequisites manually.

Before you install and configure Data Protection for Exchange Server, verify that you satisfy the hardware and software requirements.

About this task

The following examples show how to use the **msiexec** command to install the Data Protection for Exchange Server Management Console and Data Protection for Exchange Server.

Procedure

1. To install the Data Protection for Exchange Server Management Console, issue each of these **msiexec** commands on a single line.

```
msiexec /i"x:\fcm\x64\mmc\8120\enu\IBM Spectrum Protect for Mail  
- MS Exchange - Management Console.msi" RebootYesNo="No"  
Reboot="Suppress" ALLUSERS=1 INSTALLDIR="c:\program files\tivoli"  
ADDLOCAL="Client" TRANSFORM=1033.mst /qn /!v "c:\temp\DpExcMmcLog.txt"
```

Where *x*: is your DVD drive.

2. To install Data Protection for Exchange Server, issue each of these **msiexec** commands on a single line:

```
msiexec /i"x:\fcm\x64\exc\8120\enu\IBM Spectrum Protect for Mail  
- MS Exchange.msi" RebootYesNo="No" Reboot="Suppress" ALLUSERS=1  
INSTALLDIR="c:\program files\tivoli\tsm" ADDLOCAL="Client"  
TRANSFORM=1033.mst /qn /!v "c:\temp\DpExcLog.txt"
```

Where *x*: is your DVD drive.

What to do next

Important:

- You must place quotation marks around the following items:
 - A directory path that contains spaces.
 - An argument that specifies multiple features. Although you must use quotation marks around the complete argument, you must still place a backslash before each internal quotation mark.
- All features that are listed in a custom installation must be specified after the **addlocal** option.

Related concepts:

“Prerequisites” on page 33

Creating and testing a silent installation package on a DVD or a file server

The administrator can choose to make an installation package available by burning a DVD or placing the package in a shared directory on a file server.

Before you begin

Before you begin, you must choose a location for the package. If you are burning a DVD, it is convenient to use a staging directory. If you are placing the package on a file server, you can use a staging directory or build the package directly on the file server.

About this task

Typically, the installation package contains the Data Protection for Exchange Server code distribution files and a batch file for a silent installation.

Procedure

1. Issue the following commands to create the package:

Table 8. Commands for creating a silent installation package

Command	Description
<code>mkdir c:\tdpdpkg</code>	Create a staging directory for the silent-install package
<code>cd /d c:\tdpdpkg</code>	Go to the staging directory
<code>xcopy g:*.* . /s</code>	Copy the DVD distribution files to the staging directory
<code>copy c:\spinstall.bat</code>	Replace the existing <code>spinstall.bat</code> with the one created in the previous step

This example uses `c:\tdpdpkg` as a staging directory.

2. After you create the installation package, test the silent installation.
3. After you complete the test, place the package on a DVD or make it available from a shared directory.
4. After you make the package available on a DVD or from a shared directory, complete these steps to run the silent installation package on another computer.

Option	Description
From a silent installation package on DVD:	Enable the <code>autostartoption</code> to cause the silent installation to begin as soon as the DVD is inserted into the drive. If you do not enable the <code>autostartoption</code> , start the <code>spinstall.bat</code> file from the root of the DVD by issuing the following command: <code>cd /d g:\</code> <code>spinstall.bat</code>
From a distribution directory:	If the package is placed in a shared directory that is called <code>tdpdpkg</code> at <code>\\machine1\d\$</code> , another computer can run the <code>net use x: \\machine1\d\$</code> command to share the drive as drive <code>x</code> . You can issue the following command: <code>cd /d x:\tdpdpkg</code> <code>spinstall.bat</code>

In either case, the silent installation begins. Allow enough time for the unattended installation to complete. No visual cues exist to inform you when the installation is finished, although you can add visual cues to the batch file.

Silent installation error messages

The `spinstall.exe` program can produce error messages if it cannot start properly.

In most cases, administrators encounter these messages when a severe error occurs. Users rarely see these messages. When you get an error message, it displays in a message box. Every error message has a number. These messages are system error messages and there is no way to suppress them in your script.

Upgrading Data Protection for Microsoft Exchange Server

You can upgrade Data Protection for Exchange Server from an earlier version of the software.

Procedure

1. Download the updates.
2. To install the updates, run setupfcm.exe.
3. To start Microsoft Management Console (MMC), click **Start > All Programs > IBM Spectrum Protect > Data Protection for Microsoft Exchange Server > DP for Exchange Management Console**. When you start MMC after you install the updates, the configuration wizard automatically starts. The configuration wizard guides you through the process of provisioning and installing the remaining files.

Depending on the software licenses that are found on the system, the configuration process varies. The wizard provides instructions to guide you through the process.

4. If the configuration wizard does not start automatically, click **IBM Spectrum Protect** in the navigation tree, and click **Configuration**. Then, double-click **Wizards**.

Data Protection for Exchange Server migration

You can migrate data from earlier versions of Data Protection for Exchange Server.

After you upgrade from an older version of Data Protection for Exchange Server to a newer version, you can use VSS data restore operations to restore VSS backups that were originally created with the older version of the software.

Managing migrated backups to a Database Availability Group node

When you configure Data Protection for Exchange Server to back up databases in a DAG to a common DAG node, all DAG databases are backed up with the new DAG node name.

Before you begin

If you are migrating from a version that is earlier than Data Protection for Exchange Server V6.4, manage the backups from the previous versions by following these guidelines:

- Do not mix backups that are created with previous versions of Data Protection for Exchange Server with new backups that are created by using the DAG node. To separate the backups, keep the previous backups under the previous Data Protection node name that is defined in the `dsm.opt` file in the `C:\Program Files\Tivoli\tsm\TDPEExchange` directory, and use a new DAG node name to store the new backups.
- To view or restore a backup that is stored under the previous node name, you must change the Data Protection for Exchange Server configuration.
- You must manually delete backups over time assuming that the old backups are no longer useful.

Procedure

1. After you complete your migration, ensure that the first backup you do is a full backup.
2. To view and restore backups that are stored under the previous Data Protection node name, complete these steps:
 - a. Remove the **DAG Node** by using the General properties page, configuration wizard, or the **set** command on the command-line interface.
 - b. Restart or refresh Microsoft Management Console (MMC) or command-line interface.
 - c. Click the **Recover** tab in MMC, or run a `tdpexcc query tsm *` command. Because the **DAG Node** parameter is not set, Data Protection for Exchange Server lists the backups that are stored under the Data Protection for Exchange Server node.
 - d. Proceed to restore one or more of the listed backups.
3. If required, delete the backups that are expired.

Related concepts:

“Policy management with Data Protection for Microsoft Exchange Server” on page 14

Chapter 3. Configuring

You can use configuration wizards to configure Data Protection for Microsoft Exchange Server, or you can complete the steps manually. For best results, be guided by the step-by-step instructions in the configuration wizards.

About this task

The following list identifies the ways to configure Data Protection for Microsoft Exchange Server software by using the configuration wizard.

TSM Configuration

When you select the TSM Configuration configuration option, you configure Data Protection for Microsoft Exchange Server to work with IBM Spectrum Protect server.

Data Protection for Microsoft Exchange Server must be installed on your system. An IBM Spectrum Protect server must be available to communicate with Data Protection for Microsoft Exchange Server.

Mailbox Restore Only

When you select the Mailbox Restore Only configuration option, you configure Data Protection for Microsoft Exchange Server to restore mailboxes from Exchange database .EDB files. Additional data protection features are not available. This option is ideal when you only want to restore mailboxes from .EDB files and do not want the additional Data Protection for Microsoft Exchange Server software functionality. The functionality offered with this configuration option is included in the other configuration options.

Proxy node definitions for VSS backups

Data Protection for Exchange Server uses the IBM Spectrum Protect backup-archive client to implement VSS backup operations. As such, you must use two node names for VSS operations; one for the backup-archive client and the other for Data Protection for Exchange Server.

As part of the configuration procedure, a proxy relationship is defined for these node names. By default, this proxy relationship is defined when you run the configuration wizard. Follow the guidelines in this topic to manually complete the configuration.

The proxy relationship allows node names to process operations on behalf of another node name. When you register these nodes to the IBM Spectrum Protect server for VSS operations, specify the IBM Spectrum Protect `USerid=<node name>` parameter.

Two types of node names are defined in proxy node relationships:

- *Target node*: A node name that controls data backup and restore operations and also owns the data on the IBM Spectrum Protect server. This node name is specified in the Data Protection for Exchange Server `dsm.opt` file.
- *Agent node*: A node name that processes operations on behalf of a target node. This node name is specified in the backup-archive client `dsm.opt` file.

To establish the proxy relationship, on the IBM Spectrum Protect server, issue the **grant proxynode** command. For example:

```
GRANT PROXYNODE TARGET=dpexc_node_name AGENT=dsmagent_node_name
```

If you are running backups of availability databases in an AlwaysOnAvailability Group on SQL Server 2012 and later versions, a cluster node name is also required.

- *Cluster node*: A node name that stores data in a failover cluster or AlwaysOn availability group configuration.

To establish the proxy relationship, on the IBM Spectrum Protect server, issue the **grant proxynode** command. For example:

```
GRANT PROXYNODE TARGET=alwayson-node agent=tdpsql-node
GRANT PROXYNODE TARGET=alwayson-node agent=dsmagentnode
GRANT PROXYNODE TARGET=tdpsql-node agent=dsmagentnode
```

Required node names for basic VSS operations

VSS operations require specific node name settings.

To process basic VSS operations, you must have one target node and one agent node.

Table 9. Required node names for basic VSS operations

Proxy node type	Node name	Where to specify
Target node	The Data Protection for Exchange Server node name	Use the nodename option in the Data Protection for Exchange Server options file (dsm.opt)
Agent node	The Local DSMAGENT Node name that must match the backup-archive client node name	Use the localdsmagentnode parameter in the Data Protection for Exchange Server configuration file (tdpexc.cfg)

Note: For basic VSS operations, the agent node and target node are on the same system.

Required node names for basic VSS offloaded backups

VSS offloaded backups require specific node name settings.

To complete VSS offloaded backups, you must have one target node and two agent nodes:

Table 10. Required node names for basic VSS offloaded backups

Proxy node type	Node name	Where to specify
Target node	Data Protection for Exchange Server node name	Use the nodename option in the Data Protection for Exchange Server options file (dsm.opt)
Agent node	Local DSMAGENT Node	Use the localdsmagentnode parameter in the Data Protection for Exchange Server configuration file (tdpexc.cfg)
Agent node	Remote DSMAGENT Node	Use the remotedsmagentnode parameter in the Data Protection for Exchange Server configuration file (tdpexc.cfg)

Target node

This node name is where Data Protection for Exchange Server is installed. This node name (specified with the **nodename** option in the `dsm.opt` file) is referred to as the Data Protection for Exchange Server node name.

Agent node - Local DSMAGENT Node

This node name is where the backup-archive client and VSS provider are installed.

This node processes the VSS operations because Data Protection for Exchange Server do not process VSS operations.

This node name is referred to as the Local DSMAGENT Node and is specified with the **localdsmagentnode** parameter in the Data Protection for Exchange Server configuration file (`tdpexc.cfg` by default). To specify this parameter with the Properties window of Microsoft Management Console (MMC), select VSS backup. From the Properties window, you can update the Local DSMAGENT Node name. Otherwise, use the **tdpexcc set** command to specify this parameter.

Agent node - Remote DSMAGENT Node

This node name is a separate system that must also have the backup-archive client, VSS provider, and the Exchange System Management Tools installed (make sure you install the same level of the Exchange System Management Tools that is installed on your Exchange production server).

This node moves VSS snapshot data from local shadow volumes to the IBM Spectrum Protect server. It is also responsible for doing the Exchange Integrity Check.

This node name is referred to as the Remote DSMAGENT Node and is specified with the **remotedsmagentnode** parameter in the Data Protection for Exchange Server configuration file (`tdpexc.cfg` by default). To specify this parameter with the Properties window of MMC, select VSS backup. From here, you can update the Remote DSMAGENT Node name. Otherwise, use the **tdpexcc set** command to specify this parameter.

The choice of available systems depends on whether the systems have access to the local shadow volumes that contain the VSS snapshot backups. This node name is only valid for VSS environments that support shadow copies that can be transported. You cannot specify the node name if you are using the default VSS system provider.

Ensure that the **localdsmagentnode** and **remotedsmagentnode** are registered to the same IBM Spectrum Protect server that is specified in the Data Protection for Exchange Server options file (`dsm.opt`) and the backup-archive client options file (also `dsm.opt`).

Specifying configuration parameters for IBM Spectrum Protect

After Data Protection for Exchange Server is registered to IBM Spectrum Protect, you must configure the node name, password, the communications method, and the appropriate parameters to connect to the IBM Spectrum Protect server.

About this task

Parameter values are stored in an options file that is located by default in the Data Protection for Exchange Server installation directory.

Procedure

1. Use a text editor to edit the `dsm.optoptions` file.

The `dsm.opt` options file includes the following parameters, which are necessary for initial configuration:

COMMethod

Specify the communication protocol to use between the Data Protection for Exchange Server node and the IBM Spectrum Protect server.

Depending on the `commethod` option that you choose, specify one of the following connectivity parameters for the `commethod` values.

- For VSS backups, specify the **COMMethod** option in the Data Protection for Exchange Server options file.
- For VSS backups, specify the **COMMethod** option in the backup-archive client options file that is used as the local DSMAGENT node. If the environment is configured for VSS offloaded backups, also specify the **COMMethod** option in the backup-archive client options file that is used as the remote DSMAGENT node.

NODename

Specify the IBM Spectrum Protect node name that IBM Spectrum Protect uses to identify the system that runs Data Protection for Exchange Server.

PASSWORDAccess

Specify either the default `generate` value to generate a password automatically, or specify the `prompt` password to respond to a request for a password.

2. Optional: modify the default values for the following parameters:

CLUSTERnode

Leave this option blank. The default value is used.

COMPRESSION

Specify the `compression yes` option if any of the following conditions exist:

- The network adapter has a data overload
- Communications between Data Protection for Exchange Server and IBM Spectrum Protect server are over a low-bandwidth connection
- Heavy network traffic exists
- You can also use the `compressalways yes` option (with the `compression yes` setting) to specify that file compression must continue even if the file grows as a result of data compression.

Specify the compression no option if any of the following conditions exist:

- The computer that runs Data Protection for Exchange Server has a processor overload; the added processor usage might cause issues for other applications that include the server. You can monitor processor and network resource usage with the Performance Monitor program that is included with Windows.
- You are not constrained by network bandwidth; you can achieve the best performance by leaving the compression no option unchanged and enabling hardware compaction on the tape drive, which also reduces storage requirements.

For VSS backups, specify the **COMPRESSIon** option in the backup-archive client options file that is used as the local DSMAGENT node. If the environment is configured for VSS offloaded backups, specify the **COMPRESSIon** option in the backup-archive client options file that is used as the remote DSMAGENT node.

DEDUPLICATION

Specify whether the IBM Spectrum Protect API deduplicates data before the data is sent to the IBM Spectrum Protect server. Specify Yes or No. The value applies only if IBM Spectrum Protect allows client-side data deduplication.

When you specify both deduplication and **ENABLELANFree** options, the deduplication option is ignored.

You can enable client-side data deduplication by specifying DEDUPLICATION YES in the dsm.opt file.

ENABLECLIENTENCRYPTKEY

Specify this option to encrypt databases during backup and restore processing by generating one random encryption key per session.

Restriction: You can back up encrypted VSS databases only to the IBM Spectrum Protect server. You cannot back up encrypted data to a Data Protection for Exchange Server.

You can specify DES56 (56 bit), AES128 (128 bit), or AES256 (256 bit). The most secure data encryption method is AES256.

In the options file, you must also specify the databases that you want to encrypt by adding an include statement with the `include.encrypt` option.

For VSS backups, specify the encryption options in the backup-archive client options file that is used as the local DSMAGENT node. If the environment is configured for VSS offloaded backups, specify the encryption options in the backup-archive client options file that is used as the remote DSMAGENT node.

If you make changes in the backup-archive client options file, ensure that you restart the IBM Spectrum Protect Client Acceptor Daemon (CAD) service for the Exchange or SQL Server.

ENABLELANFree

If you run data backup and restore operations in a LAN-free environment, specify **ENABLELANFree** *yes* in the DSMAGENT (VSS Requestor) options file.

INCLUDE and EXCLUDE

To set policy for VSS backups, specify the VSSPOLICY statement in the Data Protection for Exchange Server configuration file.

The general **include** and **exclude** syntax is similar to the following syntax:

```
include "objectNameSpecification" [ManagementClassName]
exclude "objectNameSpecification"
```

where `objectNameSpecification` is as follows:

```
ExchangeServerName\ExchangeDatabaseName\...\backupType
```

and where `backupType` is one of the following types:

```
full, copy, incr, diff
```

This example binds mailbox history objects to management class CLASS4:

```
INCLUDE "\...\MAILBOXINFO\...\*" CLASS4
```

What to do next

You can create more Data Protection for Exchange Server options files to point to another IBM Spectrum Protect server. You can create more than one options file, where each file contains different parameters to use with a single IBM Spectrum Protect server.

Specifying Data Protection for Exchange Server DAG member name parameters

You must register the system where Data Protection for Exchange Server is installed to the IBM Spectrum Protect server with a Database Availability Group (DAG) member name.

About this task

When you configure Data Protection for Exchange Server, the IBM Spectrum Protect configuration wizard manages the creation of the IBM Spectrum Protect nodes and node attributes. You can customize the configuration template script to add additional node attributes, for example, backup compression. Alternatively, to customize IBM Spectrum Protect nodes, you can use the Administrative client options with the **DSMADMC** command.

The DAG member name owns and manages all Data Protection for Exchange Server data that is backed up to the IBM Spectrum Protect server.

Procedure

1. Specify the DAG member name with the **nodename** option in the `dsm.opt` options file. By default, the `dsm.opt` options file is in the Data Protection for Exchange Server installation directory.
2. To run VSS operations, register DAG member names for more systems if required.
3. Configure the following IBM Spectrum Protect parameters when you register your Data Protection for Exchange Server DAG member name to the IBM Spectrum Protect server:

- **MAXNUMMP** Specify the maximum number of mount points that a client node is allowed to use on the IBM Spectrum Protect server during a backup operation.
- **TXNGroupmax** Specify the number of files that are transferred as a group between Data Protection for Exchange Server and the IBM Spectrum Protect server, between transaction commit points. This parameter must have a value of 12 or greater.
- **COMPRESSIon** Specify whether the backup-archive client node compresses data before it sends the data to the IBM Spectrum Protect server during a backup operation. For VSS operations, specify `COMPRESSion=Yes` in the backup-archive client options file (`dsm.opt`) in the backup-archive client directory.

Specifying configuration and options files in non-default locations

The Data Protection for Microsoft Exchange Server software uses default configuration and options files. If you want to use non-default configuration and options files, use command-line parameters to specify alternative configuration and option files when you start Data Protection for Microsoft Exchange Server.

Before you begin

The information in this procedure does not apply to managing remote Data Protection for Microsoft Exchange Server installations.

About this task

MMC that is used for Data Protection for Microsoft Exchange Server software is started with the `flashcopymanager.exe` file. The `flashcopymanager.exe` file accepts the following parameters:

```
/mscFilename=filename # Name of the MMC snap-in control file
/author # Opens the MMC console in author mode.
```

For example:

```
flashcopymanager.exe parameter1=filename
parameter2=filename ...
```

The `flashcopymanager.exe` file accepts the following parameters to set the configuration files:

```
/EXCCONFigfile=filename # Exchange configuration file
/EXCOPTfile=filename # Exchange OPT file
```

Procedure

Start MMC with the parameters by using `flashcopymanager.exe`. You can also start and run multiple instances of MMC concurrently. With the command-line parameters, each instance operates by using a different configuration that is based on the specified configuration and option files.

Setting user preferences

Use the property pages in the Data Protection Properties window to customize Data Protection for Exchange Server configuration preferences.

Before you begin

The property pages customize preferences such as logging of activity, how languages and information are displayed, and tune performance. The information about the General property page is required to back up data, but the properties are set when you complete the configuration wizard.

When configuring preferences, consider the backup strategy, resource needs, policy settings, and hardware environment of your system.

Procedure

To configure user preferences, complete the following steps:

1. In the navigation tree of Microsoft Management Console (MMC), select the **SQL** instance for which you want to edit preferences.
2. Click **Properties** in the Actions pane.
3. Edit the property page and click **OK** to save your changes and close the window.

What to do next

Tip: You can also view or edit properties for the dashboard and the Management Console. To open the properties window, click **Dashboard** in the navigation tree, and click **Properties** in the Actions pane.

Data Protection properties

Use property pages to customize your configuration preferences.

The available property pages for a workload vary depending on whether it is configured for the stand-alone environment or the IBM Spectrum Protect environment.

You can view or edit property pages by selecting a workload from the **Protect and Recover Data** node in the navigation tree of the Management Console, and clicking **Properties** in the Actions pane.

Server Information

Use the Server Information property page to obtain information about the server that provides backup services.

The fields that display depends on whether the product is configured for a stand-alone snapshot environment or for an IBM Spectrum Protect environment.

Note: References to the stand-alone snapshot environment are specific to IBM Spectrum Protect Snapshot.

Node name

Specifies the name that is used to identify the client node for stand-alone backup operations or backup operations to IBM Spectrum Protect server.

TSM API version

Specifies the version of the IBM Spectrum Protect application programming interface (API).

Server name

For backups to IBM Spectrum Protect, specifies the name of the IBM Spectrum Protect server that you are connected to.

For a stand-alone configuration, `Virtual Server` is displayed.

Server Network Host name

Specifies the network host name for the IBM Spectrum Protect server.

For a stand-alone configuration, `FLASHCOPYMANAGER` is displayed.

Server type

For backups to IBM Spectrum Protect, specifies the type of operating system of the IBM Spectrum Protect server.

For a stand-alone configuration, `Virtual Platform` is displayed.

Server version

Specifies the version of the IBM Spectrum Protect server.

Compression mode

Specifies whether compression is used during backup operations to the IBM Spectrum Protect server. The possible values are `Yes`, `No`, and `Client Determined`.

Domain name

Specifies the policy domain that the node belongs to. A policy domain contains one or more policy sets.

For Exchange systems, the domain name, policy set, and management class are listed for the Data Protection node. To get these parameters for the DAG node, log on to the IBM Spectrum Protect server or contact your IBM Spectrum Protect server administrator.

Active Policy Set

Specifies the policy set that is active for the policy domain. A policy set contains one or more management class definitions.

Default Management Class

The default policy or management class that contains attributes. These attributes determine how long backup versions are stored, where backup versions are stored, and how many backup versions are retained.

Server Password

Use the Server Password property page to change the password for the Data Protection node that you use to access the IBM Spectrum Protect server. This property page applies only to IBM Spectrum Protect configurations.

The following fields are displayed in the property page:

Old password

Type the IBM Spectrum Protect password that you want to change.

New password

Type a new password. Follow the IBM Spectrum Protect server password policy rules.

Confirm new password

Type the new password again. Click **OK** to save your changes.

Policy Management

Use the Policy Management property page to add or update a backup policy, which controls how different backup versions are retained on local shadow volumes on stand-alone snapshot configurations.

Backup retention on local shadow volumes is determined by version and time-based policies. Ensure that sufficient local storage space is available on local shadow volumes for your VSS backup. The amount of storage space that is required depends on the VSS Provider that you use.

The following fields are displayed in the property page:

Policy Specify the unique name of a backup policy for the stand-alone configuration.

Number of Snapshots to keep

Specify the number of backup versions to retain on local shadow volumes. Enter a value from 1 to 9999. Type NL to retain as many backup versions as permitted by available storage space. The default value is 2.

This parameter does not apply to incremental backup versions of Exchange Server data. Incremental backups do not participate in expirations because of version limit because there is never more than one version of an incremental backup object. There is only one version of an incremental backup object because incremental backups are always uniquely named.

Days to keep a Snapshot

Specify the number of days to retain backup versions on local shadow volumes. Enter a value from 0 to 9999. Type NL to retain as many backup versions as permitted by available storage space. When the value is set to 0, snapshots are kept for the current day. The default value is 30.

VSS Policy Binding

Use the VSS Policy Binding property page to bind storage snapshots to back up policies or management classes. VSS policies determine how backups are managed and retained.

VSS policy statements are processed from the end to the beginning and processing stops when the first matching statement is reached. To ensure that more specific statements are processed, specify the more general specification before the more specific ones.

The policy statements do not take effect on existing or new backups until the next backup is issued.

Managed Capacity

Use the Managed Capacity property page to track the capacity of managed storage.

The information that is provided can assist you with storage capacity planning during activities such as license renewal.

Diagnostics

Use the Diagnostics property page to select the type of tracing to run on various components of Data Protection for Microsoft Exchange Server.

When you encounter a problem, open the Diagnostics property page. Select the diagnostic mode that you want to use by clicking **Normal**, **Complete**, or **Custom**. Then, click **Begin** to start the trace. Close the property page. Re-create the problem, open the Diagnostics property page, and click **End** to stop the tracing and collect the data.

If you are using this property page from the Dashboard property sheet, you can run trace only for Microsoft Management Console (MMC).

Diagnostic modes

The following diagnostic mode is available in the Diagnostics property page from the Dashboard property sheet:

MMC - use this mode to set tracing for only MMC.

The following diagnostic modes are available in the Diagnostics property page in the workload property sheets. The type of tracing that is enabled for each mode is listed in the table. Specific trace flags, and guidance on when to use each mode is also listed.

Table 11. Diagnostics modes and their usage

Mode	Components traced along with trace flags used	When to use
Normal	MMC, DP (service), API (service,api_detail)	If using legacy operations, you can use this mode as it results in small output size
Complete	MMC, DP (service), API (service,api_detail), Agent (service)	Use for VSS operations, results in large output size
Custom	Any combination	Use if specific flags are needed

Normal

Click **Normal** to collect trace and log files for legacy operations. Not applicable for Data Protection for Microsoft Exchange Server.

Complete

Click **Complete** to collect trace and log files for VSS operations.

Custom

Click **Custom**, then click the check mark icon to select the trace and log files that you want to collect. Use this mode only if specific trace flags are required.

Enable snapin tracing

Select this box to enable tracing of MMC. Click **Review** to view the trace file.

Set Default Trace Flags

Click **Set Default Trace Flags** to set the most commonly requested trace flags.

Enable Data Protection tracing

Select this box to enable tracing of Data Protection for Microsoft Exchange Server operations. Click **Review** to view the trace file. Add or update trace flags in the field.

Enable DSM Agent tracing

Select this box to enable tracing for the IBM Spectrum Protect client node. You must restart the client acceptor service before you start the trace. Click **Review** to view the trace file. Add or update trace flags in the field.

Enable API tracing

Select this box to enable tracing for the IBM Spectrum Protect API. Click **Review** to view the trace file. Add or update trace flags in the field.

Email Select diagnostic files and click **Email** to send a diagnostic email to an IBM service representative with the selected files attached. You must configure your email information before you can send the data to an IBM service representative. To configure your email information, go to the Dashboard and click **Properties**. Then, click **Email** to open the email property page.

Screenshot

This function is enabled after you click **Begin**. Click **Screenshot** to open the Diagnostic Screenshot Tool. This tool is a modeless dialog that remains open until you close it or click **End** or **Cancel**.

Click **Add New Screenshot** to add a screen capture to the FlashCopyManager\ProblemDetermination folder. The screen capture can be selected with other diagnostic data.

Tracing details for each component

All trace files are stored in the flashcopymanager folder, which is C:\Program Files\Tivoli\flashcopymanager by default. When you click **End**, these files are automatically copied, compressed, and stored in the C:\Program Files\Tivoli\flashcopymanager\problemdetermination folder along with other information.

MMC Options are stored in MMC user settings file. The following files are created as a result of the diagnostic functions:

TraceFm.trc
TraceUx.trc

Data Protection

Tracing options are stored in MMC user settings file and passed to the Data Protection component as part of the command. The following file is generated:

TraceFileExc.trc

Agent Tracing options are stored in the VSS requestor dsm.opt file. The following file is generated:

TraceFileAgent.trc

API Tracing options are stored in the respective Data Protection dsm.opt file. The following file is generated:

TraceFileExcAPI.trc

General (Exchange)

Use the General (Exchange) property page to specify general preferences for the **Exchange Server** workload. This property page applies only if your workload is configured to back up data to IBM Spectrum Protect.

Temporary log restore path

Specify the default temporary path to use when you restore logs and patch files. For best performance, specify a path that is on a different physical device than the current active logger. If you do not enter a path, the default is the value of the TEMP environment variable. When you run a full restore, copy restore, or database copy restore, all log files that are in the specified path are erased.

Back up DAG databases to common node

Specify the node name that you want to use to back up databases from a Database Availability Group (DAG) on Exchange Server 2010. With this setting, all active and passive copies of the databases are backed up to the same file space on the IBM Spectrum Protect server. The database copies are managed as a single entity, regardless of which Database Availability Group member they were backed up from.

When you use this setting, IBM Spectrum Protect applies the same policy across all DAG members, regardless of which DAG member ran the backup.

Temporary database restore path

Specify the directory where the database files that are being restored are temporarily located. Ensure that the directory provides enough space to store the entire mailbox database file. If you do not specify a directory, the database files are restored into a directory that is specified by the TEMP environment variable. This option is only available for mailbox restore operations.

Alias of temporary mailbox

Specifies the alias of a mailbox to use as a temporary storage location during mailbox restore operations. The temporary mailbox is used during restore operations of mailboxes that were deleted, re-created, or moved since the time of the backup. By default, the mailbox restore operation uses the administrator user's mailbox as a temporary storage location.

Exchange client access server

Specify the name of the Client Access Server (CAS) that you want to use. This field is available only for Microsoft Exchange Server 2010.

By default, IBM Spectrum Protect Snapshot uses the local server as the CAS if the local server has the CAS role installed. The CAS that is defined by the logon user mailbox database is used if the local server does not have the CAS role installed.

You can find the name of the current CAS, which is defined by the current logon user mailbox database, by running this Exchange Management Shell command:

```
Get-MailboxDatabase -Identity <logon user mailbox database> |  
select RpcClientAccessServer
```

To use a different CAS, you can define the CAS to be used here.

Restore mail messages as unread

Select this check box to specify that restored mail messages are marked as unread.

Backup mailbox history

Select this check box if you are using mailbox restore operations and you want the mailbox history to be backed up.

Tip: If you do not intend to run mailbox restore operations, clear this check box. This action can improve backup performance.

Logging

Use the Logging property page to specify activity log preferences.

Log File Name

Specifies the name of the file in which activities are logged.

Enable pruning

Specifies that older entries from the log are to automatically be deleted. By default, log pruning is activated and performed daily.

Number of days to keep old entries

Specifies the number of days to keep old entries in the log before they are pruned. By default, 60 days of log entries are saved in the pruning process.

Prune now

Click this option to delete older entries from the Data Protection for Exchange Server activity log when a command runs.

Regional

Use the Regional property page to set preferences that affect how languages and information are displayed and logged.

Regional and Language options

Select this option to set preferences for Microsoft Management Console (MMC). MMC uses the same regional settings as the Windows system.

Language

Select the language to use for log files and the command-line interface.

Date Format

Select a date format to use for log files and the command-line interface. The available choices represent several ways to place the month (*mm*), day (*dd*), year (*yyyy*), and period of day (*a.m.* or *p.m.*). The default date format is *mm/dd/yyyy*.

Time Format

Select a time format to use for log files and the command-line interface. The available choices represent several ways to place the hour (*hh*), minutes (*mm*), and seconds (*ss*). The default time format is *hh:mm:ss*.

Number Format

Select a number format to use for log files and the command-line interface. The available choices represent several ways to place the decimal, comma, and spaces. The default number format is *xxx,xxx.dd*.

Match MMC Language

Select this option to change MMC regional settings to match the system's regional and language options. By selecting this option, the number, date, and time formats are matched to the default formats of the selected language.

VSS Options

Use the VSS Options property page to configure preferences that are used during VSS backup and restore operations.

Default Backup Destination

Select the default storage location for your backups.

Tip: You must have the IBM Spectrum Protect Snapshot license to use the IBM Spectrum Protect software. If you have only the Data Protection license, only the IBM Spectrum Protect option is enabled.

You can select from the following storage locations:

TSM The backup is only stored on IBM Spectrum Protect server storage. For Exchange Server, IBM Spectrum Protect server is the default backup destination.

Local The backup is only stored on local disk.

Both The backup is stored on both IBM Spectrum Protect server storage and local disk.

For IBM Spectrum Protect configurations, the backups can be stored on a local disk, but managed on the IBM Spectrum Protect server. The IBM Spectrum Protect server maintains the metadata or information about where the local snapshot is stored.

Local DSMAGENT Node name

Specify the node name for the DSM Agent node of the local client system that creates the VSS backups.

Remote DSMAGENT Node name

Specify the node name of the system that moves the VSS data to IBM Spectrum Protect server storage during offloaded backups. If you do not use offloaded backups, you can leave this field blank.

Import VSS snapshots only when needed

Select the check box to have Data Protection for Exchange Server import VSS snapshots to the Windows system where the snapshots are created. The check box is selected by default. During backup processing, transportable snapshots are automatically created and imported to storage systems when the snapshots are required.

Ensure that the check box is selected if you want to take the following actions:

- Import the VSS snapshots to a local server
- Keep more than 100 backup versions
- Extend the number of LUNs that the server can use, for example, in a VMWare environment

Tip: If you work in a VMware environment and want to use VMware vMotion, ensure that the LUNs are correctly zoned to the ESX hosts. The import process maps the VSS snapshot to the ESX host where the Windows virtual machine is running.

Clear the check box if you do not want to create transportable VSS snapshots during backup processing and automatically import the snapshot to storage systems after the backup is completed.

During Instant Restore, automatically stop and restart necessary Microsoft Exchange services

When this option is selected, during instant restore operations, the following Microsoft Exchange services are, as necessary, automatically stopped and restarted:

- (DAG environments only) Exchange Replication Service
- (Exchange 2013 only) Exchange Search Host Controller Service

Mount read only

Select the check box to specify that backups are to be mounted as read-only VSS snapshots by default. However, at mount time, you can override this value and do a read/write mount. If you change the default, the corresponding update is made in your configuration file automatically.

Mount read/write (modifies backup, applies to COPY backups only)

Select the check box to specify that backups are to be mounted as read/write VSS snapshots by default. You can mount only COPY backups as read/write. After mounting, the original backup is marked as modified and while you can mount it again in the future, it can no longer be used as a restore point in future database restore operations. However, at mount time, you can override this value and do a read-only mount.

Mount read/write (without modifying backup)

Select this check box to specify that backups are to be mounted as read/write copies of the backup by default. With this option, you can mount both FULL and COPY backup types as read/write. After mounting, the original backup is not modified and you can use it again in future database restore operations. However, at mount time you can override this value and do a read-only mount.

Important:

This mount option is only available for SAN Volume Controller (SVC) devices and requires IBM System Storage Support for Microsoft Virtual Disk and Volume Shadow Copy Services version 4.12 or later. Also, you must allocate extra target volumes on your SVC storage device to accommodate the number of concurrent read/write mounts you want to do. An extra target volume that matches the size of the volume to be mounted is needed for each concurrent read/write mount of that volume.

Custom Settings

Use the Custom Settings property page to set your filtering options and control the amount of information that is returned from the server.

Select **Show Refresh Options** in the toolbar in the Recover view.

In environments where thousands or millions of backup objects are stored on the IBM Spectrum Protect server, it can be helpful to disable the automatic refresh mode. You can click **Refresh Options** and use the toolbar to switch between manual and automatic refresh mode.

Automatic and manual refresh modes differ in the following ways:

- In automatic refresh mode, a view automatically refreshes the first time that you select it. If there are thousands or millions of objects on the server, the refresh can take a long time to complete.
- In manual refresh mode, the view is not automatically refreshed. A name filter is available on the **Refresh Options** toolbar that you can use to narrow down the

number of objects selected. After you enter a name pattern, you can click **Refresh**. By using manual refresh mode and limiting your query by using filters, you can reduce the amount of information that is returned from the server. Reducing the amount of information that is returned from the server can improve query and restore performance.

You can also specify a wildcard character (*) in the name pattern to assist your filtering effort.

MAPI Settings

If you use Exchange Server 2013, use the MAPI Settings property page to verify that the user mailbox is online. You can also view and update the MAPI registry key that enables Data Protection for Exchange Server to connect to the Exchange Server.

Data Protection for Exchange Server automatically generates a default value for the registry key. Edit the registry key only if the default value is incorrect.

The values that you enter override the registry key that is in the HKEY_CURRENT_USER\Software\Microsoft\Windows NT\Current Version\Windows Messaging Subsystem directory. If you modify the registry incorrectly, the connection to the Exchange Server might fail.

You can use this property page only if you use Exchange Server 2013 or later versions.

RpcHttpProxyMap_TSM

Change the registry key values to reflect the correct domain, endpoint, and Remote Procedure Call (RPC) authentication methods for your environment. By default, the format is:

```
Domain=Proxy Server,RpcHttpAuthenticationMethod,  
RpcAuthenticationMethod,IgnoreSslCert
```

For example:

```
companyname.local=https://exchange.companyname.com,ntlm,ntlm,false
```

where:

- *Domain* value is the domain suffix of the personalized server ID, for example, *companyname.local*. Specify any domain or a substring of a domain, or the asterisk (*) and question mark (?) wildcard characters, for example, **.companyname.local*.
- *Proxy Server* value is the RPC proxy server that has the Client Access Server (CAS) role. Specify the fully qualified domain name (FQDN) of the RPC proxy server. Precede the FQDN by *http://* for an HTTP connection, or *https://* for an HTTPS connection. For example, *https://exchange.companyname.com*
- *RpcHttpAuthenticationMethod* value is the method that is used to authenticate RPC-over-HTTP connections. Specify *NTLM*, *Basic*, *Negotiate*, or *WinNT*.
- *RpcAuthenticationMethod* value is the method that is used to authenticate RPC-over-TCP connections. Specify *NTLM*, *Negotiate*, *WinNT*, *Anonymous*, or *None*.
- *IgnoreSslCert* value indicates whether the Exchange Server validates SSL certificates. For the Exchange Server to ignore invalid certificates, specify *False*.

Domain

Change the domain name to reflect the correct domain if for example, you have multiple domains, or the default domain value is incorrect. To match all domains, enter the asterisk (*) wildcard character. When you change this domain value, the *Domain* value of the registry key automatically updates in the *RpcHttpProxyMap_TSM* field.

Use HTTPS authentication

Select this check box if RPC-over-HTTPS is enabled for the Exchange Server that is hosting the MAPI profile. Otherwise, clear this check box to ensure that HTTP authentication is used for RPC-over-HTTP connections. When you change this authentication value, the *RpcAuthenticationMethod* value of the registry key automatically updates in the *RpcHttpProxyMap_TSM* field.

Related tasks:

“Ensuring successful MAPI connections” on page 98

Configuring Data Protection for Microsoft Exchange Server by using the IBM Spectrum Protect Configuration Wizard

Configuration requirements for Data Protection for Exchange Server, IBM Spectrum Protect, and other applications vary. The requirements depend on which Data Protection for Exchange Server features you want to use. For example, if you plan to use VSS operations, the IBM Spectrum Protect backup-archive client, serving as the VSS Requestor, must also be installed and configured.

Before you begin

When you are remotely configuring Data Protection for Exchange Server, you must first install IBM Spectrum Protect Snapshot for Windows on the Data Protection node server, as shown in the example that is used in the procedure. You must then run the IBM Spectrum Protect Configuration Wizard, on at least one occasion, on the Data Protection node server.

Procedure

1. Start Microsoft Management Console (MMC) by clicking **Start > All Programs > IBM Spectrum Protect > Data Protection for Microsoft Exchange Server > DP for Exchange Management Console**.
2. From the start page, click **Configuration**. Alternatively, from the navigation tree, go to the **Configuration** node. Then, double-click **Wizards**.
3. In the results pane, double-click **TSM Configuration** to open the IBM Spectrum Protect Configuration Wizard.
4. Follow the instructions on the pages of the wizard and click **Next** to move to the next page.
 - a. In the Data Protection Selection page, select **Exchange Server**. Click **Next**.
 - b. Review the results of the requirements check and ensure that you address any errors or warnings.

Click **Show Details** to view a list of individual requirement results.

- If you do not have a license for the application that you are configuring, the license requirement check fails. You must either go back to the Data Protection Selection page and clear the selected application to proceed with the configuration, or obtain the necessary license

- If you do not have all the user roles that are required for individual mailbox restore operations, click the **Warnings** link and follow the wizard prompts to add the missing Exchange Server roles. If you are a member of the Exchange Organization Management group and have sufficient role-based access control (RBAC) permissions, you can automatically add the missing roles. If you are not a member of the Exchange Organization Management group and have insufficient RBAC permissions, you must manually add the missing roles.
- c. In the TSM Node Names page, specify the IBM Spectrum Protect node names, which exist on the same system, to use for the applications that you want to protect.
- In the **VSS Requestor** field, enter the node name that communicates with the VSS service to access the Exchange Server data.
 - In the **Data Protection for Exchange Server** field, enter the node name where the Data Protection application is installed. This node name is used to store the Data Protection for Exchange Server backups. If you configure the **DAG Node** on this wizard page, Exchange Server DAG database backups are not stored under the Data Protection for Exchange Server node. The backups are stored under the DAG node.
 - In the **DAG Node** field, enter the node name that you want to use to back up databases in an Exchange Server DAG. With this setting, backups from all DAG members that are configured to use the DAG node are backed up to a common file space on the IBM Spectrum Protect server. The database copies are managed as a single entity, regardless of which DAG member they are backed up from. This setting can prevent Data Protection for Exchange Server from making too many backups of the same database.

Ensure that you configure all of your DAG members that have copies of the same database so that all use the same DAG node. On the IBM Spectrum Protect server, ensure that you register the DAG node name. All of the DAG member nodes (the Data Protection nodes) must be granted *proxynode* authority to run backups on behalf of the DAG node. All of the DSM Agent nodes (the backup-archive client nodes) must also be granted *proxynode* authority. If you do not want to manually update these properties, you can use the configuration wizard to set the properties on the IBM Spectrum Protect server.

Create a node name that can help you to distinguish the type of backup that runs. For example, if your host name is *MALTA*, you can name the VSS Requestor node name *MALTA*, and you can create a Data Protection node name that is called *MALTA_EXC*. For an Exchange configuration, the DAG node name does not need to relate to the VSS Requestor or the Data Protection for Exchange Server node name. For example, you can name it *TSM DAG*.

- d. Enter information for the IBM Spectrum Protect server that you are connecting to and click **Next** to continue.
- In the **IBM Spectrum Protect Server Address** field, enter the TCP/IP domain name or a numeric IP address for the IBM Spectrum Protect server that contains the backups. Obtain this information from your IBM Spectrum Protect server administrator.
 - In the **IBM Spectrum Protect Server Port** field, enter the port number for the IBM Spectrum Protect server that contains the backups. Obtain this information from your IBM Spectrum Protect administrator.

- Specify whether to have the wizard to configure the IBM Spectrum Protect server for you by generating a configuration macro file.

If you click **No**, the macro file is available at the final page of the wizard and can be provided to the IBM Spectrum Protect administrator as an example of one way to configure the IBM Spectrum Protect server to support application data protection.

If you click **Yes**, the wizard starts the macro during the Configuration step in the wizard. Review the macro file and update it if needed.

After you click **Yes**, enter the following information in the appropriate fields and perform the following actions:

- The name of the IBM Spectrum Protect administrator account.
 - The password for the IBM Spectrum Protect administrator.
 - Click **Test Communications** if you want to test your connection with the IBM Spectrum Protect server. This option is not available until the VSS Requestor is installed.
 - Click **Review/Edit** to review or update the IBM Spectrum Protect macro file. Alternatively, you can review the macro file and directly run the commands on the IBM Spectrum Protect server.
- e. Select the **Default** configuration setting. When you select the **Default** configuration setting, the VSS Requestor is configured in addition to the applications that you selected. The client and agent services are also registered and configured, and a schedule to support historical managed capacity is defined.
 - f. Review the results of the configuration process. Click **Show Details** to view a list of individual configuration results.

5. Click **Finish** in the Completion page.

6. Optional: For a VSS configuration, verify that the **Run VSS diagnostics when this wizard exits** option is selected. When this option is selected, a diagnostic process tests the VSS snapshots on your system after you complete the wizard.

Attention: If the configuration is for space-efficient target volumes for SAN Volume Controller or Storwize V7000, testing VSS snapshots deletes previous backups that are created for the volumes that are selected in the test wizard.

Related concepts:

“Security requirements for backup and restore operations” on page 87

Verifying the configuration

After you configure Data Protection for Exchange Server, verify that the configuration wizard automatically installs the IBM Spectrum Protect backup-archive client.

Procedure

1. In MMC, click the **Automate** tab to access the integrated command-line interface.
2. Click the **Open folder** icon, and select the `verify_exc.txt` file.
3. Click **Open**. These commands are displayed in the command-line pane:


```
query tdp
query tsm
query exchange
```
4. With the cursor in the command-line panel, press Enter to run the commands and verify your configuration. The following examples show the command output that each command generates.

- Command: **query tdp**

```
C:\Program Files\Tivoli\tsm\TDPEExchange>tdpexcc query tdp

IBM Spectrum Protect for Mail:
Data Protection for Microsoft Exchange Server
Version 8, Release 1, Level 0
(C) Copyright IBM Corporation 1998, 2016. All rights reserved.

Data Protection for Exchange Preferences
-----

BACKUPDESTINATION..... LOCAL
CLIENTAccessserver.....
DAGNODE..... FCMDAG2
DATEformat ..... 1
LANGuage ..... ENU
LOCALDSMAgentnode..... PEGUVM1
LOGFile ..... tdpexc.log
LOGPrune ..... 60
MOUNTWait ..... Yes
NUMBERformat ..... 1
REMOVEDSMAGENTnode..... PEGUVM2
STOPservicesonir ..... Yes
STOREMAILBOXInfo ..... Yes
TEMPDBRestorepath..... C:\temp\DB
TEMPLOGRestorepath..... C:\temp\LOG
TIMEformat ..... 1
IMPORTVSSSNAPSHOTSONLYWhenneeded.... Yes

The operation completed successfully. (rc = 0)
```

- Command: **query tsm**

```
C:\Program Files\Tivoli\tsm\TDPEExchange>tdpexcc query tsm

IBM Tivoli Storage Manager for Mail:
Data Protection for Microsoft Exchange Server
Version 7, Release 1, Level 2.0
(C) Copyright IBM Corporation 1998, 2015. All rights reserved.

Tivoli Storage Manager Server Connection Information
-----

Nodename ..... PEGUVM1_EXC
Network Host Name of Server ..... fvtseries11esx6
TSM API Version ..... Version 7, Release 1, Level 1.42

Server Name ..... FVTSERIES11ESX6_SERVER1
Server Type ..... Windows
Server Version ..... Version 7, Release 1, Level 0.0
Compression Mode ..... Client Determined
Domain Name ..... FCM_PDEXC
Active Policy Set ..... STANDARD
Default Management Class ..... STANDARD

The operation completed successfully. (rc = 0)
```

- Command: **query exchange**

```

C:\Program Files\Tivoli\tsm\TDPEXchange>tdpexcc query exchange

IBM Spectrum Protect for Mail:
Data Protection for Microsoft Exchange Server
Version 8, Release 1, Level 0
(C) Copyright IBM Corporation 1998, 2016. All rights reserved.

Querying Exchange Server to gather database information, please wait...

Microsoft Exchange Server Information
-----

Server Name:                PEGUVM1
Domain Name:                cvtdomain1.local
Exchange Server Version:    14.3.181.6   (Exchange Server 2010)

Databases and Status
-----

Mailbox Database 1474758353
Circular Logging - Disabled
DAG Status - None
Recovery - False
Mailbox Database 1474758353                Online

PEGUVM1_DB5G_local
Circular Logging - Disabled
DAG Status - None
Recovery - False
PEGUVM1_DB5G_local                Online

PEGUVM1_DB7L_sevsvc
Circular Logging - Disabled
DAG Status - Passive (Healthy)
Recovery - False
PEGUVM1_DB7L_sevsvc                Online

PEGUVM1_DB8K_stdsvc
Circular Logging - Disabled
DAG Status - Active
Recovery - False
PEGUVM1_DB8K_stdsvc                Online

Volume Shadow Copy Service (VSS) Information
-----

Writer Name                  : Microsoft Exchange Writer
Local DSMAgent Node         : PEGUVM1
Remote DSMAgent Node        : PEGUVM2
Writer Status                : Online
Selectable Components        : 4

The operation completed successfully. (rc = 0)

```

Configuring a Data Protection for Microsoft Exchange Server remote system to integrate with IBM Spectrum Protect

By using the TSM Configuration Wizard, you can configure a remote system to communicate with an IBM Spectrum Protect server.

Before you begin

On the local system, verify the following system requirements:

- Windows 7, Windows 8, Windows 2008, Windows 2008 R2, Windows 2012, Windows 2012 R2, or a later version is installed

- PowerShell version 3.0 or later is installed, if you are running Windows 7, Windows 8, Windows 2008, or Windows 2008 R2. On Windows 2012 and later versions, PowerShell version 4.0 is installed by default.

On the remote system, verify the following system requirements:

- Windows 2008, Windows 2008 R2, Windows 2012, Windows 2012 R2, or a later version is installed
- Windows PowerShell version 3.0 or later is installed, if you are running Windows 2008, or Windows 2008 R2. On Windows 2012 and later versions, PowerShell version 4.0 is installed by default.
- The required workload is configured.

Procedure

1. On the local system, start Microsoft Management Console (MMC).
2. From MMC, use Manage Computers to add the remote system.
3. In the navigation tree, verify that the remote system is displayed.
4. Click **Manage > Configuration > Wizards**.
5. Select **TSM Configuration**.
6. On the Data Protection Selection page, verify that the following information is entered correctly:
 - The remote computer name in the window title.
 - The correct system information.
7. Select the application to be configured and click **Next**.
8. For Exchange or SQL Server, the license check might fail. If the test fails, provide the file path and name for the location on the remote server.
9. On the TSM Node Names page, verify that the following information is entered correctly:
 - VSS Requestor
 - The Data Protection or file system name, depending on the application that is configured

For systems with a Database Availability Group (DAG) or an AlwaysOn Availability Group, the corresponding DAG node or AlwaysOn node is detected.

10. On the TSM Server Settings page, type the server name and port number.
11. For the **Would you like this wizard to configure your TSM server?** question, select **Yes**.
12. Click **Review / Edit**. If the domain is not entered correctly, update the information. Click **OK**.
13. On the Custom Configuration page, select **Default**.
14. On the Configuration page, click **Show Details**. Verify the progress and status of the configuration.
15. Click **Finish** to complete the wizard.

What to do next

To verify that the configuration is set up correctly, complete the following steps:

1. In the navigation tree, for the remote system, expand **Protect and Recover** and click the application that is configured.

2. Open the Properties and click **Server Information**. Verify that the correct information is displayed.
3. Query the components and verify that a successful backup can be completed.

Configuring Data Protection for Microsoft Exchange Server by using the Mailbox Restore Only Configuration Wizard

Configuration requirements for Data Protection for Exchange Server, IBM Spectrum Protect, and other applications vary. The requirements depend on which Data Protection for Exchange Server features you want to use. For example, if you plan on using VSS operations, the IBM Spectrum Protect backup-archive client (serving as the VSS Requestor), must also be installed and configured.

Procedure

1. Start Microsoft Management Console (MMC) by clicking **Start > All Programs > IBM Spectrum Protect > Data Protection for Microsoft Exchange Server > DP for Exchange Management Console**.
2. From the start page, click **Configuration**. Alternatively, from the tree view, go to the **Configuration** node. Then, double-click **Wizards**.
3. In the results pane, double-click **Mailbox Restore Only** to open the Mailbox Restore Only Configuration Wizard.
4. Follow the instructions on the pages of the wizard. Review the results of the requirements check. Correct any errors or warnings in the requirements check.
 - a. Click **Show Details** to view a list of individual requirement results. If you do not have all the user roles that are required for individual mailbox restore operations, click the **Warnings** link and follow the wizard prompts to add the missing Exchange Server roles. If you are a member of the Exchange Organization Management group and have sufficient role-based access control (RBAC) permissions, you can automatically add the missing roles. If you are not a member of the Exchange Organization Management group and have insufficient RBAC permissions, you must manually add the missing roles.
 - b. Click **Next** to move to the next page of the wizard.
5. Click **Finish** in the Completion page.

Related concepts:

“Security requirements for backup and restore operations” on page 87

Manually configuring Data Protection for Exchange Server for IBM Spectrum Protect Configuration

For best results, use the configuration wizards to configure Data Protection for Exchange Server. The wizards provide you with a step-by-step guide of the configuration requirements. However, if you prefer to do these steps manually, follow these configuration instructions.

Configuring the computer that runs the Exchange Server

Perform these steps on the computer where the Exchange Server is installed and running.

Before you begin

Before you begin, ensure that the Exchange Server is running.

If you configure the DSM Agent node (the backup-archive client node) manually, ensure that you set the **PASSWORDAccess** option to generate in the `dsm.opt` file for the IBM Spectrum Protect backup-archive client. Also ensure that the stored password for the DSMAGENT Node is valid.

Procedure

1. Specify your Data Protection for Exchange Server node name and communication method in the `dsm.opt` file that is located by default in the Data Protection for Exchange Server installation directory.
2. Using the **set** command, specify your Data Protection for Exchange Server preferences (date format, log file) in the `tdpexc.cfg` file in the Data Protection for Exchange Server installation directory.
3. (VSS only) If you are configuring Data Protection for Exchange Server in an Exchange Server Database Availability Group (DAG) environment, issue the **set** command to create a common node name for backing up DAG servers. For example:

```
tdpexcc set DAGNODE=TSMDAG1
```

Where TSMDAG1 is the DAG node name that is used to back up all databases in a DAG, regardless of which DAG member the database is backed up from.

Important: On the IBM Spectrum Protect server, ensure that you register the DAG node. All DAG members need proxy authority to run backups on behalf of the DAG node.

4. If you schedule more than one DAG member to back up a database, and you must prevent DAG databases from being backed up too frequently, set the minimum amount of time, in minutes, that passes before a database can be backed up again by using the **/MINimumbackupinterval**. This parameter must be specified as part of a **backup** command script that is run by the IBM Spectrum Protect scheduler.

For example, include the following statement in the `C:\BACKUP.CMD` script:

```
tdpexcc backup DB1 full /minimumbackupinterval=60
```

5. Optional: To reduce the load on a production Exchange Server in a DAG, you can specify the **/PREFERDAGPASSive** parameter. If a healthy passive database copy is not available, this parameter backs up a passive database copy. The backup is created from the active database copy. This parameter must be specified as part of a **backup** command script that is run by the IBM Spectrum Protect scheduler.

For example, include the following statement in a `C:\BACKUP.CMD` script:

```
tdpexcc backup DB1 full /minimumbackupinterval=60 /preferdagpassive
```

6. (VSS only) Specify your **VSSPOLICY** statement in your Data Protection for Exchange Server configuration file. Exchange servers that use the same DAG node name must share the VSS Policy.
7. (VSS only) Configure the IBM Spectrum Protect backup-archive client if it is not already configured. If the backup-archive client is already configured, you

can use existing client services. The backup-archive client Setup Wizard can guide you through the configuration process. In the backup-archive client GUI menu, select **Utilities > Setup Wizard > Help me configure the TSM Backup Archive Client**. The node name for this system is referred to as the **Local DSMAGENT Node** and is specified with the **localdsmagentnode** parameter in the Data Protection for Exchange Server configuration file (tdpexc.cfg).

For more information about installing the IBM Spectrum Protect backup-archive client for Windows, see *Install the UNIX and Linux backup-archive clients* (http://www.ibm.com/support/knowledgecenter/SSEQVQ_8.1.0/client/c_inst_baunix.html).

8. (VSS only) Install and configure the IBM Spectrum Protect Client Acceptor Service (CAD) if it is not already installed and configured. In the backup-archive client GUI menu, select **Utilities > Setup Wizard > Help me configure the TSM Web Client**. Make sure that the CAD service is running before you proceed to the next step.
9. (VSS only) Install and configure the IBM Spectrum Protect Remote Client Agent Service (DSMAGENT) if it is not already installed and configured. In the backup-archive client GUI menu, select **Utilities > Setup Wizard > Help me configure the TSM Web Client**. If a DSMAGENT is already installed and configured, you can use the existing one.
10. (VSS only) Install IBM Spectrum Protect Snapshot if you want to manage local persistent VSS snapshots, which are created for VSS backups to LOCAL, VSS Instant Restores, and you want to run offloaded backups.
11. (VSS only) Add the Microsoft Exchange Server binary path to the PATH statement in the system environment variables. For example:
`"C:\Program Files\Microsoft\Exchange Server\V15\Bin"`
Verify that **ESEUTIL.EXE** tool exists in this directory. This tool is used by Data Protection for Exchange Server to run automatic integrity checks on the VSS backup.
12. (VSS only) Install and configure a VSS provider. Consult the VSS provider documentation for information about configuration of that software. You do not need to install and configure any components if you use the default Windows VSS System Provider.
13. (VSS only) Define storage space to hold VSS backups that is on local shadow volumes. Define enough space to store all copies of the VSS backups as designated by your policies. Provisioning storage space to manage VSS snapshots depends on the VSS provider that you use. Consult the VSS Provider documentation for more details.

Related concepts:

"Proxy node definitions for VSS backups" on page 47

Configuring the IBM Spectrum Protect server

Complete these steps on the IBM Spectrum Protect server.

Before you begin

Ensure that the IBM Spectrum Protect server is available.

Procedure

1. Define the policy domains, policy sets, management classes, copy groups, and storage pools. Choose what is necessary to meet your Data Protection for Exchange Server backup and restore requirements. For VSS operations, IBM Spectrum Protect server authentication must be on.

2. Register your Data Protection for Exchange Server node name and password by issuing the IBM Spectrum Protect **register node** command. For example, for VSS operations, this node is the target node. When you register nodes to the IBM Spectrum Protect server specifically for VSS operations, specify the IBM Spectrum Protect **USerid**=<node name> parameter.
3. If not already defined, register your IBM Spectrum Protect backup-archive client node name and password for the system where the Exchange Server is installed. For example, this agent node is the Local DSMAGENT Node for VSS operations.
4. (VSS only) If you plan to run offloaded backups from a particular system, first register the IBM Spectrum Protect backup-archive client node name and password for the system. For example, the agent node is the Remote DSMAGENT Node. *BAOFF* is used here (and in Step 5) to differentiate between this Remote DSMAGENT Node and the Local DSMAGENT Node (Step 3). You can replace *BAOFF* with the node name of your backup-archive client, and remove the *BAOFF* from the **grant proxynode** command.
5. (VSS only) Define the proxy node relationship (for the target node and agent nodes) by issuing the IBM Spectrum Protect **grant proxynode** command. For example:

```
grant proxynode target=DAG node name agent=BAnodename
```
6. If you created a node name for backing up databases in an Exchange Server Database Availability Group (DAG), ensure that the following tasks are complete.
 - a. Register the IBM Spectrum Protect backup-archive client and DAG node names and passwords with the IBM Spectrum Protect **register node** command.
 - b. Ensure that the IBM Spectrum Protect administrator issues the **grant proxynode** command for each member server in the DAG to grant permission to the DAG member server to act as a proxy for the DAG node. If the configuration wizard is not used to configure the IBM Spectrum Protect server, the proxies must be defined. In addition, the backup archive client node and the Data Protection node require proxynode authority. The backup archive client node also requires proxynode authority to act on behalf of the Data Protection node. For example, the IBM Spectrum Protect administrator can issue the following commands on the IBM Spectrum Protect server:

```
register node backup_archive_client_node password
  userID=backup_archive_client_node
register node data_protection_node password userID=data_protection_node
grant proxynode target=data_protection_node agent=backup_archive_client_node
register node DAG_node password userID=DAG_node
grant proxynode target=DAG_node agent=backup_archive_client_node
grant proxynode target=DAG_node agent=data_protection_node
```

What to do next

If any warning messages are displayed during the configuration process, resolve the issue noted in the warning. Some warnings include a link to a macro that you can use to configure IBM Spectrum Protect. Other warnings have links to web sites where you can download the packages that you require to successfully complete the configuration process.

Configuring the system that runs offloaded backups

Perform the following steps on the computer that is running the offloaded backups: This task is for VSS operations only.

Procedure

1. Configure the IBM Spectrum Protect backup-archive client if it is not already configured. If the backup-archive client is already configured, you can use existing client services. In the backup-archive client GUI menu, select **Utilities > Setup Wizard > Help me configure the TSM Backup Archive Client**. The node name for this system is called the Remote DSMAGENT Node and is specified with the **remotedsmagentnode** parameter in the Data Protection for Exchange Server configuration file (tdpexc.cfg) on the local, not offload, system.
2. Install and configure the IBM Spectrum Protect Client Acceptor (CAD) Service and the Remote Client Agent Service (DSMAGENT) if these services are not already installed. If a client CAD Service is already installed and configured, you can use an existing one. Use the backup-archive client Setup Wizard to guide you through the CAD installation process by selecting **Utilities > Setup Wizard > Help me configure the TSM Web Client**.
3. Install the Microsoft Exchange Server management tools from the Microsoft Exchange Server installation media. Take note of the Microsoft Exchange Server Management tools binary directory (for example: C:\Program Files\Microsoft\Exchange Server\V15\Bin). Verify that the ESEUTIL.EXE tool is stored in this directory. Data Protection for Exchange Server uses this tool to run automatic integrity checking of the VSS backup. Also, the Exchange Server does not need to be installed or running on this system. Only the Microsoft Exchange Server management tools must be installed on this system. For more information about the necessary license requirements, see the Microsoft Exchange Server documentation.
4. Add the Microsoft Exchange Server binary path to the PATH statement in the system environment variables. For example:
`"C:\Program Files\Microsoft\Exchange Server\V15\Bin"`
5. Install and configure a VSS provider if you do not use the default system VSS provider. Consult the VSS provider documentation for information about the configuration of that software.

Configuring your system for mailbox restore operations

To use the Data Protection for Microsoft Exchange Server mailbox restore feature, you must complete more configuration in the configuration wizard.

About this task

Because of an Exchange Server requirement, the Data Protection for Microsoft Exchange Server configuration wizard checks that user permissions and software versions are correct.

- Ensure that you have the role-based access control (RBAC) permissions to complete individual mailbox restore operations.
- Install the correct version of Microsoft Exchange Server MAPI Client and Collaboration Data Objects on the Exchange server from which you are running the mailbox restore operations.

Download and install the Exchange MAPI and Microsoft Outlook MAPI on different servers. Do not install Microsoft Outlook 2010 or 2013 on the same

server that Data Protection for Microsoft Exchange Server uses for mailbox restore operations. Conflicts might occur in the MAPI configurations.

Procedure

1. If you are using an incorrect Microsoft Exchange Server MAPI Client and Collaboration Data Objects version, click the **Warnings** link and go to the site where you can download the correct version.
2. If you do not have all the management roles for individual mailbox restore operations, click the **Warnings** link and follow the wizard prompts to add the missing Exchange roles. If you are a member of the Exchange Organization Management group, you can automatically add the missing roles. If you are not a member of the Exchange Organization Management group, you must manually add the missing roles.
3. Configure the Client Access Server (CAS) role to run Mailbox Restore operations on Exchange Server 2010 or later versions. For more information about specifying the CAS with the **set** command, see the **Set syntax** command.

Related concepts:

“Security requirements for backup and restore operations” on page 87

Related reference:

“Set syntax” on page 219

Configuring mailbox history handling for improved performance

Mailbox history includes only the mailboxes from databases that are backed up. If you back up mailbox history with a version of Data Protection for Exchange Server earlier than version 7.1, you can manually delete the old mailbox history.

About this task

Data Protection for Exchange Server backs up a new set of mailbox history data. With the new mailbox history data, you can experience better performance when you back up mailbox history. It is also easier to find the mailbox when you restore a mailbox. Additionally, when you retrieve mailbox history, the mailbox names can be displayed in multiple languages.

Deleting the old mailbox history is not required. If you delete the old mailbox history data, you lose the location history information for the deleted and moved mailboxes in the backup copies that earlier versions of Data Protection for Exchange Server.

Procedure

1. Issue the following command to save the mailbox history to a file:

```
tdpexcc q tsm /showmailboxinfo > E:\MyMailboxHistory.txt
```

Keep this file for reference. You can use the backup copy when you need location information for the deleted and moved mailboxes
2. If you must restore a mailbox from the old backup copies, and the mailbox location changes before you delete the mailbox history, use the **/MAILBOXORIGLOCATION** parameter to restore the mailbox. After the old backup copies expire, mailbox history works without you having to specify the **/MAILBOXORIGLOCATION** parameter.

3. Complete the following steps to delete the old mailbox history from the IBM Spectrum Protect server.
 - a. Start the IBM Spectrum Protect command-line administrative interface, `dsmadm.exe`.
 - b. Log on to the IBM Spectrum Protect server.
 - c. Issue the following command to query the filespace name:

Query Filespace *node_name file_space_name*

The format of the filespace name for mailbox history is *DomainName\MAILBOXINFO*. For example, the following command queries the filespace for the mailbox history for the *CXCLAB_EXC* node. The *node_name* is the **DAGNODE** name, or the Exchange Server node name when the **DAGNODE** is not being used.

tsm: FCM>QUERY FILESPACE CXCLAB_EXC *MAILBOXINFO

The following results are displayed:

Node Name	Filespace Name	FSID	Platform	Filespace Type	Is Filespace Unicode?	Capacity	Pct Util
CXCLAB_EXC	cxcserver.-com\MAILBOXINFO	52	TDP MSE-xchg	API:ExcData	No	0 KB	0.0

4. Issue the following command to delete the filespace for the old mailbox history while bearing in mind that all previous backups, including backups of Exchange Server 2010 data, might be deleted if you do not enter the command correctly.

DElete Filespace *node_name file_space_name\MAILBOXINFO*

For example, the following command deletes the filespace for the mailbox history for the *CXCLAB_EXC* node:

tsm: FCM>DELETE FILESPACE CXCLAB_EXC cxcserver.com\MAILBOXINFO

Verifying the configuration of Data Protection for Exchange Server

Common errors might occur when a VSS operation runs. If commands complete without errors or warnings, you have verification that the Data Protection for Exchange Server configuration is correct. You can also verify that your Exchange Server is ready to run VSS operations.

Verifying the server configuration from the integrated command line

The configuration is verified as correct when these commands complete without errors or warnings.

Procedure

1. Click the **Automate** tab to access the integrated command-line interface.
2. On the lower half of the screen, click the Open folder icon, and select the `verify_exc.txt` file.
3. Click **Open**. These commands are displayed in the command-line panel:

```
query tdp
query tsm
query exchange
```

- With the cursor in the command-line panel, press Enter to run the commands and verify your configuration. The following examples show the command output that each command generates.

Command: query tdp

```
C:\Program Files\Tivoli\tsm\TDPEExchange>tdpexcc query tdp

IBM Tivoli Storage Manager for Mail:
Data Protection for Microsoft Exchange Server
Version 7, Release 1, Level 2.0
(C) Copyright IBM Corporation 1998, 2015. All rights reserved.

Data Protection for Exchange Preferences
-----

BACKUPDEStination..... LOCAL
CLIENTAccessserver.....
DAGNODE..... FCMDAG2
DATEformat ..... 1
LANGuage ..... ENU
LOCALDSMAgentnode..... PEGUVM1
LOGFile ..... tdpexc.log
LOGPrune ..... 60
MOUNTWait ..... Yes
NUMBERformat ..... 1
REMOVEDSMAGentnode..... PEGUVM2
STOPservicesonir ..... Yes
STOREMAILBOXInfo ..... Yes
TEMPDBRestorepath..... C:\temp\DB
TEMPLOGRestorepath..... C:\temp\LOG
TIMEformat ..... 1
IMPORTVSSSNAPSHOTSONLYWhenneeded.... Yes

The operation completed successfully. (rc = 0)
```

Command: query tsm

```
C:\Program Files\Tivoli\tsm\TDPEExchange>tdpexcc query tsm

IBM Tivoli Storage Manager for Mail:
Data Protection for Microsoft Exchange Server
Version 7, Release 1, Level 2.0
(C) Copyright IBM Corporation 1998, 2015. All rights reserved.

Tivoli Storage Manager Server Connection Information
-----

Nodename ..... PEGUVM1_EXC
Network Host Name of Server ..... fvtseries11esx6
TSM API Version ..... Version 7, Release 1, Level 1.42

Server Name ..... FVTSERIES11ESX6_SERVER1
Server Type ..... Windows
Server Version ..... Version 7, Release 1, Level 0.0
Compression Mode ..... Client Determined
Domain Name ..... FCM_PDEXC
Active Policy Set ..... STANDARD
Default Management Class ..... STANDARD

The operation completed successfully. (rc = 0)
```

Command: query exchange

```

C:\Program Files\Tivoli\tsm\TDPEExchange>tdpexcc query exchange

IBM Tivoli Storage Manager for Mail:
Data Protection for Microsoft Exchange Server
Version 7, Release 1, Level 2.0
(C) Copyright IBM Corporation 1998, 2015. All rights reserved.

Querying Exchange Server to gather database information, please wait...

Microsoft Exchange Server Information
-----
Server Name:                PEGUVM1
Domain Name:                cvtdomain1.local
Exchange Server Version:    14.3.181.6   (Exchange Server 2010)

Databases and Status
-----
Mailbox Database 1474758353
Circular Logging - Disabled
DAG Status - None
Recovery - False
Mailbox Database 1474758353                Online

PEGUVM1_DB5G_local
Circular Logging - Disabled
DAG Status - None
Recovery - False
PEGUVM1_DB5G_local                Online

PEGUVM1_DB7L_sevsvc
Circular Logging - Disabled
DAG Status - Passive (Healthy)
Recovery - False
PEGUVM1_DB7L_sevsvc                Online

PEGUVM1_DB8K_stdsvc
Circular Logging - Disabled
DAG Status - Active
Recovery - False
PEGUVM1_DB8K_stdsvc                Online

Volume Shadow Copy Service (VSS) Information
-----
Writer Name                  : Microsoft Exchange Writer
Local DSMAGENT Node         : PEGUVM1
Remote DSMAGENT Node        : PEGUVM2
Writer Status                : Online
Selectable Components        : 4

The operation completed successfully. (rc = 0)

```

Verifying the Exchange Server is ready to start VSS operations

Complete the following tests to verify that your Exchange Server is ready to run VSS operations.

Before you begin

- For best results, complete these tests before you install IBM Spectrum Protect.
- Test the core VSS function first. VSS function can be validated with the Windows Server-embedded **DISKSHADOW** command. The **DISKSHADOW** command is available for Windows Server 2008, Windows Server 2008 R2, or later operating systems.

About this task

The following list identifies the DISKSHADOW tests to complete before any IBM Spectrum Protect components are installed.

Procedure

1. Test non-persistent shadow copy creation and deletion as follows:

- Run DISKSHADOW in a command window
- DISKSHADOW>begin backup
- DISKSHADOW>add volume f: (Database volume)
- DISKSHADOW>add volume g: (Log volume)
- DISKSHADOW>create
- DISKSHADOW>end backup
- DISKSHADOW>list shadows all (this process might take a few minutes)
- DISKSHADOW>delete shadows all

Note: Volumes on drive F and drive G represent the Exchange Database and log volumes. Repeat this test four times, and verify that the Windows Event Log contains no errors.

2. Test persistent shadow copy creation and deletion as follows:

- Run DISKSHADOW in a command window
- DISKSHADOW>set context persistent
- DISKSHADOW>begin backup
- DISKSHADOW>add volume f: (Database volume)
- DISKSHADOW>add volume g: (Log volume)
- DISKSHADOW>create
- DISKSHADOW>end backup
- DISKSHADOW>list shadows all (This process might take a few minutes)
- DISKSHADOW>delete shadows all

Note: Volumes on drive F and drive G represent the Exchange Server database and log volumes. Repeat this test four times, verify that the Windows Event Log contains no errors.

3. Test non-persistent transportable shadow copy creation and deletion as follows:

- Run DISKSHADOW on a command window
- DISKSHADOW>set context persistent
- DISKSHADOW>set option transportable
- DISKSHADOW>begin backup
- DISKSHADOW> add volume f: (Database volume)
- DISKSHADOW> add volume g: (Log volume)
- DISKSHADOW>set metadata c:\metadata\exchangemeta.cab (specify the path where you want the metadata to be stored)
- DISKSHADOW> create
- DISKSHADOW>end backup
- Manually copy the exchangemeta.cab file from the source server to the offload server and run these two commands:
 - DISKSHADOW>LOAD METADATA *path to exchangemeta.cab*
 - DISKSHADOW>IMPORT

- DISKSHADOW>list shadows all (This process might take a few minutes)
- DISKSHADOW>delete shadows all

Note: Volumes f: and g: represent the Exchange Server database and log volumes. Repeat this test four times, and verify that the Windows Event Log contains no errors.

What to do next

When these tests complete without errors, you can install IBM Spectrum Protect. Use the DiskShadow tool for verification. The DiskShadow tool is preinstalled on the Windows Server operating system.

Note: On the last step of the configuration wizard, a VSS diagnostic check is run to verify the VSS setup. Any warnings must be fixed before you finish the configuration and start a Data Protection for Exchange Server operation.

Common errors returned from VSS operations

You can diagnose the cause of common errors that might occur when a VSS operation runs.

The following two errors are commonly returned:

ANS1017E (RC-50) Session rejected: TCP/IP connection failure

This message is displayed when the IBM Spectrum Protect backup-archive client CAD is either not running or is not configured properly.

ANS1532E (RC5722) Proxy Rejected: Proxy authority is not granted to this node.

This message is displayed when the IBM Spectrum Protect server is not configured correctly for the proxy nodes.

Transitioning Exchange Server backups from IBM Spectrum Protect Snapshot to IBM Spectrum Protect

Configure IBM Spectrum Protect Snapshot so that you can access both a local and IBM Spectrum Protect server. Use this approach if you move to the IBM Spectrum Protect environment and want to continue interacting with locally managed snapshots until policy marks them for expiration.

Before you begin

If you use the **Standalone** and IBM Spectrum Protect server configuration wizards to configure IBM Spectrum Protect Snapshot, you do not need to manually implement the following procedures. To interact with IBM Spectrum Protect, run the **TSM** configuration wizard. To interact with the IBM Spectrum Protect Snapshot server, run the **Standalone** configuration wizard. You can move from one type of server to another by running the corresponding configuration wizard at any time.

About this task

If you do not use the configuration wizards, coordinate efforts with your IBM Spectrum Protect server administrator to complete the following manual tasks. Some of the following command examples are formatted on multiple lines. Issue each command on a single line.

Configuring the IBM Spectrum Protect server Procedure

1. Select or create the policy definitions that are used for each type of backup you plan to use. You can provide the administrator with the existing local-defined policy settings in your IBM Spectrum Protect Snapshot stand-alone environment. Use the GUI or the command-line interface of Data Protection for Microsoft Exchange Server to retrieve this information.
2. Register your Data Protection for Microsoft Exchange Server node name and password with the IBM Spectrum Protect **register node** command. The `userid` option must also be specified with the **register node** server command. For example:

```
register node DPnodename DPpassword userID=DPnodename
```
3. If not already defined in IBM Spectrum Protect, register the IBM Spectrum Protect backup-archive client node name and password for the workstation where the Exchange server is installed. For example:

```
register node BAnodename BApasword userID=BAnodename
```
4. Define the proxy node relationship for the Target Node and agent nodes with the IBM Spectrum Protect **grant proxynode** command. For example:

```
grant proxynode target=DPnodename agent=BAnodename
```

Configuring the computer that runs the Exchange Server Procedure

1. In the directory where the Data Protection for Microsoft Exchange Server is installed, make a copy of the options file named `dsm.opt`. After you begin by using the IBM Spectrum Protect server, the copy is used for access to the IBM Spectrum Protect Snapshot stand-alone environment. One method of making the copy is to start the Exchange command-line prompt from the IBM Spectrum Protect Snapshot Snapin: In the IBM Spectrum Protect Snapshot Snapin Tree view, an Exchange server node is displayed for each Exchange server instance on the computer.
 - a. Select an Exchange server instance in the tree view. The integrated command line and an Actions pane are displayed.
 - b. Start the Data Protection for Microsoft Exchange Server command line from the Actions pane. Select:
Launch Command Line
 - c. To make a copy of the options file, enter:

```
copy dsm.opt dsm_local.opt
```
2. In the same directory, make a copy of the Data Protection for Microsoft Exchange Server configuration file. For example:

```
copy tdpexc.cfg tdpexc_local.cfg
```

Preserve the contents of the local configuration file under these conditions:

 - You specified policy bindings during the use of IBM Spectrum Protect Snapshot.
 - You are updating the policy bindings to reflect changes in your policy specifications for your IBM Spectrum Protect server usage.
3. In the IBM Spectrum Protect backup-archive client installation directory, make a copy of the VSS requestor options file named `dsm.opt`. Use the Windows **copy** command. For example:

```
C:\Program Files\Tivoli\TSM\baclient>copy dsm.opt dsm_local.opt
```

4. In all of the files named `dsm.opt`, modify the `TCPSEVERADDRESS` line. Replace `FLASHCOPYMANAGER` with the IP address of the IBM Spectrum Protect server. For example:

```
TCPServeraddress 9.52.170.67
```

To accomplish this task, use a text editor like Notepad or Wordpad.

5. To access the IBM Spectrum Protect Snapshot stand-alone environment during the transition period, open a Windows command prompt and change the directory to the IBM Spectrum Protect backup-archive client installation directory. This path is the default:

```
C:\Program Files\Tivoli\TSM\baclient
```

Create an alternative Windows service for the IBM Spectrum Protect Client Acceptor service by using the **dsmcutil** command. For example:

```
dsmcutil install cad /name:tsmcad4local  
/node:my_backup-archive_client_node  
/password:my_TSM_server_password  
/optfile:"C:\Program Files\Tivoli\TSM\baclient\dsm_local.opt"  
/httpport:1583
```

For more information about using the **dsmcutil** command, see `dsmcutil` command (http://www.ibm.com/support/knowledgecenter/SSEQVQ_8.1.2/client/c_cfg_dsmcutil_usewin.html).

6. Create an alternate Windows service for the IBM Spectrum Protect remote agent service. For example:

```
dsmcutil install cad /name:tsmcad4remote  
/node:my_backup-archive_client_node  
/password:my_TSM_server_password  
/optfile:"C:\Program Files\Tivoli\TSM\baclient\dsm_remote.opt"  
/httpport:1583
```

7. Edit the `dsm_local.opt` file in the Data Protection for Microsoft Exchange Server installation directory. Add this line:

```
HTTPPORT 1583
```

8. Start the alternate IBM Spectrum Protect Client Acceptor service:

```
dsmcutil start /name:tsmcad4local
```

9. Stop and restart the original IBM Spectrum Protect Client Acceptor service so that the new values in the `dsm.opt` file are activated. You can do this action through the Windows Services GUI or by using the **dsmcutil** command:

```
dsmcutil stop /name:"TSM Remote Client Agent"  
dsmcutil stop /name:"TSM Client Acceptor"  
dsmcutil start /name:"TSM Client Acceptor"
```

10. As backups start occurring and are managed in the IBM Spectrum Protect server environment, phase out the remaining backups that are created in the IBM Spectrum Protect Snapshot stand-alone environment. You can choose between two ways of achieving the phase-out:

- a. In the IBM Spectrum Protect Snapshot stand-alone environment, define a time-based policy that automatically causes the old backups to expire and delete. For example, if you want to expire each backup after it is 30 days old, update the time-based policy by using the command:

```
tdpexcc update policy mypolicy /daysretain=30  
/tsmoptfile=dsm_local.opt  
/configfile=tdpexc_local.cfg
```

You can also change these parameters by using the Local Policy Management dialog that is accessed from the **Utilities** menu of the Data

Protection for Microsoft Exchange Server Backup/Restore GUI. For information about how to start the GUI, see the section that describes how to access the IBM Spectrum Protect Snapshot stand-alone environment.

The process of expiring backups when their age exceeds the **daysretain** limit depends upon a basic function that is run in the stand-alone environment. The function must include an operation that queries the backups. If you do not regularly use the stand-alone environment client, you can use a scheduler to periodically start a command such as:

```
tdpexcc query tsm * /all
/tsmoptfile=dsm_local.opt
/configfile=tdpexc_local.cfg
```

For example, if your backups are created each week, then you can schedule the **query** command to run weekly to cause the expiration of out-of-date backups.

The last backup that is created while you run the stand-alone environment, is not automatically deleted by the process of expiring the backups. For that result, use the explicit delete operation, as described next.

- b. Alternatively, you can explicitly delete each backup when you determine that it is no longer needed. Use the Data Protection for Microsoft Exchange Server **delete backup** command, or the **Delete Backup** (right mouse-click menu option) in the GUI **Restore** tab.
11. To access the IBM Spectrum Protect Snapshot stand-alone environment:
 - a. Open the Automate tab to access the integrated command-line prompt.
 - b. Start IBM Spectrum Protect Snapshot stand-alone commands by appending the **/tsmoptfile** option, for example:

```
tdpexcc query tsm * /all
/tsmoptfile=dsm_local.opt
/configfile=tdpexc_local.cfg
```

- c. Start the IBM Spectrum Protect Snapshot GUI by issuing this command at the command prompt.

```
flashcopymanager.exe /tsmoptfile=dsm_local.opt
/configfile=tdpexc_local.cfg
```

12. If necessary, start the IBM Spectrum Protect Snapshot stand-alone environment to restore from a backup that was created in that environment.
13. When the transition is complete and you no longer need access to the IBM Spectrum Protect Snapshot stand-alone environment, you can remove the alternate services. To remove the services, use the Windows Services GUI or the **dsmcutil** command:

```
dsmcutil remove /name:tsmagent4local
dsmcutil remove /name:tsmcad4local
```

Examples of IBM SAN Volume Controller and IBM Storwize V7000 configuration scenarios

Configuration examples are scenarios that you can use to help you plan your data backup and recovery solutions.

Production application data is on standard volumes. Keep 14 snapshot backup versions. Use minimum storage space for snapshot backup versions. A full physical copy is not required. Complete two VSS backups per day.

SAN Volume Controller and Storwize V7000 settings

Create 14 space-efficient target volumes for each source volume to

be protected. Enable the autoexpand option for the space-efficient target volumes. Add the space-efficient target volumes to the VSS_FREE pool.

VSS Provider settings

Set the background copy rate to 0.

settings

Set the policy to retain 14 local backup versions. Schedule snapshot backups as required by setting the backup destination option to LOCAL.

After 14 VSS backups are completed, the 15th VSS backup causes the oldest backup to be deleted and reuses that target set.

Production application data is on standard volumes. Keep one snapshot backup version. Use minimum storage space for snapshot backup versions. A full physical copy is not required. Perform one VSS backup per day and send the backup to IBM Spectrum Protect.

SAN Volume Controller and Storwize V7000 settings

Create two space-efficient target volumes for each source volume to be protected. Enable the autoexpand option for the space-efficient target volumes. Add the space-efficient target volumes to the VSS_FREE pool .

VSS Provider settings

Set the background copy rate to 0.

settings

Set the policy to retain two local backup versions. Schedule snapshot backups as required by setting the backup destination to BOTH

Set the policy for local snapshot backups to retain n+1 backup versions so that n snapshot backups are available for restore. Otherwise, a local backup version might not be available if a VSS backup fails after the prior backup was deleted.

Production application data is on standard volumes. Keep one snapshot backup version. A full physical copy is required. Minimize space usage of background copies. Perform one VSS backup per day and send the backup to IBM Spectrum Protect.

SAN Volume Controller and Storwize V7000 settings

Create one standard target volume for each source volume to be protected. Add standard target volumes to the VSS_FREE pool.

VSS Provider settings

Use the default background copy rate of 50. Configure a custom value to use incremental FlashCopy.

settings

Set the policy to retain one local backup version. Schedule snapshot backups as required by setting the backup destination to BOTH.

When you use incremental FlashCopy backup processing, the VSS provider does not delete the single snapshot target set even though FlashCopy Manager software deletes the prior VSS snapshot before it creates a new snapshot.

Production application data is on standard volumes. Keep two snapshot backup versions. Full physical copies are required for local backup versions. Begin VSS backups every 12 hours with one backup sent to IBM Spectrum Protect daily.

SAN Volume Controller and Storwize V7000 settings

Create three standard target volumes for each source volume to be protected. Add standard target volumes to the VSS_FREE pool .

VSS Provider settings

Use the default background copy rate of 50.

settings

Set the policy to retain three local backup versions. Schedule VSS backups as follows: set the backup destination to LOCAL at 11:00, set the backup destination to BOTH at 23:00.

Set the policy for local snapshot backups to retain n+1 backup versions so that you can restore n snapshot backups.

Production application data is on standard volumes. Keep four snapshot backup versions. Use minimum storage space for snapshot backup versions. A full physical copy is not required. Perform VSS backups every six hours with one backup daily sent to IBM Spectrum Protect.

SAN Volume Controller and Storwize V7000 settings

Create five space-efficient target volumes for each source volume to be protected. Enable the autoexpand option for the space-efficient target volumes. Add space-efficient target volumes to the VSS_FREE pool .

VSS Provider settings

Use the default background copy rate of 0.

settings

Set the policy for local snapshot backups to retain five local backup versions. Schedule VSS backups as follows: set the backup destination to LOCAL at 06:00, 12:00, and 18:00, set the backup destination to BOTH at 00:00.

- Set policy to retain n+1 backup versions so that n snapshot backups are available for restore

Production application data is on space-efficient volumes. Keep two snapshot backup versions. A full physical copy is required for local backup versions. Perform VSS backups every six hours with one backup daily sent to IBM Spectrum Protect.

SAN Volume Controller and Storwize V7000 settings

Create three space-efficient target volumes for each source volume to be protected. Allocate the same percentage of real storage as for source volumes. Add space-efficient target volumes to the VSS_FREE pool .

VSS Provider settings

Use the default background copy rate of 50.

settings

Set the policy to retain three local backup versions. Schedule VSS backups as follows: set the backup destination to LOCAL at 06:00, 12:00, and 18:00, set the backup destination to BOTH at 00:00.

Set the policy for local snapshot backups to retain $n+1$ backup versions so that n snapshot backups are available for restore operations. This setting allows thin provisioning for both source and target volumes, and allows them to grow together.

Chapter 4. Protecting data

You can back up and restore your Microsoft Exchange Server data by using Microsoft Management Console (MMC) or the command-line interface.

About this task

If required, you can manage your installations remotely.

Note: For information about protecting Microsoft Exchange Server data in VMware environments, see chapter *Protection for in-guest applications* in the *IBM Spectrum Protect for Virtual Environments: Data Protection for VMware User's Guide*.

Prerequisites

With Data Protection for Microsoft Exchange Server, you can back up and restore Exchange Server data and protect your Exchange Server environment.

You can use Data Protection for Exchange Server to run backup and restore operations in a Database Availability Group (DAG) environment. A DAG consists of mailbox servers that provide recovery from database, server, or network failures. DAGs provide continuous replication and continuous mailbox availability.

Security requirements for backup and restore operations

For Data Protection for Exchange Server security, users who are logged on to the Exchange Server must have role-based access control (RBAC) permissions to access mailboxes and to complete mailbox restore tasks.

If your user name is authorized by the security policy in your organization, you can add user names in the Exchange Organization Management role group or subgroups. Users whose name is in the Exchange Organization Management role group or subgroups can complete mailbox restore operations. Users whose name is not in the Exchange Organization Management role group or subgroups might experience slower performance when completing restore operations.

You must define a minimum set of management roles and role scope for the Exchange user.

- Set the role and scope:

Management roles

"Active Directory Permissions", "Databases", "Disaster Recovery", "Mailbox Import Export", "View-Only Configuration", and "View-Only Recipients".

To restore an Exchange 2013 public folder mailbox, the Exchange user must also have the Public Folders management role. To restore mail to a Unicode PST file, the Exchange user must have the Mailbox Import Export management role.

The following Exchange Powershell cmdlet sets RBAC permissions:

```
New-RoleGroup -Name "My Admins" -Roles "Active Directory Permissions", "Databases",  
"Disaster Recovery", "Mailbox Import Export", "Public Folders",  
"View-Only Configuration", "View-Only Recipients" -Members operator1
```

The preceding example creates a new group, My Admins, with minimum roles to run Data Protection for Exchange Server, and assigns user operator1 to this group. The operator1 user can run Data Protection for Exchange Server but with limited Exchange privileges, for example, the user cannot create or remove a user mailbox.

Management role scope

Ensure that the following Exchange objects are in the management role scope for the user name who is logged on to the Exchange Server:

- The Exchange Server that contains the required data
 - The recovery database that Data Protection for Exchange Server creates
 - The database that contains the active mailbox
 - The database that contains the active mailbox of the user who completes the restore operation
- Verify that the Exchange user name is a member of a local Administrators group, and has an active Exchange mailbox in the domain.

By default, Windows adds the Exchange Organization Administrators group to other security groups, including the local Administrators group. For Exchange users who are not members of the Exchange Organization Management group, you must manually add the user account to the local Administrators group. By using the Local Users and Groups tool on the computer of the domain member, select **Administrative tools > Computer Management > Local Users and Groups tool**. On a domain controller computer that does not have a local Administrators group or Local Users and Groups tool, manually add the user account to the Administrators group in the domain by selecting **Administrative tools > Active Directory Users and Computers tool**.

Software requirements for backup and restore operations

To protect Microsoft Exchange Server 2010 and 2013 data, verify that your environment is set up correctly.

Ensure that your environment is set up to meet the following requirements.

Microsoft Exchange Server 2010 and 2013 requirements

Data Protection for Exchange Server requires that you have local Administrator privileges.

Membership in the Organization Management group is not required because you might not want to grant Organization Management group permissions to all Exchange Server backup and restore operators. Instead, you can define customized role-based access control (RBAC) roles and management role scope so that Exchange Server users can run only limited operations within a limited scope.

Microsoft Exchange Server 2013 requirements

In Exchange Server 2013 mailbox restore operations, the MAPI clients must use the Remote Procedure Call over HTTP protocol (RPC over HTTPS, also known as Outlook Anywhere). You cannot use the RPC over TCP because Microsoft does not use that protocol.

Use Exchange Server 2013 CU2 or later versions, and download the correct MAPI. These software requirements are documented in the Hardware and Software Requirements technote at this location: All Requirements(<http://www.ibm.com/support/docview.wss?uid=swg21219345>). Follow the link to the requirements technote for your specific release or update level.

After you configure your environment, mailbox restore operations work in the same way as with previous versions of Microsoft Exchange Server.

Software requirements for mailbox restore operations

When you restore mailboxes and mailbox data, you can choose where to restore the mail and how to restore the mail. You can restore mailbox data from the GUI or command-line interface.

From these interfaces, you can restore data interactively by using the Mailbox Restore Browser or directly from Exchange Server database files. When you restore mailboxes and mailbox data on Exchange Server 2010 or later, ensure that your environment is set up to meet the following requirements:

- Ensure that the administrator account that is used to perform the mailbox restore operation has an active Exchange mailbox in the domain.
- Ensure that the user name who is logged in has role-based access control (RBAC) permissions to complete individual mailbox restore operations.
- Ensure that the directory where you restore a mailbox has enough temporary disk space to store the entire mailbox database and log files. To specify the restore directory path, use the following settings on the General property page for the Exchange Server workload:
 - **Temporary Log Restore Path**
 - **Temporary Database Restore Path**

If you do not specify a directory, the database files are restored into a directory that is specified by the TEMP environment variable.

- Ensure that the correct version of Microsoft Exchange Server MAPI Client and Collaboration Data Objects is installed on the Exchange Server that you use to run the mailbox restore operations. The correct version is identified in the Hardware and Software Requirements technote that is associated with the level of your software. This technote is available at this web page: All Requirements(<http://www.ibm.com/support/docview.wss?uid=swg21219345>)
Follow the link to the requirements technote for your specific release or update level.

The amount of time that is needed to complete the restore process depends on the size of the mailbox databases, the network speed, and the number of mailboxes to process.

VSS backup methods

Depending on your Exchange Server environment, you can run only full backups, full plus incremental backups, or full plus differential backups. Your backup strategy might also include backing up data to IBM Spectrum Protect or local shadow volumes.

As you consider Exchange Server backup methods, understand all aspects of Exchange Server disaster recovery and the backup recommendations that Microsoft provides.

Follow these guidelines:

- Do not implement incremental and differential backups together.
- If you choose a strategy that involves incremental or differential backups, you must disable circular logging on the databases of the Exchange Server.

Full backup method

Use the full backup method during low usage times because a full backup can take a long time to run. However, the restore process is the most efficient because only the most recent full backup is restored.

Full backup plus incremental backup method

Use the full backup plus incremental backup method when the normal backup schedule or network capacity cannot support a full backup.

To minimize the effect on the backup schedule and network traffic during peak times, you can run a periodic full backup, followed by a series of incremental backups. For example, you can schedule full backups on the weekend and incremental backups during the week. You can run full backups during low usage times and when increased network traffic can be tolerated.

If you use this backup strategy, modify the IBM Spectrum Protect storage management policies to ensure that all incremental backups are collocated on the IBM Spectrum Protect server. In this way, you can improve data restore performance by reducing the number of media mounts that are necessary to restore a series of incremental backups.

Full backup plus differential backup method

Use the full backup plus differential backup method if your backup schedule and network capacity can facilitate backing up all transaction logs that accumulate between full backup operations. This strategy requires that only one differential backup plus the last full backup be transferred to complete a restore operation. However, the same amount of data must be transferred in the differential image, as in the series of incremental backup operations.

Therefore, a full backup plus differential backup policy increases network traffic and IBM Spectrum Protect storage usage. This policy assumes that the differential backups are processed as often as the incremental backups.

Consider the potential advantages and whether you can justify the additional resources that are necessary to resend all prior transaction logs with each subsequent differential backup.

IBM Spectrum Protect backups versus local shadow volumes backups

When you create a policy for your backups, you must choose whether to back up data to IBM Spectrum Protect storage versus VSS disks. Data backups to IBM Spectrum Protect typically takes longer to process than backups to local shadow volumes.

Backing up Exchange Server data to IBM Spectrum Protect is necessary when long-term storage is required. For example, saving Exchange Server data on tape for archival purposes requires long-term storage. IBM Spectrum Protect backups are also necessary for disaster recovery situations when the disks that are used for local backups are unavailable.

By maintaining multiple backup copies on IBM Spectrum Protect server storage, a point-in-time copy is available if backups on the local shadow volumes become corrupted or deleted.

Local shadow volumes

When you back up data to local shadow volumes, ensure that sufficient local storage space is assigned to the local shadow volumes. Create different sets of policies for backups to both local shadow volumes and to IBM Spectrum Protect server storage. If you use a VSS provider other than the Windows VSS System Provider, follow the backup recommendations of the VSS provider.

You can run backups to local shadow volumes by time and backup versions. It is more effective to base policy for local backups on version limits because local snapshots are created more frequently and VSS storage provisioning and space limitations apply. In Database Availability Group (DAG) environments, all of the DAG members must use the same local VSS policy.

Environment and storage resources also impact how many backup versions you can maintain on local shadow volumes for VSS fast restore and VSS instant restore operations, and on IBM Spectrum Protect server for VSS restore operations.

Database Availability Group backup and restore operations

To optimize use of available server resources, Database Availability Group (DAG) members often store a subset of the Exchange Server databases in a combination of active and passive copies.

Typical DAG configuration

In the following example, three copies of five databases span five servers in a DAG. This configuration ensures that two servers in the DAG never have the same set of database copies. The configuration also provides greater resilience to failures. Specifically, three servers must fail before the servers lose access to a database.




Figure 2. Typical DAG configuration

Typical data protection deployments in DAG environments

You can back up data from any DAG member and restore the data to any DAG member. You can also back up data from either the active or passive copy. Full and incremental database backups do not have to be completed from the same DAG member. All databases included in a VSS type backup are integrated.

The following figure illustrates a deployment of a backup task that is distributed across DAG members.




Figure 3. Example of backups that are distributed across DAG members

To specify a backup of all DAG nodes, issue the same backup command on each node. The command file contains separate backup commands per database. For example:

```
tdpexcc backup DB1 full /minimumbackupinterval=60 /preferdagpassive
tdpexcc backup DB2 full /minimumbackupinterval=60 /preferdagpassive
tdpexcc backup DB3 full /minimumbackupinterval=60 /preferdagpassive
```

In this deployment, one schedule applies to all nodes. The same backup command file is used for each node.

The following figure illustrates another possible backup task distribution across DAG members.




Figure 4. Another example of backups distributed across DAG members

In this deployment, one schedule applies to all nodes. The same backup command file is used for for all nodes. The command file contains separate backup commands per database on that node. For example:

```
tdpexcc backup DB1 full /minimumbackupinterval=60 /preferdagpassive
tdpexcc backup DB2 full /minimumbackupinterval=60 /preferdagpassive
tdpexcc backup DB3 full /minimumbackupinterval=60 /preferdagpassive
```

Best practices for backing up a Database Availability Group

When you back up data, distribute the backup workload for scalability and isolate backup activity to a dedicated backup node. When you isolate backup activity, it minimizes the impact to production databases.

As a best practice, identify all replica copies of the same database and eliminate redundant backups of the same databases. You can apply retention policies to databases. Back up databases from any node in the availability group and run restore operations from any node in the availability group.

Complete backups for replicated database copies from the same Exchange Server. Additionally, complete backups on the passive database copies. When you backup passive database copies, you do not increase the load on the production Exchange Server.

When you back up databases, follow these guidelines:

- Use a DAG member to store DAG database backups.
- Ensure that the same VSS policy applies to all DAG members.
- Ensure that the first backup is a FULL backup when you move backups to DAG member backups.
- Ensure that previous backups are manually deleted after you move backups to DAG member backups, assuming that those backups are no longer needed.
- Run backups from a passive database copy to avoid increasing the load on the active databases.
- Schedule all DAG members that have a copy of the database to back up the database at the same time. To set the minimum amount of time before a backup of another DAG copy of the same database is allowed, specify the **MINIMUMBACKUPINTERVAL** parameter. When you specify this parameter, only one backup is taken per backup cycle.
- If the Exchange Server database belongs to a DAG and is an active database copy, specify the **/EXCLUDEAGACTIVE** parameter to exclude the databases from the backup.
- If the Exchange Server database belongs to a DAG and is a passive database copy, specify the **/EXCLUDEDAGPASSIVE** parameter to exclude the databases from the backup.
- If the Exchange Server database does not belong to a DAG, specify the **/EXCLUDENONDAGDBS** parameter to exclude the databases from the backup.
- To a backup is to be taken from a passive copy unless no valid passive copy is available, specify the **/PREFERDAGPASSIVE** parameter.
- To bypass an integrity check if two or more valid database copies exist in a DAG, specify the **/SKIPINTEGRITYCHECK** parameter.

Best practices for restoring a Database Availability Group

In a DAG environment, you must restore databases on an active database copy. To restore to a passive database copy, you must first move the copy to the active state. After the restore operation is complete, you can move the active database copy to the passive state.

If you back up data to a local system, you can complete data restore operations only on the Exchange Server where the backup is taken.

Starting Microsoft Management Console

After you complete the configuration process, start Microsoft Management Console (MMC) to protect your Exchange Server data.

Before you begin

If you try to use Data Protection for Exchange Server before you complete the configuration process, the software does not function correctly.

About this task

Data Protection for Exchange Server software is displayed in MMC as a plug-in. MMC uses a navigation tree to organize the computer data that is registered. Each computer icon that is followed by the word *Dashboard* represents a physical computer.

When you register a computer, information about the computer is collected and stored. Password information is encrypted and stored separately. The computers that are registered are tracked with a globally unique identifier (GUID). The GUID is used when you back up and restore data.

You can create groups of computers. These groups consolidate information when you view the dashboard, prepare reports, and run group commands. By default, the computers in a group are selected when you complete tasks for the group, but you can select more computers in the tree to include in an operation.

Procedure

To start MMC, click **Start > All Programs > Data Protection for Microsoft Exchange Server > DP for Exchange Management Console**.

Starting the Data Protection for Exchange Server command-line interface

You can start the Data Protection for Exchange Server command-line interface by using a Windows command prompt with administrative privileges. Alternatively, you can start the command-line interface from Microsoft Management Console (MMC).

Procedure

1. Start MMC.
2. In the navigation tree, select the computer node where you want to run the commands.
3. Expand the **Protect and Recover Data** node.
4. In the navigation tree, select an Exchange Server node.
5. Click the **Automate** tab. An integrated command line is available in the task window. You can use the interface to enter PowerShell cmdlets or command-line interface commands. The output is displayed in the main window.
6. From the drop-down list, change **PowerShell** to **Command Line**.

Related information:

“Command-line overview” on page 152

Managing Data Protection for Exchange Server installations remotely

From a single Data Protection for Exchange Server installation you can manage all of the Data Protection for Exchange Server installations in your enterprise.

Before you begin

Your system must run Microsoft Windows 2008 or later versions, PowerShell 3.0 or later, and Data Protection for Exchange Server. On Windows 2012 and later versions, PowerShell version 4.0 is installed by default. For information about downloading, installing, and enabling Windows PowerShell, see this web page: Microsoft Windows Management Framework 3.0 Downloads (<http://www.microsoft.com/en-us/download/details.aspx?id=34595>)

Procedure

Enable remote management for Data Protection for Exchange Server installations by issuing the following Windows PowerShell command.

```
Enable-PSRemoting -force
```

This command enables remote management in most environments. If you use Microsoft Exchange, complete the following steps:

1. On the primary system, issue the following command:

```
enable-wsmancredssp -role client -delegatecomputer remote_computername
```
2. On each remote system that runs Microsoft Exchange, issue the following command:

```
enable-wsmancredssp -role server
```
3. Add the Data Protection for Exchange Server servers to the trusted hosts list by issuing the following command on each remote system:

```
Set-Item WSMan:\localhost\Client\TrustedHosts -Value remote_server_name -Force
```
4. Verify that Windows PowerShell Remoting is configured correctly by issuing the following cmdlet.

```
invoke-command -computername remote_computername  
-scriptblock {pwd} -Credential administrator
```
5. After you make configuration changes, restart the winrm service by entering the following command:

```
Restart-Service winrm
```

Adding remote systems

You can add remote systems in MMC.

Procedure

1. From Microsoft Management Console (MMC), in the Actions pane, click **Manage Computers**.
2. Verify that the local system is listed in both the Tree Nodes and Computers panes.
3. From the Tree Nodes pane, click the add icon. The icon is green and resembles the symbol for addition.
4. Type the name and description for the new tree node.
5. From the Computers pane, click the add icon. The computers that you add are associated with the tree node that you are creating. If you add only one computer, the tree node type can be either **Dashboard** or **Group**. If you add

more than one computer, the tree node type is **Group**. If you add only one computer, from the Tree Nodes pane, you can toggle between the **Dashboard** and **Group** types.

6. Type the system name and a description. For systems that are not in the domain, provide the fully qualified address. Alternatively, to select a system that is based on another system in the domain or to read a list of computers from a file, on the Computers pane, click **Import**. Clicking **Import** displays a dialog called Add Computers. From the Add Computers dialog, there are two tabs: Active Directory and Import. To complete the Add Computers dialog window entries, complete the following steps:

- a. For the Active Directory tab, complete these fields

Domain

The current domain is displayed. The domain cannot be changed.

Location

The organizational unit that is used to search for computers. The default value is displayed.

Name By default, the wildcard character (*) is displayed. You can leave the default value or enter a specific name.

Account

The current account is displayed. If you want to use a different account, click **Search** to search the domain for other computers. The Search capability is enabled only when the **Location** and **Name** fields have values.

- b. For the Import tab, browse to find a comma-separated values (.CSV) file that contains computer entries. After you find a .CSV file and click **Import**, the contents of the .CSV file are read as entries and are added to the list. The following .CSV file is an example of a valid .CSV file for the import activity:

```
NewNode1,Group1,CurrentUser,Test node 1
NewNode2
NewNode3,,Description of NewNode3
NewNode4,Group2,CurrentUser,Test node 4
```

The first column (the node name) is required. The other data columns are optional. The list is processed by position. For the group, if a group does not exist, the group is created.

7. From the Computers pane, click **Test Connection**. The test status is reported in the Message column of the Computers table.
8. Click **OK** to close the Manage Computers window.
9. Verify that the tree node is listed in the navigation tree. The remote node does not have all of the functionality available for local systems. For example, entries for learning, online support, and favorite links are not displayed. For tree node type **Dashboard**, the main window displays the Protect, Recover, and Automate tabs. For tree node type **Group**, the main window displays the Group Dashboard, Group Reports, and Group Commands tabs.
10. After you add systems, you can remove (delete) the systems. You can also select the system to edit the properties, including tree node type, that you entered when you added the system. If you want to change the order of the systems that are displayed in the navigation tree, use the GUI controls from the Manage Computers window.

Determining managed storage capacity

You can track the capacity of managed storage assets. This information can be useful when you are calculating storage requirements for license renewal.

About this task

Typically, the capacity that is used by server data differs from the capacity of the volume that contains that data. For example, a set of databases might require a capacity of 1 GB and be on a 10 GB volume. When a snapshot of the volume is created, the Data Protection for Exchange Server managed capacity measurement is 10 GB.

Procedure

1. From Microsoft Management Console (MMC), select an Exchange Server instance.
2. On the **Protect** tab, click **Properties** in the Action pane.
3. Select **Managed Capacity** from the list of available property pages. The managed capacity is calculated and displayed.
4. To view a list of the volumes that contain backups and their respective managed capacities, click **Show Details**.

Backing up Exchange Server data

By using Microsoft Volume Shadow Copy Service (VSS), you can back up Exchange Server data and mount the backup if required.

About this task

Data Protection for Exchange Server tracks and stores mailbox location history, which is used to automate mailbox restore operations. This action causes a delay to occur before each backup. In small or centralized Active Directory environment, the delay might be a few seconds or minutes. In large or geographically dispersed environments, the delay might take more time.

If you do not plan to use mailbox restore, you can safely disable mailbox history.

Related concepts:

“Data protection in VSS environments” on page 3

“Data backup processing” on page 7

Ensuring successful MAPI connections

If you use Exchange Server 2013, use the MAPI Settings property page to verify that the user mailbox is online. You can also view and update the MAPI registry key that enables Data Protection for Exchange Server to connect to the Exchange Server.

Before you begin

Ensure that the correct version of Microsoft Exchange Server MAPI Client and Collaboration Data Objects is installed on the Exchange Server. The correct version is identified in the Hardware and Software Requirements technote that is associated with the level of your software.

About this task

For mailbox restore operations to succeed in Exchange Server 2013 environments, the MAPI client must use Remote Procedure Call over HTTPS (RPC over HTTPS), also known as Outlook Anywhere. You cannot use RPC over TCP.

Procedure

1. From Microsoft Management Console (MMC), select an Exchange Server instance.
2. On the **Protect** tab, click **Properties** in the Action pane.
3. Select **MAPI Settings** from the list of property pages.
4. Verify that the following information is correct in the Exchange Server environment:
 - The **mailbox alias** field shows the mailbox that you are logged in to. Verify that you can open the mailbox in Microsoft Outlook or Outlook Web Access (OWA).
 - The **Exchange Profile Server** field shows the correct mailbox endpoint on the Exchange Server that has the Client Access Server (CAS) role. Verify that you can open the target mailbox in Outlook or OWA.
5. Edit the registry key only if the default value is incorrect. Use one of the following methods.
 - Enter the registry key value in the `RpcHttpProxyMap_TSM` field.
 - Enter the `Domain` field value and select or clear the **Use HTTPS authentication** check box. When you change either of these values, the values of the registry key automatically updates in the `RpcHttpProxyMap_TSM` field.

Consider that the values that you enter override the registry key that is in the `HKEY_CURRENT_USER\Software\Microsoft\Windows NT\Current Version\Windows Messaging Subsystem` directory. If you modify the registry incorrectly, the connection to the Exchange Server might fail.

RpcHttpProxyMap_TSM

Change the registry key values to reflect the correct domain, endpoint, and Remote Procedure Call (RPC) authentication methods for your environment. By default, the following format is used.

```
Domain=Proxy Server,RpcHttpAuthenticationMethod,  
RpcAuthenticationMethod,IgnoreSslCert
```

For example:

```
companyname.local=https://exchange.companyname.com,ntlm,ntlm,false
```

where:

- *Domain* value is the domain suffix of the personalized server ID, for example, `companyname.local`. Specify any domain or a substring of a domain, or the asterisk (*) and question mark (?) wildcard characters, for example, `*.companyname.local`.
- *Proxy Server* value is the RPC proxy server that has the Client Access Server (CAS) role. Specify the fully qualified domain name (FQDN) of the RPC proxy server. Precede the FQDN by `http://` for an HTTP connection, or `https://` for an HTTPS connection. For example, `https://exchange.companyname.com`

- *RpcHttpAuthenticationMethod* value is the method that is used to authenticate RPC-over-HTTP connections. Specify *NtLm*, *Basic*, *Negotiate*, or *WinNt*.
- *RpcAuthenticationMethod* value is the method that is used to authenticate RPC-over-TCP connections. Specify *NtLm*, *Negotiate*, *WinNt*, *Anonymous*, or *None*.
- *IgnoreSslCert* value indicates whether the Exchange Server validates SSL certificates. For the Exchange Server to ignore invalid certificates, specify *False*.

Domain

Change the domain name to reflect the correct domain if for example, you have multiple domains, or the default domain value is incorrect. To match all domains, enter the asterisk (*) wildcard character. When you change this domain value, the *Domain* value of the registry key automatically updates in the *RpcHttpProxyMap_TSM* field.

Use HTTPS authentication

Select this check box if RPC-over-HTTPS is enabled for the Exchange Server that is hosting the MAPI profile. Otherwise, clear this check box to ensure that HTTP authentication is used for RPC-over-HTTP connections. When you change this authentication value, the *RpcAuthenticationMethod* value of the registry key automatically updates in the *RpcHttpProxyMap_TSM* field.

Related tasks:

“Specifying configuration parameters for IBM Spectrum Protect” on page 50

“Troubleshooting MAPI connection issues” on page 136

Related reference:

“MAPI Settings” on page 63

Backing up Exchange Server data by using VSS

By using Microsoft Volume Shadow Copy Service (VSS), you can back up Exchange Server data and mount the backup if required.

Before you begin

- You must have a VSS provider that is configured for your environment.
- If you back up Exchange Server databases in a Database Availability Group (DAG) environment, and you want to back up your databases to a common node, ensure that you set up a DAG node name (DAGNODE).

Tip: Backing up DAG databases to a common node is helpful when you want to manage backups with a single policy, regardless of which DAG server completes the backup.

You can set up the DAG node name in the **DAG Node** field in the TSM Node Names page of the IBM Spectrum Protect configuration wizard, or in the **Back up DAG databases to common node** field in the General properties page for your Exchange Server workload.

Procedure

1. Start Microsoft Management Console (MMC) and click **Exchange Server** in the navigation tree.
2. On the **Protect** tab, select one or more databases to back up. Alternatively, click the **Protect Data** shortcut in the start page of MMC.

- a. Filter the list of available databases in the results pane by entering a keyword in the **Search** field.
 - b. If you are running backup operations in an Exchange Server DAG environment, you can back up an active database copy or passive database copy. View the copy status in the **DAG Status** column on the **Protect** tab.
3. Specify the backup options. If the backup options are not displayed, click **Show Backup Options**.
- To use offloaded backups, set the **Offload** option to **True**.
If you use offloaded backups, specify the remote client node, **RemoteDSMAGENTNode**, that runs the VSS offloaded backups on a remote computer. This option applies only to VSS backups.
 - Select **Skip Integrity Check** and choose one of the following options.

Table 12. Options for integrity checking

Task	Action
Bypass integrity checking for all database and log files	Select Yes
Run integrity checking to verify that all database and log files are free of errors	Select No This option is the default.
Bypass integrity checking for database files only if at least two valid copies of a database (one active and one passive copy) exist in a DAG	Select Skip Database Check If Healthy
Bypass integrity checking for database and log files only if at least two valid copies of a database (one active and one passive copy) exist in a DAG	Select Skip Database And Log Check If Healthy

- If you are scheduling the backup of databases in an Exchange Server DAG, set the **Minimum Backup Interval** value to the minimum amount of time, in minutes, before a backup of another copy of the same DAG database can begin. The default value is 0, which means that you can back up the database again immediately after a backup operation of that database is complete. The time of the last database backup is determined from the Exchange Server and not the IBM Spectrum Protect server.
This option specifies that only one database copy can be backed up within a time frame. This option prevents all members in a DAG from backing up the database. Specify this setting for tasks that are scheduled to run when you click **Run Scheduled**.
 - If you are scheduling the backup of databases in an Exchange Server DAG, set **PreferDAGPassive** option to **True** to skip the backup for an active database copy unless no valid passive copy is available. If no valid passive copy is available, the backup is created from the valid active database copy. Specify this setting for tasks that are scheduled to run when you click **Run Scheduled**.
4. In the Actions pane, click **Backup Destination** to specify whether you want the data to be backed up to your local server, an IBM Spectrum Protect server, or both.
 5. Optional: Choose a mode for the current task:
 - **Run Interactively**: Click this item to run the current task interactively. This selection is the default.

- **Run Scheduled:** Click this item to convert the current action into a scheduled task. When you select this item, the schedule wizard runs the command that is required to complete the task.
6. To create the backup, select the backup action in the Actions pane. You can run a full, copy, incremental, or differential backup with the VSS backup method.

Related concepts:

“VSS backup requirements” on page 4

“Offloaded VSS backups” on page 12

Related tasks:

“Restoring a Database Availability Group database backup” on page 107

Related reference:

“Backup command” on page 153

Mounting Exchange Server backups

To see a copy of Exchange Server data from a specific point in time, mount a snapshot backup.

About this task

A copy of data from a specific time is also known as a point-in-time consistent copy or online snapshot.

Restriction: You cannot use Microsoft Management Console (MMC) to mount a backup to a different server. To mount a VSS snapshot to a remote server, enter the **mount backup** command at the command line, or use the **Mount-DpExcBackup** cmdlet.

When you submit a mount request, all of the volumes that are contained in the original snapshot set are imported. If the number of volumes that are imported exceed the maximum number of allowable mapped volumes for the environment, the mount operation can fail.

You can mount VSS snapshot backups either as read-only or read/write. When a snapshot backup is mounted as read/write, you can do individual mailbox or mail item restores without needing to copy the Exchange database file from the snapshot backup into the recovery database (RDB); which greatly reduces the restore time. There are two variations of the mount read/write option:

- **Mount read/write (modifies backup, applies to COPY backups only)**

For VSS providers that support transportable shadow copies, you can mount a COPY type backup as read/write. After mounting, your COPY backup is marked as modified and while you can mount it again in the future, this backup can no longer be used as a restore point in future full database restore operations. It can be used for mailbox restore operations only. All databases on the snapshot volume that are mounted as read/write are marked as modified.

- **Mount read/write (without modifying backup)**

This mount option is only available for SAN Volume Controller (SVC) devices and requires IBM System Storage Support for Microsoft Virtual Disk and Volume Shadow Copy Services version 4.12 or later. With this option, you can mount writable VSS snapshots of FULL or COPY backups.

Note: You can override your default mount options as specified in the configuration file by using either the `/MOUNTRW` parameter on the `mount backup` command, or the `Mount-DpExcBackup` cmdlet.

Procedure

1. Start MMC.
2. Click **Recover Data** in the welcome page of MMC.
3. In the Actions pane on the **Recover** tab, click **Mount Backup**.
4. Either type the path to the empty NTFS or ReFS folder where you want to mount the backup or browse to find the path. Click **OK**. On the **Recover** tab, the backup that you mounted is displayed.
5. Optional: If required, select the **Mount the snapshots in read/write mode** option.
6. In the Actions pane, select the **Explore** and **Unmount Backup** tasks for the backup that you mounted.

Deleting Exchange Server backups

You can remove an Exchange Server backup that you created with the VSS backup method. Use this procedure only for deletions that are outside the scope of your standard policy management deletions.

Before you begin

Typically, backups are deleted automatically based on user-defined policy management settings. This procedure is necessary only if you must delete backups that are outside the scope of IBM Spectrum Protect Snapshot policy definitions.

If you back up Exchange Server Database Availability Group (DAG) databases to IBM Spectrum Protect, you can delete the database backup from the DAG member to a local shadow volume only from the Exchange Server on which the backup is created.

If you delete a remotely-mounted backup, the snapshots and the relationship between the source and target volumes on the storage device are also deleted. However, the target volume that is imported and mounted might continue to exist. In addition, the target volume might not be available to the server where the remote mount occurred. The operations to the target volume depend on the VSS hardware provider and the storage device implementation.

After the maximum number of remotely-mounted backup versions or the maximum number of days to retain a backup is exceeded, the associated backup is expired and deleted.

Procedure

1. Start Microsoft Management Console (MMC).
2. Click **Recover Data** in the welcome page of MMC.
3. On the **Recover** tab for the Exchange Server instance, select **View: Database Restore**. In the Results pane, browse to and select one or more database backups to delete.
4. In the Actions pane, click **Delete Backup**. While a backup is being deleted, two tasks are displayed in the task window to show that the deletion is in progress, and that the view is being refreshed.

Restoring Exchange Server data

You can restore an Exchange Server database backup to a recovery database or to an alternate (or relocated) database. You can also restore a replicated database copy in a Database Availability Group (DAG).

Related concepts:

“Data backup processing” on page 7

Setting data restore options in Microsoft Management Console

To optimize the data restore process for your environment, modify the default options that are available in Microsoft Management Console (MMC).

Procedure

1. On the Recover tab, select **Database Restore**.
2. Click **Show Restore Options** to modify the default restore options as follows:

Table 13. Database restore options

Option	Action
Auto Select	<p>For this option, specify a value of Yes (default) to quickly select the backup objects to restore. With automatic selection, when you select the most recent backup to restore, all associated backups are automatically selected, up to the previous full backup. When you specify Yes, the automatic selection option applies to full backups, differential backups, and incremental backups, but not to copy backups. This option affects backups in the following ways:</p> <ul style="list-style-type: none">• When you click a differential backup, the associated full backup is also selected.• When you click an incremental backup, the associated full backup and all associated earlier incremental backups are also selected.• For VSS backup, automatically selects all databases that were backed up together to the local destination. However, databases that were backed up to IBM Spectrum Protect are not automatically selected.
From Server	<p>Enter the name of the server where the original backup is completed. The default value is the local server.</p>

Table 13. Database restore options (continued)

Option	Action
<p>Instant Restore</p>	<p>For this option, specify a value of Yes to use volume-level snapshot restore (instant restore) for local VSS backups if the backup exists on SAN-attached volumes. Specify a value of No to disable instant restore, which bypasses volume-level copy and uses file-level copy (fast restore) to restore the files from a local VSS backup. The default value is Yes, which uses volume-level snapshot restore if it is available.</p> <p>This option is available for VSS operations only. If you use instant restore for SAN Volume Controller earlier than version 5.1 or DS8000, ensure that any previous background copies that involve the volumes that are being restored are completed before you initiate the instant restore.</p> <p>This option is automatically set to No during <i>restore into</i> operations.</p> <p>In an instant restore operation, files on the destination file system are overwritten. Incremental and differential backups are automatically converted to file-level restores. An instant restore operation requires that the drive or volume where the mailbox database is located must be available. Any other process or application must not have access to the drive or volume.</p>
<p>Mount Databases After Restore</p>	<p>For this option, specify a value of Yes to automatically mount databases after backups are recovered. No is the default value for this option.</p>
<p>Run Database Recovery</p>	<p>If more backup files need to be restored before you run database recovery, select False.</p> <p>If you want to run database recovery after the restore is done, select True and specify whether only restored logs or both restored and current logs should be used for the recovery.</p> <ul style="list-style-type: none"> • Replay Restored and Current Logs Replays all transaction log entries both in the restored log files and in the log files on the server that have not been backed up. You cannot specify this option for instant restore operations. • Replay Restored Only Replays only restored logs. <p>Note: From the MMC, the ERASEexistinglogs parameter can be applied to erase the existing transaction log files for the database that is being restored before you restore it.</p>

Restoring an Exchange Server database

You can use the *restore into* function to restore an Exchange Server database backup to a recovery database or alternate database. You can also restore a DAG active or passive database copy to a recovery database or alternate database.

Before you begin

- Ensure that your system is set up to use the DAG node name (DAGNODE). You can specify the DAG node name in the **DAG Node** field in the TSM Node Names page of the IBM Spectrum Protect configuration wizard or in the **Backup DAG databases to common node** field in the General properties page for your Exchange Server workload.
- You can restore mailboxes with the Mailbox Restore Browser or Mailbox Restore functions. In some rare cases, however, you might want to restore data into a recovery database or alternate database. Ensure that a recovery database or alternate database exists before you attempt the restore operation.

About this task

- For database backups in the Exchange Server Database Availability Group (DAG) environment, you can restore a database regardless of which DAG member the database was backed up from because all database copies are backed up by using a single DAG node. Local backups must be restored on the node where the backup was completed.
- Running any type of *restore into* function automatically disables VSS instant restore capability.

When you restore a database by using instant restore processing, data that exists in the destination database is overwritten, and is no longer available after restore processing is complete. When you restore a database by using the *restore into* function, you restore data to an alternate target destination. The data is not restored to the original source destination. For the restore operation to be successful, the alternate target destination must be of equal or greater size as the original source volume.

- To complete restore operations, backups must be taken on the same version of Exchange Server.
- You cannot use multiple instances of Data Protection for Microsoft Exchange Server to restore databases into the recovery database simultaneously.

Procedure

1. From Microsoft Management Console (MMC), create the recovery database if one does not exist. You can also use PowerShell commands (cmdlets) to do this step.
2. Use Data Protection for Exchange Server to restore the mailbox database.
3. From MMC, right-click the backup that you want to restore, and click **Restore Into**, or select the backup and click **Restore Into** in the Actions pane. For example, if Maildb1 is the name of the relocated database that you are restoring, the command-line entry is as follows:

```
TDPEXC RESTore Maildb1 FULL /INTODB=Maildb1
```
4. When you restore data to a recovery database, specify the option to replay restored logs only; otherwise the restore can fail.
 - a. Issue **/recover=applyrestoredlogs** at the command prompt.
 - b. On the **Restore** tab, select **Replay Restored Logs ONLY**.

Only transaction logs that are contained in the backup are applied to the mailbox database when a recovery database restore operation is processing.

Related concepts:

“VSS restore requirements” on page 5

“Software requirements for mailbox restore operations” on page 89

Related reference:

“Restore command” on page 186

Restoring a Database Availability Group database backup

You can restore a replicated database copy in a Database Availability Group (DAG).

About this task

You can use Exchange Management Shell commands, which are provided in parentheses.

Procedure

1. Make the database that you want to restore active (**Move-ActiveMailboxDatabase**).
2. Suspend replication of all passive copies of the database (**Suspend-MailboxDatabaseCopy**).
3. Unmount the active mailbox database (**Dismount-Database**).
4. If you are using VSS instant restore, and the **During Instant Restore, automatically stop and restart necessary Microsoft Exchange services** option is not selected in Microsoft Management Console (MMC), or the **STOPSERVICESONIR** parameter is set to NO at the command line, stop the following replication services on all copies of the database.
 - (DAG environments only) Exchange Replication Service
 - (Exchange Server 2013 only) Exchange Search Host Controller Service
5. Restore the database and logs by using the command line or MMC.

Restriction: The database must not be mounted automatically after the restore. If you use MMC, ensure that the **Mount Databases After Restore** option is set to **No** in the Restore pane. If you use the command line, set the **/mountdatabases** parameter to NO.

However, if the **During Instant Restore, automatically stop and restart necessary Microsoft Exchange services** option is selected, or the **STOPSERVICESONIR** parameter is set to YES, you can set the **Mount Databases After Restore** option to YES.

6. If the service is stopped, start the replication service before you mount the active mailbox database. Otherwise, the database mount fails (**Mount-Database**).
7. Verify the health of the database before you update or reseed to replicated database copies. (**Get-MailboxDatabaseCopyStatus**)
8. Update or reseed all replicas (**Update-MailboxDatabaseCopy**). By completing this step, you can help to avoid potential transaction log synchronization problems that might arise if replication is resumed directly.
9. Move the active database to the server that you want (**Move-ActiveMailboxDatabase**).

Related reference:

“Restore command” on page 186

Complete restore or replacement of Exchange Server

You can completely recover Exchange Server 2010 or later versions.

For more information, see this article: Backup, Restore, and Disaster Recovery (<http://technet.microsoft.com/en-us/library/dd876874.aspx>)

Restoring mailbox data

In environments that run Microsoft Exchange Server 2010 or later versions, you can use the Data Protection for Exchange Server mailbox restore feature to run individual mailbox and item-level recovery operations. For information about features delivered to support Microsoft Exchange 2016, see related reference below.

Related reference:

“Updates for V8.1.2” on page xiii

Individual mailbox recovery

Data Protection for Exchange Server backs up at the database level, and also restores individual items from the database backup.

Backing up Exchange servers at the mailbox item-level can cause the following issues:

- Insufficient scalability as item-level backups that are run hourly on each day of the week still prove to be an inadequate solution.
- More resource strain is added to the production servers.
- Since database backups are still done, the Exchange data is duplicated as item-level backups. The same data is backed up a second time.

To address these issues, Microsoft provides these features in Exchange:

- “Deleted Item Restore” can be configured to keep items within the Exchange Server databases, even after they are deleted. This option enables the items to be restored later.
- “Deleted Mailbox Restore” can be configured to keep mailboxes within the Exchange Server databases, even after they are deleted. This option enables the items to be restored or reconnected later.
- The recovery database enables a database to be restored to a special database. Wizards and tools are provided by Exchange to extract data from this database. This process can be done without disrupting the production servers.

Restoring mailbox data

Data Protection for Exchange Server backs up mailbox data at the database level, and also restores individual mailbox items from the database backup.

Before you begin

You must have role-based access control (RBAC) permissions to complete individual mailbox restore operations.

If you plan to restore mail or folders by using a Simple Mail Transfer Protocol (SMTP) server, ensure that you configure the SMTP server before you start a restore operation. To set the configuration in the Management Console, right-click **Dashboard** in the tree view and select **Properties**. From the E-mail property page, enter the SMTP server and port.

About this task

- You can restore an entire mailbox or selected mailbox items to the original mailbox, another online mailbox on the same server, or to a Unicode .pst file. When you restore a mailbox to the mailbox restore destination, Data Protection for Exchange Server automatically restores the mail items in the Recoverable Items folder.
 - You cannot restore the Recoverable Items folder and subfolder hierarchy to a mailbox restore destination. You can restore only the mail items in the folders.
 - The mail items that you can restore depends on whether the mailbox is enabled for mailbox restore operations.
 - You can restore the Recoverable Items content for a public folder mailbox but not for each public folder in the public folder mailbox.
 - You can exclude the mail items in the Recoverable Items folder in mailbox restore operations.
 - You cannot create a subfolder in the Recoverable Items folder in a mailbox.
- In Exchange Server 2013, you can restore a public folder mailbox database, a public folder mailbox, or only a part of the mailbox, for example, a specific public folder.
 - To restore an Exchange 2013 public folder mailbox, the Exchange user must have the Public Folders management role.
 - You can restore a public folder mailbox only to an existing public folder mailbox that is on the Exchange Server.
 - You can restore a public folder only to an existing public folder. The public folder on the Exchange Server must have the same folder path as the public folder to be restored. If the public folder is deleted from the public folder mailbox on the Exchange Server, you must re-create the public folder with the same folder path as the public folder to be restored, before you start the restore operation.
 - As a best practice, restore public folder mailboxes separately from user mailboxes. Select only one public folder mailbox to restore at a time if you want to restore a specific public folder in the mailbox, or if you want to restore to a different public folder mailbox than the original mailbox.

If you restore multiple mailboxes in a single restore operation, and at least one of the mailboxes is a public folder mailbox, the mailboxes are restored only to their original mailbox locations. You cannot specify a filter or an alternate mailbox destination.
 - You might restore to a different public folder mailbox than the original mailbox if, for example, the public folder is relocated after the time of the backup. Before you complete the public folder restore operation, ensure that the public folder exists with the same folder path in the alternate mailbox location.
- In Exchange Server 2010 or later, you can restore an archive mailbox or a part of the mailbox, for example, a specific folder. You can restore archive mailbox messages to a mailbox that is on the Exchange Server, to an archive mailbox, or to an Exchange Server .pst file.

If you enable a user mailbox to be archived, ensure that the user is logged on to that mailbox at least once before you complete a backup and restore operation on the mailbox.
- If you restore multiple mailboxes, and you want to retain the recovery database after the restore operation is complete, ensure that all the mailboxes are in the same recovery database.

- By default, Data Protection for Exchange Server restores the latest backup that is available for the specified mailbox.

The amount of time that it takes to complete the restore process depends on the size of the mailbox databases, the network speed, and the number of mailboxes to process.

Procedure

1. Start Microsoft Management Console (MMC) and select **Exchange Server** in the navigation tree.
2. On the **Recover** tab for the Exchange Server instance, select the **Mailbox Restore** view.
3. Select one or more mailboxes to restore. A list of mailboxes that are backed up is displayed. If you restore mail to a Unicode personal folder (.pst) file, or you restore a mailbox that is deleted or re-created after the time of the backup, Data Protection for Exchange Server requires a temporary mailbox to store the mailbox messages. Create a temporary mailbox by setting the **Alias** of temporary mailbox option on the Properties page, under the **General** tab.

Tip: Ensure that the temporary mailbox that you create is on a database with enough disk storage capacity to accommodate all of the mailbox items that you are restoring.

4. Optional: By default, the entire mailbox is restored. You can use the **Item-Level Mailbox Filters** to identify individual messages to restore:
 - a. Click **Show Filter Options** and **Add Row**.
 - b. In the **Column Name** field, click the down arrow and select an item to filter.
 - You can filter public mailbox folders only by the **Folder Name** column.
 - You can filter Unicode .pst files only by **Backup Date**, **Folder Name**, and **All Content** filters.
 - You can filter by backup date, and click the default date and time to edit the table cell. To change the date, click the arrow at the end of the cell. The calendar date selection tool is displayed. After you select a date, to display the date in the field, press **Enter**. To edit the time, enter the time by using the 12-hour clock time convention such as 2 p.m.
When you specify a backup date, Data Protection for Exchange Server searches for a backup that corresponds to that exact date. If a backup with that exact date is not found, Data Protection for Exchange Server selects the first backup after that date.
 - c. In the **Operator** field, select an operator.
 - d. In the **Value** field, specify a filter value.
 - e. If you want to filter on more items, click **Add Row**.
5. Specify the restore options by clicking **Show Restore Options**.

Table 14. Database restore options

Task	Action
Keep Recovery Database After Restore	Use this option to retain a recovery database after a mailbox restore operation is complete. The default value is No . If you set the value to Yes , Data Protection for Exchange Server automatically retains the recovery database after mailbox restore processing.
Mailbox	If the alias of the mailbox to restore is not displayed in the list of mailboxes, specify the alias. This option overrides any selected mailboxes.
Mailbox Original Location	Use this option only if the mailbox was deleted or re-created since the time of the selected backup, and mailbox history is disabled. Specify the Exchange Server and the database where the mailbox was at the time of the backup. Use the following format: server-name,db-name, for example, serv1,db1.
Mark Restored Messages As Unread	Use this option to automatically mark the mailbox messages as unread after the restore operation is completed. The default value is Yes .
Use Existing Recovery Database	Use this option to restore the mailbox from an existing recovery database. The default value is Yes . If you set the value to No and a recovery database is mounted on the server before you restore the mailbox, Data Protection for Exchange Server automatically removes the recovery database during mailbox restore processing.
Enter Mount Point or Directory for MountRW Mailbox Restore	Select this option to specify either a directory path or a mount point for a read/write mount of a local VSS snapshot backup, that you want to use for a mailbox restore. If you do not use read/write mounts, no entry is necessary. Ensure that Use Existing Recovery Database value is set to No .

6. To complete the restore operation, click one of the following **Restore** options.

Table 15. Restore options

Task	Action
Restore Mail to Original Location	Select this action to restore mail items to their location at the time of the backup operation.

Table 15. Restore options (continued)

Task	Action
Restore Mail to Alternate Location	<p>Select this action to restore the mail items to a different mailbox.</p> <p>Note: If deleted mail items or tasks are flagged in the Recoverable Items folder of a mailbox, the items are restored with the flag attribute to the Flagged Items and Tasks view in the target mailbox.</p>
Restore Mail to non-Unicode PST file	<p>Select this action to restore the mail items to a non-Unicode personal folders (.pst) file.</p> <p>When you restore mail items to a .pst file with one selected mailbox, you are prompted for a file name. When you restore mail items to a .pst file with more than one selected mailbox, you are prompted for a directory location. Each mailbox is restored to a separate .pst file that reflects the name of the mailbox at the specified directory.</p> <p>If the .pst file exists, the file is used. Otherwise, the file is created.</p> <p>Restriction: The contents of each folder cannot exceed 16,383 mail items.</p>
Restore Mail to Unicode PST file	<p>Select this action to restore the mail items to a Unicode .pst file.</p> <p>When you restore mail items to a .pst file with one selected mailbox, you are prompted for a file name. When you restore mail items to a .pst file with more than one selected mailbox, you are prompted for a directory location.</p> <p>You can enter a standard path name (for example, c:\PST\mailbox.pst) or a UNC path (for example, \\server\c\$\PST\mailbox.pst). When you enter a standard path, the path is converted to a UNC path. If the UNC is a non-default UNC path, enter the UNC path directly.</p> <p>Each mailbox is restored to a separate .pst file that reflects the name of the mailbox at the specified directory. If the .pst file exists, the file is used. Otherwise, the file is created.</p>

Table 15. Restore options (continued)

Task	Action
<p>Restore Public Folder Mailbox</p>	<p>Select this action to restore a public folder mailbox to an existing online public folder mailbox.</p> <p>You can filter the mailbox and restore a specific public folder to an existing online public folder. In the Folder to be restored field, enter the name of the public folder that you want to restore. If you are restoring a subfolder in a parent folder, specify the full folder path in this format: <i>parent_folder_name/sub_folder_name</i>. To restore all subfolders in a parent folder, use <i>parent_folder_name/*</i>. If the full folder path includes spaces, enclose the folder path in double quotation marks, and do not append a backslash character (\) at the end of the folder path.</p> <p>You can also restore all or part of a public folder mailbox to a different public folder mailbox than the original mailbox. In the Target public folder mailbox field, specify the destination public folder mailbox that you want to restore to.</p>
<p>Restore Mail to Archive Mailbox</p>	<p>This action applies to a primary mailbox or an archive mailbox. Select this action to restore all or part of either type of mailbox to the original archive mailbox or to an alternate archive mailbox.</p> <p>You can filter the archive mailbox and restore a specific mailbox folder. In the Folder to be restored field, enter the name of the folder in the archive mailbox that you want to restore. If you are restoring a subfolder in a parent folder, specify the full folder path in this format: <i>parent_folder_name/sub_folder_name</i>. To restore all subfolders in a parent folder, use <i>parent_folder_name/*</i>. If the full folder path includes spaces, enclose the folder path in double quotation marks, and do not append a backslash character (\) at the end of the folder path.</p> <p>In the Target archive mailbox field, specify the archive mailbox destination that you want to restore to.</p>
<p>Exclude recoverable mail items while restoring the mailbox</p>	<p>Apply this action if you are restoring an online, public folder, or archive mailbox to an original mailbox, alternate mailbox, or to a Unicode .pst file.</p> <p>Specify a value of Yes to exclude the mail items in the Recoverable Items folder in mailbox restore operations. No is the default value.</p>

Related concepts:

“Security requirements for backup and restore operations” on page 87

Related tasks:

“Troubleshooting a MAPI error that prevents multiple mailboxes restoring in a Microsoft Exchange 2013 environment” on page 136

Related reference:

“Restoremmailbox command” on page 199

Restoring mailbox messages interactively with the Mailbox Restore Browser

You can use the Mailbox Restore Browser to interactively restore a mailbox or items from a mailbox on an Exchange Server.

Before you begin

You must have role-based access control (RBAC) permissions to complete individual mailbox restore operations.

If you plan to restore mail or folders by using a Simple Mail Transfer Protocol (SMTP) Server, ensure that you configure the SMTP Server before you start a restore operation. Set the configuration in Microsoft Management Console (MMC) by right-clicking **Dashboard** in the navigation tree and selecting **Properties**. Then, in the E-mail property page, enter the SMTP server and port.

Download and install the Exchange MAPI and Microsoft Outlook MAPI on different servers. Do not install Microsoft Outlook 2010 or 2013 on the same server that Data Protection for Exchange Server uses for mailbox restore operations. Conflicts might occur in the MAPI configurations.

About this task

- You can restore an entire mailbox or selected mailbox items to the original mailbox, another online mailbox on the same server, or to a Unicode .pst file. When you restore a mailbox to the mailbox restore destination, Data Protection for Exchange Server automatically restores the mail items in the Recoverable Items folder.
 - You cannot restore the Recoverable Items folder and subfolder hierarchy to a mailbox restore destination. You can restore only the mail items in the folders.
 - The mail items that you can restore depends on whether the mailbox is enabled for mailbox restore operations.
 - You can restore the Recoverable Items content for a public folder mailbox but not for each public folder in the public folder mailbox.
 - You can exclude the mail items in the Recoverable Items folder in mailbox restore operations.
 - You cannot create a subfolder in the Recoverable Items folder in a mailbox.
 - The Mailbox Restore Browser displays folders that are normally hidden from view, for example, in the Recoverable Items folder. Folder names in the Recoverable Items folder are internal to Microsoft Exchange and are not translated by Microsoft. Therefore, if you use a language other than English, the folder names still display in English.
- In Exchange Server 2013, you can restore a public folder mailbox database, a public folder mailbox, or only a part of the mailbox, for example, a specific

public folder. However, you cannot restore individual messages in a public folder by using the Mailbox Restore Browser interface.

- To restore an Exchange 2013 public folder mailbox, the Exchange user must have the Public Folders management role.
- You can restore a public folder mailbox only to an existing public folder mailbox that is on the Exchange Server.
- You can restore a public folder only to an existing public folder. The public folder on the Exchange Server must have the same folder path as the public folder to be restored. If the public folder is deleted from the public folder mailbox on the Exchange Server, you must re-create the public folder with the same folder path as the public folder to be restored, before you start the restore operation.
- As a best practice, restore public folder mailboxes separately from user mailboxes. Select only one public folder mailbox to restore at a time if you want to restore a specific public folder in the mailbox, or if you want to restore to a different public folder mailbox than the original mailbox.
If you restore multiple mailboxes in a single restore operation, and at least one of the mailboxes is a public folder mailbox, the mailboxes are restored only to their original mailbox locations. You cannot specify a filter or an alternate mailbox destination.
- You might restore to a different public folder mailbox than the original mailbox if, for example, the public folder is relocated after the time of the backup. Before you complete the public folder restore operation, ensure that the public folder exists with the same folder path in the alternate mailbox location.
- If you restore multiple mailboxes, and you want to retain the recovery database after the restore operation is complete, ensure that all the mailboxes are in the same recovery database.
- By default, Data Protection for Exchange Server restores the latest backup that is available for the specified mailbox.

Restriction: Only mailboxes within the same database can be restored in a single mailbox restore action.

Procedure

1. Start MMC.
2. Under the **Protect and Recover Data** node in the navigation tree, select **Exchange Server**.
3. On the Recover panel, click **View > Mailbox Restore Browser**.
4. In the Select Source window, specify the mailbox that you want to restore. Choose from the actions in the following table:

Table 16. Selecting mailboxes to restore

Task	Action
Browse mailboxes and select one to restore	<ol style="list-style-type: none"> 1. From the drop-down list, select Browse Mailboxes. 2. Select a mailbox. 3. Click OK. <p>Tip: Use the Search field to filter the mailboxes. You can also sort the mailboxes by columns.</p>

Table 16. Selecting mailboxes to restore (continued)

Task	Action
Specify a mailbox to restore by name	<ol style="list-style-type: none"> 1. In the Mailbox Name field, enter the name of the mailbox to restore. 2. Click OK.
Restore a mailbox backup that was created at a specific time	<ol style="list-style-type: none"> 1. In the Backup Date/Time field, click the default date and time to edit the details. 2. To change the date, click the calendar icon, select a date, and press Enter. 3. To change the time of day, use the 12-hour system convention such as 2 p.m. 4. Click OK.
Review the mailbox backups that are available to restore before you complete the restore operation	<ol style="list-style-type: none"> 1. From the drop-down list, select Browse Mailboxes. 2. Select a mailbox for which backups exist. 3. From the Available Database Backups list, review the backups that are available for the mailbox and select a backup version to restore. 4. Ensure that the Backup Date/Time field reflects the time stamp for the selected mailbox backup. 5. Click OK.
Restore a mailbox that was deleted or re-created after the time of the backup	<p>In the Actions pane, click Properties, and on the General page, enter the temporary mailbox alias.</p> <p>Tip: If you do not enter the alias, the mailbox restore operation uses the administrator mailbox as a temporary storage location.</p>
Browse all databases in a backup	<ol style="list-style-type: none"> 1. From the drop-down list, select Browse Databases. 2. Select a database. 3. Click OK. <p>Tip: Use the Search field to filter the databases. You can also sort the mailboxes by columns.</p>

After the selected mailbox is restored to the recovery database, the restored mailbox and folders are displayed in the results pane.

5. In the results pane, browse the folders and messages that are contained within the selected mailbox. Choose from the following actions to select the mailbox, folder, or message to restore:

Table 17. Previewing and filtering mailbox items

Task	Action
Preview mailbox items	<ol style="list-style-type: none"> 1. Select a mailbox item to display its contents in the preview pane. 2. When an item contains an attachment, click the attachment icon to preview its contents. Click Open or save the item by clicking Save.
Filter mailbox items	<p>Use the filter options to narrow the list of folders and messages in the result pane.</p> <ol style="list-style-type: none"> 1. Click Show Filter Options and Add Row. 2. Click the down arrow in the Column Name field and select an item to filter. You can filter by folder name, subject text, and so on. You can filter public mailbox folders only by the Folder Name column. When you select All Content, the mailbox items are filtered by attachment name, sender, subject, and message body. 3. In the Operator field, select an operator. 4. In the Value field, specify a filter value. 5. If you want to filter on more items, click Add Row. 6. Click Apply Filter to filter the messages and folders.

6. In the Actions pane, click the folder or messages restore task that you want to run. If you click **Save Mail Message Content**, which becomes available only when a message is selected in the preview pane, a Windows Save File window is displayed. Specify the location and message name and click **Save**. The Restore Progress window opens and shows the progress of the restore operation. Data Protection for Exchange Server restores the mailbox backup to its original mailbox location.
7. To restore a mailbox or mailbox item to either of the following locations, complete the following steps. Choose from the actions in the following table:

Table 18. Restoring a mailbox to another mailbox or .pst file

Task	Action
Restore a mailbox or mailbox item to a different mailbox	<ol style="list-style-type: none"> 1. On the Actions pane, click Open Exchange Mailbox. 2. Enter the alias of the mailbox to identify it as the restore destination. 3. Drag the source mailbox to the destination mailbox on the results pane. <p>Restriction: You cannot drag mail items or subfolders in the Recoverable Items folder to a destination mailbox.</p>

Table 18. Restoring a mailbox to another mailbox or .pst file (continued)

Task	Action
Restore a mailbox to an Outlook personal folders (.pst) file	<ol style="list-style-type: none"> 1. On the Actions pane, click Open non-Unicode PST File. 2. When the Windows File window opens, select an existing .pst file or create a .pst file. 3. Drag the source mailbox to the destination .pst file on the results pane. <p>Restriction: You can use the Mailbox Restore Browser only with non-Unicode .pst files.</p>
Restore Public Folder Mailbox	<p>Select this action to restore a public folder mailbox to an existing online public folder mailbox.</p> <p>You can filter the mailbox and restore a specific public folder to an existing online public folder. In the Folder to be restored field, enter the name of the public folder that you want to restore. If you are restoring a subfolder in a parent folder, specify the full folder path in this format: <i>parent_folder_name/sub_folder_name</i>. To restore all subfolders in a parent folder, use <i>parent_folder_name/*</i>. If the full folder path includes spaces, enclose the folder path in double quotation marks, and do not append a backslash character (\) at the end of the folder path.</p> <p>You can also restore all or part of a public folder mailbox to a different public folder mailbox than the original mailbox. In the Target public folder mailbox field, specify the destination public folder mailbox that you want to restore to.</p>

In the Actions pane, the **Close Exchange Mailbox** and **Close PST File** tasks are displayed only when a destination mailbox or .pst file is opened.

8. Optional: Remove the recovery database by clicking **Close Mailbox to Restore**. This option is displayed only after a recovery database is created. Data Protection for Exchange Server removes the recovery database and cleans up the restored files. If you do not select **Close Mailbox to Restore**, the recovery database is not removed even if you exit MMC.

If MMC also detects a recovery database that is created outside of Data Protection for Exchange Server, it automatically connects to it. When you complete your mailbox restore tasks, you must manually remove the recovery database. You cannot use the **Close Mailbox to Restore** option.

Related concepts:

“Mailbox restore operations” on page 26

Related tasks:

“Troubleshooting mailbox restore errors” on page 134

Restoring mailboxes directly from Exchange database files

If the backup database (EDB) file and log files are available on the disk of a supported Microsoft Exchange Server, you can restore an individual mailbox directly from the EDB file.

Before you begin

If you use Spectrum Protect for Virtual Environments software, review the following guidelines before you restore the mailbox:

- You can use Spectrum Protect for Virtual Environments to back up an Exchange Server in a virtual machine. For more information about the **backup** command, see Backup command (http://www.ibm.com/support/knowledgecenter/SSEB6_8.1.0/ve.user/r_ve_vmcli_backup.html).
- To restore mailboxes from the backups that are created by Spectrum Protect for Virtual Environments, mount the virtual volumes that contain the EDB file and log files with read/write access. You can obtain read/write access by clearing the **Mount virtual volume as read only** check box.
- If the log files are on a different volume than the EDB file, mount the volume that contains the log files on an unused drive letter. In this way, you can apply the transaction logs to the EDB file.

If you use Data Protection for Exchange Server to back up the Exchange Server, you can enter the following command to restore the database files to a local disk:
tdpexcc RESTOREFILES

Verify that read/write access to the EDB file is available.

Verify that the Exchange Server transaction log files are available.

Procedure

1. From the Exchange server, start Data Protection for Exchange Server.
2. After you log on to Data Protection for Exchange Server, in the navigation area, select the **Exchange Server** node and the **Recover** tab. The Mailbox Restore Browser view opens.
3. In the Actions pane, click **Open EDB File on Disk**.
4. In the window, enter or browse to the location of the EDB file.
5. In the window, enter or browse to the location of the log file directory. Specifying a path to the log file directory is not required. However, the amount of time that is necessary to complete the restore operation is reduced when you provide the log file directory path.
6. Click **OK**. The EDB file is opened and the mailboxes are displayed.
7. Select the mailbox that you want to restore and the type of restore that you want to complete. For example, you can restore a mailbox to a PST file.
8. When the restore operation is complete, click **Close Mailbox to Restore**. You are prompted to save or delete the recovery database folder.

Related reference:

“Restoremmailbox command” on page 199

Restoring a deleted mailbox or items from a deleted mailbox

You can use the Data Protection for Exchange Server mailbox restore operation to restore a mailbox (or items from a mailbox) that was deleted from an Exchange Server.

Before you begin

If you are restoring a mailbox that was deleted or re-created since the time of the backup, you must specify a temporary mailbox with enough storage capacity to accommodate all the mailbox items that you are restoring. Specify a temporary mailbox by setting the `/TEMPMAILBOXAlias` parameter. If the `/TEMPMAILBOXAlias` parameter is not set, the default mailbox is the logon user mailbox.

Procedure

Decide where the mailbox data from the deleted mailbox is to be restored. With the mailbox restore operation, you have three options as follows:

1. Restore the deleted mailbox data to the original location. Before you run the mailbox restore operation, re-create the mailbox that is using Exchange.
2. Restore the deleted mailbox data into an active alternative mailbox in an online Exchange Server.
3. Restore the deleted mailbox data into an Exchange Server personal folders (.pst) file.

Related reference:

“Restoremailbox optional parameters” on page 203

“Restoremailbox command” on page 199

Viewing, printing, and saving reports

You can access reports on recent activity and historical managed capacity. You can determine which licenses and software are installed.

Procedure

1. Select **Reporting** in the **Manage** section. A list of available reports is displayed. Each report provides a summary of the report contents.
2. Select a report from the list. The selected report displays.
3. To print or save the current report, click the appropriate icon at the top of the report.

Chapter 5. Automating

With Data Protection for Exchange Server *automation* capability, you can run commands from the command line, create scripts, schedule tasks, and use Microsoft Management Console (MMC) to start tasks. The tasks that you can automate are based on the scripts and schedules that you create.

About this task

Data Protection for Exchange Server supports you automating tasks from the command-line interface or Microsoft Windows PowerShell command prompt (Version 3.0 or later). You can also use the **Automate** tab in the MMC.

Preparing to use Windows PowerShell cmdlets with Data Protection for Exchange Server

Data Protection for Exchange Server includes a set of Windows PowerShell cmdlets to help you manage Data Protection for Exchange Server components in your environment.

About this task

You can issue the cmdlets that are provided with Data Protection for Exchange Server in Windows environments.

Data Protection for Exchange Server cmdlets help support a seamless management environment and greatly improve remote management and automation capabilities. You can aggregate cmdlets together to form commands and use the large volume of existing cmdlets from other vendors.

Before you use the cmdlets, complete the following steps.

Procedure

1. Log on to the system as an administrator.
2. From a Windows PowerShell command prompt, issue the following command:

```
set-executionpolicy remotesigned
```

3. During installation of Data Protection for Exchange Server, the following Windows PowerShell modules are imported automatically from the TDPEXchange folder.

- FmModuleExc.dll
- FmModuleMMC.dll

If you wish to import the Windows PowerShell modules manually, from the Windows PowerShell command prompt, import modules, with the administrator credentials, as follows:

- a. Navigate to the TDPEXchange folder.
- b. Issue the following commands:

```
import-module .\FmModuleExc.dll  
import-module .\FmModuleMMC.dll
```

- c. (Optional) To use the cmdlets in these modules any time that you start Windows PowerShell, add the following lines to your profile. The following path is the default profile path.

```
$path = (get-itemproperty -path "HKLM:\SOFTWARE\IBM\TDPEExchange\  
currentversion\mmc" -ea SilentlyContinue).path  
if ($null -ne $path)  
{  
    dir "$path\fmmodule*.dll" | select -expand fullname | import-module  
    -force -Global  
}
```

What to do next

For information about creating, running, monitoring, and troubleshooting scripts with cmdlets, see Windows PowerShell 3.0 or later documentation. For more information about Windows PowerShell cmdlets, consistent naming patterns, parameters, arguments, and syntax, see this web page as a starting point: Microsoft TechNet: Getting Started with Windows PowerShell (<http://technet.microsoft.com/en-us/library/hh857337.aspx>).

Cmdlets for Microsoft Management Console

The following list identifies the cmdlets that you can use when interacting with Microsoft Management Console (MMC).

- **Clear-FcmMmcManagedCapacityHistory**
- **Clear-FcmMmcScheduledActivityHistory**
- **Disable-FcmMmcSchedule**
- **Enable-FcmMmcSchedule**
- **Get-FcmMmcActivity**
- **Get-FcmMmcComputerInformation**
- **Get-FcmMmcManagedCapacityHistory**
- **Get-FcmMmcReport**
- **Get-FcmMmcSchedule**
- **Get-FcmMmcScheduledActivity**
- **New-FcmMmcSchedule**
- **Remove-FcmMmcSchedule**
- **Set-FcmMmcSchedule**
- **Start-FcmMmcSchedule**

To view the details about a specific cmdlet, run the **Get-Help** cmdlet with the cmdlet name. For example:

```
Get-Help New-FcmMmcSchedule
```

To continue the example, to see examples for the cmdlet, enter:

```
get-help New-FcmMmcSchedule -examples
```

For more information, enter:

```
get-help New-FcmMmcSchedule -detailed
```

For technical information, enter:

```
get-help New-FcmMmcSchedule -full
```

For online product information, enter:

```
get-help New-FcmMmcSchedule -online
```

For information about a specific parameter, enter:

```
help New-FcmMmcSchedule -Parameter backupdestination
```

To display the help in a separate window, include the **-showwindow** parameter with the **help** command.

Cmdlets for protecting Microsoft Exchange Server data

The following table identifies the cmdlets that you can use to protect Microsoft Exchange Server data.

Table 19. Cmdlets to protect Microsoft Exchange Server data. The following table identifies the cmdlets that you can use to protect Microsoft Exchange Server data.

Cmdlet name	Related command-line interface command	Short description
Add-DpExcPolicy	tdpexcc create policy	Create a policy for Data Protection for Exchange Server.
Backup-DpExcComponent	tdpexcc backup	Back up a Microsoft Exchange Server database.
Copy-DpExcPolicy	tdpexcc copy policy	Copy an existing policy.
Dismount-DpExcBackup	tdpexcc unmount backup	Dismount a backup.
Get-DpExcBackup	tdpexcc query tsm *	Query backups.
Get-DpExcComponent	tdpexcc query exchange	Query the Exchange Server for all databases that are available for backup.
Get-DpExcConfig	tdpexcc query tdp	Display configuration information.
Get-DpExcConnection	tdpexcc query tsm	Query a list of the current values set in the configuration file for IBM Spectrum Protect.
Get-DpExcInformation	tdpexcc query exchange	Query general local Exchange Server information.
Get-DpExcMailboxLocationHistory	tdpexcc q tsm /showMailboxInfo	Query the mailbox location history.
Get-DpExcManagedCapacity	tdpexcc query managedcapacity	Query managed capacity for Microsoft Exchange Server.
Get-DpExcPolicy	tdpexcc query policy	Display policy information.
Mount-DpExcBackup	tdpexcc mount backup	Mount a backup to provide access to the files that the backup contains. You can mount a backup as read-only or read/write.
Remove-DpExcBackup	tdpexcc delete backup	Remove the backup.
Remove-DpExcPolicy	tdpexcc delete policy	Delete the policy.
Reset-DpExcTsmPassword	tdpexcc changetsmpassword	Change the IBM Spectrum Protect password used by Data Protection for Exchange Server.
Restore-DpExcBackup	tdpexcc restore	Restore a backup.
Restore-DpExcMailbox	tdpexcc restore mailbox	Restore a mailbox.
Set-DpExcConfig	tdpexcc set paramname	Set the application configuration parameters in a configuration file.
Set-DpExcPolicy	tdpexcc update policy	Update a policy.

To view the details about a specific cmdlet, run the **Get-Help** cmdlet with the cmdlet name. For example:

```
Get-Help Backup-DpExcComponent
```

To continue the example, to see examples for the cmdlet, enter:

```
get-help Backup-DpExcComponent -examples
```

For more information, enter:

```
get-help Backup-DpExcComponent -detailed
```

For technical information, enter:

```
get-help Backup-DpExcComponent -full
```

For online product information, enter:

```
get-help Backup-DpExcComponent -online
```

For information about a specific parameter, enter:

```
help Backup-DpExcComponent -Parameter backupdestination
```

To display the help in a separate window, include the **-showwindow** parameter with the **help** command.

Automating tasks

You can automate a workload by entering Windows PowerShell cmdlets or command-line interface commands in the integrated command line.

About this task

You use the Automate view to work with commands. You can create, save, store, and schedule commands to run at the scheduled time.

Procedure

1. To open the Automate view, select a workload that you want to work with and click **Automate**. An integrated command line is available in the task window from which you can enter PowerShell cmdlets or command-line interface commands.

2. Change **PowerShell** to **Command Line**.

3. To run a command, type a command in the details pane and click the **Execute** icon. You can issue the commands with or without specifying `tdpexcc`.

For example, for each selected workload instance, you can enter a single command or multiple commands, such as:

```
q tsm  
q exc
```

You can also run a saved task by clicking the **Open** icon, selecting the command file, and clicking the **Execute** icon. The output is displayed in the main window.

4. Click the **Save** icon and follow the prompts to save a command for future use.
5. To schedule a command, click the **Schedule this command** icon to open the scheduling wizard. Follow the prompts in the wizard to create a schedule for the command. The output of the command is displayed in the results pane.
6. (Optional) Save or send the command output to an email address.

What to do next

You can automate commands from the Protect, Recover, Schedule, and Task List views in Microsoft Management Console (MMC):

1. Start MMC and select the **Exchange Server** instance in the navigation tree.
2. Click the tab for the task you want to do (**Protect** or **Recover**).
3. Automate the command by using one of the following methods:

Result pane

Select the item for your task in the result pane, and select **Run Scheduled** in the toolbar menu. Click the appropriate task in the **Action** pane. When the scheduling wizard starts, enter the information for each prompt to create a scheduled task.

You can select the type of scheduler you want to use to manage your scheduled operations. Click the relevant radio button to select either the local Windows Scheduler or the TSM scheduler.

Task List pane

When a task is submitted, it displays in the task list pane. Select the appropriate task, then click **Schedule command script** in the task list toolbar. When the schedule wizard starts, enter the information for each prompt to create a scheduled task.

You can also right-click a task in the Task List pane and click **Copy**. Then, click the **Automate** tab and paste the command in the field.

IBM Spectrum Protect task scheduler

Review these guidelines when you define an IBM Spectrum Protect schedule.

- If you use the IBM Spectrum Protect-prompted scheduling mode, ensure that the Data Protection for Exchange Server option file specifies the `tcpclientaddress` and `tcpclientport` options. If you want to run more than one scheduler service, use the same `tcpclientaddress`. However, you must use different values for `tcpclientport` in addition to the different node names. As an example, you might want to run more than one scheduler service when you schedule Data Protection for Exchange Server and the regular Windows backup client.

You can use server-prompted scheduling only when TCP/IP communication is used. By default, Data Protection for Exchange Server uses the client polling schedule mode.

- If you make any changes that affect the scheduler in the Data Protection for Exchange Server options file, restart the scheduler to activate the changes. For example, the IBM Spectrum Protect server address, the schedule mode, or the client TCP address or port can affect the scheduler. To restart the scheduler, issue the following commands:

```
net stop "Data Protection for Exchange Server Scheduler"  
net start "Data Protection for Exchange Server Scheduler"
```

- The default IBM Spectrum Protect scheduler log file (`dsmsched.log`) contains status information for the IBM Spectrum Protect scheduler. In this example, the file is in this path:

```
d:\Program Files\Tivoli\TSM\TDPEExchange\dsmsched.log
```

You can override this file name by specifying the `schedlogname` option in the Data Protection for Exchange Server options file.

- Data Protection for Exchange Server creates a log file with statistics about the backed up database objects when the `/logfile` parameter is specified during the

tdpexcc command. Outputs from the scheduled commands are sent to the scheduler log file (`dsmsched.log`). After scheduled work is completed, check the log to verify that the work is completed successfully.

When a scheduled command is processed, the scheduler log might contain the following entry:

```
Scheduled event eventname completed successfully
```

This result indicates that IBM Spectrum Protect successfully issued the scheduled command that is associated with the *eventname*. No attempt is made to determine whether the command succeeded or failed. To assess the success or failure of the command, evaluate the return code from the scheduled command in the scheduler log. The scheduler log entry for the command return code is prefaced with the following text:

```
Finished command. Return code is:
```

If any scheduled backups fail, the scheduler script exits with the same error code as the failed backup command. A non-zero error code means that the backup failed.

- If `passwordaccess generate` is not specified in the `dsm.opt` file, then the IBM Spectrum Protect password must be specified on the **tdpexcc** command. To specify the password, use the `/tsmpassword` parameter in the command file that is run by the scheduler (`excfull.cmd`). You can also specify the password on the Data Protection for Exchange Server command line. For example:

```
tdpexcc query tsm /tsmnode=mars1 /tsmpassword=newpassword
```

Chapter 6. Troubleshooting

Data Protection for Exchange supports you in protecting Microsoft Exchange databases.

About this task

If you encounter a problem, you typically start with a symptom, or set of symptoms, and trace the root cause. Problem determination, however, is not the same as problem solving. During the process of problem determination, you might obtain sufficient information to enable you to solve the problem. In some cases, you cannot solve a problem even after you determine its cause. For example, a performance problem might be caused by a limitation of your hardware

Diagnosing problems

One of the most difficult challenges of troubleshooting in a client-server environment is determining which component is the origin of the problem. VSS diagnostic wizards are available to help you test VSS snapshots on your system. You can determine whether the source of the problem is a general VSS issue or a IBM Spectrum Protect issue.

Diagnosing VSS issues

You can test VSS persistent, non-persistent, and resync snapshots on your system with the assistance of a VSS diagnostics wizard.

Before you begin

Attention: Do not run these tests if you are already using SAN Volume Controller or Storwize V7000 space-efficient snapshots on your computer. If you do so, existing snapshots might be removed.

Procedure

1. Start Microsoft Management Console (MMC).
2. To open the diagnostics wizard, complete these steps:
 - a. Click **Diagnostics** in the results pane of the welcome page.
 - b. In the Actions pane, click **VSS Diagnostics**.

A list of volumes are displayed, and the status of each test is displayed when it is completed.

3. To view the results of the persistent and non-persistent snapshot testing, complete these steps:
 - a. Select the volumes or mount points to test and click **Next**.
 - b. Click **Show VSS Information** to view details about the VSS providers, writers, and snapshots that are available on your system.

The results of the persistent and non-persistent snapshot testing displays as Passed or Failed.

4. To view the results of the resync snapshot testing, complete these steps:

Note: VSS ResyncLUNs API instant restore tests will revert the data on the volume to an earlier time. Do not enable these instant restore tests on production volumes as data loss may occur.

- a. To test if the selected volumes support the VSS ResyncLuns API, select a volume and then click **Next**.
- b. Verify that the **Testing resync snapshot** field indicates a successful result.

The results of the resync snapshot testing display as Passed or Warning.

Note: On non-IBM storage devices, resync snapshots are necessary only for instant restore. Resync snapshots have no impact on backup and fast restore on non-IBM storage devices.

5. Review the results of the snapshot testing and click **Next**. The final results of the persistent and non-persistent snapshot testing display as Success or Unsuccessful.
6. Depending on the results, complete these steps:
 - If the testing status is a success, click **Finish** and exit the wizard.
 - If the testing status is not successful, click **Previous** and review information in the Rule dialog.
7. Return to the Management window and begin backup operations.

Determining that the problem is a Data Protection for Exchange issue or a general VSS issue

The Data Protection client interacts closely with the backup-archive client (DSMAGENT). The client completes all of the Virtual Shadow Copy Service (VSS) operations. You can test the connectivity between the Data Protection client and the IBM Spectrum Protect and determine whether the source of the problem is the Microsoft VSS service or with the IBM Spectrum Protect.

About this task

- The vssadmin and diskshadow tools are applications that can run backups that use the Microsoft Exchange VSS APIs.

vssadmin

A utility that is installed with your operating system. It can show current volume shadow copy backups and all installed shadow copy writers and providers in the command window.

diskshadow

The diskshadow tool is available on Windows 2008 server and 2008 R2.

With these tools, you can determine the following items:

- Verify VSS provider configurations
 - Rule out any possible VSS problems before you run the IBM Spectrum Protect VSS functions
 - That you might have a VSS configuration problem or a real hardware problem if an operation does not work with diskshadow or vssadmin
 - That you might have an IBM Spectrum Protect problem if an operation works with diskshadow or vssadmin but not with the IBM Spectrum Protect
- For VSS operations, you can re-create the problem with the Microsoft diskshadow tool. If you are able to re-create the problem with the diskshadow tool, the source of the problem is likely to be within the VSS provider or the Exchange server.

Procedure

1. Test the connectivity between the Data Protection client and the IBM Spectrum Protect DSMAgent.
 - a. Select the Exchange workload that you want to work with and click the **Automate** tab to open the **Automate** view.
 - b. To verify that your installation and configuration is correct, issue the **Query Exchange** command in the lower details pane and click **Execute** (or **Enter**). Alternatively, issue the **TDPEXCC QUERY EXCHANGE** command on the computer where the Exchange server is installed. The results are displayed in the pane.

The **TDPEXCC QUERY EXCHANGE** command returns information about the following items:

- Exchange server status
- Circular logging
- VSS components

The following example shows a sample of the output that is generated by the **TDPEXCC QUERY EXCHANGE** command:

```
Volume Shadow Copy Service (VSS) Information
-----
Writer Name           : Microsoft Exchange Writer
Local DSMAgent Node   : SERVERA
Writer Status         : Online
Selectable Components : 4
```

If the **TDPEXCC QUERY EXCHANGE** command does not return all of this information, you might have a proxy configuration problem. Contact the IBM Spectrum Protect server administrator to have the correct server **GRANT PROXY** commands that are issued to enable proxy authority for nodes. If all of the information returned to you seems correct, proceed to the next step.

2. To determine whether the problem is with the Microsoft VSS service or a problem within the IBM Spectrum Protect code, use the `vssadmin` and `diskshadow` tools to re-create the error as follows:
 - a. Issue **VSSADMIN** commands, as shown in this example:

```
VSSADMIN LIST WRITERS
VSSADMIN LIST PROVIDERS
VSSADMIN LIST SHADOWS
```

The **VSSADMIN LIST SHADOWS** command does not list shadow copies of SAN-attached volumes.
The `vssadmin` tool uses Microsoft Software Shadow Copy provider to list the shadow copies that are created.
 - b. Before you install IBM Spectrum Protect for Mail, test the core VSS function. Do the following `diskshadow` testing before you install any IBM Spectrum Protect components:
 - Test non-persistent shadow copy creation and deletion by issuing the following **DISKSHADOW** commands:

```
diskshadow>set verbose on
diskshadow>begin backup
diskshadow>add volume f: (database volume)
diskshadow>add volume g: (log volume)
diskshadow>create
diskshadow>end backup
diskshadow>list shadows all
diskshadow>delete shadows all
diskshadow>list shadows all
```

Volumes *f:* and *g:* represent the Exchange database and log volumes. Repeat issuing the **DISKSHADOW** commands four times and verify that the Windows event log file contains no errors.

- Test persistent shadow copy creation and deletion by issuing the following **DISKSHADOW** commands:

```
diskshadow>set context persistent
diskshadow>set verbose on
diskshadow>begin backup
diskshadow>add volume f: (database volume)
diskshadow>add volume g: (log volume)
diskshadow>create
diskshadow>end backup
diskshadow>list shadows all (this might take a few minutes)
diskshadow>delete shadows all
diskshadow>list shadows all
```

Volumes *f:* and *g:* represent the Exchange database and log volumes. Repeat issuing the **diskshadow** commands four times and verify that the Windows event log file contains no errors.

- Test persistent transportable shadow copy creation and deletion by issuing the following **DISKSHADOW** commands:

```
diskshadow>set context persistent
diskshadow>set option transportable
diskshadow>add volume f: (database volume)
diskshadow>add volume g: (log volume)
diskshadow>set metadata c:\metadata\exchangemeta.cab
(the path where you want the metadata stored)
diskshadow>create
```

You must copy the *exchangemeta.cab* file from the source server to the offload server. After you copy the file, issue the following commands:

```
diskshadow>load metadata newpath/exchangemeta.cab
diskshadow>import
diskshadow>list shadows all (this might take a few minutes)
diskshadow>delete shadows all
```

Volumes *f:* and *g:* represent the Exchange database and log volumes. Repeat issuing the **diskshadow** commands four times and verify that the Windows event log file contains no errors.

3. Perform the following tests to ensure that VSS is working correctly:
 - a. Test nonpersistent shadow copy creation and deletion:
 - Run “**DISKSHADOW k: l:**” where *k:* and *l:* are the Exchange Server database and log volumes.
 - Repeat the previous step 4 times.
 - Inspect the Windows Event Log to ensure that the results are appropriate.
 - b. Test persistent shadow copy creation and deletion:
 - Run “**DISKSHADOW -p k: l:**” where *k:* and *l:* are the Exchange Server database and log volumes. Run “**DISKSHADOW -da**” if you do not have enough space.
 - Repeat the previous step 4 times.
 - Inspect the Windows Event Log to ensure that the results are appropriate.
 - c. Test nonpersistent transportable shadow copy creation and deletion (VSS Hardware Provider environments only):
 - Run “**DISKSHADOW -p -t=export.xml k:l:**” where *k:* and *l:* are the Exchange Server database and log volumes.

- Copy the resultant “export.xml” file from computer 1 to computer 2 before you continue to the next step.
- On the computer you have set aside for offload, run “DISKSHADOW -i=export.xml”
- Inspect the Windows Event Log to ensure that things look appropriate.

If any of these tests fail repeatedly, you have a hardware configuration problem or a real VSS Problem. Consult your hardware documentation for known problems or search Microsoft Knowledge Database for any information.

If all tests pass, continue to Step 4.

4. Re-create your specific problem by using diskshadow. If you can re-create your problem, only through a series of steps (for example: a backup fails only when you perform two consecutive local backups), try to perform those same tests by using diskshadow.
 - Exchange VSS backups to Local are simulated by running a diskshadow persistent snapshot.
 - Exchange VSS backups to the IBM Spectrum Protect are simulated by running a diskshadow nonpersistent snapshot.
 - Exchange VSS backups to Local and to the IBM Spectrum Protect are simulated by running a diskshadow persistent snapshot.
 - Offloaded Exchange VSS backups to the IBM Spectrum Protect are simulated by running a diskshadow nonpersistent, transportable snapshot.

See the diskshadow documentation for the specific commands for performing backups.

If you can re-create the problem, it most likely is a general VSS issue. See the Microsoft Knowledge Database for information. If your operation passes successfully with diskshadow, it most likely is an IBM Spectrum Protect or Data Protection for Exchange client problem.

What to do next

For more information, see this technote: Verifying VSS functionality for the Data Protection Exchange backup (<https://www.ibm.com/support/docview.wss?uid=swg21403456>)

Resolving reproducible problems

When a component fails to operate as designed, try to reproduce the problem and capture information about the current operating environment at the time of the error. You can troubleshoot VSS backup and restore operations, mailbox restore errors, and VSS and SAN Volume Controller, Storwize V7000, or DS8000 problems.

Troubleshooting VSS backup and restore operations

If you encounter a problem during VSS backup and restore processing, attempt to reproduce the problem in your environment.

Before you begin

If a VSS backup fails, verify that sufficient disk space is available to store the snapshot.

Procedure

1. Try the operation that failed again.
2. Restart the IBM Spectrum Protect services, including the TSM Client Acceptor and the TSM Remote Client Agent.
3. If the problem still exists, close other applications, especially those applications that interact with Exchange, for example, antivirus applications. Retry the operation that failed.
4. If the problem persists, look for information in the event logs: `tdpexc.log` and `dsmerror.log`. You can also review the messages in the Windows event log. Log entries might exist to help you identify the VSS event that triggers the issue.
5. If you do not find a resolution to the problem in the log files, complete the following steps:
 - a. Shut down the Exchange server or the computer.
 - b. Restart the Exchange server or the computer.
 - c. Run the operation that failed.

Failovers from VSS instant restore processing to VSS fast restore processing

If an error occurs early in a VSS instant restore operation, the error might cause the system to fail over to VSS fast restore processing. However, if an error occurs later in the instant restore operation, instant restore processing might fail without failing over to fast restore processing.

About this task

Errors in VSS instant restore operations might occur, for example, if the volume where the restored database is stored is used by another process.

Procedure

Check the error message in the `dsmerror.log` file.

Troubleshooting VSS limitations with IBM SAN Volume Controller and IBM Storwize V7000

When you run IBM Spectrum Protect Snapshot for Exchange Server VSS backups (non-offloaded) to a backup destination of IBM Spectrum Protect server, the IBM SAN Volume Controller or IBM Storwize V7000 LUNs can sometimes remain mapped to the Windows host even though the backup is complete.

Procedure

Use a backup destination other than IBM Spectrum Protect server (BOTH or LOCAL).

Results

When you run two IBM Spectrum Protect Snapshot for Exchange Server VSS backups and if the volumes are large, or the background copy rate is set to a low number, or both conditions occur, the second VSS backup might be presented to be in a hang state. Typically, the Exchange Server data is on IBM SAN Volume Controller or IBM Storwize V7000 disks. However, the second backup is waiting for the IBM SAN Volume Controller or IBM Storwize V7000 background copy of the first backup to complete before proceeding. IBM SAN Volume Controller or IBM Storwize V7000 does not allow two background copies of the same volume to occur at the same time. You might not know that the second backup is waiting for the first background copy to complete.

You might also see timeout errors if the previous IBM SAN Volume Controller or IBM Storwize V7000 background copy takes too long.

What to do next

To resolve timeout issues, schedule VSS backups so that enough time elapses between backups, or increase the copy rate of the IBM SAN Volume Controller or IBM Storwize V7000 background copy.

Troubleshooting VSS limitations with IBM N-series and NetAppFAS series

If you plan to run VSS backups with backup destination set to LOCAL, understand the limitations in the VSS Provider for NetApp FAS series and IBM N-series, and in SnapDrive 4.2 and earlier versions, that affect the way in which you can run your VSS backup operations. You must configure your environment correctly to avoid snapshot deletions, backup failure, and out-of-space conditions on the production volumes.

Before you begin

- Ensure that a NAS file server LUN that is used by Exchange Server databases is fully dedicated to the database. Exchange Server databases cannot share LUNs.
- Ensure that a NAS filer LUN that is used by Exchange Server databases is the only LUN on the filer volume. For example, if the Exchange Server uses four LUNs, four corresponding filer volumes must exist, where each volume contains one LUN.
- If the NetApp volume type is Traditional, ensure that VSS backups with backup destination set to LOCAL are bound to a management class that has `verExists=1`. This setting is not required if flexible volumes are used.
- Ensure that VSS backups with backup destination set to LOCAL are either a full or copy backup. You cannot mix local backups of type FULL and COPY.
- Ensure that VSS backups with backup destination set to TSM are a full or copy backup. There are no restrictions on IBM Spectrum Protect backups.
- When you run VSS backups, ensure that previous backups finish completely before you start a new backup. To avoid issues on Exchange Server, the VSS service, and, the NAS filer, avoid backup overlaps.

About this task

The following backup procedure is an example that shows how to optimally run VSS backups by using both IBM Spectrum Protect and local backup destinations. The following assumptions apply to this example backup procedure:

- Stated configuration requirements are in place.
- Daily VSS full backups to a local destination occurs every four hours - 12 a.m., 4 a.m., 8 a.m., 12 p.m., 4 p.m., 8 p.m.
- The VSS backup to IBM Spectrum Protect takes one hour to complete.
- The VSS backup to a local destination takes five minutes to complete.

Procedure

Set your daily VSS full back schedule to run in one of the following ways:

- Run daily VSS full backups to a local destination every four hours - 12 a.m., 4 a.m., 8 a.m., 12 p.m., 4 p.m., 8 p.m.
- Run daily VSS full backups to IBM Spectrum Protect storage by one of the following two methods:
 - Set **backupdestination** to BOTH to run at 12 a.m. Because this setting runs a 12 a.m. backup to a local destination, do not separately schedule a 12 a.m. backup to a local destination.
 - Set full offloaded-backup to run at 1 a.m. No VSS local backup is available to restore VSS backups between 1 a.m. and 4 a.m., when the next VSS backup to a local destination occurs.
- Set weekly VSS full backups to run to IBM Spectrum Protect, as an offloaded backup, at 5 a.m.

Troubleshooting mailbox restore errors

If you encounter a mailbox restore error, determine whether the problem is reproducible on other Exchange servers.

About this task

Some of the mailbox restore errors that you might encounter include MAPI connection issues to the mailbox, insufficient role-based access control (RBAC) permissions to complete the restore operation, or issues with the Mailbox Restore Browser feature.

Troubleshooting insufficient RBAC roles and permissions

For the following mailbox restore errors, ensure that the RBAC roles and management role scope are set on the Exchange objects for the Exchange user.

Procedure

1. If a mailbox fails to open and the error message indicates a missing RBAC permission, ensure that the user who is logged on to the mailbox has the required RBAC roles, and the management scope for those roles includes the database that contains the mailbox. Then, open the mailbox again.
2. If a mailbox restore operation fails and the error message indicates a missing RBAC permission, ensure that the user who is logged on to the mailbox has the required RBAC roles, and the management scope for those roles includes the source and target databases. Then, restart the restore operation.

Related concepts:

“Security requirements for backup and restore operations” on page 87

Troubleshooting mailbox permissions, authentication methods, and registry key settings in a Microsoft Exchange 2013 environment

To resolve mailbox restore errors in an Exchange Server 2013 environment, ensure that the Exchange Server mailbox permissions, authentication methods, registry key settings, and the Client Access Server (CAS) role are configured correctly.

Procedure

1. Grant full access permission to the user who is logged on to the target mailbox. When the administrator mailbox is used, Exchange Server 2013 usually blocks full access permission for the administrator by default.
2. To restore an Exchange 2013 public folder mailbox, ensure that the Exchange user has the Public Folders management role.
3. Log on to an Exchange Server 2013 mailbox as the Exchange Server administrator and ensure that sufficient storage space is available in the administrator mailbox.
4. Ensure that you can access the mailbox that you logged on to and the target mailbox in either Microsoft Outlook or Outlook Web Access.
5. Specify an Exchange Server 2013 CAS by setting the **CLIENTACCESSServer**=*servername* parameter. If you are using a load balancer, set the **CLIENTACCESSServer** parameter to point to the CAS instead of the load balancer.
6. Open the administrator mailbox and the target mailbox. On the Actions pane in the Mailbox Restore Browser interface, click **Open Exchange Mailbox**.
7. Verify that the MAPI registry key, `RpcHttpProxyMap_TSM`, is correct to enable Data Protection for Exchange Server to connect to the Exchange Server. Use one of the following methods:
 - Check the registry key that is in the `HKEY_CURRENT_USER\Software\Microsoft\Windows NT\Current Version\Windows Messaging Subsystem` directory. Change the registry key values to reflect the correct domain, endpoint, and Remote Procedure Call (RPC) authentication methods for your environment. For example, you might specify HTTPS as the authentication method if RPC-over-HTTPS connections are enabled for the Exchange Server that is hosting the MAPI profile. Otherwise, you might use HTTP authentication for RPC-over-HTTP connections.
 - Use the MAPI Settings property page in Microsoft Management Console (MMC) to ensure that the MAPI registry key is correct. Change the registry key values to reflect the correct domain, endpoint, and Remote Procedure Call (RPC) authentication methods for your environment.

By default, the following registry key format is used.

```
Domain=Proxy Server,RpcHttpAuthenticationMethod,  
RpcAuthenticationMethod,IgnoreSslCert
```

where:

- *Domain* value is the domain suffix of the personalized server ID, for example, `companyname.local`. Specify any domain or a substring of a domain, or the asterisk (*) and question mark (?) wildcard characters, for example, `*.companyname.local`.
- *Proxy Server* value is the RPC proxy server that has the Client Access Server (CAS) role. Specify the fully qualified domain name (FQDN) of the RPC proxy server. Precede the FQDN by `http://` for an HTTP connection, or `https://` for an HTTPS connection. For example, `https://exchange.companyname.com`

- *RpcHttpAuthenticationMethod* value is the method that is used to authenticate RPC-over-HTTP connections. Specify NTLM, Basic, Negotiate, or WinNT.
- *RpcAuthenticationMethod* value is the method that is used to authenticate RPC-over-TCP connections. Specify NTLM, Negotiate, WinNT, Anonymous, or None.
- *IgnoreSslCert* value indicates whether the Exchange Server validates SSL certificates. For the Exchange Server to ignore invalid certificates, specify `False`.

The default registry key looks like the following example:

```
contoso.com=https://mail.contoso.com,ntlm,ntlm,false
```

Troubleshooting MAPI connection issues

Procedure

To diagnose MAPI-to-mailbox connection issues, enter the **TDPMAPI TESTMAPI** command with these parameters:

/MAILBOXALIAS

This parameter is the alias name for the mailbox that you are logged on to. The parameter refers to the email alias for the user and is the portion of the email address before the @ symbol. Run this command for the mailbox to be restored and the mailbox that you are logged on to.

/EXCSERVER

(Exchange Server 2010 environments) This parameter is the name of the Exchange Server that has the Client Access Server (CAS) role. By default, the parameter points to the local server. Use the Exchange PowerShell command, **get-ExchangeServer | fl**, to determine the Exchange Server that defined the CAS role for the mailbox database. You must specify this parameter when a CAS Load Balancer exists within the environment.

(Exchange Server 2013 environments) This parameter is the name of the mailbox endpoint of the user who is logged in. Use the Exchange PowerShell command, **whoami | Get-Mailbox | fl ExchangeGUID**, to determine the value. You must specify this parameter for Exchange Server 2013.

/TRACEFILE

This parameter is the file name that is used to store the output from tracing operations. By default, tracing is turned off. You can qualify the file name by specifying a drive and a full directory path. You must have write permissions for the user that runs the command.

Troubleshooting a MAPI error that prevents multiple mailboxes restoring in a Microsoft Exchange 2013 environment

When you restore multiple mailboxes on a server that is running Exchange Server 2013, the mailbox restore operation might partially fail and report a MAPI error.

About this task

In Exchange Server 2013, Client Throttling Policy (the **RcaMaxConcurrency** parameter), specifies how many concurrent connections you can maintain at one time. If you attempt to make more concurrent requests than the **RcaMaxConcurrency** parameter allows, the new connection attempt fails. However, the existing connections remain valid.

Procedure

Increase the **RcaMaxConcurrency** value for the logon user mailbox. For more information about this setting, see Microsoft documentation: Exchange 2013 Client Throttling ([http://technet.microsoft.com/en-us/library/bb232205\(v=exchg.150\).aspx](http://technet.microsoft.com/en-us/library/bb232205(v=exchg.150).aspx))

Troubleshooting issues with mailbox restore operations or Mailbox Restore Browser operations on remote systems

If you run complex mailbox restore operations on a remote system and need to query many mailboxes on the remote system, an out-of-memory exception can occur if there is not enough Microsoft Windows PowerShell memory to run the operation. To resolve the out-of-memory exception, increase the default memory value for the remote Microsoft Windows Power Shell session. You must increase both the machine-wide memory setting and the plug-in memory setting. After doing so, restart the WinRM service, and run the operation again.

About this task

You might get an out-of-memory exception when you attempt to run the following tasks:

- If you restore multiple mailboxes across multiple databases, you might see the following message:
Specified method is not supported.
- If you complete a mailbox restore task on the remote system, the list of mailboxes might not be displayed in the **Source** mailbox navigation tree of the MMC. You might see the following message:
Error: Processing data for a remote command failed
with the following error message:
The WSMAN provider host process did not return a proper response.
A provider in the host process may have behaved improperly.
For more information, see the about_Remote_Troubleshooting Help topic.
OperationStopped: (<Machine_Name>:String) [],
PSRemotingTransportExceptionJobFailure

Procedure

1. Increase the machine-wide memory setting.
 - a. At the Microsoft Windows PowerShell command line, navigate to `WSMan:\localhost\Shell\MaxMemoryPerShell1MB`.
 - b. Increase the value of **MaxMemoryPerShell1MB**.
2. Increase the memory setting for plug-ins.
 - a. At the Microsoft Windows PowerShell command line, navigate to `WSMan:\localhost\Plugin\Microsoft.PowerShell\Quotas\MaxMemoryPerShell1MB`.
 - b. Increase the value of **MaxMemoryPerShell1MB**.
3. Restart the WinRM service, and run the operation that you require again.

Example

To increase the maximum of memory that is allocated per shell to 4 GB, enter the following cmdlets at the Microsoft Windows PowerShell command line.

1. **Set-Item WSMAN:\localhost\Shell\MaxMemoryPerShell1MB 4096**

2. **Set-Item** WSMan:\localhost\Plugin\Microsoft.PowerShell\Quotas\
MaxMemoryPerShellMB 4096
3. **Restart-Service** winrm

Troubleshooting an SMTP restore issue that occurs when you restore email with large attachments in the Mailbox Restore Browser interface

If you restore an email with an attachment that is larger than 3 MB to an SMTP server, a Microsoft fix is required.

About this task

You might see the following error message:

```
QFD: System.Net.Mail - SmtplibClient class throws exceptions if file attachment  
is over 3 MB
```

Procedure

Resolve the issue by applying the fix that is available at this web page: Microsoft Connect Visual Studio and .NET Framework Downloads (<http://support.microsoft.com/kb/2183292>)

Troubleshooting a limitation with deleted mailbox history in the Mailbox Restore Browser interface

Data Protection for Exchange Server does not record the time when mailboxes are deleted.

About this task

After a mailbox is deleted, the **Available Database Backups** list in the Mailbox Restore Browser continues to list database backups that contained the mailbox prior to its deletion.

From the **Available Database Backups** list, ensure that the backup version that you select, for the restore task, contains a copy of the mailbox. If the database backup is completed after the mailbox was deleted, the mailbox is not available for the restore.

Troubleshooting VSS and SAN Volume Controller, Storwize V7000, or DS8000

If you experience VSS and SAN Volume Controller, Storwize V7000, or DS8000 problems, use these troubleshooting tips to help you discount some common configuration issues.

Procedure

1. Verify connectivity to the CIMOM (Common Information Model Object Manager) as follows:
 - a. Refer to your SAN Volume Controller, Storwize V7000, or DS8000 documentation.
 - b. Run the **IBMVCFG LIST** command. The default location is %Program Files%\IBM\Hardware Provider for VSS-VDS.
 - c. Issue the **IBMVCFG SHOWCFG** command to view the provider configuration information.

2. Verify CIMOM operational issues as follows:
 - a. If your backup or restore operation fails, check the IBMVSS.log file.
If the backup or restore failure is from a CIMOM failure, the log displays output similar to the following example:


```
Wed Jan 13 17:34:34.793 - Calling AttachReplicas
Wed Jan 13 17:34:35.702 - AttachReplicas: 909ms
Wed Jan 13 17:34:35.702 - returnValue: 34561
Wed Jan 13 17:34:35.718 - AttachReplicas returned: 34561
...
...
Wed Jan 13 17:34:35.779 - IBMVSS: AbortSnapshots
```

A return value of 0 means that it was successful.
 - b. To determine why a backup or restore operation failed, look at the log files. The files are generated by the CLI or graphical user interface (GUI), depending on how you run your operation. The log files might provide more information about the failure.
3. If the failure seems to be for a different reason than a CIMOM failure, verify your host configuration. Run the latest support levels of the software for SAN Volume Controller, Storwize V7000, or DS8000.
4. If you are unable to resolve these problems, provide the following information to IBM Support:
 - Information that is listed in the IBM Spectrum Protect diagnostic information section
 - HBA type, firmware, and driver levels
 - SDD version
 - SAN Volume Controller microcode version (if applicable)
 - DS8000 microcode version (if applicable)
 - Storwize V7000 microcode version (if applicable)
 - SAN Volume Controller or Storwize V7000 Master Console version (if applicable)
 - For DS8000, the CIM Agent version (if applicable)
 - IBMVSS.log
 - IBMVDS.log
 - Application Event Log
 - System Event Log

Resolving problems with IBM Support

Contact IBM Support for further assistance if you have a problem that you are unable to solve by applying maintenance fixes, reproducing the issue, or using the information in previous topics. IBM Support might request to see some or all of the trace and log files while investigating a problem that you report.

About this task

Also, you might be asked to set a trace the Data Protection client when using VSS technology, and then collect the log. IBM Support uses the information that is captured in the log file to trace a problem to its source or to determine why an error occurred.

Gathering trace and log files

Data Protection for Exchange Server uses several components. Each component is in its own directory along with its respective troubleshooting files. By using the Trace and Log Files view, you can easily view these files in a central location.

About this task

The following files are examples of the files that you can view, including default log and trace files:

Examples of Data Protection for Exchange Server default log and trace files:

- Installation directory: C:\Program Files\Tivoli\TSM\TDPEExchange
- dserror.log
- tdpexc.log
- *TraceFileExc.trc*

If the tdpexc.log is defined in a path other than the default C:\Program Files\Tivoli\TSM\TDPEExchange\tdpexc.log, the reports do not include the following information for scheduled backup and restore operations:

- Task completion
- Type of data protection activity
- Amount of data protection activity

The charts and reports display only information that is present in the default log file tdpexc.log.

Examples of VSS requestor default log and trace files:

- Installation directory: C:\Program Files\Tivoli\TSM\baclient
- dsmerror.log

Examples of IBM VSS provider for SAN Volume Controller, Storwize V7000, and DS8000 log files

- IBMVDS.log
- IBMVss.log

Procedure

1. When you encounter a problem in the Management Console, create trace files by using the Diagnostics property page.
 - a. Click **Properties > Diagnostics**, and click **Begin**.
 - b. Close the property page and reproduce the problem.
 - c. Open the Diagnostics property page and click **Stop**. Clicking the **Diagnostics** button is the preferred method for gathering information to send to your service representative. This method gathers all the information that is needed. Even if a problem occurs only on the command-line interface, command, you can always gather information by using the Automate tab. The log files are displayed in the Trace and Log Files view.
2. Click the trace or log file that you want to view. The contents of the file are displayed in the results pane.

Gathering installation log files to debug installation problems

If a problem occurs during the installation process, gather details about the installation process. You can create a detailed log file of the failed installation that can help IBMSupport to analyze and evaluate the problem.

About this task

The installation wizard collects log files for the installation process.

To help you quickly resolve problems, IBMSupport needs the following information:

- Operating system level
- Service pack
- Description of the hardware that is installed and operating in the production environment
- Installation package (from the DVD or downloaded) and level
- Any Windows event log that is relevant to the failed installation
- Windows services that were active during the failed installation (for example, antivirus software)
- Whether you are logged on to the local system console (not through a terminal server)
- Whether you are logged on as a local administrator, not a domain administrator (cross-domain installations are not supported)

Procedure

1. To create a detailed log file (setup.log) of the failed installation, enter the following command to run the setup program (setup.exe):

```
setup /v"1*v setup.log"
```
2. To view the log files, go to **Manage > Diagnostics > Trace And Log Files** on the navigation pane. The log files are listed in the upper window pane. When you select the log file, the file is displayed in the lower window pane.

Gathering traces for the Data Protection client when using VSS technology

You must gather traces for Data Protection for Exchange Server, the IBM Spectrum Protect application programming interface (API), and the DSMAGENT processes to ensure a good diagnosis of the Volume Shadow Copy Service (VSS) operation.

About this task

To diagnose Data Protection for Exchange VSS operational problems, gather these traces:

- Data Protection for Exchange trace
- IBM Spectrum Protect API trace
- DSMAGENT trace
- Exchange VSS Writer tracing

Procedure

1. Gather the Data Protection for Exchange trace as follows:

- a. To create the trace flag, issue the “/TRACEFILE” and “/TRACEFLAGS” command-line options with the following example command:

```
TDPEXCC BACKUP SG1 FULL /TRACEFILE=DPTRACE.TXT /TRACEFLAG=SERVICE
```
- b. Enable tracing for FlashCopy Manager. For information about how to enable tracing, see Viewing trace and log files for IBM Spectrum Protect Snapshot for Windows (http://www.ibm.com/support/knowledgecenter/SSERFV_8.1.0/fcm.win/t_fcm_diag_traceandlogfiles.html).
2. Gather the IBM Spectrum Protect API trace as follows: Enable tracing with the DP/Exchange dsm.opt file and the “TRACEFILE” and “TRACEFLAGS” keywords. The following text is an example of the entry in the DP/Exchange dsm.opt file:

```
TRACEFILE APITRACE.TXT  
TRACEFLAG SERVICE
```
3. Gather the DSMAGENT trace as follows: Enable tracing with the dsmagent (baclient) dsm.opt file and the “TRACEFILE” and “TRACEFLAGS” keywords. The following text is an example of the entry in the dsmagent (baclient) dsm.opt file:

```
TRACEFILE AGTTRACE.TXT  
TRACEFLAG SERVICE PID TID ENTER ALL_VSS SBRM RESTORE
```

The trace flag, in this instance, is ALL_VSS (you might need different traceflags, depending on the circumstance).

4. Gather the Exchange VSS Writer trace. Event logging is the only extra tracing that can be turned on. Complete these steps to modify the level of event logging for the Exchange Store Writer:
 - a. Open Microsoft Management Console (MMC).
 - b. Find the server object.
 - c. Right-click the server on which you want to increase the logging level and click **Properties** or **Manage Diagnostic Logging Properties**, depending on the Exchange version.
 - d. Click the **Diagnostics Logging** tab.
 - e. Expand the **MSExchangeIS** node in the **Services** pane and click **System**.
 - f. Click **Exchange writer** in the **Categories** pane and select the logging level.
 - g. Click **Apply** and then **OK** to close the Properties window.
5. Enable the Volume ShadowCopy service debug trace features in Windows. For information about enabling debug tracing, see the following web pages:
 - How to enable the Volume Shadow Copy service's debug tracing features in Microsoft Windows Server 2003 and Windows 2008 (<http://support.microsoft.com/kb/887013>)
 - Using Tracing Tools with VSS (<http://msdn.microsoft.com/en-us/library/windows/desktop/dd765233%28v=vs.85%29.aspx>)

Gathering information about Exchange with VSS before calling IBM

The Data Protection client depends on the operating system and the Exchange application. Collecting all the necessary information about the environment can significantly assist Support in determining the source of problem.

Procedure

Gather as much of the following information as possible before you contact IBM Support:

- The exact level of the Windows operating system, including all service packs and hotfixes that were applied.
- The exact level of the Exchange Server, including all service packs and hotfixes that were applied.
- The exact level of Data Protection for Exchange with Volume Shadow Copy Service (VSS) Backup/Restore support.
- The exact level of the IBM Spectrum Protect API.
- The exact level of the IBM Spectrum Protect server.
- The exact level of the IBM Spectrum Protect backup-archive client.
- The exact level of the IBM Spectrum Protect storage agent (if LAN-free environment).
- The IBM Spectrum Protect server platform and operating system level.
- The output from the IBM Spectrum Protect server **QUERY SYSTEM** command.
- The output from the Data Protection for Exchange **TDPEXCC QUERY EXCHANGE** command.
- The device type (and connectivity path) of the Exchange databases and logs.
- (SAN only) The specific hardware that is being used. For example: HBA, driver levels, microcode levels, SAN Volume Controller or Storwize V7000 levels, DS8000 hardware details.
- Permissions and the name of the user ID being used to run backup and restore operations.
- The name and version of antivirus software.
- (SAN only) The VSS hardware provider level.
- The VSS hardware provider log files. See the documentation of the specific VSS hardware provider on how to enable tracing and collect the trace log files.
- (SAN only) The IBM CIM agent level for DS8000, SAN Volume Controller, or Storwize V7000.
- A list of vendor-acquired Exchange applications that are running on the system.
- A list of other applications that are running on the system.
- A list of the steps that are needed to re-create the problem (if the problem can be re-created).
- If the problem cannot be re-created, list the steps that caused the problem.
- Does the problem occur on other Exchange servers?

Gathering information about Exchange Server with VSS before you call IBM

You can collect several log files and other data for Data Protection for Exchange Server server diagnosis.

About this task

The Management Console (MMC) automatically collects information in a package file, which you can send to Support. To collect this information manually, refer to the following file list.

Procedure

1. Gather as many of the following files as possible before you contact IBM Support.

- The contents of the C:\Program Files\Tivoli\tsm\baclient\adsm.sys\vss_staging directory and subdirectories. Gather the appropriate directories if you are using the VSSALTSTAGINGDIR option.
 - The Data Protection for Exchange Server configuration file. The default configuration file is tdpexc.cfg.
 - The Data Protection for Exchange Server IBM Spectrum Protect API options file. The default options file is dsm.opt.
 - The IBM Spectrum Protect registry hive export.
 - The Exchange Server registry hive export.
 - The IBM Spectrum Protect Server activity log. The Data Protection client logs information to the server activity log. An IBM Spectrum Protect administrator can view this log for you if you do not have an IBM Spectrum Protect administrator user ID and password.
 - If the Data Protection client is configured for LAN-free data movement, also collect the options file for the IBM Spectrum Protect storage agent. The default name for this file is dsmsta.opt.
 - Any screen captures or command-line output of failures or problems.
2. Gather the following IBM Spectrum Protect log files, which can indicate the date and time of a backup, the data that is backed up, and any error messages or completion codes that might help to determine your problem:
- The Data Protection for Exchange Server log file. The default location of this file is C:\Program Files\Tivoli\TSM\TDPEExchange\tdpexc.log
 - The IBM Spectrum Protect API Error log file. The default location of this file is C:\Program Files\Tivoli\TSM\TDPEExchange\dsierror.log
 - The DSMAGENT error log file. The default location of this file is C:\Program Files\Tivoli\TSM\baclient\dsmerror.log
 - The dsmcrash.dmp and DSMAGENT crash log file, if requested. The default location is C:\Program Files\Tivoli\TSM\baclient\dsmcrash.log.

Important: The Windows event log receives information from the Exchange Server and many different components that are involved during a Volume Shadow Copy Service (VSS) operation. Export the event log to a text file format.

3. Use the Data Protection for Exchange console to list the events that originate by Data Protection for Exchange. Select **Dashboard - ServerName > Diagnostics > System Information** and double-click the dpevents.ps1 script in the PowerShell section of the **System Information** page. On Windows Server 2008 or later, You can use PowerShell scripting to list the events information. You can also use the export function from within the Event Viewer to do this function. The utility, by default, produces a tabular listing of all event log records in three sections (one section per event log type). Specify the type of event log you require by using one of the following /L parameters:

```

/L Application
/L Security
/L System

```

The following example generates output only for the application and system event logs:

```

cscript c:\windows\system32\eventquery.vbs /L Application >eq_app.out
cscript c:\windows\system32\eventquery.vbs /L System >eq_sys.out

```

You can use the /V parameter to receive detailed (verbose) output:

```
cscript c:\windows\system32\eventquery.vbs /V >eq.out
cscript c:\windows\system32\eventquery.vbs /L System /V >eq_sys.out
```

You can use the /FO parameter to specify tabular, list, or comma-separated (CSV) output. You can use the following methods to specify the output:

```
/FO TABLE
/FO LIST
/FO CSV
```

The default format is TABLE. The LIST output puts each column of the record on a separate line. This technique is similar to how the IBM Spectrum Protect administrator's command-line interface (CLI) displays output when it is too wide for tabular display. The CSV output can be loaded into a spreadsheet or database tool for easier viewing. The following example generates a detailed CSV file of the application log:

```
cscript c:\windows\system32\eventquery.vbs /L Application /FO CSV /V >eq_app.out
```

You can get more help information for the tool by using the following example:

```
cscript c:\windows\system32\eventquery.vbs /?
```

4. To increase the number of events that are logged by the Microsoft Exchange Writer, use the **Set-EventLogLevel** PowerShell cmdlet command. For more information about the **Set-EventLogLevel** PowerShell cmdlet command, see the Microsoft documentation.

The following VSS provider log files can also be helpful, if applicable:

- System Provider - (Windows Event Log)
- IBM System Storage SAN Volume Controller, IBM Storwize V7000, or IBM System Storage DS8000 series - %Program Files%\IBM\Hardware Provider for VSS\IBMVss.log
- NetApp - %Program Files%\SnapDrive*.log
- XIV - zip up all of the files in the C:\Windows\Temp\xProvDotNet directory

Viewing and modifying system information

You can view and edit scripts that provide information about system components including, for example, Windows-related services for Data Protection for Microsoft Exchange Server, Windows event log entries, and Volume Shadow Copy Service (VSS) information.

About this task

The System Information view is extensible. You can take advantage of this flexibility to add and share customized scripts.

Procedure

1. Open the System Information view as follows:
 - a. Click **Diagnostics** in the results pane of the welcome page.
 - b. Double-click **System Information** in the results pane. A list of scripts is displayed in the results pane of the System Information view. The types of scripts that are displayed are PowerShell scripts, Windows Management Instrumentation scripts, and IBM Spectrum Protect scripts.
2. Add, update, or delete your scripts, as follows:

Action	Steps
Add your own scripts	<ol style="list-style-type: none"> 1. Click New in the Actions pane. 2. If you want to copy your scripts directly to the ProgramFiles\Tivoli\FlashCopyManager\Scripts directory, make sure that your scripts follow these extension requirements: <ul style="list-style-type: none"> • PowerShell scripts: <i>filename.ps1</i> • Windows Management Instrumentation (WMI) scripts: <i>filename.wmi</i> • IBM Spectrum Protect scripts: <i>filename.tsm</i> IBM Spectrum Protect Snapshot uses the file type extension to determine how to run the script.
View or edit an existing script	<ol style="list-style-type: none"> 1. From the list of script files in the results pane, select the name of a script that you want to view or edit. Tip: The name of the script is displayed in the Actions pane. Click the name of the script in the Actions pane to reveal or hide a list of actions to process. 2. To open the script file for viewing or editing, click Command Editor in the Actions pane. 3. View or edit the script. 4. Click OK to save your changes, or click Cancel to exit the System Information Command Editor without saving any changes.
Delete a script	<ol style="list-style-type: none"> 1. From the list of script files in the results pane, select the name of a script that you want to delete. Tip: The name of the script is displayed in the Actions pane. Click the name of the script in the Actions pane to reveal or hide a list of actions to process. 2. Click Delete in the Actions pane.

Emailing files to IBM Support

You can send diagnostic information to IBM Support.

Before you begin

About this task

The Email Support files feature collects all detected configuration, option, system information, trace, and log files. It also collects information about services, operating systems, and application versions. These files are compressed and then attached in an email.

Procedure

1. Start the Microsoft Management Console (MMC).
2. Click **Diagnostics** in the results pane of the welcome page.
3. Click the **E-Mail Support files** icon in the Action pane.
4. Enter the required information in the various fields and click **Done**. The information is sent to the designated support personnel and the dialog closes.

Results

Files are collected, compressed, and stored in the flashcopymanager\problemdetermination folder. The files are deleted and replaced each time that you email the support files. If the Email feature is not configured, or is blocked by a

firewall, or if the files are large, use another method to transfer them. You can copy the files directly from the `flashcopymanager\problemdetermination` folder and transfer the files to another site by using another method such as FTP.

Online IBM support

Multiple online support resources are available for you to use.

The following list identifies where you can find information online:

- Tivoli Storage Manager wiki (<https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Storage%20Manager>).
- Storage Management community on Service Management Connect (<https://www.ibm.com/developerworks/servicemanagement/sm/index.html>).
- IBM Spectrum Protect for Mail (www.ibm.com/software/products/en/spectrum-protect-for-mail). Enter the search term to narrow the search criteria for your support requirements. Examples of search terms that you might use include an authorized program analysis report (APAR) number, release level, or operating system.

Chapter 7. Performance tuning

Many factors can affect the backup and restore performance of your Exchange Server.

Some of these factors, such as hardware configuration, network type, and capacity, are not within the scope of Data Protection for Exchange Server. Some options that are related to Data Protection for Exchange Server can be tuned for optimum performance. In addition, the following issues affect performance:

- Database backups to local shadow volumes eliminates the transfer of data to the IBM Spectrum Protect server.
- During VSS backup processing, the consistency of the database backup is verified. Therefore, backup processing time can be significant. You can improve backup processing performance when you specify the **/SKIPINTEGRITYCHECK** options with the **backup** command to bypass integrity checking.

Restriction: If you bypass integrity checking, the backup that is stored on IBM Spectrum Protect server might not be valid and data loss can occur, unless the database that you are backing up is in a Database Availability Group (DAG) environment and has at least two valid copies (one active and one passive copy).

- The time that is required to complete a snapshot, ranges from seconds to minutes, depending on the type of VSS provider used. Depending on the size of the database and log files, integrity checking can delay the completion of the backup.
- Backup-archive client settings can affect performance when you back up data to the IBM Spectrum Protect server.
- Performing Data Protection for Exchange Server VSS backups from an Exchange Server DAG passive copy can offload I/O and possibly processor resources from the production server.

If the update for IBM Spectrum Protect server APAR IC86558 is not applied, apply the update.

For VSS backups, the **RESOURCEUTILIZATION** client option is also important. This option increases or decreases the capacity of the client to create multiple sessions. The higher the value, the more sessions the client can start. The range for the option is from 1 to 10.

If you run multiple backups in parallel, stagger the backup times by several minutes. The staggered backup times ensure that the snapshots are not created at the same time. When you use VSS, only one snapshot set can be created at a time.

Related tasks:

“Specifying configuration parameters for IBM Spectrum Protect” on page 50

“Managing Exchange Database Availability Group members by using a single policy” on page 23

“Binding backups to a policy” on page 21

LAN-free data movement

Running Data Protection for Exchange Server in a LAN-free environment means that data can be directly sent to storage devices.

When you implement a LAN-free environment, data bypasses potential network congestion. However, you must be properly equipped to operate in a LAN-free environment. For more information about setting up a LAN-free environment, see LAN-free client-data backup: Scenario (http://www.ibm.com/support/knowledgecenter/SSSQZW_7.1.1/com.ibm.itsm.sta.doc/c_scenario_lanfree.html).

In addition to specific LAN-free requirements, you must specify the following options. For VSS backups, specify these options in the backup-archive client options file.

enablelanfree yes

This option specifies whether to enable an available LAN-free path.

lanfreecommmethod

Specifies a communication protocol.

lanfreetcport

Specifies the TCP/IP port number where the IBM Spectrum Protect Storage Agent is listening.

lanfreetcpserveraddress

Specifies the TCP/IP address for the IBM Spectrum Protect Storage Agent.

For more information about these options, see Installing and configuring the client (https://www.ibm.com/support/knowledgecenter/SSSQZW_7.1.1/com.ibm.itsm.sta.doc/t_extlib_inst_client.html).

Chapter 8. Reference

Reference topics provide information related to Data Protection for Microsoft Exchange Server. Topics include the backup and restore commands that you can issue at the command-line interface as an alternative to using Microsoft Management Console (MMC) and frequently asked questions.

Support for Microsoft Exchange 2016

With IBM Spectrum Protect for Mail: Data Protection for Microsoft Exchange Server version 7.1.4.2, features that support Microsoft Exchange 2016 were added and you can now protect and manage your Microsoft Exchange 2016 environment.

Mailbox filter options

When you restore an individual mailbox, you can use mailbox filters to identify individual messages to restore. With Microsoft Exchange 2016, the Folder Name filter option is supported.

For example, to restore a folder that is named "folder A" in the mailbox "MailboxA", run the following command:

```
tdpexcc restoremailbox "MailboxA" /MailboxFilter="folder, folderA"
```

Individual mailbox restore options

You can restore individual mailbox items from database backups. The following table describes the differences between the mailbox restore features that are supported with Microsoft Exchange 2010 and 2013 and the features that are supported with Microsoft Exchange 2016.

Table 20. Mailbox restore options

Feature	Description	Exchange 2010, 2013	Exchange 2016
Mailbox restore	Mailbox restore browser	The Mailbox Restore Browser is available under the Protect and Recover Data node on the MMC	Not supported with Microsoft Exchange 2016.
	Non-Unicode mailbox restore	The Restore Mailbox to non-Unicode PST file option is available for selection in the Actions pane.	Not supported with Microsoft Exchange 2016. Instead, the mailbox is automatically restored to a Unicode PST file.

Temporary mailbox folder cleanup

When a mailbox is successfully restored, with Data Protection for Microsoft Exchange Server version 7.1.4.2, the temporary mailbox folder that was created during the restore operation can be deleted automatically.

Note: To enable automatic temporary folder deletion with Microsoft Exchange 2016, you must log on as an Exchange Server administrator and ensure that the **ApplicationImpersonation** role is assigned to your user. This role is not enabled by default.

Message application programming interface (MAPI) client and collaboration data objects (CDO)

The MAPI/CDO library is not supported with Microsoft Exchange 2016. The MAPI/CDO Changes table describes the impact of this change in your Data Protection for Microsoft Exchange Server solution.

Table 21. MAPI/CDO changes

Feature	Description	Exchange 2010, 2013	Exchange 2016
MAPI Settings	MAPI Settings property page	The MAPI Settings property page is available under the Protect and Recover Data node on the MMC.	Not supported with Microsoft Exchange 2016.
	MAPI configuration checks that use the configuration wizard	When you use the configuration wizard to configure IBM Spectrum Protect Snapshot for Exchange Server on the MMC, the system automatically runs a number of checks to verify that the Microsoft Exchange Server MAPI client and CDO are correctly installed.	Not supported with Microsoft Exchange 2016, therefore MAPI configuration checks are skipped.

Command-line overview

The name of the Data Protection for Exchange Server command-line interface is **tdpexcc.exe**. This program is in the directory where Data Protection for Exchange Server is installed.

Command-line parameter characteristics

The command-line parameters have the following characteristics:

- Positional parameters do not include a leading slash (/) or dash (-).
- Optional parameters can display in any order after the required parameters.
- Optional parameters begin with a forward slash (/) or a dash (-).
- Minimum abbreviations for keywords are indicated in uppercase text.
- Some keyword parameters require a value.
- For those keyword parameters that require a value, the value is separated from the keyword with an equal sign (=).
- If a parameter requires more than one value after the equal sign, the values are separated with commas.
- Each parameter is separated from the others by using spaces.
- If a parameter value includes spaces, the value must be enclosed in double quotation marks.
- A positional parameter can display only once per command invocation.

Command-line interface help

Issue the **tdpexcc ?** or **tdpexcc help** command to display help for the command-line interface. You can see more specific help for commands by entering a command like the following example: **tdpexcc help backup**, where **backup** is an example of a command.

Related tasks:

“Starting the Data Protection for Exchange Server command-line interface” on page 95

Backup command

Use the **backup** command to run Exchange Server database backups from the Exchange Server to IBM Spectrum Protect server storage.

Microsoft Exchange Server considers the wildcard character (*) to be an invalid character when used in database names. Databases that contain the wildcard character (*) in their name are not backed up. When a full VSS snapshot backup (created for back up to local shadow volumes) is run, the backup remains active until the backup version is expired on the IBM Spectrum Protect server according to the defined server policy. As a result, different active backups can exist at the same time:

- VSS local (full)
- VSS local (copy)
- VSS IBM Spectrum Protect server (full)
- VSS IBM Spectrum Protect server (copy)

The Exchange database file size might increase as a result of increased database commitments that are triggered by backup operations. This behavior is standard for the Microsoft Exchange server.

For IBM SAN Volume Controller and IBM Storwize V7000 storage systems, only one backup is allowed to occur while the background copy process is pending. A new backup is not started until the background copy process for the previous backup is completed. As a result, local backups for IBM SAN Volume Controller and IBM Storwize V7000 storage systems must be initiated at a frequency that is greater than the time required for the background copy process to complete.

Data Protection for Exchange Server supports the following types of backup:

Full Back up the entire database and transaction logs. If a successful integrity check and backup are obtained, the Exchange Server truncates the committed log files.

Incremental

Back up the transaction logs. If a successful integrity check and backup are obtained, the Exchange Server deletes the committed log files.

Differential


Back up the transaction logs but do not delete them.

Copy Back up the entire database and transaction logs. Do not delete the transaction logs.

Backup syntax

To view available options and truncation requirements, use the **backup** command.

TDPEXCC command



Backup positional parameters

Positional parameters immediately follow the **backup** command and precede the optional parameters.

The following positional parameters specify the object to back up:

*** | *db-name1, ...,db-nameN backuptype***

***** Back up all databases sequentially.

db-name

Back up the specified database. If separated by commas, ensure that there is no space between the comma and the database name. If any database name contains blanks, enclose the database name in double quotation marks. The database name is case sensitive.

The following positional parameters specify the type of backup to run:

FULL | COPY | INCRemental | DIFFerential

FULL Back up the entire database and transaction logs, and if a successful backup is obtained, truncate the transaction logs.

COPY Back up the entire database and transaction logs, do not truncate the transaction logs.

INCRemental

Back up the transaction logs, and if a successful backup is obtained, truncate the transaction logs.

DIFFerential

Back up the transaction logs but do not truncate them.

Backup optional parameters

Optional parameters follow the **backup** command and positional parameters.

/BACKUPDESTination=TSM | LOCAL | BOTH

Use the ***/BACKUPDESTination*** parameter to specify the location where the backup is stored.

You can specify:

TSM The backup is stored on IBM Spectrum Protect server storage only. This option is the default value.

LOCAL The backup is stored on local shadow volumes only.

BOTH The backup is stored on IBM Spectrum Protect server storage and local shadow volumes.

/CONFIGfile=configfilename

Use the ***/CONFIGfile*** parameter to specify the name (*configfilename*) of the Data Protection for Exchange Server configuration file that contains the values to use for a **backup** operation.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the Data Protection for Exchange Server installation directory is used. If the ***/CONFIGfile*** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `tdpexc.cfg`.

If the *configfilename* variable includes spaces, enclose the entire ***/CONFIGfile*** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\file.cfg"
```

DAGNODE=nodename

Specify the node name that you want to use to back up the databases in an Exchange Server Database Availability Group. With this setting, backups from all Database Availability Group members that are configured to use the DAG node are backed up to a common file space on the IBM Spectrum Protect server. The database copies are managed as a single entity, regardless of which Database Availability Group member they are backed up from. This setting can prevent IBM Spectrum Protect Snapshot for Microsoft Exchange Server from making too many backups of the same database.

/EXCLUDEDAGACTive

Use the **/EXCLUDEDAGACTive** parameter to exclude databases from the backup if they belong to a Database Availability Group and are an active database copy.

/EXCLUDEDAGPASsive

Use the **/EXCLUDEDAGPASsive** parameter to exclude the databases from backup if they belong to a Database Availability Group and are a passive database copy.

/EXCLUDENONDAGDbs

Use the **/EXCLUDENONDAGDbs** parameter to exclude the databases from backup if they do not belong to a Database Availability Group.

/EXCLUDEDB=db-name,...

Use the **/EXCLUDEDB** parameter to exclude the specified databases from the backup operation.

/LOGFile=logfilename

Use the **/LOGFile** parameter to specify the name of the activity log file that is generated by Data Protection for Exchange Server.

The *logfilename* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfilename* variable can include a fully qualified path. However, if no path is specified, the log file is written to the Data Protection for Exchange Server installation directory.

If the *logfilename* variable includes spaces, enclose the entire **/LOGFile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\mytdpexchange.log"
```

If the **/LOGFile** parameter is not specified, log records are written to the default log file, *tdpexc.log*.

The **/LOGFile** parameter cannot be turned off, logging always occurs.

When you use multiple simultaneous instances of Data Protection for Exchange Server to complete operations, use the **/LOGFile** parameter to specify a different log file for each instance used. This function directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=numdays | No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the

option that you set for the **/LOGPrune** parameter, some days of data are saved. By default, 60 days of log entries are saved. The option **No** can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify **no**, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/MINimumbackupinterval=numberOfMinutes

If you are scheduling the backup of databases in an Exchange Server Database Availability Group, specify the minimum amount of time, in minutes, before a backup of another copy of the same Database Availability Group database can begin. The range is 1 - 9999.

Setting this parameter specifies that only one database copy can be backed up within a time frame. This option prevents all of the members in a Database Availability Group from backing up the database, which would be redundant and invalidate the IBM Spectrum Protect storage management policy.

/MOUNTwait=Yes|No

Use the **/MOUNTwait** parameter to specify whether Data Protection for Exchange Server is to wait for removable media to mount (such as tapes or CDs) or to stop the current operation. This situation occurs when the IBM Spectrum Protect server is configured to store backup data on removable media and waits for a required storage volume to be mounted.

You can specify:

Yes Wait for tape mounts. This option is the default.

No Do not wait for tape mounts.

/OFFLOAD

Specify this parameter to complete the integrity check and backup of files to IBM Spectrum Protect on the system that is specified by the **remotedsmagentnode** instead of the local system. This parameter is only valid when **/backupdestination=TSM**. This parameter requires a VSS

provider that supports transportable shadow copies. You cannot specify the parameter with the default Windows VSS System Provider.

/PREFERDAGPASSive

If you are scheduling the backup of databases in an Exchange Server Database Availability Group, set this parameter to back up a passive database in an Exchange Server Database Availability Group unless no valid passive copy is available. If no valid passive copy is available, the backup is created from the active database copy.

/Quiet This parameter prevents status information from being displayed. The level of information that is written to the activity log is not affected.

/SKIPINTEGRITYCHECK

Use the **/SKIPINTEGRITYCHECK** parameter to specify whether Data Protection for Exchange Server bypasses the integrity checking of databases and log files, or automatically runs the integrity checking of databases and log files.

You can specify the following values:

- No** Run integrity checking to verify that all database and log files do not contain integrity issues. This option is the default.
- Yes** Bypass integrity checking of all database and log files during backup processing.

SKIPDBCHECKIFHEALTHY

Bypass integrity checking of database files only if at least two valid copies of a database (one active and one passive copy) exist in a Database Availability Group (DAG).

SKIPDBANDLOGCHECKIFHEALTHY

Bypass integrity checking of all database and log files during backup processing only if at least two valid copies of a database (one active and one passive copy) exist in a DAG.

Attention: If you do not specify a value with the **/SKIPINTEGRITYCHECK** parameter, integrity checking of database and log files is bypassed. If you bypass integrity checking, the backup that is stored on IBM Spectrum Protect server might not be valid, or data loss can occur.

/TSMNODE=*tsmnode*

Use the *tsmnode* variable to refer to the IBM Spectrum Protect node name that Data Protection for Exchange Server uses to log on to the IBM Spectrum Protect server.

You can store the node name in the IBM Spectrum Protect options file (*dsm.opt*). This parameter overrides the value in the IBM Spectrum Protect options file if **PASSWORDACCESS** is set to PROMPT. This parameter is not valid when **PASSWORDACCESS** is set to GENERATE in the options file.

/TSMOPTFile=*tsmoptfilename*

Use the *tsmoptfilename* variable to identify the Data Protection for Exchange Server options file.

The file name can include a fully qualified path name. If no path is specified, the directory where Data Protection for Exchange Server is installed is searched.

If the *tsmoptfilename* variable includes spaces, enclose the entire **/TSMOPTFile** parameter entry in double quotation marks. For example:

```
/TSMOPTFile="c:\Program Files\file.opt"
```

The default is `dsm.opt`.

`/TSMPassword=tsmpassword`

Use the `tsmpassword` variable to refer to the IBM Spectrum Protect password that Data Protection for Exchange Server uses to log on to the IBM Spectrum Protect server.

If you specified **PASSWORDACCESS GENERATE** in the Data Protection for Exchange Server options file (`dsm.opt`), you do not need to supply the password here because the one that is stored in the registry is used. However, to store the password in the registry, you must specify the IBM Spectrum Protect password the first time that Data Protection for Exchange Server connects to the IBM Spectrum Protect server.

If you do specify a password with this parameter when **PASSWORDACCESS GENERATE** is in effect, the command-line value is ignored unless the password for this node is not yet stored in the registry. In that case, the specified password is stored in the registry and used when you run this command.

If **PASSWORDACCESS PROMPT** is in effect, and you do not specify a password value on the command line, then you are prompted for a password.

The IBM Spectrum Protect password that Data Protection for Exchange Server uses to log on to the IBM Spectrum Protect server can be up to 63 characters in length.

`/UpdateMailboxInfoOnly`

Specify the `/UpdateMailboxInfoOnly` parameter to update only the mailbox history information in Microsoft Exchange Server database backups, for example:

```
tdpexcc backup DB1 full /UpdateMailboxInfoOnly
```

where `DB1` is the database name, and `full` is the type of database backup.

Restriction: This parameter does not back up the Exchange Server database.

Examples: backup command

The examples in this topic show you how to use the **backup** command.

This example shows how to run a full VSS backup of exactly one copy of a database that contains multiple copies in an Exchange Server Database Availability Group (DAG). The command instructs Data Protection for Exchange Server to back up only the database `KEENV1_M_DB1` if a minimum of 60 minutes passes since the latest backup of the database, and if no other member in the `FCMDAG2` Database Availability Group is backing it up. Include this command in a command script (for example, `c:\backup.cmd`). Then, define an IBM Spectrum Protect schedule that starts this command script, and associate all DAG nodes to this schedule.

```
tdpexcc backup KEENV1_M_DB1 full /minimumbackupinterval=60
```

This example shows how to run a full VSS backup of one valid passive copy of a database that contains multiple copies in an Exchange Server Database Availability Group (DAG). If a valid passive copy is not available, the active database copy is backed up. The command instructs Data Protection for Exchange Server to back up only the passive copy of database `KEENV1_M_DB1` if a minimum of 60 minutes passes since the latest backup of the database, and if no other member in the `FCMDAG2` Database Availability Group is backing it up. If no passive database copy

is available, back up the active database copy. Include this command in a command script (for example, c:\backup.cmd). Define an IBM Spectrum Protect schedule that starts this command script, and associate all DAG nodes to this schedule.

```
tdpexcc backup KEENVMI_M_DB1 full /minimumbackupinterval=60 /preferdagpassive
```

Changetsmpassword command

To change the IBM Spectrum Protect password that is used by Data Protection for Exchange Server, use the **changetsmpassword** command. The password is used to log on to the IBM Spectrum Protect server.


If you do not enter the old and new passwords, Data Protection for Exchange Server prompts you for the old and new passwords. Data Protection for Exchange Server does not display the password on the screen.

The IBM Spectrum Protect password that Data Protection for Exchange Server uses to log on to the IBM Spectrum Protect server can be up to 63 characters in length.

Changetsmpassword syntax

Use the **changetsmpassword** command syntax diagrams as a reference to view available options and truncation requirements.

TDPEXCC command



Changetsmpassword positional parameters

Positional parameters immediately follow the **changetsmpassword** command and precede the optional parameters.

oldpassword newpassword verifypassword

oldpassword

Specifies the current password that is used by Data Protection for Exchange Server.

newpassword

Specifies the new password that is used by Data Protection for Exchange Server.

verifypassword

Specifies the new password again for verification.

Changetsmpassword optional parameters

Optional parameters follow the **changetsmpassword** command and positional parameters.

/CONFIGfile=*configfilename*

Use the **/configfile** parameter to specify the name of the Data Protection for Exchange Server configuration file that contains the values for the Data Protection for Exchange Server configuration options.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the Data Protection for Exchange Server installation directory is used. If the **/configfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `tdpexc.cfg`.

If the *configfilename* variable includes spaces, enclose the entire **/configfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\file.cfg"
```

/LOGFile=*logfile*

Use the **/logfile** parameter to specify the name of the activity log file that is generated by Data Protection for Exchange Server.

The *logfile* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfile* variable can include a fully qualified path. However, if no path is specified, the log file is written to the Data Protection for Exchange Server installation directory.

If the *logfile* variable includes spaces, enclose the entire **/logfile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\mytdpexchange.log"
```

If the **/logfile** parameter is not specified, log records are written to the default log file, `tdpexc.log`.

The **/logfile** parameter cannot be turned off, logging always occurs.

When you use multiple simultaneous instances of Data Protection for Exchange Server to perform operations, use the **/logfile** parameter to specify a different log file for each instance used. This function directs logging for each instance to a different log file and prevents interspersed log file records.

Attention: Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=*numdays* | **No**

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days

of data are saved. By default, 60 days of log entries are saved. The option `No` can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify `no`, the log file is not pruned.
- If you do not specify `/LOGPrune`, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify `/LOGPrune`, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify `/LOGPrune` without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the `/LOGFile` parameter or **logfile** setting.

`/TSMNODE=tsmnodename`

Use the *tsmnodename* variable to refer to the IBM Spectrum Protect node name that Data Protection for Exchange Server uses to log on to the IBM Spectrum Protect server.

You can store the node name in the IBM Spectrum Protect options file (`dsm.opt`). This parameter overrides the value in the IBM Spectrum Protect options file if **PASSWORDACCESS** is set to `PROMPT`. This parameter is not valid when **PASSWORDACCESS** is set to `GENERATE` in the options file.

`/TSMOPTFile=tsmoptfilename`

Use the *tsmoptfilename* variable to identify the Data Protection for Exchange Server options file.

The file name can include a fully qualified path name. If no path is specified, the directory where Data Protection for Exchange Server is installed is searched.

If the *tsmoptfilename* variable includes spaces, enclose the entire `/tsmoptfile` parameter entry in quotation marks. For example:

```
/TSMOPTFile="c:\Program Files\file.opt"
```

The default is `dsm.opt`.

Example: changetsmpassword command

The following example changes the IBM Spectrum Protect password that is used by Data Protection for Exchange Server:

```
tdpexc changetsmpassword oldpw newpw newpw
```

Delete backup command

To delete a VSS backup of an Exchange Server database, use the **delete backup** command.

To complete a Data Protection for Exchange Server delete backup, you must have local registry rights for all versions of Exchange Server. When a full VSS snapshot backup is completed, the backup remains active until the backup version is deleted by issuing the **delete backup** command, or until the backup version is expired by VSS according to the defined policy. Two different active backups can exist at the same time:


- Full backup, and any associated incremental backups and differential backups.
- Copy backup, and any associated incremental backups and differential backups.


When you delete an active full or copy backup, the state of the previous active full or copy backup changes from inactive to active. However, the current active incremental or differential backup is not deleted. The backup seems to be erroneously associated with the newly active full or copy backup. Also, the incremental or differential backup (associated with the previous inactive full or copy backup that changed to active) remains inactive. This inactive incremental or differential backup might not display in the query output unless the **/all** parameter is specified with the **query fcm** command.

Delete Backup syntax

Use the **delete backup** command syntax diagrams as a reference to view available options and truncation requirements.

TDPEXCC command





Delete Backup positional parameters

Positional parameters immediately follow the **delete backup** command and precede the optional parameters.

The following positional parameters specify the backup to delete:

* | *database_name*

* Delete the active backups of all databases.

database_name

Delete a backup of the specified database name. The active backup is deleted unless you specify a different backup with the **/object** parameter. When multiple active incremental backups exist, the **/object** parameter must be specified with the **delete** command.

Multiple entries are separated by commas. If separated by commas, ensure that there is no space between the comma and the database name. If any database name contains blanks, enclose the database name in double quotation marks.

Attention:

- Be careful to delete only the backups you want.
- Deleting incremental or differential backups can cause loss of recovery points.
- Deleting a full backup might cause incremental or differential backups to remain in a suspended state and are considered useless without a corresponding full backup.

The following positional parameters specify the type of delete backup to perform:

FULL | **COPY** | **INCREMENTAL** | **DIFFERENTIAL**

FULL Delete full type backups.

COPY Delete copy type backups.

INCREMENTAL

Delete incremental type backups.

DIFFerential

Delete differential type backups.

Delete Backup optional parameters

Optional parameters follow the **delete backup** command and positional parameters.

/BACKUPDESTination=TSM|LOCAL

Use the **/backupdestination** parameter to specify the location from where the backup is to be deleted. The default is the value (if present) specified in the Data Protection for Exchange Server preferences file (*tdpexc.cfg*). If no value is present, the backup is deleted from IBM Spectrum Protect server storage.

You can specify:

TSM The backup is deleted from IBM Spectrum Protect server storage. This option is the default value.

LOCAL The backup is deleted from the local shadow volumes.

/CONFIGfile=configfilename

Use the **/configfile** parameter to specify the name (*configfilename*) of the Data Protection for Exchange Server configuration file that contains the values to use for a **delete backup** operation.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the Data Protection for Exchange Server installation directory is used. If the **/configfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is *tdpexc.cfg*.

If the *configfilename* variable includes spaces, enclose the entire **/configfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\file.cfg"
```

DAGNODE=nodename

Specify the node name that you want to use to back up the databases in an Exchange Server Database Availability Group. With this setting, backups from all Database Availability Group members that are configured to use the DAG node are backed up to a common file space on the IBM Spectrum Protect server. The database copies are managed as a single entity, regardless of which Database Availability Group member they are backed up from. This setting can prevent IBM Spectrum Protect Snapshot for Microsoft Exchange Server from making too many backups of the same database.

/FROMEXCSERVer=server-name

Use the **/fromexcserver** parameter to specify the name of the Exchange Server where the original backup was performed.

The default is the local Exchange Server. However, you must specify the name if the Exchange Server is not the default.

/LOGFile=logfilename

Use the **/logfile** parameter to specify the name of the activity log file that is generated by Data Protection for Exchange Server.

The *logfilename* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The

logfile variable can include a fully-qualified path. However, if no path is specified, the log file is written to the Data Protection for Exchange Server installation directory.

If the *logfile* variable includes spaces, enclose the entire **/logfile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\mytdpexchange.log"
```

If the **/logfile** parameter is not specified, log records are written to the default log file, *tdpexc.log*.

The **/logfile** parameter cannot be turned off, logging always occurs.

When using multiple simultaneous instances of Data Protection for Exchange Server to perform operations, use the **/logfile** parameter to specify a different log file for each instance used. This directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=*numdays* | No

Use the **/logprune** parameter to disable log pruning or to explicitly request that the log be pruned for one command run. By default, log pruning is enabled and performed once per day. The *numdays* variable represents the number of days to save log entries. By default, 60 days of log entries are saved in the pruning process. You can use the Data Protection for Exchange Server GUI or the **set** command to change the defaults so that log pruning is disabled, or so that more or less days of log entries are saved. If you use the command line, you can use the **/logprune** parameter to override these defaults. When the value of the **/logprune** variable *numdays* is a number in the range 0 - 9999, the log is pruned even if log pruning has already been performed for the day.

Changes to the value of the **timeformat** or **dateformat** parameter can result in the log file being pruned unintentionally. If the value of the **timeformat** or **dateformat** parameter has changed, prior to issuing a Data Protection for Exchange Server command that might prune the log file, perform one of the following actions to prevent the log file from being pruned:

- Make a copy of the existing log file.
- Specify a new log file with the **/logfile** parameter or *logfile* setting.

/OBJECT=*objectname*

Use the **/object** parameter to specify the name of the backup object that you want to delete. The object name uniquely identifies each backup object and is created by Data Protection for Exchange Server.

Use the Data Protection for Exchange Server query `tsm * /all` command to view the names of all available backup objects.

The **/object** parameter is used to delete only one incremental backup at a time. When multiple active incremental backups exist, the **/object** parameter must be specified with the **delete** command. If it is not specified, the **delete** command fails.

/OLDERthan=*days*

Use the **/olderthan** parameter to specify how old backup files can be before they are deleted. The *days* variable can range from 0 - 9999. There is no default value for **/olderthan**.

/Quiet This parameter prevents status information from being displayed. The level of information that is written to the activity log is not affected.

/TSMNODE=tsmnode

Use the *tsmnode* variable to refer to the IBM Spectrum Protect node name that Data Protection for Exchange Server uses to log on to the IBM Spectrum Protect server.

You can store the node name in the IBM Spectrum Protect options file (*dsm.opt*). This parameter overrides the value in the IBM Spectrum Protect options file if **PASSWORDACCESS** is set to PROMPT. This parameter is not valid when **PASSWORDACCESS** is set to GENERATE in the options file.

/TSMOPTFile=tsmoptfilename

Use the *tsmoptfilename* variable to identify the Data Protection for Exchange Server options file.

The file name can include a fully qualified path name. If no path is specified, the directory where Data Protection for Exchange Server is installed is searched.

If the *tsmoptfilename* variable includes spaces, enclose the entire **/tsmoptfile** parameter entry in double quotation marks. For example:

```
/TSMOPTFile="c:\Program Files\file.opt"
```

The default is *dsm.opt*.

/TSMPassword=tsmpassword

Use the *tsmpassword* variable to refer to the IBM Spectrum Protect password that Data Protection for Exchange Server uses to log on to the IBM Spectrum Protect server.

If you specified **PASSWORDACCESS** GENERATE in the Data Protection for Exchange Server options file (*dsm.opt*), you do not need to supply the password here because the one that is stored in the registry is used. However, to store the password in the registry, you must specify the IBM Spectrum Protect password the first time Data Protection for Exchange Server connects to the IBM Spectrum Protect server.

If you do specify a password with this parameter when **PASSWORDACCESS** GENERATE is in effect, the command-line value is ignored unless the password for this node has not yet been stored in the registry. In that case, the specified password is stored in the registry and used when you run this command.

If **PASSWORDACCESS** PROMPT is in effect, and you do not specify a password value on the command line, then you are prompted for a password.

The IBM Spectrum Protect password that Data Protection for Exchange Server uses to log on to the IBM Spectrum Protect server can be up to 63 characters in length.

Help command


To display help for Data Protection for Exchange Server commands, use the **tdpexcc help** command.

This command lists one or more commands and their parameters. If you cannot see all of the help on a screen, set the width of your screen display to a value greater than 80 characters. For example, set the screen width to 100 characters.

Help syntax

Use the **help** command syntax diagrams as a reference to view available options and truncation requirements.

TDPEXCC command



Help optional parameters

Optional parameters follow the Data Protection for Exchange Server **help** command.

The following optional parameters specify the help to be displayed:

*| *command*

Identifies the specific Data Protection for Exchange Server command that is to be displayed. If you use the asterisk (*) wildcard character, help for all Data Protection for Exchange Server commands is displayed.

The valid command names are shown:

```
BACKup
CHANGETSMPassword
HELP
MOUNT
Query
RESTore
RESTOREFiles
RESTOREMailbox
SET
```

*| *subcommand*

Help can be displayed for commands that have several subcommands, for example, the **query** command. If you do not specify a subcommand or asterisk (*) wildcard character, help for all Data Protection for Exchange Server **query** commands is displayed.

The valid subcommand names for the **query** command are shown:

```
EXCHange
managedcapacity
policy
TDP
TSM
```


Mount backup command

To mount backups, use the **mount backup** command.

Mount Backup syntax

Use the **mount backup** command syntax diagrams as a reference to view available options and truncation requirements.

TDPEXCC command



Mount backup positional parameter

The positional parameters immediately follow the **mount backup** command and precede the optional parameters.

The following positional parameters specify the objects to mount:

component name[(object-id)]=mount point root dir[,component name=mount point root dir]

component name[(object-id)]

Specify the backup of a local Exchange database.

mount point root dir

Specify the absolute path to the directory where the snapshots are going to be displayed as mount point directories. The directory must be empty. If not empty, an error is reported.

The list must contain all non-qualified objects or all qualified objects. The list cannot contain a combination of non-qualified objects and qualified objects. Specify the list by using the following syntax:

```

mount backup object-1[(object-1-id)]= mount-point-1[,object-2[(object-2-id)]
=mount-point-2...]

```

For example:

```

tdpexc mount backup excdb(20120815064316)=f:\emptyfolder

```

Mount Backup optional parameters

Optional parameters follow the **mount backup** command and positional parameters.

/CONFIGfile=*configfilename*

Use the **/configfile** parameter to specify the name (*configfilename*) of the Data Protection for Exchange Server configuration file that contains the values to use for a **mount backup** operation.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the Data Protection for Exchange Server installation directory is used. If the **/configfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `tdpexc.cfg`.

If the *configfilename* variable includes spaces, enclose the entire **/configfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\tdpexc.cfg"
```

/LOGFile=*logfile*

Use the **/logfile** parameter to specify the name of the activity log file that is generated by Data Protection for Exchange Server. The *logfile* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfile* variable can include a fully qualified path. However, if no path is specified, the log file is written to the Data Protection for Exchange Server installation directory.

If the *logfile* variable includes spaces, enclose the entire **/logfile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\tdpexc.log"
```

If the **/logfile** parameter is not specified, log records are written to the default log file, `tdpexc.log`.

The **/logfile** parameter cannot be turned off, logging always occurs.

/LOGPrune=*numdays* | **No**

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option **No** can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify **no**, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.

- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, *60*, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/MOUNTRW=Yes | No

You can mount a read/write copy of your IBM Spectrum Protect backup so that you can modify the copy without invalidating the backup. You use this option to indicate whether a snapshot backup is mounted as read/write. The default value is as specified in the configuration file with the **/MOUNTRW** parameter. If a default value is not specified in the configuration file, the default value is No. By specifying the **/MOUNTRW** option, you override the default value.

The following values are available:

No Perform a read-only mount operation.

Yes Perform a read/write mount operation. The behavior of the read/write mount is controlled by the **USESNAPOFASNAPMount** parameter in the configuration file.

- If **USESNAPOFASNAPMount** is set to No, you can mount only COPY backups as read/write. After mounting, the original backup is modified and can no longer be used as a restore point in future database restore operations (on the **VSS Options** properties page, the **Mount read/write (modifies backup, applies to COPY backups only)** check box is selected).
- If **USESNAPOFASNAPMount** is set to Yes, you can mount both FULL and COPY backup types as read/write (on the **VSS Options** properties page, the **Mount read/write (without modifying backup)** check box is selected). In this instance, the backups are not modified and can be used in future restore operations.

Important:

This mount option is only available for SAN Volume Controller (SVC) devices and requires IBM System Storage Support for Microsoft Virtual Disk and Volume Shadow Copy Services version 4.12 or later. Dynamic target allocation is not supported. You must allocate more target volumes on your SVC storage device to accommodate the number of concurrent read/write mounts you want to do. An extra target volume, that matches the size of the volume to be mounted, is needed for each concurrent read/write mount of that volume.

/REMOTECOMPUTER=computername

Enter the computer name or IP address of the remote system where the backup was created.

/REMOTECOMPUTERUser=user

Enter the user name that is used to log on to the server specified with the

REMOTECOMPUTER parameter. If a domain is required to log on with the domain account, enter *Domain\User*. To log on to the local account, the domain is not required. There is no default value.

/REMOTECOMPUTERPassword=passwd

Enter the password for the user name that is specified with the **REMOTECOMPUTERUser** parameter. There is no default value.

/TSMNODE=tsmnodename

Use the *tsmnodename* variable to refer to the IBM Spectrum Protect node name that Data Protection for Exchange Server uses to log on to the IBM Spectrum Protect server.

You can store the node name in the IBM Spectrum Protect options file (*dsm.opt*). This parameter overrides the value in the IBM Spectrum Protect options file if **PASSWORDACCESS** is set to PROMPT. This parameter is not valid when **PASSWORDACCESS** is set to GENERATE in the options file.

/TSMOPTFile=tsmoptfilename

Use the *tsmoptfilename* variable to identify the IBM Spectrum Protect options file.

The file name can include a fully qualified path name. If no path is specified, the directory where Data Protection for Exchange Server is installed is searched.

If the *tsmoptfilename* variable includes spaces, enclose the entire **/tsmoptfile** parameter entry in double quotation marks. For example:

```
/TSMOPTFile="c:\Program Files\dsm.opt"
```

The default is *dsm.opt*.

/TSMPassword=tsmpassword

Use the *tsmpassword* variable to refer to the IBM Spectrum Protect password that Data Protection for Exchange Server uses to log on to the IBM Spectrum Protect server.

If you specified **PASSWORDACCESS** GENERATE in the Data Protection for Exchange Server options file (*dsm.opt*), supplying the password is not necessary here because the one that is stored in the registry is used. However, to store the password in the registry, you must specify the IBM Spectrum Protect password the first time Data Protection for Exchange Server connects to the IBM Spectrum Protect server.

If you do specify a password with this parameter when **PASSWORDACCESS** GENERATE is in effect, the command-line value is ignored unless the password for this node is not yet stored in the registry. In that case, the specified password is stored in the registry and used when you run this command.

If **PASSWORDACCESS** PROMPT is in effect, and you do not specify a password value on the command line, then you are prompted for a password.

The IBM Spectrum Protect password that Data Protection for Exchange Server uses to log on to the IBM Spectrum Protect server can be up to 63 characters in length.

Query Exchange command

To query the local Exchange Server for general information, use the **query exchange** command.


The **query exchange** command returns the following information:

- Exchange Server name and version
- Domain name
- Names of all databases
- Status (online, offline) of all databases
- Circular logging status (enabled, disabled)
- The following VSS information:
 - Writer name
 - Local DSMAgent node
 - Remote DSMAgent node
 - Writer status (online, offline)
 - Number of selectable components

Query Exchange syntax

Use the **query exchange** command syntax diagrams as a reference to view available options and truncation requirements.

TDPEXCC command



Query Exchange optional parameters

Optional parameters follow the **query exchange** command.

/CONFIGfile=filename

Use the **/CONFIGfile** parameter to specify the name (*filename*) of the Data Protection for Exchange Server configuration file that contains the values to use for a **query exchange** operation.

The *filename* variable can include a fully qualified path. If the *filename* variable does not include a path, the Data Protection for Exchange Server installation directory is used. If the **/CONFIGfile** parameter is not specified, or if the *filename* variable is not specified, the default value is `tdpexc.cfg`.

If the *filename* variable includes spaces, enclose the entire **/CONFIGfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\file.cfg"
```

/LOGFile=logfilename

Use the **/LOGFile** parameter to specify the name of the activity log file that is generated by Data Protection for Exchange Server. The *logfilename*

variable identifies the name of the activity log file. If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfile* variable can include a fully qualified path. However, if no path is specified, the log file is written to the Data Protection for Exchange Server installation directory. If the *logfile* variable includes spaces, enclose the entire **/LOGFile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\mytdpexchange.log"
```

If the **/LOGFile** parameter is not specified, log records are written to the default log file, *tdpexc.log*. The **/LOGFile** parameter cannot be turned off, logging always occurs.

When you use multiple simultaneous instances of Data Protection for Exchange Server to perform operations, use the **/LOGFile** parameter to specify a different log file for each instance used. This function directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=numdays | No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option **No** can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify **no**, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

Query Managedcapacity command

To determine the amount of managed capacity in use, and to assist with your storage planning, use the **query managedcapacity** command.

Purpose

The **query managedcapacity** command displays capacity that relates to the volumes that are represented in local inventory that is managed by Data Protection for Exchange Server. You can issue this command on all Windows operating systems that are supported by Data Protection for Exchange Server.

The capacity that is displayed includes deactivated backups, that is, backups which have not expired, on the IBM Spectrum Protect server. Once a deleted backup has been expired by the IBM Spectrum Protect server, the capacity that is displayed no longer contains capacity for the deleted backup.

TDPEXCC command: Query MANAGEDCAPacity

►► TDPEXCC—Query MANAGEDCAPacity —————►
 |_/Detailed|

Parameters

/Detailed

Results in a detailed listing of snapped volumes. If this option is not specified, then only the total capacity is displayed.

In this example, the **tdpexcc query managedcapacity** command displays the total amount of managed capacity in use in the local inventory. The following output is displayed:

```
Total Managed Capacity : 124.65 GB (133,844,426,752 bytes)
```

In this example, the **tdpexcc query managedcapacity /detailed** command displays a detailed listing of total amount of managed capacity and the snapped volumes in use. The following output is displayed:

```
IBM Spectrum Protect Snapshot for Mail
Snapshot for Microsoft Exchange Server
Version 8, Release 1, Level 0
(C) Copyright IBM Corporation 1998, 2016. All rights reserved.

Total Managed Capacity : 31.99 GB (34,353,438,720 bytes)

Volume           : M:
Managed Capacity : 16.00 GB (17,176,719,360 bytes)

Volume           : F:
Managed Capacity : 16.00 GB (17,176,719,360 bytes)
Total Managed Capacity : 1,019.72 MB (1,069,253,632 bytes)
```

Query policy command

To query local policy information, use the **query policy** command.

Query Policy

This command is used to list the attributes of a policy.

TDPEXCC command

►► TDPEXCC—Query POLicy—*polycyname*—►

└─*─┘

When you specify *, all policies are queried. The results of the query are displayed:

Connecting to Exchange Server, please wait...		
Policy	Number of snapshots to keep	Days to keep a snapshot
-----	-----	-----
EXCPOL	3	60
STANDARD	2	30

Query TDP command

To query a list of the current values that are set in the configuration file for Data Protection for Exchange Server, use the **query tdp** command.

Query TDP syntax

Use the **query TDP** command syntax diagrams as a reference to view available options and truncation requirements.

TDPEXCC command: Query TDP

►► TDPEXCC—Query TDP—►

└─/CONFIGfile=—*tdpexc.cfg*—┘

└─*configfilename*—┘

►

└─/LOGFile=—*tdpexc.log*—┘

└─*logfilename*—┘

└─/LOGPrune=—*60*—┘

└─*numdays*—┘

└─No—┘

Query TDP optional parameters

Optional parameters follow the **query TDP** command.

/CONFIGfile=*configfilename*

Use the **/CONFIGfile** parameter to specify the name (*configfilename*) of the Data Protection for Exchange Server configuration file that contains the values to use for a **query TDP** operation.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the Data Protection for Exchange Server installation directory is used. If the **/CONFIGfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `tdpexc.cfg`.

If the *configfilename* variable includes spaces, enclose the entire **/CONFIGfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\file.cfg"
```

/LOGFile=*logfile*

Use the **/LOGFile** parameter to specify the name of the activity log file that is generated by Data Protection for Exchange Server.

The *logfile* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfile* variable can include a fully qualified path. However, if no path is specified, the log file is written to the Data Protection for Exchange Server installation directory.

If the *logfile* variable includes spaces, enclose the entire **/LOGFile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\mytdpexchange.log"
```

If the **/LOGFile** parameter is not specified, log records are written to the default log file, `tdpexc.log`.

The **/LOGFile** parameter cannot be turned off, logging always occurs.

When you use multiple simultaneous instances of Data Protection for Exchange Server to perform operations, use the **/LOGFile** parameter to specify a different log file for each instance used. This function directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=*numdays* | **No**

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option **No** can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify **no**, the log file is not pruned.

- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is *60*.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, *60*, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

Examples: query tdp command

This output example provides a sample of the text, messages, and process status that displays when the **query tdp** command is used.

The following code sample includes output for a VSS configuration:

```
IBM Tivoli Storage Manager for Mail:
Data Protection for Microsoft Exchange Server
Version 7, Release 1, Level 3.0
(C) Copyright IBM Corporation 1998, 2015. All rights reserved.

Data Protection for Exchange Preferences
-----

BACKUPDEStination..... LOCAL
CLIENTAccessserver.....
DAGNODE..... DAG1
DATEformat ..... 1
LANGuage ..... ENU
LOCALDSMAgentnode..... TIVVM400
LOGFile ..... tdpexc.log
LOGPrune ..... 60
MOUNTWait ..... Yes
NUMBERformat ..... 1
REMOVEDSMAgentnode.....
TEMPDBRestorepath.....
TEMPLOGRestorepath.....
TIMEformat ..... 1
IMPORTVSSSNAPSHOTSONLYWhenneeded.... Yes

Completed
```

Query TSM command

To display IBM Spectrum Protect server information, use the **query tsm** command.

This command displays the following information:

- IBM Spectrum Protect node name
- Network host name of the server
- IBM Spectrum Protect API version
- Server name, type, and version
- Compression mode

- Domain
- Active policy set
- Default management class

This command also displays a list of backups that are stored on the IBM Spectrum Protect server that match the databases that are entered. Active and inactive objects are displayed. However, only the active backup objects are displayed by default. To include inactive backup versions in the list, use the **/all** optional parameter.

Query TSM syntax

To view available options and truncation requirements, use the **query TSM** command.

TDPEXCC command



Query TSM positional parameters

Positional parameters immediately follow the **query TSM** command and precede the optional parameters.

The following positional parameters specify the object to query. If none of these positional parameters are specified, only the IBM Spectrum Protect API and IBM Spectrum Protect server information is displayed:

* | *dbname*

* Query all backup objects for all databases.

dbname

Query all backup objects for the specified database. Multiple entries are separated by commas.

The following positional parameters specify the type of backup to query. If this parameter is not specified, all backup types are displayed:

FULL | **COPY** | **INCRemental** | **DIFFerential**

FULL Query only full backup types

COPY Query copy backup types only

INCRemental

Query only incremental backup types

DIFFerential

Query only differential backup types

Query TSM optional parameters

Optional parameters follow the **query TSM** command and positional parameters.

/Active

Use the **/Active** parameter to display active backup objects only. This setting is the default setting.

/All Use the **/All** parameter to display both active and inactive backup objects. If the **/All** parameter is not specified, only active backup objects are displayed.

/CONFIGfile=filename

Use the **/CONFIGfile** parameter to specify the name of the Data Protection for Exchange Server configuration file that contains the values for the Data Protection for Exchange Server configuration options.

The *filename* variable can include a fully qualified path. If the *filename* variable does not include a path, the Data Protection for Exchange Server installation directory is used. If the **/CONFIGfile** parameter is not specified, or if the *filename* variable is not specified, the default value is *tdpexc.cfg*.

If the *filename* variable includes spaces, enclose the entire **/CONFIGfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\file.cfg"
```

/DAGNode=nodename

Specify the node name that you want to use to back up and restore the databases in an Exchange Server Database Availability Group (DAG). With this setting, backups from all DAG members that are configured to use the DAG node are backed up to a common file space on the IBM Spectrum Protect server. The database copies are managed as a single entity,

regardless of which DAG member they were backed up from. This setting can prevent Data Protection for Exchange Server from making too many backups of the same database.

/DEtail

Use the **/DEtail** parameter to display comprehensive information about the status of the IBM Spectrum Protect server, including the following details:

- Backup compressed

Table 22. Backup compressed values

Value	Status
Yes	One or more objects are compressed.
No	No objects are compressed.
Unknown	It is not known whether the backup is compressed, which includes backups before Data Protection for Exchange Server version 6.3.

- Backup encryption type

Table 23. Backup encryption type values

Value	Status
None	None of the objects are encrypted.
AES-256	The objects are encrypted with AES-256 encryption. The most secure data encryption method is AES256.
AES-128	The objects are encrypted with AES-128 encryption.
DES-56	The objects are encrypted with DES-56 encryption.
Unknown	It is not known whether the objects in the database are encrypted, which includes backups before Data Protection for Exchange Server version 6.3.

- Backup client-deduplicated

Table 24. Backup client-deduplicated values

Value	Status
Yes	One or more objects are client-side deduplicated.
No	No objects are client-side deduplicated.
Unknown	It is not known whether the backup is client-side deduplicated, which includes backups before Data Protection for Exchange Server version 6.3.

- Backup supports instant restore

Table 25. Backup supports instant restore values

Value	Status
Yes	One or more objects support instant restore.
No	No objects support instant restore.

Table 25. Backup supports instant restore values (continued)

Value	Status
Unknown	The backup might not support instant restore. This setting applies to backups before Data Protection for Exchange Server version 6.3.

/FROMEXCSERVER=*servername*

Use the **/fromexcserver** parameter to specify the name of the Exchange Server where the original backup is completed.

The default is the local Exchange Server. However, you must specify the name if the Exchange Server is not the default.

If a DAG node is specified by using the **dagnode** parameter, Data Protection for Exchange Server uses this node name instead of the Data Protection for Exchange Server node to back up databases in an Exchange Server Database Availability Group. Therefore, the **query tsm** command automatically displays the backups that were created by the other DAG members, without having to specify the **/fromexcserver** parameter.

/LOGFile=*logfile*

Use the **/LOGFile** parameter to specify the name of the activity log file that is generated by Data Protection for Exchange Server.

The *logfile* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfile* variable can include a fully qualified path. However, if no path is specified, the log file is written to the Data Protection for Exchange Server installation directory.

If the *logfile* variable includes spaces, enclose the entire **/LOGFile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\mytdpexchange.log"
```

If the **/LOGFile** parameter is not specified, log records are written to the default log file, *tdpexc.log*.

The **/LOGFile** parameter cannot be turned off; logging always occurs.

When you use multiple simultaneous instances of Data Protection for Exchange Server to complete operations, use the **/LOGFile** parameter to specify a different log file for each instance used. This function directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=*numdays* | No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option **No** can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify *no*, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/REMOTECOMPUTER=computername

Enter the IP address or host name for the remote system where you want to query the data backups.

/REMOTECOMPUTERUser=user

Enter the user name that is used to log on to the server specified with the **REMOTECOMPUTER** parameter. If a domain is required to log on with the domain account, enter *Domain\User*. To log on to the local account, the domain is not required. There is no default value.

/REMOTECOMPUTERPassword=password

Enter the password for the user name that is specified with the **REMOTECOMPUTERUser** parameter. There is no default value.

/SHOWMAILBOXInfo

Use the **/SHOWMAILBOXInfo** parameter to display mailbox history information in backup databases.

/TSMNODE=nodename

Use the *tsmnodename* variable to refer to the IBM Spectrum Protect node name that Data Protection for Exchange Server uses to log on to the IBM Spectrum Protect server.

You can store the node name in the IBM Spectrum Protect options file (*dsm.opt*). This parameter overrides the value in the IBM Spectrum Protect options file if **PASSWORDACCESS** is set to **PROMPT**. This parameter is not valid when **PASSWORDACCESS** is set to **GENERATE** in the options file.

/TSMOPTFile=optfilename

Use the *optfilename* variable to identify the Data Protection for Exchange Server options file.

The file name can include a fully qualified path name. If no path is specified, the directory where Data Protection for Exchange Server is installed is searched.

If the *optfilename* variable includes spaces, enclose the entire **/TSMOPTFile** parameter entry in double quotation marks. For example:

```
/TSMOPTFile="c:\Program Files\file.opt"
```

The default is `dsm.opt`.

`/TSMPassword=password`

Use the *password* variable to refer to the IBM Spectrum Protect password that Data Protection for Exchange Server uses to log on to the IBM Spectrum Protect server.

If you specified `PASSWORDACCESS GENERATE` in the Data Protection for Exchange Server options file (`dsm.opt`), supplying the password here is not necessary because the one that is stored in the registry is used. However, to store the password in the registry, you must specify the IBM Spectrum Protect password the first time that Data Protection for Exchange Server connects to the IBM Spectrum Protect server .

If you do specify a password with this parameter when `PASSWORDACCESS GENERATE` is in effect, the command-line value is ignored unless the password for this node is not yet stored in the registry. In that case, the specified password is stored in the registry and used when you run this command.

If `PASSWORDACCESS PROMPT` is in effect, and you do not specify a password value on the command-line interface, you are prompted for a password.

The IBM Spectrum Protect password that Data Protection for Exchange Server uses to log on to the IBM Spectrum Protect server can be up to 63 characters in length.

Examples: query tsm command

The **query tsm** command displays information about the IBM Spectrum Protect API and IBM Spectrum Protect server.

The following example includes output from the **tdpexcc query tsm** command:

```
IBM Tivoli Storage Manager for Mail:
Data Protection for Microsoft Exchange Server
Version 7, Release 1, Level 3.0
(C) Copyright IBM Corporation 1998, 2015. All rights reserved.

Tivoli Storage Manager Server Connection Information
-----
Nodename ..... MALTA_EXC
Network Host Name of Server ..... GIJOE
TSM API Version ..... Version 7, Release 1, Level 2.0

Server Name ..... GIJOE_SERVER1_230
Server Type ..... Windows
Server Version ..... Version 7, Release 1, Level 2.0
Compression Mode ..... Client Determined
Domain Name ..... FCM_PDEXC
Active Policy Set ..... STANDARD
Default Management Class ..... STANDARD

Completed
```

For backups that are created from an Exchange Server Database Availability Group (DAG) member, both the DAG node name and the name of the server on which that backup was run are displayed with the **query tsm** command. The following example queries the IBM Spectrum Protect server for the backup objects that were backed up under the DAG node name `DAG2` from DAG member server `TIVVM483`:

Command:

```
tdpexcc query tsm *
```

Output:

```
IBM Tivoli Storage Manager for Mail:
Data Protection for Microsoft Exchange Server
Version 7, Release 1, Level 3.0
(C) Copyright IBM Corporation 1998, 2015. All rights reserved.

Querying Tivoli Storage Manager server for a list of database backups, please wa
it...

Connecting to TSM Server as node 'TIVVM483_EXC'...
Connecting to Local DSM Agent 'TIVVM483'...
Using backup node 'DAG2'...

DAG           : DAG2

Database      : RATTEST_DAGDB

Backup Date   Size      S Fmt  Type  Loc      Object Name
-----
03/27/2014 16:11:14 149.07MB A VSS  full  Srv 20140327161114
                13.01MB
                136.06MB
03/27/2014 18:02:01 14.00MB A VSS  incr  Srv 20140327180201
                14.00MB
                Logs
```

The following example queries the IBM Spectrum Protect server in detail for the backup objects that were backed up under the DAG node name DAG2 from DAG member server TIVVM483:

Command:

```
tdpexcc query tsm * /detail
```

Output:

```

IBM Tivoli Storage Manager for Mail:
Data Protection for Microsoft Exchange Server
Version 7, Release 1, Level 3.0
(C) Copyright IBM Corporation 1998, 2015. All rights reserved.

Querying Tivoli Storage Manager server for a list of database backups, please wa
it...

Connecting to TSM Server as node 'TIVVM483_EXC'...
Connecting to Local DSM Agent 'TIVVM483'...
Using backup node 'DAG2'...

Backup Object Information
-----

Exchange Server Name ..... TIVVM483
Database Availability Group ..... DAG2
Backup Database Name ..... RATTEST_DAGDB
Backup Method ..... VSS
Backup Location ..... Srv
Backup Object Type ..... full
Mounted as .....
Backup Object State ..... Active
Backup Creation Date / Time ..... 03/27/2014 16:11:14
Backup Compressed ..... No
Backup Encryption Type ..... None
Backup Client-deduplicated ..... No
Backup Supports Instant Restore ..... No
Backup Object Size / Name ..... 149.07MB / 20140327161114 (From DBC
opy)
Backup Object Size / Name ..... 13.01MB / Logs
Backup Object Size / Name ..... 136.06MB / File

```

Restore command

To restore a backup from IBM Spectrum Protect storage to an Exchange server, use the **restore** command.

When you use the **restore** command, remember the following guidelines:

- When you restore inactive backups or active incremental backups, use the **/object** parameter to specify the name of the backup object to restore. This object name uniquely identifies the backup instance in IBM Spectrum Protect storage. You can issue a **tdpexcc query tsm *** command to obtain a list of the object names.

If the **tdpexcc restore dbname incr** command is entered (without the **/object** parameter) to restore multiple active incremental backups, all multiple active incremental backups are restored sequentially. The **/object** parameter is used to restore only one incremental backup at a time.

- To initiate recovery, you must use the **/recover** parameter when you restore the last backup object. In addition, the value of **/templogrestorepath** is not the same value as the current location. If the value is the same, the database can become corrupted.
 - Specify **/recover=applyalllogs** to replay the restored-transaction log entries and the current active-transaction log entries.
 - Specify **/recover=applyrestoredlogs** to replay only the restored-transaction log entries. The current active-transaction log entries are not replayed.

When you choose this option for a restore, your next backup must be a full or copy backup.

Failure to use the **/recover** parameter when you restore the last backup set leaves the databases unmountable.

- Specify `/mountdatabases=yes` if you are restoring the last backup set and you want the database mounted automatically after the recovery is completed.

With Microsoft Exchange Server, you cannot specify the asterisk (*) wildcard character in database names. Databases that contain the asterisk (*) wildcard character in their name are not backed up.

The GUI provides an easy-to-use, flexible interface to help you complete a restore operation. The interface presents information in a way that allows multiple selection and, in some cases, automatic operation.

Important:

If the Windows event log, Data Protection for Exchange Server log file, or a command error indicates that a restore operation failed, this failure might be caused by the `restore.env` file that remains behind. This file is created by the Microsoft restore interface and is used for debugging the restore failure. This file is named `Ennrestore.env` where *Enn* is the base name of the restored transaction log files. After the restore error is resolved, remove any remaining `restore.env` files before you attempt the next restore operation. For more information, see the following Microsoft web page: [Restore.env file documentation](#)

Data Protection for Exchange Server supports the following types of restore:

Full Restore a full backup.

Copy Restore a copy backup.

Incremental

Restore an incremental backup.

Differential

Restore a differential backup.

Related tasks:


“Restoring an Exchange Server database” on page 106


“Restoring a Database Availability Group database backup” on page 107

Restore syntax

Use the **restore** command syntax diagrams as a reference to view available options and truncation requirements.

TDPEXCC command





Restore positional parameters

Positional parameters immediately follow the **restore** command and precede the optional parameters.

The following positional parameters specify the object to restore:

* | *database_name1*, ..., *database_nameN*

* Restore all database names sequentially.

database_name

Restore the specified database. Multiple entries are separated by commas. If separated by commas, ensure that there is no space

between the comma and the database name. If any database name contains blanks, enclose the database name in double quotation marks.

The following positional parameters specify the type of restore to perform:

FULL | COPY | INCRemental | DIFFerential

FULL Restore a Full type backup.

COPY Restore a Copy type backup.

INCRemental
Restore an Incremental type backup.

DIFFerential
Restore a Differential type backup

Restore optional parameters

Optional parameters follow the **restore** command and positional parameters.

/BACKUPDESTination=TSM | LOCAL

Use the **/BACKUPDESTination** parameter to specify the location from where the backup is to be restored. The default is the value (if present) specified in the Data Protection for Exchange Server preferences file (*tdpexc.cfg*). If no value is present, the backup is restored from IBM Spectrum Protect server storage.

You can specify:

TSM The backup is restored from IBM Spectrum Protect server storage. TSM is the default value.

LOCAL The backup is restored from the local shadow volumes.

/CONFIGfile=configfilename

Use the **/CONFIGfile** parameter to specify the name of the Data Protection for Exchange Server configuration file that contains the values for the Data Protection for Exchange Server configuration options.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the Data Protection for Exchange Server installation directory is used. If the **/CONFIGfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is *tdpexc.cfg*.

If the *configfilename* variable includes spaces, enclose the entire **/CONFIGfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\file.cfg"
```

DAGNODE=nodename

Specify the node name that you want to use to back up the databases in an Exchange Server Database Availability Group. With this setting, backups from all Database Availability Group members that are configured to use the DAG node are backed up to a common file space on the IBM Spectrum Protect server. The database copies are managed as a single entity, regardless of which Database Availability Group member they are backed up from. This setting can prevent IBM Spectrum Protect Snapshot for Microsoft Exchange Server from making too many backups of the same database.

/ERASEexistinglogs=YES|NO

Use the **/ERASEexistinglogs** parameter to erase the existing transaction log files for the database that is being restored before you restore the specified databases. If you do not erase existing data, existing transaction logs can be reapplied when the Exchange databases are mounted. This parameter is valid only when you restore a VSS FULL or VSS COPY backup of Exchange Server databases.

/FROMEXCServer=server-name

Use the **/FROMEXCServer** parameter to specify the name of the Exchange Server where the original backup was done.

The default is the local Exchange Server. However, you must specify the name if the Exchange Server is not the default.

If a DAG node is specified by using the **/dagnode** parameter, Data Protection for Exchange Server uses this node name instead of the Data Protection for Exchange Server node to back up databases in an Exchange Server Database Availability Group. Therefore, the **restore** command automatically restores the backups that were created by the other DAG members, without having to specify the **/fromexcserver** parameter.

/INSTANTRestore=YES|NO

Use the **/INSTANTRestore** parameter to specify whether to use volume level snapshot or file level copy to restore a VSS backup on local shadow volumes. A SAN Volume Controller, Storwize family, or IBM System Storage DS8000 series storage system is required to perform VSS instant restore operations.

You can specify:

YES Use volume level snapshot restore for a VSS backup on local shadow volumes if the backup exists on volumes that support it. YES is the default.

NO Use file level copy to restore the files from a VSS backup on local shadow volumes. Bypassing volume-level copy means that Exchange database files, log files, and the checkpoint file are the only data overwritten on the source volumes.

When you run VSS instant restore operations, you must make sure that any previous background copies (that involve the volumes that are being restored) are completed before initiating the VSS instant restore. The **/INSTANTRestore** parameter is ignored and VSS instant restore capabilities are automatically disabled when doing any type of VSS restore into operation.

When you run a VSS instant restore in a Database Availability Group (DAG) environment, do not choose the option that automatically mounts the databases after the recovery is completed. As described in the Database Availability Group considerations section, to run the VSS instant restore for databases in a DAG environment, you must stop the Microsoft Exchange Replication service before doing the VSS instant restore or the restore fails. In this case, after the VSS instant restore is completed, start the Microsoft Exchange Replication service and then finally mount the database.

/INTODB=into-db-name

Use the **/INTODB** parameter to specify the name of the Exchange Server database into which the VSS backup is restored. The database name must

be specified with the *into-db-name* variable. For example, if RDB is the name of the database into which the VSS backup is restored, the command-line entry would be as follows:

```
TDPEXCC RESTore Maildb1 FULL /INTODB=RDB
```

However, when you restore a database that is relocated (system file path, log file path, or database file path), you must specify the same database name as the one you are restoring. For example, if Maildb1 is the name of the relocated database that is being restored, the command-line entry would be as follows:

```
TDPEXCC RESTore Maildb1 FULL /INTODB=Maildb1
```

Considerations:

- There is no default value.
- To restore into a Recovery Database (RDB) or alternate database, an RDB or alternate database must exist before you attempt the restore operation.

/LOGFile=logfilename

Use the **/LOGFile** parameter to specify the name of the activity log file that is generated by Data Protection for Exchange Server.

The *logfilename* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfilename* variable can include a fully qualified path. However, if no path is specified, the log file is written to the Data Protection for Exchange Server installation directory.

If the *logfilename* variable includes spaces, enclose the entire **/LOGFile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\mytdpexchange.log"
```

If you do not specify the **/LOGFile** parameter, log records are written to the default log file, *tdpexc.log*.

The **/LOGFile** parameter cannot be turned off, logging always occurs.

When you using multiple simultaneous instances of Data Protection for Exchange Server to perform operations, use the **/LOGFile** parameter to specify a different log file for each instance used. This function directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=numdays | No

Use the **/LOGPrune** parameter to disable log pruning or to explicitly request to prune the log for one command run. By default, log pruning is enabled and done once per day. The *numdays* variable represents the number of days to save log entries. By default, 60 days of log entries are saved in the pruning process. You can use the Data Protection for Exchange Server GUI or the **set** command to change the defaults so that log pruning is disabled, or so that more or less days of log entries are saved. If you use the command line, you can use the **/LOGPrune** parameter to override these defaults. When the value of the **/LOGPrune** variable *numdays* is a number in the range 0 to 9999, the log is pruned even if log pruning is already done for the day.

Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in pruning the log unintentionally. If the value of the **TIMEformat** or **DATEformat** parameter is changed, before you issue a Data Protection for Exchange Server command that might prune the log file, do one of the following actions to prevent the log file from being pruned:

- Make a copy of the existing log file.
- Specify a new log file with the **/logfile** parameter or **logfile** setting.

/MOUNTDatabases=No|Yes

Use the **/MOUNTDatabases** parameter to specify whether to mount the databases after the restore operation is completed. Specify one of the following values:

- Yes** Mount the databases after the restore operation is completed.
- No** Do not mount the databases after the restore operation is completed. No is the default.

/MOUNTWait=Yes|No

Use the **/MOUNTWait** parameter to specify whether Data Protection for Exchange Server is to wait for removable media to mount (such as tapes or CDs) or to stop the current operation. This situation occurs when the IBM Spectrum Protect server is configured to store backup data on removable media and waits for a required storage volume to be mounted.

You can specify:

- Yes** Wait for tape mounts. This option is the default.
- No** Do not wait for tape mounts.

/OBJECT=object-name

Use the **/OBJECT** parameter to specify the name of the backup object you want to restore. The object name uniquely identifies each backup object and is created by Data Protection for Exchange Server.

Use the Data Protection for Exchange Server **query tsm** command to view the names of the backup objects.

If the **tdpexcc restore dbname incr** command is entered (without the **/OBJECT** parameter) to restore multiple active incremental backups, all multiple active incremental backups are restored sequentially. The **/OBJECT** parameter is used to restore only one incremental backup at a time.

/Quiet This parameter prevents status information from being displayed. The level of information that is written to the activity log is not affected.

/RECOVER=APPLYRESToredlogs|APPLYALLlogs

Use this parameter to specify whether or not you want to run recovery after you restore an object. This parameter is specified on the last backup object restored. To initiate recovery, use the **/RECOVER** parameter when you restore the last backup object. In addition, the value of **/TEMPLOGRESTorepath** is not to be the same value as the current location. If the value is the same, the database can become corrupted. Failure to use the **/RECOVER** parameter when you restore the last backup set leaves the databases unmountable. If this situation occurs, you can either restore the last backup again and specify the **/RECOVER=value** option or you can use the Microsoft **ESEUTIL /cc** command to run recovery manually.

Specify one of the following values when you use this parameter:

APPLYALLlogs

Specify `/recover=applyalllogs` to replay the restored-transaction log entries AND the current active-transaction log entries. Any transaction logs entries that are displayed in the current active-transaction log are replayed. This option is the default.

APPLYRESToredlogs

Specify `/recover=applyrestoredlogs` to replay only the restored-transaction log entries. The current active-transaction log entries are not replayed.

When you choose this option for a restore, your next backup must be a full or copy backup.

Considerations:

- When you restore multiple backup objects, the **/RECOVER** option is to be used on the restore of the last object.
If you specify **/RECOVER=APPLYRESToredlogs** when doing a restore, the next backup must be a full backup.

/TEMPLOGRESTorepath=path-name

Use the **/TEMPLOGRESTorepath** parameter to specify the default temporary path to use when you are restoring logs and patch files. For best performance, this logger is to be on a different physical device than the current active-transaction logger.

If you do not specify the **/TEMPLOGRESTorepath** parameter, the default value is the value that is specified by the **/TEMPLOGRESTorepath** option in the Data Protection for Exchange Server configuration file. The default Data Protection for Exchange Server configuration file is `tdpexc.cfg`.

If you do not specify the **/TEMPLOGRESTorepath** parameter, and the **/TEMPLOGRESTorepath** value does not exist in the Data Protection for Exchange Server configuration file, the `TEMP` environment variable value is used.

When you run a `FULL` or `COPY` restore operation, all log files in the path that is specified by the **/TEMPLOGRESTorepath** parameter are erased.

In addition, the value of **/TEMPLOGRESTorepath** is not to be the same value as the current location. If the value is the same, the database can become corrupted.

Do not specify double-byte characters (DBCS) within the temporary log path.

/TSMNODE=tsmnodename

Use the `tsmnodename` variable to refer to the IBM Spectrum Protect node name that Data Protection for Exchange Server uses to log on to the IBM Spectrum Protect server.

You can store the node name in the IBM Spectrum Protect options file (`dsm.opt`). This parameter overrides the value in the IBM Spectrum Protect options file if **PASSWORDACCESS** is set to `PROMPT`. This parameter is not valid when **PASSWORDACCESS** is set to `GENERATE` in the options file.

/TSMOPTfile=tsmoptfilename

Use the `tsmoptfilename` variable to identify the Data Protection for Exchange Server options file.

The file name can include a fully qualified path name. If no path is specified, the directory where Data Protection for Exchange Server is installed is searched.

If the *tsmoptfilename* variable includes spaces, enclose the entire **/TSMOPTfile** parameter entry in double quotation marks. For example:
`/TSMOPTfile="c:\Program Files\file.opt"`

The default is `dsm.opt`.

/TSMPassword=*tsmpassword*

Use the *tsmpassword* variable to refer to the IBM Spectrum Protect password that Data Protection for Exchange Server uses to log on to the IBM Spectrum Protect server.

If you specified **PASSWORDACCESS GENERATE** in the Data Protection for Exchange Server options file (`dsm.opt`), supplying the password here is not necessary because the one that is stored in the registry is used. However, to store the password in the registry, you must specify the IBM Spectrum Protect password the first time Data Protection for Exchange Server connects to the IBM Spectrum Protect server.

If you do specify a password with this parameter when **PASSWORDACCESS GENERATE** is in effect, the command-line value is ignored unless the password for this node is not yet stored in the registry. In that case, the specified password is stored in the registry and used when you run this command.

If **PASSWORDACCESS PROMPT** is in effect, and you do not specify a password value on the command line, then you are prompted for a password.

The IBM Spectrum Protect password that Data Protection for Exchange Server uses to log on to the IBM Spectrum Protect server can be up to 63 characters in length.

Restorefiles command

To restore the flat files from a specified Data Protection for Exchange Server backup into a specified directory, use the **restorefiles** command.

When working with the **restorefiles** command, follow these guidelines:

- This command does not require an Exchange Server to be installed on, or accessible from the system where the **restorefiles** command is run.
- Files can be restored to an alternative system or to an alternative directory on the same system as the Exchange Server.
- The command continues until it succeeds, or until the destination volume does not contain enough space for the operation.
- When you restore files from an inactive backup or an active incremental backup, use the **/object** parameter to specify the name of the backup object. The object name uniquely identifies the backup instance in IBM Spectrum Protect server storage. A list of backup object names is obtained by issuing the **query tsm** command.
- This command is only available on the command-line interface. It is not available in the Data Protection for Exchange Server graphical user interface.
- The **restorefiles** command overwrites files that exist and have the same name.

- If a log file from an incremental backup has the same name as the log file from the full backup operation, you can run two consecutive **restorefiles** commands to the same directory:

For a full backup, enter the following command:

```
tdpexcc restorefiles DB1 FULL /into=d:\temprestore
```

For an incremental restore with backed up log files, enter the following command:

```
tdpexcc restorefiles DB1 INCR /into=d:\temprestore
```

The **/into** parameter can be a directory on any mounted local disk. It is not possible to use a mapped network drive as a restore destination.

- When you use the **restorefiles** command to restore local VSS backups, the command must be issued from the system on which the snapshot was created.


Related concepts:

“Mailbox restore operations” on page 26

Restorefiles syntax

Use the **restorefiles** command syntax diagrams as a reference to view available options and truncation requirements.

TDPEXCC command



Restorefiles positional parameters

Positional parameters immediately follow the **restorefiles** command and precede the optional parameters.

The following positional parameters specify the object to restore:

|*database_name1*,...,*database_nameN

***** Sequentially restore all flat files for the database.

dbname

Restore the specified database files. Multiple entries are separated by commas.

The following positional parameters specify the type of backup from which the files are restored:

FULL | COPY | INCRemental | DIFFerential

FULL Restore the files from a Full type backup for VSS.

COPY Restore the files from a Copy type backup for VSS.

INCRemental

Restore the files from an Incremental type backup for VSS.

DIFFerential

Restore the files from a Differential type backup for VSS.

Restorefiles optional parameters

Optional parameters follow the **restorefiles** command and positional parameters.

/BACKUPDESTINATION

VSS backups can have a backup destination of TSM or LOCAL.

/CONFIGfile=*configfilename*

Use the **/configfile** parameter to specify the name of the Data Protection for Exchange Server configuration file that contains the values for the Data Protection for Exchange Server configuration options.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the Data Protection for Exchange Server installation directory is used. If the **/configfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `tdpexc.cfg`.

If the *configfilename* variable includes spaces, enclose the entire **/configfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\file.cfg"
```

DAGNODE=*nodename*

Specify the node name that you want to use to back up the databases in an Exchange Server Database Availability Group. With this setting, backups from all Database Availability Group members that are configured to use the DAG node are backed up to a common file space on the IBM Spectrum Protect server. The database copies are managed as a single entity, regardless of which Database Availability Group member they are backed up from. This setting can prevent IBM Spectrum Protect Snapshot for Microsoft Exchange Server from making too many backups of the same database.

/FROMEXCServer=server-name

Use the **/fromexcserver** parameter to specify the name of the Exchange Server where the original backup was done. The default is the local Exchange Server name.

If a DAG node is specified by using the **/dagnode** parameter, Data Protection for Exchange Server uses this node name instead of the Data Protection for Exchange Server node to back up databases in an Exchange Server Database Availability Group. Therefore, the **restore** command automatically restores the backups that were created by the other DAG members, without having to specify the **/fromexcserver** parameter.

/INTO=pathname

Use the **/into** parameter to specify the root directory where files must be restored. The **restorefiles** operation creates a subdirectory under the root directory that contains the name of the database. Restored files are placed in that subdirectory. If the **/into** parameter is not specified, the files are restored into the directory where the **restorefiles** command is issued. The **/into** parameter can be a directory on any mounted local disk. It is not possible to use a mapped network drive as a restore destination.

/LOGFile=logfilename

Use the **/logfile** parameter to specify the name of the activity log file that is generated by Data Protection for Exchange Server.

The *logfilename* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfilename* variable can include a fully qualified path. However, if no path is specified, the log file is written to the Data Protection for Exchange Server installation directory.

If the *logfilename* variable includes spaces, enclose the entire **/logfile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\mytdpexchange.log"
```

If the **/logfile** parameter is not specified, log records are written to the default log file, *tdpexc.log*.

The **/logfile** parameter cannot be turned off, logging always occurs.

When you use multiple simultaneous instances of Data Protection for Exchange Server to run operations, use the **/logfile** parameter to specify a different log file for each instance used. This parameter directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=numdays | No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option **No** can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify *no*, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/MOUNTWait=Yes | No

Use the **/mountwait** parameter to specify whether Data Protection for Exchange Server waits for removable media to mount (such as tapes or CDs) or to stop the current operation. This situation occurs when the IBM Spectrum Protect server is configured to store backup data on removable media and waits for a required storage volume to be mounted.

You can specify:

Yes Wait for tape mounts. This option is the default.

No Do not wait for tape mounts.

/OBJECT=object name

Use the **/object** parameter to specify the name of the backup object files that you want to restore. The object name uniquely identifies each backup object and is created by Data Protection for Exchange Server.

Use the Data Protection for Exchange Server **query tsm *** command to view the names of the backup objects.

/Quiet This parameter prevents status information from being displayed. The level of information that is written to the activity log is not affected.

/TSMNODE=tsmnodename

Use the *tsmnodename* variable to refer to the IBM Spectrum Protect node name that Data Protection for Exchange Server uses to log on to the IBM Spectrum Protect server.

You can store the node name in the IBM Spectrum Protect options file (*dsm.opt*). This parameter overrides the value in the IBM Spectrum Protect options file if **PASSWORDACCESS** is set to PROMPT. This parameter is not valid when **PASSWORDACCESS** is set to GENERATE in the options file.

/TSMOPTFile=tsmoptfilename

Use the *tsmoptfilename* variable to identify the Data Protection for Exchange Server options file.

The file name can include a fully qualified path name. If no path is specified, the directory where Data Protection for Exchange Server is installed is searched.

If the *tsmoptfilename* variable includes spaces, enclose the entire **/tsmoptfile** parameter entry in double quotation marks. For example:
`/TSMOPTfile="c:\Program Files\file.opt"`

The default is `dsm.opt`.

/TSMPassword=tsmpassword

Use the *tsmpassword* variable to refer to the IBM Spectrum Protect password that Data Protection for Exchange Server uses to log on to the IBM Spectrum Protect server.

If you specified **PASSWORDACCESS GENERATE** in the Data Protection for Exchange Server options file (`dsm.opt`), supplying the password here is not necessary because the one that is stored in the registry is used. However, to store the password in the registry, you must specify the IBM Spectrum Protect password the first time Data Protection for Exchange Server connects to the IBM Spectrum Protect server.

If you do specify a password with this parameter when **PASSWORDACCESS GENERATE** is in effect, the command-line value is ignored unless the password for this node is not yet stored in the registry. In that case, the specified password is stored in the registry and used when you run this command.

If **PASSWORDACCESS PROMPT** is in effect, and you do not specify a password value on the command line, then you are prompted for a password.

The IBM Spectrum Protect password that Data Protection for Exchange Server uses to log on to the IBM Spectrum Protect server can be up to 63 characters in length.

Restoremailbox command

To restore mailbox-level data or mailbox-item-level data from Data Protection for Exchange Server backups, use the **restoremailbox** command.

The **restoremailbox** command applies to any Data Protection for Exchange Server VSS backups:

- VSS backups that are stored on IBM Spectrum Protect server
- VSS backups that are stored on local shadow volumes

When you use the **restoremailbox** command, follow these guidelines:

- Ensure that you have the required role-based access control (RBAC) permissions to complete individual mailbox restore operations.
- You can restore multiple mailboxes in a single mailbox restore operation.
- You can use the **restoremailbox** command to restore data to a mailbox on the Exchange Server or to an Exchange Server.pst file.

When you restore to a Unicode .pst file, except for the **Folder Name** and **All Content** filters, the filters are ignored. The amount of time that is needed to complete the restore process depends on the size of the mailbox databases, the network speed, and the number of mailboxes to process.

For non-Unicode .pst files, you can limit the range of the mailbox data to restore by using the **/mailboxfilter** parameter to specify filters that are based on the following mailbox message elements:

- Sender name
- Folder name

- Message body
- Subject line
- Attachment name
- Range of the message delivery date and time
- You can use the **restoremailbox** command on the primary Exchange Server or on an alternate Exchange Server that is in the same domain.
- In Exchange 2013, you can use the **restoremailbox** command to restore a public folder mailbox database, a public folder mailbox, or only a part of the mailbox, for example, a specific public folder.
 - To restore an Exchange 2013 public folder mailbox, you must have the Public Folders management role.
 - You can restore a public folder mailbox only to an existing public folder mailbox that is on the Exchange server.
 - You can restore a public folder only to an existing public folder. The public folder on the Exchange server must have the same folder path as the public folder to be restored. If the public folder is deleted from the public folder mailbox on the Exchange server, you must re-create the public folder with the same folder path as the public folder to be restored, before you start the restore operation.
 - As a best practice, restore public folder mailboxes separately from user mailboxes. Select only one public folder mailbox to restore at a time if you want to restore a specific public folder in the mailbox, or if you want to restore to a different public folder mailbox than the original mailbox.

If you restore multiple mailboxes in a single restore operation, and at least one of the mailboxes is a public folder mailbox, the mailboxes are restored only to their original mailbox locations. You cannot specify a filter or an alternate mailbox destination.
 - You can restore to a different public folder mailbox than the original mailbox if, for example, the public folder was relocated since the time of the backup. Before you complete the public folder restore operation, ensure that the public folder exists with the same folder path in the alternate mailbox location.
- In Exchange Server 2010 or later, you can use the **restoremailbox** command to restore an archive mailbox or only a part of the mailbox, for example, a specific folder. You can restore archive mailbox messages to an existing mailbox on the Exchange server, to an archive mailbox, or to an Exchange Server .pst file.

If you enable a user mailbox to be archived, ensure that the user is logged on to that mailbox at least once before you complete a backup and restore operation on the mailbox.
- You can use the **restoremailbox** command with the following parameter and options:
 - Set the **/KEEPADB** parameter option to Yes to retain a recovery database after one or more mailboxes are restored. Set the parameter value to No to automatically remove the recovery database after mailbox restore processing. Regardless of the option that you set, Data Protection for Exchange Server retains the recovery database if the mailbox restore operation fails after the recovery database is successfully restored.

If you restore multiple mailboxes, and you want to retain the recovery database after the restore operation is complete, ensure that all mailboxes are in the same recovery database.

- Set the **/USEEXISTINGRDB** parameter option to **Yes** to restore a mailbox from an existing recovery database. Set the parameter value to **No** to automatically remove the existing recovery database during mailbox restore processing.
- Set the **/mailboxoriglocation** parameter to specify the server and the database where the mailbox is at the time of backup. You set this option when the mailbox history is disabled and when the mailbox that you are restoring is either moved or deleted since the time of the backup.
- If a mailbox is deleted or re-created since the time of the backup, you must use a temporary mailbox with enough capacity to contain all of the mailbox items that you are restoring. The mailbox of the user who is logged in is used as temporary mailbox by default.

You can set the **/tempmailboxalias** optional parameter by selecting **Properties** from the Actions pane. In the **Data Protection Properties** window, select the **General** page, where you can specify the temporary log restore path, the temporary database restore path and the alias of the temporary mailbox.

- You can use the **restoremailbox** command to recover and restore different types of mail items in the Recoverable Items folder.
 - The mail items that you can restore depends on whether the mailbox is enabled for mailbox restore operations.
 - You cannot restore the Recoverable Items folder and subfolder hierarchy to a mailbox restore destination. You can restore only the contents of the email folders.
 - You cannot add a subfolder to the Recoverable Items folder in a mailbox.

Related tasks:


“Restoring mailboxes directly from Exchange database files” on page 119


“Restoring mailbox data” on page 108

Restoremailbox syntax


Use the **restoremailbox** command syntax diagram as a reference to view available options and truncation requirements.


TDPEXCC command





/MAILBOXRESTOREDESTINATION options:





Notes:

- 1 You can specify the **/MAILBOXFILTER** parameter multiple times however, you must specify each **/MAILBOXFILTER** subparameter only once.

Restoremailbox positional parameters

Positional parameters immediately follow the **restoremailbox** command and precede the optional parameters.

original-mailboxname

Use this parameter to specify the name of the mailbox to restore from. The mailbox name can be either the mailbox-alias, the mailbox-display name, or the mailbox globally unique identifier (GUID). The *original-mailboxname* parameter is required.

To specify more than one name, separate them by commas.

If any mailbox name contains blanks, enclose the entire mailbox name in double quotation marks.

Restoremailbox optional parameters

Optional parameters follow the **restoremailbox** command and positional parameters.

/CONFIGfile=configfilename

Use the **/configfile** parameter to specify the name of the Data Protection for Microsoft Exchange Server configuration file that contains the values for the Data Protection for Microsoft Exchange Server configuration options.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the Data Protection for Microsoft Exchange Server installation directory is used. If the **/configfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `tdpexc.cfg`.

If the *configfilename* variable includes spaces, enclose the entire **/configfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\file.cfg"
```

DAGNODE=nodename

Specify the node name that you want to use to back up the databases in an Exchange Server Database Availability Group. With this setting, backups from all Database Availability Group members that are configured to use the DAG node are backed up to a common file space on the IBM Spectrum Protect server. The database copies are managed as a single entity, regardless of which Database Availability Group member they are backed up from. This setting can prevent IBM Spectrum Protect Snapshot for Microsoft Exchange Server from making too many backups of the same database.

/FROMArchive

Use the **/FROMArchive** parameter only if you are restoring an archive mailbox and you specify the mailbox alias of the primary mailbox. If you

specify the primary mailbox alias and you do not specify this parameter option, by default, the primary mailbox is restored.

To restore an archive mailbox to another archive mailbox, specify both the `/FROMArchive` and the `/MAILBOXRESTOREDESTINATION=ARCHIVEMAILBOX,target-mailboxname` parameters. For example:

```
tdpexcc restoremailbox "OriginalArchiveMailboxName" /FROMArchive  
/MAILBOXRESTOREDESTINATION=ARCHIVEMAILBOX,"TargetArchiveMailboxName"
```

/EXCLUDEDUMPster=No|Yes

Use the `/EXCLUDEDUMPster` parameter to specify whether Data Protection for Exchange Server includes or excludes the mail items in the Recoverable Items folder in mailbox restore operations.

You can specify the following values:

- No** Restore the mail items in the Recoverable Items folder to a mailbox restore destination. This option is the default.
- Yes** Do not restore the mail items in the Recoverable Items folder to a mailbox restore destination.

If you are restoring the mailbox of George Clark, for example, you can exclude the Recoverable Items folder contents as shown in the following example:

```
tdpexcc restoremailbox "George Clark" /EXCLUDEDUMPster=YES  
/USEEXISTINGRDB=NO /KEEPRDB=NO
```

/KEEPRDB=No|Yes

Use the `/KEEPRDB` parameter to specify whether Data Protection for Microsoft Exchange Server retains a recovery database for reuse in mailbox restore operations, or automatically removes the recovery database after mailbox restore operations.

You can specify the following values:

- No** Do not retain a recovery database for mailbox restore operations. Remove the recovery database after mailbox restore processing. This option is the default.
- Yes** Retain the recovery database for mailbox restore operations.

/LOGFile=logfilename

Use the `/logfile` parameter to specify the name of the activity log file that is generated by Data Protection for Microsoft Exchange Server.

The `logfilename` variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The `logfilename` variable can include a fully qualified path. However, if no path is specified, the log file is written to the Data Protection for Microsoft Exchange Server installation directory.

If the `logfilename` variable includes spaces, enclose the entire `/logfile` parameter in double quotation marks. For example:

```
/LOGFile="c:\Program Files\mytdpexchange.log"
```

If you do not specify the `/logfile` parameter, log records are written to the default log file, `tdpexc.log`.

The `/logfile` parameter cannot be turned off, logging always occurs.

When you use multiple simultaneous instances of Data Protection for Microsoft Exchange Server to process operations, use the **/logfile** parameter to specify a different log file for each instance that is used. This parameter directs logging for each instance to a different log file and prevents interspersed log file records.

Note: Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=numdays | No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, some days of data are saved. By default, 60 days of log entries are saved. The option **No** can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify **no**, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/MAILBOXFILTER=ATTACHMENTNAME | ENDDATETIME | FOLDER | MESSAGEBODY | SENDER | STARTDATETIME | SUBJECT | ALLCONTENT

Use the **/MAILBOXFILTER** parameter to specify filters to restrict the mailbox data that is restored. When you are restoring to a Unicode .pst file, you can apply only the **FOLDER** and **ALLCONTENT** filters. You can apply only the **FOLDER** filter when you are restoring a public folder to an online public folder, or when you are restoring an archive mailbox folder.

You can specify multiple filters; however, you must specify each filter one time. For each filter that you specify, a separate **/MAILBOXFILTER** parameter must be used. For example:

```
tdpexc.exe restoremailbox dchang /MAILBOXFILTER=STARTDATETIME,07/01/2013 /MAILBOXFILTER=ENDDATETIME,07/31/2013
```

Mailbox data that matches a combination of all filters that are specified is restored. If no filters are specified, by default, all data in the mailbox is restored.

Specify one of the following filters when you use this parameter:

ATTACHMENTNAME,*attachmentname-search-text*

Use `/MAILBOXFILTER=attachmentname attachmentname-search-text` to restore only the mailbox messages that contain a match of the specified text within a message attachment name. The match is not case-sensitive. For example, an *attachmentname-search-text* of Rob matches the attachment name: Rob, robert.txt, PROBE, and prObE.pdf.

Enclose the *attachmentname-search-text* variable in double quotation marks.

Note: The ATTACHMENTNAME filter does not match the attachment names of encrypted mailbox messages. If a mailbox message is encrypted, it is skipped by the ATTACHMENTNAME filter.

ENDDATETIME,*end-date[,end-time]*

Use `/MAILBOXFILTER=enddatetime, end-date, end-time` to restore only the mailbox messages that are sent or received earlier than the specified date and time.

The *end-date* variable is required. Use the same date format for the *end-date* that you selected with the DATEFORMAT option in the Data Protection for Exchange options file.

The *end-time* variable is optional. Use the same time format for the *end-time* variable that you selected with the TIMEFORMAT option in the Data Protection for Exchange options file.

The ENDDATETIME filter date and time must be later than the STARTDATETIME filter date and time. If no time is specified, all messages that are sent or received on that date are restored.

FOLDER,*folder-name*

Use `/MAILBOXFILTER=folder, original-folder-name` to restore only the mailbox messages that are in the specified folder within the original mailbox. The match is not case-sensitive.

Enclose the *original-folder-name* variable in double quotation marks.

- To filter a public folder to restore, ensure that you are restoring the folder to an existing public folder that has the same folder path as the public folder to be restored. If the original public folder is deleted after the time of the backup, re-create the public folder. Specify the full path to the folder. If the full directory path includes spaces, enclose the directory path in double quotation marks, and do not append a backslash character (\) at the end of the directory path.

For example, to restore a folder that is named "SubFolder" under "ParentFolder", specify "ParentFolder/SubFolder" as the folder path. To restore all folders in a parent folder, use *ParentFolder/**.

- To restore a specific folder in an archive mailbox, ensure that you specify the full directory path to the folder.

To restore an archive mailbox to another archive mailbox, you must specify both the `/MAILBOXFILTER=folder, original-folder-name` parameter and the

`/MAILBOXRESTOREDESTINATION=ARCHIVEMAILBOX,target-mailboxname` parameter. For example:

```
tdpexcc restoremailbox "OriginalArchiveMailboxName"  
/MailboxFilter=folder,"folderA" /MAILBOXRESTOREDESTINATION=  
ARCHIVEMAILBOX,"TargetArchiveMailboxName"
```

- To restore the folder of a mailbox to a Unicode .pst file, ensure that you specify the `/MAILBOXFILTER=FOLDER` parameter with the `/MAILBOXRESTOREDESTINATION=UNICODEPST` parameter. Specify the full directory path to the folder. For example, to restore a folder that is named "SubFolder" under "ParentFolder", specify "ParentFolder/SubFolder" as the folder path. To restore all folders in a parent folder, use `ParentFolder/*`.
- To restore only the mail items in the Deletions subfolder of the Recoverable Items/ folder, specify the `/MAILBOXFILTER=FOLDER` parameter with the correct folder value for the target destination.

As shown in the following example, if you are restoring mail items to the original mailbox, specify the Deletions folder.

```
tdpexcc restoremailbox "george clark" /MailboxFilter=folder,  
"Deletions"
```

If you are restoring mail items to a Unicode .pst file, specify the full folder path to the Deletions folder.

```
tdpexcc restoremailbox "george clark" /MailboxFilter=folder,  
"Recoverable Items/Deletions" /KEEPADB=NO /USEExistingrdb=NO  
/MAILBOXRESTOREDESTINATION=UNICODEPST,c:\gclark.pst
```

MESSAGEBODY,*messagebody-search-text*

Use `/MAILBOXFILTER=messagebody,messagebody-search-text` to restore only the mailbox messages that contain a match of the specified text within the message body. The match is not case-sensitive. For example, a *messagebody-search-text* of Rob matches the message body text: Rob, robert, PROBE, and pr0be.

Enclose the *messagebody-search-text* variable in double quotation marks.

The **MESSAGEBODY** filter does not match the message body of encrypted mailbox messages. If a mailbox message is encrypted, it is skipped by the **MESSAGEBODY** filter.

SENDER,*sender-name*

Use `/MAILBOXFILTER=sender,sender-name` to restore only the mailbox messages that are received from the specified message sender.

Enclose the *sender-name* variable in double quotation marks.

STARTDATETIME,*start-date[,start-time]*

Use `/MAILBOXFILTER=startdatetime,start-date,start-time` to restore only the mailbox messages that are sent or received after the specified date and time.

The *start-date* variable is required. Use the same date format for the *start-date* that you selected with the `DATEFORMAT` option in the Data Protection for Exchange options file.

The *start-time* variable is optional. Use the same time format for the *start-time* variable that you selected with the `TIMEFORMAT` option in the Data Protection for Exchange options file."

The **STARTDATETIME** filter date and time must be earlier than the **ENDDATETIME** filter date and time. If no time is specified, all messages that are sent or received on that date are restored.

SUBJECT,*subject-search-text*

Use `/MAILBOXFILTER=subject,subject-search-text` to restore only the mailbox messages that contain a match of the specified text within the message subject line. The match is not case-sensitive. For example, a *subject-search-text* of Rob matches the subject text: Rob, robert, PROBE, and prObE.

Enclose the *subject-search-text* variable in double quotation marks.

ALLCONTENT,*allcontent-search-text*

Use `/MAILBOXFILTER=allcontent,allcontent-search-text` to restore only the mailbox messages that contain a match of the specified text that is contained within the message sender, message subject line, message body, or message attachment. The match is not case-sensitive. For example, an *allcontent-search-text* of Rob matches Rob, robert, PROBE, and prObE contained within the message sender, the subject line, or the message body.

Enclose the *allcontent-search-text* variable in double quotation marks.

The **ALLCONTENT** filter does not match the message body of encrypted mailbox messages. If a mailbox message is encrypted, the **ALLCONTENT** filter matches only text that is contained within the message sender or the subject line.

/MAILBOXORIGLOCATION=server-name,db-name

Use the `/mailboxoriglocation` parameter to specify the Exchange Server and the database where the mailbox is at the time of backup.

If you do not specify the `/mailboxoriglocation` parameter, the default value is the location (found in the mailbox location history) of the mailbox to restore from, for the backup time specified. If no mailbox location history is available, the default value is the current active location of the mailbox.

server-name

The name of the Exchange Server where the mailbox is at the time of backup.

db-name

The name of the database where the mailbox is at the time of backup.

The `/mailboxoriglocation` parameter is only necessary if the mailbox to be restored from is moved or deleted after the time of the backup, and no mailbox location history is available. This parameter is case-sensitive. Data Protection for Microsoft Exchange Server 6.1 (or later versions) maintains mailbox location history.

A **restoremailbox** operation from a backup that is processed by Data Protection for Microsoft Exchange Server before version 6.1 fails if the `/mailboxoriglocation` parameter is not specified for mailboxes that meet one or both of the following conditions:

- The mailbox to be restored is moved. The mailbox is not in the same server and the same database where the mailbox is at the time of the backup.
- The mailbox to be restored is deleted and the restore destination is to an alternate mailbox or to a .pst file.

For example:

```
TDPEXCC RESTOREMAILBOX johngrimshawe /MAILBOXORIGLOCATION=serv1,mdb1
/MAILBOXRESTOREDate=03/06/2013
/MAILBOXRESTOREDESTINATION=PST,c:\team54\rcvr.pst
```

/MAILBOXRESTOREDate=restore-date

Use the **/mailboxrestoredate** parameter with or without the **/mailboxrestorettime** parameter to establish a date and time to restore mailbox data from. A mailbox is restored from the earliest backup that completed after the date and time is established by the **/mailboxrestoredate** and the **/mailboxrestorettime** parameters.

The backup after the specified time is selected because, if an earlier backup is selected, Data Protection for Microsoft Exchange Server misses the emails that are sent between the time of the backup and the specified time. By selecting the first backup after the specified time, Data Protection for Microsoft Exchange Server ensures that all of the emails, up to the specified time, are restored. Specify the appropriate date in the *restore-date* variable; use the same format that you selected with the DATEFORMAT option in the Data Protection for Microsoft Exchange Server options file.

If *restore-date* or *restore-time* values are not specified, no date and time is established. By default the mailbox is restored from the most recent available backup.

If either *restore-date* or *restore-time* is specified, then the mailbox is restored from the earliest backup that is taken after the established restoration date and time. If no backup of the mailbox after the established date and time is found, by default the mailbox is restored from the most recent available backup.

- If you specify both *restore-date* and *restore-time*, this selection establishes the mailbox restoration period.
- If you specify *restore-date*, and you do not specify *restore-time*, *restore-time* defaults to a value of 23:59:59. This selection establishes the *restore-date* at the specified date.
- If you specify *restore-time* without *restore-date*, then *restore-date* defaults to the current date. This selection establishes the restoration date and time as the current date at the specified *restore-time*.
- Only default time and date formats are accepted. If you use a format other than the default format to specify the time and date, the date and time is ignored.

/MAILBOXRESTORETime=restore-time

Use the **/mailboxrestorettime** parameter with or without the **/mailboxrestoredate** parameter to establish a date and time to restore a mailbox from. A mailbox is restored from the earliest backup that completed after the date and time is established by the **/mailboxrestoredate** and the **/mailboxrestorettime** parameters.

The backup after the specified time is selected because, if an earlier backup is selected, Data Protection for Microsoft Exchange Server misses the emails that are sent between the time of the backup and the specified time.

By selecting the first backup after the specified time, Data Protection for Microsoft Exchange Server ensures that all of the emails, up to the specified time, are restored. Specify the appropriate time in the *restore-time* variable; use the same format that you selected with the TIMEFORMAT option in the Data Protection for Microsoft Exchange Server options file.

If *restore-date* and *restore-time* values are not specified, no date and time is established. By default the mailbox is restored from the most recent available backup.

If either *restore-date* or *restore-time* is specified, the mailbox is restored from the earliest backup that completed after the established date and time. If no backup of the mailbox after the established date and time is found, by default the mailbox is restored from the most recent available backup.

- If you specify both *restore-date* and *restore-time*, this selection establishes the mailbox restoration period.
- If you specify *restore-date*, and you do not specify *restore-time*, *restore-time* defaults to a value of 23:59:59. This selection establishes the *restore-date* at the specified date.
- If you specify *restore-time* without *restore-date*, then *restore-date* defaults to the current date. This selection establishes the restoration date and time as the current date at the specified *restore-time*.

/MAILBOXRESTOREDESTINATION=EXCHANGE | PST | UNICODEPST | ARCHIVEMAILBOX

Use the **/mailboxrestoredestination** parameter to specify the destination to restore the mailbox data to.

If you do not specify the **/mailboxrestoredestination** parameter, by default, the EXCHANGE option is used and the **/mailboxrestoredestination** is not required. The default system behavior is to restore mailbox data to the original location in the original active mailbox. When you restore multiple mailboxes with the same **restoremailbox** command, the default system behavior is to restore mailbox data into each original active mailbox.

Mailbox items are merged into the mailbox destination. If a mailbox item exists in the mailbox destination, that item is not restored.

You must specify one of the following values when you use this parameter:

EXCHANGE,[target-mailboxname,target-foldername]

Use the **/mailboxrestoredestination=EXCHANGE** option to restore mailbox messages into a live Exchange Server.

The EXCHANGE option is the default option. If you specify the **/mailboxrestoredestination=EXCHANGE** option without specifying any variables, the result is the same as not specifying the **/mailboxrestoredestination** parameter. The mailbox data is restored to the original location in the original active mailbox.

Use **/mailboxrestoredestination=EXCHANGE,target-mailboxname,target-foldername** to restore mailbox messages into a destination other than the original location in the original active mailbox. The mailbox messages are restored into a subfolder of the specified folder within the target mailbox. The target mailbox can be the original mailbox or an alternate mailbox.

When you restore multiple mailboxes with the same **restoremailbox** command, this option restores the mailbox data into a subfolder (designated by each original mailbox-alias) of the

specified target folder in the active mailbox. The folders from the corresponding original mailbox, which contain the restored mailbox messages, are in each subfolder. The specified folder in the target mailbox contains a subfolder that is designated by the original mailbox alias name. Subfolders that contain the restored mailbox messages are in each parent subfolder. These child subfolders have the folder structure of the original mailbox.

target-mailboxname

Specify the target mailbox-alias or the target mailbox-display name. The target mailbox must be an active mailbox.

If the *target-mailboxname* variable includes spaces, enclose the entry in double quotation marks.

To restore a specific public folder to an alternate public folder mailbox, specify both the `/MAILBOXFILTER=folder,original-folder-name` parameter and the `/MAILBOXRESTOREDESTINATION=EXCHANGE,target-publicfolder-mailboxname` parameter. For example:

```
tdpexcc restoremailbox "OriginalPublicFolderMailbox"  
/MailboxFilter=folder,"folderA" /MAILBOXRESTOREDESTINATION=  
EXCHANGE,"TargetPublicFolderMailbox"
```

You can restore a public folder only to an existing public folder on the Exchange server. If the public folder is relocated to an alternate mailbox destination after the time of the backup, ensure that it exists in the alternate mailbox location with the same folder path as the folder to be restored. The restore operation does not automatically re-create the public folder in the destination mailbox.

target-foldername

The *target-foldername* variable specifies the mailbox folder in the target mailbox to restore mailbox messages to.

If you restore a mailbox to a different destination than the original mailbox, the mailbox folders are restored in the destination mailbox under a folder that is named *original-mailbox_mailbox-GUID*. In the process, the Recoverable Items folders are restored.

If you specify the *target-mailboxname* variable and the target mailbox is not the original mailbox, you must specify a folder name. However, when you restore to a mailbox in a target public folder, do not specify a target folder name. A folder name is not required for public folder restore operations.

If the mailbox folder specified by the *target-foldername* variable does not exist in the target mailbox, a folder with the target folder name is created in the target mailbox except for public folder mailboxes.

The target folder contains one subfolder for each original-mailbox that is restored (designated by each original-mailbox alias). The folders from the corresponding original mailbox, which contain the restored mailbox messages, are in each subfolder. If you did not specify the

/mailboxfilter parameter, the target folder that you specified contains, within the subfolder that is designated by the original mailbox alias, all the folders that are in the mailbox that you are restoring from. If you specified the **/mailboxfilter** parameter, the subfolder within the folder that you specified contains only the folders with messages that match the filter criteria.

If the *target-foldername* variable includes spaces, enclose the entire *target-foldername* variable entry in double quotation marks. For example:

```
/MAILBOXRESTOREDESTINATION=EXCHANGE,Kerry,"temp folder"
```

When you restore multiple mailboxes with the same **restoremailbox** command, and you specify a target folder, each original-mailbox is restored to the target folder in the target mailbox. The target folder contains one subfolder for each original-mailbox that is restored (designated by each original mailbox alias). The folders from the corresponding original mailbox, which contain the restored mailbox messages, are in each subfolder.

For example, this **restoremailbox** operation restores mailboxes "andrew baker" and "sally wood" to the folder "previous_acctmng" in the target mailbox "mary brown":

```
restoremailbox "andrew baker","sally wood"  
/mailboxrestoredest=exchange,"mary brown",previous_acctmng
```

The restored mailbox messages are placed in folders that are copied from the original mailboxes that use the following folder structure:

```
mary brown (target mailbox)
  >-previous_acctmng (specified folder)
    >-abaker (original-mailbox1 alias)
      >-Inbox (restored folder from mailbox1)
      >-Outbox (restored folder from mailbox1)
      >-My Accts (restored folder from mailbox1)
    >-swood (original-mailbox2 alias)
      >-Inbox (restored folder from mailbox2)
      >-Outbox (restored folder from mailbox2)
      >-New Accts (restored folder from mailbox2)
```

PST,non-Unicode-pst-filename-path

Use **/mailboxrestoredestination=PST,non-Unicode-pst-filename-path** to restore mailbox data to an Exchange Server personal folders (.pst) file. The mailbox data that is restored is in non-Unicode format.

You can include the *non-Unicode-pst-filename-path* variable to specify the destination where the **restoremailbox** operation writes the .pst file. The *non-Unicode-pst-filename-path* can be either a fully qualified path to a .pst file or a directory path. If you do not specify a path, the .pst file is written to the current directory.

- You can specify *non-Unicode-pst-filename-path* as a fully qualified path to a .pst file to restore all mail to that .pst file.

```
TDPEXCC RESTOREMAILBOX gclark  
/mailboxrestoredestination=PST,c:\mb\dept54\vpo.pst
```

The .pst directory must exist before you use the **restoremailbox** command. The .pst file is created if it does not exist.

If you are restoring more than one mailbox and you specify a fully qualified path to a .pst file, all the mailbox data is restored to the one .pst file specified. Inside the .pst file, the parent-level folder name is the mailbox-alias-name, followed by the rest of the mailbox folders.

- You can specify *non-Unicode-pst-filename-path* as a directory path to have IBM Spectrum Protect Snapshot for Exchange Server create a .pst file by using the mailbox-alias-name of the mailbox that is being restored, and store the .pst file in the specified directory. For example, the .pst file name of the restored mailbox "George Clark" (gclark) is gclark.pst.

```
TDPEXCC RESTOREMAILBOX "george clark"  
/mailboxrestoredestination=PST,c:\mb\dept54\
```

The .pst directory must exist before you use the **restoremailbox** command. If the .pst file does not exist, the file is created.

If you restore multiple mailboxes with the same **restoremailbox** command, and you specify a directory path, each mailbox is restored into a separate .pst file. For example, if mailboxes John (john1), John Oblong (oblong), and Barney Olaf (barneyo) are restored and the specified directory path is c:\finance, all mailboxes are restored into the c:\finance directory as shown:

```
c:\finance\john1.pst  
c:\finance\oblong.pst  
c:\finance\barneyo.pst
```

The .pst directory must exist before you use the **restoremailbox** command. The mailbox data that is restored by using */mailboxrestoredestination=PST,non-Unicode-pst-filename-path* must be less than 2 GB.

If the *non-Unicode-pst-filename-path* variable includes spaces, enclose the entire *non-Unicode-pst-filename-path* variable entry in double quotation marks and do not append a backslash character (\) at the end of folder path. For example:

```
TDPEXCC RESTOREMAILBOX "george clark"  
/mailboxrestoredestination=PST,"c:\mb\dept54\access group"
```

UNICODEPST,Unicode-pst-filename-path

Use */mailboxrestoredestination=UNICODEPST,Unicode-pst-filename-path* to restore mailbox data to an Exchange Server personal folders (.pst) file. The mailbox data that is restored is in Unicode format.

You can include the *Unicode-pst-filename-path* variable to specify the destination where the **restoremailbox** operation writes the .pst file. The *Unicode-pst-filename-path* can be either a fully qualified UNC path to a .pst file or a directory path. If you do not specify a path, the .pst file is written to the current directory. If you specify a non-UNC path (such as c:\dir\mailbox.pst), IBM Spectrum Protect Snapshot for Exchange Server tries to convert it to a UNC path for you, but it might not work for custom UNC paths or shares.

- You can specify *Unicode-pst-filename-path* as a fully qualified path to a .pst file to restore all mail to that .pst file.

```
TDPEXCC RESTOREMAILBOX gclark
/mailboxrestoredestination=UNICODEPST,c:\mb\dept54\vp0.pst
```

Important: The .pst directory must exist before you use the **restoremailbox** command. The .pst file is created if it does not exist.

If you are restoring more than one mailbox and you specify a fully qualified path to a .pst file, all the mailbox data is restored to the one .pst file specified. Inside the .pst file, the parent-level folder name is the mailbox-alias-name, followed by the rest of the mailbox folders.

- You can specify *Unicode-pst-filename-path* as a directory path to have IBM Spectrum Protect Snapshot for Exchange Server create a .pst file by using the mailbox-alias-name of the mailbox that is being restored, and store the .pst file in the specified directory. For example, the .pst file name of the restored mailbox "George Clark" (gclark) is gclark.pst.

```
TDPEXCC RESTOREMAILBOX "george clark"
/mailboxrestoredestination=UNICODEPST,c:\mb\dept54
```

The .pst directory must exist before you use the **restoremailbox** command. If the .pst file does not exist, the file is created.

If you restore multiple mailboxes with the same **restoremailbox** command, and you specify a directory path, each mailbox is restored into a separate .pst file. For example, if mailboxes John (john1), John Oblong (oblong), and Barney Olef (barneyo) are restored and the specified directory path is c:\finance, all mailboxes are restored into the c:\finance directory as shown:

```
c:\finance\john1.pst
c:\finance\oblong.pst
c:\finance\barneyo.pst
```

- To restore only the mail items in the Deletions subfolder of the Recoverable Items/ folder, specify the /MAILBOXFILTER=FOLDER parameter with the correct folder value for the target destination.

As shown in the following example, if you are restoring mail items to a Unicode .pst file, specify the full folder path to the Deletions folder.

```
tdpexcc restoremailbox "george clark" /MailboxFilter=folder,
"Recoverable Items/Deletions" /KEEPADB=NO /USEExistingrdb=NO
/MAILBOXRESTOREDESTINATION=UNICODEPST,c:\gclark.pst
```

If the *Unicode-pst-filename-path* variable includes spaces, enclose the entire *Unicode-pst-filename-path* variable entry in double quotation marks and do not append a backslash character (\) at the end of folder path. For example:

```
TDPEXCC RESTOREMAILBOX "george clark"
/mailboxrestoredestination=UNICODEPST,"c:\mb\dept54\access group"
```

ARCHIVEMAILBOX,[target-mailboxname,target-foldername]

Use /MAILBOXRESTOREDESTINATION with the ARCHIVEMAILBOX and /FROMARCHIVE parameters to restore archive mailbox messages to its original archive mailbox or to an alternate archive mailbox.

Use /MAILBOXRESTOREDESTINATION=ARCHIVEMAILBOX,target-mailboxname to specify the archive mailbox destination that you want to restore to. You can also specify a target folder name in the archive mailbox.

To restore an archive mailbox into a specific folder of an archive mailbox, specify both the `/FROMArchive` parameter and the `/MAILBOXRESTOREDESTINATION=ARCHIVEMAILBOX,target-mailboxname,target-foldername` parameters. For example:

```
tdpexcc restoremailbox "OriginalFolderName" /FROMArchive
/MAILBOXRESTOREDESTINATION=ARCHIVEMAILBOX,"TargetFolderName"
"folderA"
```

If you specify the `/MAILBOXRESTOREDESTINATION=ARCHIVEMAILBOX` parameter without specifying a target mailbox destination, the mailbox messages are restored to the original location in the original archive mailbox.

/MOUNTWait=Yes | No

Use the `/mountwait` parameter to specify whether Data Protection for Microsoft Exchange Server waits for removable media to mount (such as tapes or CDs) or to stop the current operation. This situation occurs when the IBM Spectrum Protect server is configured to store backup data on removable media and waits for a required storage volume to be mounted.

You can specify:

- Yes** Wait for tape mounts. This option is the default.
- No** Do not wait for tape mounts.

/MOUNTRW=Yes | No

You can mount a read/write copy of your IBM Spectrum Protect backup so that you can modify the copy without invalidating the backup. You use this option to indicate whether a snapshot backup is mounted as read/write. The default value is as specified in the configuration file with the `/MOUNTRW` parameter. If a default value is not specified in the configuration file, the default value is No. By specifying the `/MOUNTRW` option, you override the default value.

The following values are available:

- No** A backup is mounted as read-only, which results in a file copy of the Microsoft Exchange database file to the RDB to complete the mailbox restore operation.
- Yes** A backup is mounted as read/write to do the mailbox restore operation. The backup is mounted on the directory you specify with the `/MOUNTPath` parameter. If a `/MOUNTPath` value is not specified, a temporary directory is used (system environment TEMP variable).

Note: When you specify the `/MOUNTRW` parameter for mailbox restore operations, the following parameters also apply:

- If both the `/MOUNTRW` and `/USEEXISTINGRDB` parameters are set to Yes and a recovery database (RDB) exists on the system, the existing RDB is used for mailbox restore operations and the `/MOUNTRW` parameter is ignored.
- If the `/KEEPRDB` parameter is specified, the snapshot RDB remains mounted on the system after the mailbox restore operation is complete (you must remove the snapshot RDB manually later). If you also specified the `/MOUNTRW` parameter, you must unmount the RDB by using the **unmount backup** command or the Windows Powershell cmdlet **Dismount-DpExcBackup**.

/TEMPDBRESTorepath=path-name

Use the **/tempdbrestorepath** parameter to specify the default temporary path to use when you restore mailbox database files.

If you do not specify the **/tempdbrestorepath** parameter, the default value is the value that is specified by the TEMPDBRESTOREPATH option in the Data Protection for Exchange configuration file. The default Data Protection for Microsoft Exchange Server configuration file is `tdpexc.cfg`. If the TEMPDBRESTOREPATH value does not exist in the Data Protection for Microsoft Exchange Server configuration file, the TEMP environment variable value is used.

If the *path-name* variable includes spaces, enclose the entire **/tempdbrestorepath** parameter entry in double quotation marks. For example:

```
TDPEXC RESTOREMAILBOX richgreene  
/tempdbrestorepath="h:\Exchange Restore Directory"
```

- Do not specify a value of **/TEMPDBRESTorepath** that is the same value as the location of the active database. If the value is the same, the database might become corrupted.
- Choose a temporary database-restore location that has enough space to hold the entire restore for the database.

For better performance, the current active-transaction logger is to be on a different physical device from the paths that are specified by the values of the **/TEMPDBRESTorepath** parameter and the **/TEMPLOGRESTorepath** parameter. The paths that are specified by the values of the **/TEMPDBRESTorepath** parameter and the **/TEMPLOGRESTorepath** parameter can be on the same or separate physical devices from each other.

Do not specify double-byte characters (DBCS) within the temporary database-restore path.

/TEMPLOGRESTorepath=path-name

Use the **/templogrestorepath** parameter to specify the default temporary path to use when you restore logs and patch files.

If you do not specify the **/templogrestorepath** parameter, the default value is the value that is specified by the TEMPLOGRESTOREPATH option in the Data Protection for Exchange configuration file. The default Data Protection for Microsoft Exchange Server configuration file is `tdpexc.cfg`. If you do not specify the **/templogrestorepath** parameter and the TEMPLOGRESTOREPATH value does not exist in the Data Protection for Microsoft Exchange Server configuration file, the TEMP environment variable value is used.

- Do not specify a value of **/TEMPLOGRESTorepath** that is the same value as the current location for the database that is used for recovery. If the value is the same, the database might become corrupted.
- Choose a temporary log-restore location that has enough space to hold all the log and patch files.

For better performance, the current active-transaction logger is to be on a different physical device from the paths that are specified by the values of the **/TEMPLOGRESTorepath** parameter and the **/TEMPDBRESTorepath** parameter. The paths that are specified by the values of the **/TEMPLOGRESTorepath** parameter and the **/TEMPDBRESTorepath** parameter can be on the same or separate physical devices from each other.

Do not specify double-byte characters (DBCS) within the temporary log-restore path.

/TSMNODE=tsmnodename

Use the *tsmnodename* variable to refer to the IBM Spectrum Protect node name that Data Protection for Microsoft Exchange Server uses to log on to the IBM Spectrum Protect server.

You can store the node name in the IBM Spectrum Protect options file (*dsm.opt*). This parameter overrides the value in the IBM Spectrum Protect options file if **PASSWORDACCESS** is set to PROMPT. This parameter is not valid when **PASSWORDACCESS** is set to GENERATE in the options file.

/TSMOPTFile=tsmoptfilename

Use the *tsmoptfilename* variable to identify the Data Protection for Microsoft Exchange Server options file.

The file name can include a fully qualified path name. If no path is specified, the directory where Data Protection for Microsoft Exchange Server is installed is searched.

If the *tsmoptfilename* variable includes spaces, enclose the entire **/tsmoptfile** parameter entry in double quotation marks. For example:

```
/TSMOPTFile="c:\Program Files\file.opt"
```

The default is *dsm.opt*.

/TSMPassword=tsmpassword

Use the *tsmpassword* variable to refer to the IBM Spectrum Protect password that Data Protection for Microsoft Exchange Server uses to log on to the IBM Spectrum Protect.

If you specified **PASSWORDACCESS** GENERATE in the Data Protection for Microsoft Exchange Server options file (*dsm.opt*), supplying the password here is not necessary because the one that is stored in the registry is used. However, to store the password in the registry, you must specify the IBM Spectrum Protect password the first time that Data Protection for Microsoft Exchange Server connects to the IBM Spectrum Protect server.

If you do specify a password with this parameter when **PASSWORDACCESS** GENERATE is in effect, the command-line value is ignored unless the password for this node is not yet stored in the registry. In that case, the specified password is stored in the registry and used when you run this command.

If **PASSWORDACCESS** PROMPT is in effect, and you do not specify a password value on the command line, then you are prompted for a password.

The IBM Spectrum Protect password that Data Protection for Microsoft Exchange Server uses to log on to the IBM Spectrum Protect server can be up to 63 characters in length.

/USEEXISTINGRDB=Yes | No

Use the **/USEEXISTINGRDB** parameter to specify whether Data Protection for Microsoft Exchange Server restores mailboxes from an existing recovery database, or automatically removes an existing recovery database during mailbox restore operations.

You can specify the following values:

Yes Use an existing recovery database for mailbox restore operations. This option is the default option.

- No** Do not use an existing recovery database for mailbox restore operations. Remove the recovery database during mailbox restore processing.

Examples: restoremailbox command

You can combine the use of the `/KEEPRDB` and `/USEEXISTINGRDB` parameter options with the `restoremailbox` command.

Example: Use an existing recovery database for mailbox operations

Use an existing recovery database for restore mailbox operations so that you do not have to restore the recovery database again.

```
tdpexcc restoremailbox <MB> /USEEXISTINGRDB=Yes
```

Example: Retain a recovery database for mailbox operations

Retain a recovery database after a mailbox restore operation so that you can use the recovery database for other restore operations.

```
tdpexcc restoremailbox <MB> /KEEPRDB=YES
```

Example: Retain a recovery database for multiple mailbox restore operations, and then remove it

Because you restore multiple mailboxes at different times, you want to retain the recovery database after the first mailbox restore operation and use it for subsequent restore operations. When you restore the final mailbox, you remove the recovery database.

```
tdpexcc restoremailbox <MB_1> /KEEPRDB=YES  
tdpexcc restoremailbox <MB_2> /USEEXISTINGRDB=YES  
tdpexcc restoremailbox <MB_n> /KEEPRDB=NO
```

Example: Restore multiple mailboxes simultaneously

Simultaneously restore multiple mailboxes and ensure that the recovery database is automatically removed after each mailbox is restored.

```
tdpexcc restoremailbox <MB_1>,<MB_2> /KEEPRDB=NO
```

Example: Restore multiple mailboxes from an existing recovery database

Simultaneously restore multiple mailboxes from an existing recovery database.

Tip: Mailboxes that are not in the recovery database are bypassed during restore processing, and are indicated in the console output.

Restore the remaining mailboxes that are not in the recovery database.

```
tdpexcc restoremailbox <MB_1>,<MB_2>...<MB_n>  
/USEEXISTINGRDB=YES /KEEPRDB=NO  
tdpexcc restoremailbox <MB_1>,<MB_2>...<MB_n>  
/USEEXISTINGRDB=NO /KEEPRDB=NO
```

Set command

To set the Data Protection for Exchange Server configuration parameters in a configuration file, use the `set` command.

The values that you set are saved in a Data Protection for Exchange Server configuration file. The default file is `tdpexc.cfg`. Configuration values can also be set in the Data Protection Properties window in Microsoft Management Console (MMC).


For command invocations other than this command, the value of a configuration parameter that is specified in a command overrides the value of the configuration

parameter that is specified in the Data Protection for Exchange Server configuration file. When you use this command, if you do not override a value for the configuration file parameter, the values in the default configuration file are used.

Set syntax

Use the **set** command syntax diagrams as a reference to view available options and truncation requirements.

TDPEXCC command: SET



Related tasks:

“Configuring your system for mailbox restore operations” on page 74

Set positional parameters

Positional parameters immediately follow the **set** command and precede the optional parameters.

The following positional parameters specify the values in the Data Protection for Exchange Server configuration file. You can set one value for each **tdpexcc set** command that you run:

BACKUPDESTINATION=TSM | LOCAL | BOTH

Use the **BACKUPDESTINATION** positional parameter to specify the storage location for your backup. You can specify these options:

TSM The backup is stored on IBM Spectrum Protect server storage only. This option is the default.

LOCAL

The backup is stored only on local shadow volumes.

BOTH The backup is stored on both IBM Spectrum Protect server storage and local shadow volumes.

CLIENTACCESSserver=servername

The *servername* variable refers to the name of the server you use to access the client.

DAGNODE=nodename

Specify the node name that you want to use to back up the databases in an Exchange Server Database Availability Group. With this setting, backups from all Database Availability Group members that are configured to use the DAG node are backed up to a common file space on the IBM Spectrum Protect server. The database copies are managed as a single entity, regardless of which Database Availability Group member they are backed up from. This setting can prevent IBM Spectrum Protect Snapshot for Microsoft Exchange Server from making too many backups of the same database.

DATEformat=dateformatnum

Use the **DATEformat** positional parameter to select the format that you want to use to display dates.

The *dateformatnum* variable displays the date in one of the following formats. Select the format number that corresponds to the format that you want to use.

- 1 (Default) MM/DD/YYYY
- 2 DD-MM-YYYY
- 3 YYYY-MM-DD
- 4 DD.MM.YYYY
- 5 YYYY.MM.DD
- 6 YYYY/MM/DD
- 7 DD/MM/YYYY

Changes to the value of the **dateformat** parameter can result in an undesired pruning of the Data Protection for Exchange Server log file (tdpexc.log by default). You can avoid losing existing log file data by doing one of the following actions:

- After you change the value of the **dateformat** parameter, make a copy of the existing log file before you run Data Protection for Exchange Server.
- Specify a new log file with the **/logfile** parameter.

IMPORTVSSSNAPSHOTSONLYWhenneeded

Use the **/IMPORTVSSSNAPSHOTSONLYWhenneeded** parameter to specify whether Data Protection for Exchange Server automatically imports VSS snapshots to the Windows system where the snapshots are created.

Specify one of the following values:

- Yes** Import VSS snapshots to the Windows system where the snapshots are created. The option is the default. During backup processing, transportable snapshots are automatically created and imported to storage systems when the snapshots are required. This option is the default.
- No** Do not create transportable VSS snapshots during backup processing, and do not automatically import the snapshot to storage systems after the backup is completed.

LOCALDSMAgentnode=nodename

Specify the node name of the local system that runs the VSS backups. This positional parameter must be specified for VSS operations to be processed.

LOGFile=logfilename

Use the **LOGFile** positional parameter to specify the name of the activity log file that is generated by Data Protection for Exchange Server. The Data Protection for Exchange Server activity log records significant events, such as completed commands and error messages.

The *logfilename* variable identifies the name of the activity log file. If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfilename* variable can include a fully qualified path. However, if no path is specified, the log file is assigned to the Data Protection for Exchange Server installation directory.

LOGPrune=numdays | No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. In the configuration file, the default value for the **LOGPrune** is that specified by the **logprune** configurable option. The default value is *60*, which means 60 days of log entries are saved. The option **No** can be specified to disable log pruning.

Regardless of the option that is set in the configuration file for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify the **LOGPrune** parameter, that value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.

- You can specify **LOGPrune** without specifying *numdays* or *no*; in this case, the default value, *60*, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **LOGFile** parameter or **logfile** setting.

/MOUNTRW=Yes | No

You can mount a read/write copy of your IBM Spectrum Protect backup so that you can modify the copy without invalidating the backup. You use this option to indicate whether a snapshot backup is mounted as read/write. The default value is as specified in the configuration file with the **/MOUNTRW** parameter. If a default value is not specified in the configuration file, the default value is No. By specifying the **/MOUNTRW** option, you override the default value.

The following values are available:

- No** Perform a read-only mount operation.
- Yes** Perform a read/write mount operation. The behavior of the read/write mount is controlled by the **USESNAPOFASNAPMount** parameter in the configuration file.
- If **USESNAPOFASNAPMount** is set to No, you can mount only COPY backups as read/write. After mounting, the original backup is modified and can no longer be used as a restore point in future database restore operations (on the **VSS Options** properties page, the **Mount read/write (modifies backup, applies to COPY backups only)** check box is selected).
 - If **USESNAPOFASNAPMount** is set to Yes, you can mount both FULL and COPY backup types as read/write (on the **VSS Options** properties page, the **Mount read/write (without modifying backup)** check box is selected). In this instance, the backups are not modified and can be used in future restore operations.

Important:

This mount option is only available for SAN Volume Controller (SVC) devices and requires IBM System Storage Support for Microsoft Virtual Disk and Volume Shadow Copy Services version 4.12 or later. Dynamic target allocation is not supported. You must allocate more target volumes on your SVC storage device to accommodate the number of concurrent read/write mounts you want to do. An extra target volume, that matches the size of the volume to be mounted, is needed for each concurrent read/write mount of that volume.

MOUNTWait=Yes | No

Use the **MOUNTWait** positional parameter to specify whether Data Protection for Exchange Server waits for removable media to mount (such as tapes or CDs) or to stop the current operation. This situation occurs when the IBM Spectrum Protect server is configured to store backup data on removable media and waits for a required storage volume to be mounted.

Specify **Yes** for Data Protection for Exchange Server to wait until all initial volumes of any required removable media are made available to the IBM Spectrum Protect server before the command completes.

Specify **No** for Data Protection for Exchange Server to end the command (if removable media are required). An error message displays.

NUMberformat=*fmtnum*

Use the **NUMb**erformat positional parameter to specify the format you want to use to display numbers.

The *fmtnum* variable displays numbers by using one of the following formats. Select the format number that corresponds to the format you want to use.

- | | |
|---|--------------------|
| 1 | (Default) n,nnn.dd |
| 2 | n,nnn,dd |
| 3 | n nnn,dd |
| 4 | n nnn.dd |
| 5 | n.nnn,dd |
| 6 | n'nnn,dd |

REMOTEDSMAgentnode=*nodename*

Specify the node name of the system that moves the VSS data to IBM Spectrum Protect server storage during offloaded backups.

STOREMAILBOXInfo=**Yes** | **No**

The **STOREMAILBOX**Info parameter is used to track mailbox history for moved and deleted mailboxes. By default, this parameter is set to **Yes**. If you do not plan to use mailbox restore, you can set this option to **No**. When the option is set to **No**, Data Protection for Exchange Server does not back up the mailbox history.

In large or geographically dispersed domains, more time is required to complete the backup mailbox history task. In this scenario, you can reduce the amount of time that is required to complete the backup mailbox history task by setting the option for **STOREMAILBOX**Info to **No**. When you set the option for **STOREMAILBOX**Info to **No**, mailboxes that are not moved or are not deleted can be restored normally. Moved and deleted mailbox restores can use the **/MAILBOXORIGLOCATION** parameter (of the **Restorem**ailbox command) to specify the mailbox location at the time of the backup.

TEMPDBREStorepath=*pathname*

To specify the default temporary path to use with mailbox database files, use the **TEMPDBRE**Storepath positional parameter.

If you do not enter a path, the default value is the value of the **TEMP** environment variable.

If the path name includes spaces, you must enclose the entire **TEMPDBRE**Storepath positional parameter entry in double quotation marks. For example:

```
TDPEXCC SET TEMPDBREStorepath="h:\Exchange Restore Directory"
```

Attention: Do not specify a value of **TEMPDBRESTorepath** that is the same value as the location of the active database. If the value is the same, the database might become corrupted.

Choose a temporary database-restore location that has enough space to hold the entire restore for the database.

Tip: For better performance, ensure that the current active-transaction logger is on a different physical device from the paths that are specified by the values of the **templogrestorepath** parameter setting and the **tempdbrestorepath** parameter setting. The paths that are specified by the values of the **templogrestorepath** parameter setting and the **tempdbrestorepath** parameter setting can be on the same or separate physical devices from each other.

Restriction: Do not specify double-byte characters (DBCS) within the temporary database-restore path.

TEMPLOGRESTorepath=pathname

To specify the default temporary path to use when you are restoring logs and patch files, use the **TEMPLOGRESTorepath** positional parameter.

If you do not enter a path, the default value is the value of the **TEMP** environment variable.

If the path name includes spaces, you must enclose the entire **TEMPDBRESTorepath** positional parameter entry in double quotation marks. For example:

```
TEMPLOGRESTorepath="c:\Program Files\templog"
```

Attention: Do not specify a value of **TEMPDBRESTorepath** that is the same value as the current location for the database that is used for recovery. If the value is the same, the database might become corrupted.

Choose a temporary log-restore location that has enough space to hold all the log and patch files.

Tip: For better performance, the current active-transaction logger is to be on a different physical device from the paths that are specified by the values of the **templogrestorepath** parameter setting and the **tempdbrestorepath** parameter setting. The paths that are specified by the values of the **templogrestorepath** parameter setting and the **tempdbrestorepath** parameter setting can be on the same or separate physical devices from each other.

Restriction: Do not specify double-byte characters (DBCS) within the temporary log-restore path.

TIMEformat=formatnumber

Use the **TIMEformat** positional parameter to specify the format in which you want system time that is displayed.

The *formatnumber* variable displays time in one of the following formats. Select the format number that corresponds to the format you want to use.

- 1 (Default) HH:MM:SS
- 2 HH,MM,SS
- 3 HH.MM.SS

USESNAPOFASNAPMount=Yes | No

During mount operations, you can specify that you want to do a read/write mount by setting `/MOUNTRW=Yes`. When you set the `/MOUNTRW=Yes`, the **USESNAPOFASNAPMount** parameter applies and you can further specify whether you want to mount an existing backup or to create a snapshot of an existing backup. You can only set the **USESNAPOFASNAPMount** parameter in your configuration file.

- If **USESNAPOFASNAPMount** is set to No, the **Mount read/write (modifies backup, applies to COPY backups only)** check box is selected on the **VSS Options** properties page. After mounting, the original COPY backup can be modified and so can no longer be used as a restore point for future database restore operations.
- If **USESNAPOFASNAPMount** is set to Yes, the **Mount read/write (without modifying backup)** check box is selected on the **VSS Options** properties page. This option is only available for SAN Volume Controller (SVC) devices.

Important: You can set **USESNAPOFASNAPMount=Yes** only for SAN Volume Controller (SVC) devices with IBM System Storage Support for Microsoft Virtual Disk and Volume Shadow Copy Services version 4.12 or later. Also, you must allocate more target volumes on your SVC storage device to accommodate the number of concurrent read/write mounts you want to do. An extra target volume matching the size of the volume to be mounted is needed for each concurrent read/write mount of that volume.

Set optional parameters

Optional parameters follow the **set** command and the positional parameters.

/CONFIGfile=*configfilename*

Use the **/configfile** parameter to specify the name of the Data Protection for Exchange Server configuration file in which these values are set.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the Data Protection for Exchange Server installation directory is used. If the **/configfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `tdpexc.cfg`.

If the *configfilename* variable includes spaces, enclose the entire **/configfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\file.cfg"
```

Examples: set command

The following examples provide a sample of the text, messages, and process status that displays when you use the **set** command.

The following command specifies the file `exchange.log`, in the `d:\tsm\tdpexchange` directory, as the Data Protection for Exchange Server log file instead of the default Data Protection for Exchange Server log file, `tdpexc.log`, in the directory where Data Protection for Exchange Server is installed. An example of the output is displayed.

Command

```
tdpexcc set logfile=d:\tsm\tdpexchange\exchange.log
```

Output

```
IBM Tivoli Storage Manager for Mail:  
Data Protection for Microsoft Exchange Server  
Version 7, Release 1, Level 3.0  
(C) Copyright IBM Corporation 1998, 2015. All rights reserved.  
  
AC05054I The preference has been set successfully.
```

The following example sets *FCMDAG2* as the DAG node name in the configuration file.

Command

```
tdpexcc set dagnode=FCMDAG2
```

Output

```
IBM Tivoli Storage Manager for Mail:  
Data Protection for Microsoft Exchange Server  
Version 7, Release 1, Level 3.0  
(C) Copyright IBM Corporation 1998, 2015. All rights reserved.  
  
ACN5054I The preference has been set successfully.
```

The following statement is added to the *tdpexc.cfg* configuration file:
DAGNODE FCMDAG2


Unmount backup command


To unmount backups that are previously mounted and are managed by IBM Spectrum Protect Snapshot for Exchange Server, use the **unmount backup** command.

Unmount Backup syntax

Use the **unmount backup** command syntax diagrams as a reference to view available options and truncation requirements.

TDPEXCC command





Unmount Backup positional parameter

The positional parameter immediately follows the **unmount backup** command and precedes the optional parameters.

mount points root directory

Absolute path to the directory where the snapshots are displayed as mount point directories. For example:

mount points root dir

Unmount Backup optional parameters

Optional parameters follow the **unmount backup** command and positional parameters.

/CONFIGfile=*configfilename*

Use the **/configfile** parameter to specify the name (*configfilename*) of the configuration file that contains the values to use for an **unmount backup** operation.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the installation directory is used. If the **/configfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `tdpexc.cfg`.

If the *configfilename* variable includes spaces, enclose the entire **/configfile** parameter entry in double quotation marks. For example:

`/CONFIGfile="c:\Program Files\tdpexc.cfg"`

/LOGFile=*logfile*

Use the **/logfile** parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot for Exchange Server. The *logfile* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfile* variable can include a fully qualified path. However, if no path is specified, the log file is written to the IBM Spectrum Protect Snapshot for Exchange Server installation directory.

If the *logfile* variable includes spaces, enclose the entire **/logfile** parameter entry in double quotation marks. For example:

`/LOGFile="c:\Program Files\tdpexc.log"`

If the **/logfile** parameter is not specified, log records are written to the default log file, `tdpexc.log`.

The **/logfile** parameter cannot be turned off, logging always occurs.

/LOGPrune=*numdays* | **No**

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option **No** can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify *no*, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/REMOTECOMPUTER=computername

Enter the computer name or IP address of the remote system where the backup was created.

/REMOTECOMPUTERUser=user

Enter the user name that is used to log on to the server specified with the **REMOTECOMPUTER** parameter. If a domain is required to log on with the domain account, enter *Domain\User*. To log on to the local account, the domain is not required. There is no default value.

/REMOTECOMPUTERPassword=password

Enter the password for the user name that is specified with the **REMOTECOMPUTERUser** parameter. There is no default value.

/TSMNODE=tsmnodename

Use the *tsmnodename* variable to refer to the IBM Spectrum Protect node name that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

You can store the node name in the IBM Spectrum Protect options file (*dsm.opt*). This parameter overrides the value in the IBM Spectrum Protect options file if **PASSWORDACCESS** is set to PROMPT. This parameter is not valid when **PASSWORDACCESS** is set to GENERATE in the options file.

/TSMOPTFile=tsmoptfilename

Use the *tsmoptfilename* variable to identify the IBM Spectrum Protect options file.

The file name can include a fully qualified path name. If no path is specified, the directory where IBM Spectrum Protect Snapshot is installed is searched.

If the *tsmoptfilename* variable includes spaces, enclose the entire **/tsmoptfile** parameter entry in double quotation marks. For example:
`/TSMOPTfile="c:\Program Files\file.opt"`

The default is `dsm.opt`.

/TSMPassword=tsmpassword

Use the *tsmpassword* variable to refer to the IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

If you specified **PASSWORDACCESS GENERATE** in the IBM Spectrum Protect Snapshot options file (`dsm.opt`), supplying the password is not necessary here because the one that is stored in the registry is used. However, to store the password in the registry, you must specify the IBM Spectrum Protect password the first time IBM Spectrum Protect Snapshot connects to the IBM Spectrum Protect server.

If you do specify a password with this parameter when **PASSWORDACCESS GENERATE** is in effect, the command-line value is ignored unless the password for this node is not yet stored in the registry. In that case, the specified password is stored in the registry and used when you run this command.

If **PASSWORDACCESS PROMPT** is in effect, and you do not specify a password value on the command line, then you are prompted for a password.

The IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server can be up to 63 characters in length.

Frequently asked questions

Review the answers to the following frequently asked questions about Data Protection for Microsoft Exchange Server.

How do I compress my Data Protection for Microsoft Exchange Server backups?

Use the **compression** option to instruct the IBM Spectrum Protect API to compress data before the data is sent to the IBM Spectrum Protect server. Compression reduces traffic and storage requirements.

For VSS backups, specify the **compression** option in the backup-archive client options file that is used as the Local DSMAGENT Node. If the environment is configured for VSS offloaded backups, you must also specify the **compression** option in the backup-archive client options file that is used as the Remote DSMAGENT Node. Review the compression information in the client documentation before you compress your data.

How do I encrypt my Data Protection for Microsoft Exchange Server backups?

Use the **enableclientencryptkey** and **encryptiontype** options to encrypt Microsoft Exchange databases during backup and restore processing.

For VSS backups, specify the encryption options in the backup-archive client options file that is used as the Local DSMAGENT Node. If the environment is configured for VSS offloaded backups, you must also specify the encryption options in the backup-archive client options file that

is used as the Remote DSMAGENT Node. Review the encryption information in the client documentation before you encrypt your databases.

How do I deduplicate my Data Protection for Microsoft Exchange Server backups?

Use the **deduplication** option to enable client-side data deduplication. Client-side data deduplication is used by the IBM Spectrum Protect API to remove redundant data during backup processing before the data is transferred to the IBM Spectrum Protect server.

For VSS backups, specify the **deduplication** option in the backup-archive client options file that is used as the Local DSMAGENT Node. If the environment is configured for VSS offloaded backups, you must also specify the **deduplication** option in the backup-archive client options file that is used as the Remote DSMAGENT Node. Review the deduplication information in the client documentation before you encrypt your databases.

How do I verify that I have Microsoft Exchange Server MAPI Client and Collaboration Data Objects correctly installed to complete Data Protection for Microsoft Exchange Server mailbox restore operations on my Exchange Server?

When you use the configuration wizard in the Microsoft Management Console (MMC) to configure Data Protection for Microsoft Exchange Server, the wizard completes a requirements check. This check verifies whether the Microsoft Exchange Server MAPI Client and Collaboration Data Objects is correctly installed.

You can also issue the `tdpmapi.exe testmapi` command to verify whether the MAPI is installed correctly.

How does a Data Protection for Microsoft Exchange Server mailbox restore operation really do mailbox-level and mailbox item-level restores?

When a mailbox restore operation is initiated, Data Protection for Microsoft Exchange Server completes the following actions:

1. Starts a session with the IBM Spectrum Protect server.
2. Queries the IBM Spectrum Protect server for a list of available backups.
3. Selects an appropriate backup that is based on user input.
4. When necessary, create an Exchange recovery database.
5. Restores the selected backup into the Exchange recovery database.
6. Copies individual mailboxes or individual mailbox items from the Exchange recovery database into the specified destination.
7. Removes the Exchange recovery database and the associated files.

Can I back up and restore a Database Availability Group (DAG) copy?

Exchange Server DAG replica copies can be backed up and restored by using the VSS method.

What is a VSS restore into operation?

A VSS restore into operation can be completed on VSS backups. A VSS restore into operation allows a VSS backup of data to be restored into the recovery database, an alternate database, or a relocated database.

Are VSS restores restored into the recovery database?

Yes, VSS restores can be restored into the recovery database or into an alternate database.

Why does the VSS instant restore fail over to a VSS fast restore?

A failover can occur if the Exchange data is on storage systems that are not supported for VSS instant restore.

How does VSS instant restore work?

VSS instant restore is a volume-level hardware-assisted copy where target volumes (that contain the snapshot) are copied back to the original source volumes. A SAN Volume Controller, Storwize family, XIV, or IBM System Storage DS8000 series storage system is required to complete VSS instant restores.

Now that I am completing VSS operations, why are there so many active backups?

IBM Spectrum Protect policy manages VSS backups on local shadow volumes and on IBM Spectrum Protect server storage. This management allows for different policies, which can lead to an increase in the number of active backups.

Can I use UNC drive letters with VSS offloaded backups?

No, Data Protection for Microsoft Exchange Server VSS offloaded backups do not process correctly if the Exchange database or log location are specified with UNC-based drive letters. For example, the following path uses UNC drive letters and is not supported in a VSS offloaded backup:

```
\\host_srv1\c$\Program Files\Exchsrvr\First Database
```

The following path is specified correctly:

```
C:\Program Files\Exchsrvr\First Database
```

Drive-based names are supported when you use a volume mount point. For example:

```
X:\Exch_Mount_Point\Program Files\Exchsrvr\First Database
```

However, UNC-based naming (as shown in the following example) is not supported when you use a volume mount point:

```
\\host_srv1\x$\Exch_Mount_Point\Program Files\Exchsrvr\First Database
```

Why do I receive a TCP/IP timeout failure when I have Windows internal VSS tracing turned on?

Data Protection for Microsoft Exchange Server VSS operations might timeout with a TCP/IP failure when Windows internal VSS tracing is turned on because of the additional time that is required to write entries to the trace file. You can avoid this issue by increasing the values for the IBM Spectrum Protect server `commtimeout` and `idletimeout` options or by decreasing the amount of Windows internal VSS tracing.

How do I complete a mailbox-level and an item-level backup and restore for Exchange?

With the Data Protection for Microsoft Exchange Server mailbox restore feature, you can complete individual mailbox recovery and item-level recovery operations in Microsoft Exchange Server environments on Data Protection for Microsoft Exchange Server backups.

Can I restore a Data Protection for Microsoft Exchange Server database backup to flat files without using an Exchange Server? Can I restore a Data Protection for Microsoft Exchange Server database backup to a flat file without interrupting the Data Protection for Microsoft Exchange Server Server?

Yes, use the `restorefiles` command. For more information, see “Restorefiles command” on page 194.

How do I schedule Data Protection for Microsoft Exchange Server backups?

You can schedule Data Protection for Microsoft Exchange Server backups by using the IBM Spectrum Protect backup-archive client scheduler or the MMC scheduler.

How do I know whether my backup ran successfully?

A message displays that states the backup completed successfully. In addition, the TDPEXchange service for backup starts and finishes displays messages in the Event Viewer. The Task Manager in the MMC provides centralized information about the status of your tasks. Processing information is also available in the following files:

- Data Protection for Microsoft Exchange Server log file (default: `tdpexc.log`)
This file indicates the date and time of a backup, data backed up, and any error messages or completion codes.
- IBM Spectrum Protect server activity log
Data Protection for Microsoft Exchange Server logs information about backup and restore commands to the IBM Spectrum Protect server activity log. If you do not have an administrator user ID and password for IBM Spectrum Protect, your IBM Spectrum Protect administrator can view this log for you.
- IBM Spectrum Protect API error log file (default: `dsierror.log`)

To prevent unsuccessful backups, refer to the following facts:

- An incremental backup of an Exchange Server database can fail if a previous full backup attempt of the same database that ended prematurely. If you receive Data Protection for Microsoft Exchange Server errors ACN3025E or ACN4226E, complete a full backup of the database.
- A backup can fail if necessary transaction logs are deleted or truncated. An error message is displayed stating that log files or patch files are missing. Perform the following steps to recover from this type of backup failure:
 1. Verify that only one product is completing backups on your system.
 2. Perform a full backup.
 3. If an error is still encountered, shut down and restart the Exchange Server, then complete a full backup.
 4. If an error persists, restart the system and complete a full backup.

How do the Exchange Server transaction logs get truncated?

The log truncation can seem delayed because Exchange must make sure all log updates are sent and committed in all copies (active and passive) before it truncates the logs. A backup product, for example, IBM Spectrum Protect, completes a full backup and reports the backup is successful to Exchange. The Exchange server processes the actual log file truncation. You see evidence of this notification to truncate logs in the Windows Event log.

What do I do when the following IBM Spectrum Protect server error message is displayed: "ANR9999D smmode.c(xxxx): Error validating inserts, and so on"

You do not need to do anything because this message can be ignored. Installing a later version of IBM Spectrum Protect server prevents this message from being displayed.

Do I use the same nodename as used by my backup-archive client?

No, you must use different node names.

What happens if I bypass integrity checking of database backups?

In a Database Availability Group (DAG) environment, you can bypass integrity checking only when the database that you are backing up has at least two healthy copies of a database (one active and one passive copy).

However, if you directly bypass integrity checking by setting the **/SKIPINTEGRITYCHECK=YES** parameter, the backup that is stored on IBM Spectrum Protect server might not be valid, or data loss can occur. To ensure that at least one active and one passive database copy are available before you skip integrity checking, set the **/SKIPINTEGRITYCHECK=SkipDbCheckIfHealthy** or **/SKIPINTEGRITYCHECK=SkipDbAndLogCheckIfHealthy** parameters.

If you bypass integrity checking and the database is not recoverable because of errors, you must contact the software vendor to resolve the issue.

Appendix. Accessibility features for the IBM Spectrum Protect product family

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Overview

The IBM Spectrum Protect family of products includes the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

The IBM Spectrum Protect family of products uses the latest W3C Standard, WAI-ARIA 1.0 (www.w3.org/TR/wai-aria/), to ensure compliance with US Section 508 (www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards) and Web Content Accessibility Guidelines (WCAG) 2.0 (www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the product.

The product documentation in IBM Knowledge Center is enabled for accessibility. The accessibility features of IBM Knowledge Center are described in the Accessibility section of the IBM Knowledge Center help (www.ibm.com/support/knowledgecenter/about/releasenotes.html?view=kc#accessibility).

Keyboard navigation

This product uses standard navigation keys.

Interface information

User interfaces do not have content that flashes 2 - 55 times per second.

Web user interfaces rely on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

Web user interfaces include WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Vendor software

The IBM Spectrum Protect product family includes certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for accessibility information about its products.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service
800-IBM-3383 (800-426-3383)
(within North America)

For more information about the commitment that IBM has to accessibility, see IBM Accessibility (www.ibm.com/able).

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those

websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows: © (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_.

Trademarks

IBM, the IBM logo, and ibm.com[®] are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Linear Tape-Open, LTO, and Ultrium are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java[™] and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

SoftLayer[®] is a registered trademark of SoftLayer, Inc., an IBM Company.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

Glossary

A glossary is available with terms and definitions for the IBM Spectrum Protect family of products.

See the IBM Spectrum Protect glossary.

To view glossaries for other IBM products, see IBM Terminology.

Index

Special characters

- /BACKUPDESTINATION parameter
 - and backup command 155
 - and restore command 189
- /CONFIGFILE parameter
 - and backup command 155
 - and restore command 189
- /ERASEEXISTINGLOGS parameter
 - and restore command 190
- /EXCLUDEDAGACTIVE parameter
 - and backup command 156
- /EXCLUDEDAGPASSIVE parameter
 - and backup command 156
- /EXCLUDEDB parameter
 - and backup command 156
- /EXCLUDENONDAGDBS parameter
 - and backup command 156
- /FROMEXCSEVER parameter
 - and restore command 190
- /INSTANTRESTORE parameter
 - and restore command 190
- /INTODB parameter
 - and restore command 190
- /LOGFILE parameter
 - and backup command 156
 - and restore command 191
- /LOGPRUNE parameter
 - and restore command 191
- /MOUNTDATABASES parameter
 - and restore command 192
- /MOUNTWAIT parameter
 - and backup command 157
 - and restore command 192
- /OBJECT parameter
 - and restore command 192
- /OFFLOAD parameter
 - and backup command 157
- /QUIET parameter
 - and backup command 158
 - and restore command 192
- /RECOVER parameter
 - and restore command 192
- /TEMPLOGRESTOREPATH parameter
 - and restore command 193
- /TSMNODE parameter
 - and backup command 158
 - and restore command 193
- /TSMOPTFILE parameter
 - and backup command 158
 - and restore command 193
- /TSMPASSWORD parameter
 - and backup command 159
 - and restore command 194
- /USESNAPOFASNAPTOMOUNT parameter
 - and set command 225

A

- accessibility features 235
- APAR 147

- auto select option, GUI 104
- automated failover
 - overview 30

B

- backing up Exchange Server data 100
- backing up Exchange Server data in a DAG environment 100
- backup
 - command line 155
 - copy
 - description 7
 - database copy
 - description 7
 - differential
 - description 7
 - full 89
 - description 7
 - full plus differentials 89
 - full plus incremental 89
 - incremental
 - description 7
 - storage group
 - command line 164
- backup command 154
 - and /BACKUPDESTINATION parameter 155
 - and /CONFIGFILE parameter 155
 - and /EXCLUDEDAGPASSIVE parameter 156
 - and /EXCLUDEDB parameter 156
 - and /EXCLUDENONDAGDBS parameter 156
 - and /logfile parameter 156, 165
 - and /logprune parameter 166
 - and /MINIMUMBACKUPINTERVAL parameter 157
 - and /MOUNTWAIT parameter 157
 - and /OFFLOAD parameter 157
 - and /preferdagpassive parameter 158
 - and /quiet parameter 166
 - and /Quiet parameter 158
 - and /SKIPINTEGRITYCHECK parameter 158
 - and /TSMNODE parameter 158
 - and /TSMOPTFILE parameter 158
 - and /TSMPASSWORD parameter 159
 - and /UpdateMailboxInfoOnly parameter 159
 - and /EXCLUDEDAGACTIVE parameter 156
- DAG backup
 - and /MINIMUMBACKUPINTERVAL parameter 157
 - and /preferdagpassive parameter 158
- example 159
- overview 153
- scheduling backup command
 - and /MINIMUMBACKUPINTERVAL parameter 157
 - and /preferdagpassive parameter 158

- backup methods 11
- backup strategy 89
- full backup 89
- full plus differentials 89
- full plus incremental 89
- IBM Spectrum Protect versus local shadow volumes 89
- VSS cluster 133
- backupdestination parameter
- and delete backup command 165

backupdestination parameter (*continued*)
and set command 220

C

capacity
determining managed storage 98

changetsmppassword command
and /configfile parameter 161
and /logfile parameter 161
and /tsmnode parameter 162
and /tsmoptfile parameter 162
example 163
syntax diagram 160

circular logging 7

cluster
VSS 133

command 179
policy 176

command line parameters
/backupdestination
and set 220
/configfile
and query exchange 173
and query tdp 177
and set 225
/FROMArchive
and restoremailbox 203
/localdsmagentnode
and set 221
/logfile
and set 221
/mailboxoriglocation
and restoremailbox 208
/mailboxrestoredate
and restoremailbox 209
/mailboxrestoredestination
and restoremailbox 210
/mailboxrestoretime
and restoremailbox 209
/mountwait
and set 222
/quiet
and backup 166
/Quiet
and backup 158
and restore 192
/remotedsmagentnode
and set 223
/tempdbrestorepath
and restoremailbox 216
and set 223
/templogrestorepath
and restoremailbox 216
and set 224

command-line interface
overview 152

command-line parameters
/backupdestination
and restore 165
and restorefiles 196
/BACKUPDESTINATION
and backup 155
and restore 189
/configfile
and changetsmppassword 161
and delete backup 165

command-line parameters (*continued*)

/configfile (*continued*)
and mount backup 170
and restorefiles 196
and restoremailbox 203
and unmount backup 227

/CONFIGfile
and backup 155
and restore 189

/dateformat
and set 220

/ERASEexistinglogs
and restore 190

/EXCLUDEDAGACTive
and backup 156

/EXCLUDEDAGPASsive
and backup 156

/EXCLUDEDDB
and backup 156

/EXCLUDEDUMPster
and restoremailbox 204

/EXCLUDENONDAGDbs
and backup 156

/fromexcserver
and delete backup 165
and restorefiles 197

/FROMEXCSErVer
and restore 190

/IMPORTVSSNAPSHOTSONLYWhenneeded
and set 221

/INSTANTREStore
and restore 190

/into
and restorefiles 197

/INTODB
and restore 190

/KEEPRDB
and restoremailbox 204

/logfile
and backup 165
and changetsmppassword 161
and mount backup 170
and query exchange 173
and query tdp 177
and restorefiles 197
and restoremailbox 204
and unmount backup 227

/LOGFile
and backup 156
and restore 191

/logprune
and backup 166

/LOGPrune
and restore 191

/mailboxfilter
and restoremailbox 205

/MINimumbackupinterval
and backup 157

/MOUNTDatabases
and restore 192

/mountrw
and mount backup 171, 215, 222

/mountwait
and restorefiles 198
and restoremailbox 215

/MOUNTWait
and backup 157

- command-line parameters (*continued*)
 - /MOUNTWait (*continued*)
 - and restore 192
 - /numberformat
 - and set 223
 - /object
 - and delete backup 166
 - and restorefiles 198
 - /OBJECT
 - and restore 192
 - /OFFLOAD
 - and backup 157
 - /olderthan
 - and delete backup 166
 - /preferdagpassive
 - and backup 158
 - /quiet
 - and restorefiles 198
 - /RECOVER
 - and restore 192
 - /SHOWMAILBOXInfo
 - and query tsm 183
 - /SKIPINTEGRITYCHECK
 - and backup 158
 - /TEMPLOGRESTorepath
 - and restore 193
 - /timeformat
 - and set 224
 - /tsmnode
 - and changetsmpassword 162
 - and mount backup 172
 - and restore 167
 - and restorefiles 198
 - and restoremailbox 217
 - and unmount backup 228
 - /TSMNODE
 - and backup 158
 - and restore 193
 - /tsmoptfile
 - and changetsmpassword 162
 - and mount backup 172
 - and restore 167
 - and restorefiles 198
 - and restoremailbox 217
 - and unmount backup 228
 - /TSMOPTFile
 - and backup 158
 - and restore 193
 - /tsmpassword
 - and mount backup 172
 - and restore 167
 - and restorefiles 199
 - and restoremailbox 217
 - and unmount backup 229
 - /TSMPassword
 - and backup 159
 - and restore 194
 - /UpdateMailboxInfoOnly
 - and backup 159
 - /USEEXISTINGRDB
 - and restoremailbox 217
 - /USESNAPOFASNAPTomount
 - and set backup 225
 - and local 196
 - and tsm 196
 - and vss 196
 - dagnode 156, 165, 189, 196, 203, 220

- commands
 - query exchange 173
 - query managedcapacity 175
 - query tdp 176
 - query tsm 178
 - set 218
- communication protocol option 50
- compressalways option 50
- compression option 50
- configfile parameter
 - and changetsmpassword command 161
 - and delete backup command 165
 - and mount backup command 170
 - and query exchange command 173
 - and query tdp command 177
 - and restorefiles command 196
 - and restoremailbox command 203
 - and set command 225
 - and unmount backup command 227
- configuration
 - manual procedure
 - Exchange Server 71
 - IBM Spectrum Protect server 72
 - options 50
 - procedure
 - offloaded backups 74
 - wizard 68
- configuration files
 - non-default locations 53
- configuration preferences
 - Data Protection for Microsoft Exchange Server 54
- configure Data Protection for Exchange Server with DAG
 - node 68
- configuring
 - binding
 - policy 20
 - Data Protection for Microsoft Exchange Server 47, 64
 - Data Protection for Microsoft Exchange Server with IBM Spectrum Protect 68
 - policy 20
 - quick instructions 34
- continuous replication 12, 91
- copy backup
 - description 7
- custom settings 62

D

- DAG 1, 23, 87
- DAG node
 - configuration 68
- dagnode parameter 156, 165, 189, 196, 203, 220
- data protection
 - Exchange with VSS backup-restore support
 - gathering information before calling IBM 142
 - Exchange with VSS backup/restore support
 - determining the issue 128
 - gathering files before calling IBM 143
 - general help 127
 - tracing when using VSS 141
 - troubleshooting 138
- Data Protection for Exchange Server
 - backups 11
 - commands 152
 - configuring options 50
 - creating an installation package on a DVD 42
 - exclude processing 52

- Data Protection for Exchange Server *(continued)*
 - include processing 52
 - LAN-free
 - description 150
 - operating environment 1
 - overview 1
 - policy settings 15
 - registering 52
 - VSS planning 3
- Data Protection for Exchange Server GUI
 - starting 87, 94
- Data Protection for Exchange Server silent installation
 - playing back the installation 42
- Data Protection for Exchange Server tasks
 - automating 124
 - Data Protection for Exchange Server tasks 124
- Data Protection for Microsoft Exchange Server
 - configuring 47, 64
 - features 1
 - installing on a local system 38
 - overview 87
 - performance 149
 - quick configuration 34
 - quick installation 34
 - reference 151
 - restore types 25
 - silent installation 39
 - silent installation with batch file 41
- Data Protection for Microsoft Exchange Server scripts
 - adding 145
 - editing 145
 - viewing 145
- Data Protection for Microsoft Exchange Server silent installation
 - setup error messages 43
- Data Protection for Microsoft Exchange Server trace and log files
 - viewing 140
- Data Protection for Microsoft Exchange Server VSS backup
 - policy binding 21
- Database Availability Group 1
 - backup and restore 12, 91
- database availability groups 87
- database copy backup
 - description 7
- dateformat parameter
 - and set command 220
- delete backup command
 - and /backupdestination parameter 165
 - and /configfile parameter 165
 - and /fromexcsrv parameter 165
 - and /object parameter 166
 - and /olderthan parameter 166
 - syntax diagram 163
- deleting Exchange Server VSS backups 103
- developerWorks wiki 147
- diagnosing VSS issues for Data Protection for Exchange Server 127
- diagnostics properties 57
- differential backup
 - description 7
- disability 235
- dsm.opt file 50
 - clusternode 50
 - communication protocol 50
 - compressalways 50
 - compression 50

- dsm.opt file *(continued)*
 - enableclientcryptkey 51
 - enablelanfree 51
 - encryptiontype 51
 - include.encrypt 51
 - nodename 50

E

- email support files 146
- enableclientcryptkey option 51
- enablelanfree option 51
- encryption 51
- encryptiontype option 51
- example
 - backup command 159
 - changetempspassword command 163
 - query tdp command 178
 - query tsm command 184
 - restoremailbox command 218
 - set command 225
- excfull.log 125
- Exchange backup
 - DAG environment 100
 - VSS
 - GUI 100
- Exchange Database Availability Group
 - managing with single policy 23
- Exchange restore
 - DAG environment 104
 - VSS
 - GUI 104
- Exchange Server backup
 - DAG environment 98
 - VSS
 - GUI 98
- Exchange Server VSS backup
 - deleting 103
 - mounting 102
- Exchange VSS
 - restore considerations 5
- Exchange VSS restore considerations 5
- exclude processing 52
- EXCLUDEDUMPster parameter
 - and restoremailbox command 204
- excsched.log 125
- expiring VSS backup s
 - policy 14

F

- failover
 - overview 30
- FAQ 229
- features 1
- files
 - Data Protection for Exchange Server options 159, 167, 172, 194, 199
 - Data Protection for Microsoft Exchange Server options 217
 - dsm.opt 50
 - excfull.log 125
 - excsched.log 125
 - IBM Spectrum Protect options file 162
 - IBM Spectrum Protect Snapshot options 229

- files (*continued*)
 - tdpexc.cfg
 - and backup command 155
 - and changetempspassword command 161
 - and delete backup command 165
 - and mount backup command 170
 - and query exchange command 173
 - and query tdp command 177
 - and restore command 189, 193
 - and restorefiles command 196
 - and restoremailbox command 203, 216
 - and set command 225
 - and unmount backup command 227
 - tdpexc.log 232
 - and backup command 156
 - and changetempspassword command 161
 - and delete backup command 165
 - and mount backup command 170, 171, 215, 222
 - and query exchange command 173
 - and query tdp command 177
 - and restore command 191
 - and restorefiles command 197
 - and restoremailbox command 204
 - and set command 221
 - and unmount backup command 227
 - tdpexcc.exe 152
- flat files 231
- Frequently asked questions 229
- from server option, GUI 104
- FROMArchive parameter
 - and restoremailbox command 203
- fromexcsrv parameter
 - and delete backup command 165
 - and restorefiles command 197
- full backup
 - description 7
 - strategy 89
- full plus differential backup
 - strategy 89
- full plus incremental backup
 - strategy 89

G

- general properties for Exchange Server 59
- graphical user interface (GUI)
 - restore options 104
- GUI
 - DAG Exchange backup 98, 100
 - DAG Exchange restore 104
 - Exchange VSS backup 98, 100
 - Exchange VSS restore 104
 - individual mailbox restore 108
 - restore options 105
- guidelines
 - VSS restore 5

H

- help command
 - syntax diagram 168

I

- IBM Knowledge Center x

- IBM SAN Volume Controller and IBM Storwize V7000
 - using Data Protection for Exchange Server 28, 83
- IBM Spectrum Protect
 - policy settings 15
- IBM Spectrum Protect options file 162
- IBM Spectrum Protect Snapshot
 - transitioning backups 80
- IBM System Storage DS8000 series
 - requirements 29
- include processing 52
- include.encrypt option 51
- incremental backup
 - description 7
- individual mailbox
 - restoremailbox
 - command line 203
- individual mailbox restore
 - GUI 108
- installation
 - configuring options 50
 - registering Data Protection for Exchange Server 52
- installing
 - creating an installation package on a DVD 42
 - on a local system 38
 - quick instructions 34
 - silently with batch file 41
- installing Data Protection for Microsoft Exchange Server
 - on multiple servers (silent) 39
 - unattended (silent) 39
- into parameter
 - and restorefiles command 197

K

- KEEPADB parameter
 - and restoremailbox command 204
- keyboard 235
- Knowledge Center x

L

- LAN-free
 - description 150
- local backup policy
 - setting 20
- localsmagentnode parameter
 - and set command 221
- logfile parameter
 - and changetempspassword command 161
 - and delete backup command 165
 - and mount backup command 170
 - and query exchange command 173
 - and query tdp command 177
 - and restorefiles command 197
 - and restoremailbox command 204
 - and set command 221
 - and unmount backup command 227
- logging
 - circular 7
- logging properties 60
- logprune parameter 156, 161, 170, 174, 177, 182, 197, 205, 221, 228
 - and delete backup command 166

M

- mailbox
 - restoremailbox
 - command line 203
- mailbox history handling 75
- mailbox restore
 - guidelines 89
 - overview 26
- mailbox restore operations
 - permissions 87
 - security 87
- mailboxfilter parameter
 - and restoremailbox command 205
- mailboxoriglocation parameter
 - and restoremailbox command 208
- mailboxrestoredate parameter
 - and restoremailbox command 209
- mailboxrestoredestination parameter
 - and restoremailbox command 210
- mailboxrestorettime parameter
 - and restoremailbox command 209
- managed storage
 - determining capacity 98
- managing with single policy
 - Exchange Database Availability Group 23
- MAPI
 - ensuring successful connections 98
- MAPI settings for Exchange Server 63, 98
- migration 44
 - mailbox history handling 75
- MINimumbackupinterval parameter
 - and backup command 157
- MMC GUI
 - starting 87, 94
- mount backup command
 - and /configfile parameter 170
 - and /logfile parameter 170
 - and /tsmnode parameter 172
 - and /tsmoptfile parameter 172
 - and /tsmpassword parameter 172
 - and mountrw parameter 171, 215, 222
 - syntax diagram 169
- mounting Exchange Server VSS backups 102
- mountrw parameter
 - and mount backup command 171, 215, 222
- mountwait parameter
 - and restorefiles command 198
 - and restoremailbox command 215
 - and set command 222
- msiexec.exe
 - used for silent installation 41

N

- New in Data Protection for Exchange Server Version 8.1.2 xiii
- node name
 - Data Protection for Exchange Server 52
 - offloaded backup 48
 - proxy nodes 47
 - VSS 48
- nodename option 50
- numberformat parameter
 - and set command 223

O

- object parameter
 - and delete backup command 166
 - and restorefiles command 198
- offloaded backup
 - configuration procedure 74
 - node names 48
- offloaded VSS backup
 - overview 12
- olderthan parameter
 - and delete backup command 166
- operating environment 1
- optional parameters 196
- options
 - GUI restore
 - mountdatabases 105
 - run recovery 105
- options files
 - non-default locations 53
- overview 1, 87
 - offloaded VSS backup 12
 - VSS backup 11

P

- parameters
 - /backupdestination
 - and delete backup command 165
 - and restorefiles command 196
 - and set command 220
 - /BACKUPDESTination
 - and backup command 155
 - and restore command 189
 - /configfile
 - and changetsmpassword command 161
 - and delete backup command 165
 - and mount backup command 170
 - and query exchange command 173
 - and query tdp command 177
 - and restorefiles command 196
 - and restoremailbox command 203
 - and set command 225
 - and unmount backup command 227
 - /CONFIGfile
 - and backup command 155
 - and restore command 189
 - /dateformat
 - and set command 220
 - /ERASEexistinglogs
 - and restore command 190
 - /EXCLUDEDAGACTive
 - and backup command 156
 - /EXCLUDEDAGPASSive
 - and backup command 156
 - /EXCLUDEDDB
 - and backup command 156
 - /EXCLUDEDUMPster
 - and restoremailbox command 204
 - /EXCLUDENONDAGDBs
 - and backup command 156
 - /FROMArchive
 - and restoremailbox command 203
 - /fromexcserver
 - and delete backup command 165
 - and restorefiles command 197

parameters (*continued*)

- /FROMEXCSErVer
 - and restore command 190
- /IMPORTVSSSNAPSHOTSONLYWhenneeded
 - and set command 221
- /INSTANTREStore
 - and restore command 190
- /into
 - and restorefiles command 197
- /INTODB
 - and restore command 190
- /KEEPRDB
 - and restoremailbox command 204
- /localdsmagentnode
 - and set command 221
- /logfile
 - and changetsmppassword command 161
 - and delete backup command 165
 - and mount backup command 170
 - and query exchange command 173
 - and query tdp command 177
 - and restorefiles command 197
 - and restoremailbox command 204
 - and set command 221
 - and unmount backup command 227
- /LOGFile
 - and backup command 156
 - and restore command 191
- /logprune
 - and delete backup command 166
- /LOGPrune
 - and restore command 191
- /mailboxfilter
 - and restoremailbox command 205
- /mailboxoriglocation
 - and restoremailbox command 208
- /mailboxrestoredat
 - and restoremailbox command 209
- /mailboxrestoredestination
 - and restoremailbox command 210
- /mailboxrestoretim
 - and restoremailbox command 209
- /MINimumbackupinterval
 - and backup command 157
- /MOUNTDatabases
 - and restore command 192
- /mountrw
 - and mount backup command 171, 215, 222
- /mountwait
 - and restorefiles command 198
 - and restoremailbox command 215
 - and set command 222
- /MOUNTWait
 - and backup command 157
 - and restore command 192
- /numberformat
 - and set command 223
- /object
 - and delete backup command 166
 - and restorefiles command 198
- /OBject
 - and restore command 192
- /OFFLOAD
 - and backup command 157
- /olderthan
 - and delete backup command 166

parameters (*continued*)

- /preferdagpassive
 - and backup command 158
- /quiet
 - and delete backup command 166
 - and restorefiles command 198
- /Quiet
 - and backup command 158
 - and restore command 192
- /RECOVER
 - and restore command 192
- /remotedsmagentnode
 - and set command 223
- /SHOWMAILBOXInfo
 - and query tsm command 183
- /SKIPINTEGRITYCHECK
 - and backup command 158
- /tempdbrestorepath
 - and restoremailbox command 216
 - and set command 223
- /templogrestorepath
 - and restoremailbox command 216
 - and set command 224
- /TEMPLOGRESTorepath
 - and restore parameter 193
- /timeformat
 - and set command 224
- /tsmnode
 - and changetsmppassword command 162
 - and mount backup command 172
 - and restore command 167
 - and restorefiles command 198
 - and restoremailbox command 217
 - and unmount backup command 228
- /TSMNODE
 - and backup command 158
 - and restore command 193
- /tsmoptfile
 - and changetsmppassword command 162
 - and mount backup command 172
 - and restore command 167
 - and restorefiles command 198
 - and restoremailbox command 217
 - and unmount backup command 228
- /TSMOPTFile
 - and backup command 158
 - and restore command 193
- /tsmpassword
 - and mount backup command 172
 - and restore command 167
 - and restorefiles command 199
 - and restoremailbox command 217
 - and unmount backup command 229
- /TSMPassword
 - and backup command 159
 - and restore command 194
- /UpdateMailboxInfoOnly
 - and backup command 159
- /USEEXISTINGRDB
 - and restoremailbox command 217
- /USESNAPOFASNAPTomount
 - and set command 225

dagnode 156, 165, 189, 196, 203, 220

parameters, described

- optional
 - /logprune 156, 161, 170, 174, 177, 182, 197, 205, 221, 228

- performance 149
- policy 22
 - binding 20
 - binding Data Protection for Microsoft Exchange Server VSS
 - backups 21
 - configuring 20
 - expiring VSS backup 14
 - setting local policy 20
- policy command
 - overview 176
- policy management properties 56
- policy settings
 - Data Protection for Exchange Server and IBM Spectrum Protect 15
- preferdagspassive parameter
 - and backup command 158
- preferences 54
- printing reports 120
- product support 147
- properties
 - custom settings 62
 - diagnostics 57
 - general Exchange Server 59
 - logging 60
 - MAPI settings 63, 98
 - policy management 56
 - regional settings 60
 - VSS 61
- property pages 54
- proxy nodes 47
- publications x

Q

- query exchange command
 - and /configfile parameter 173
 - and /logfile parameter 173
 - overview 173
 - syntax diagram 173
- query managedcapacity command
 - overview 175
- query tdp command
 - and /configfile parameter 177
 - and /logfile parameter 177
 - example 178
 - overview 176
 - syntax diagram 176
- query tsm 179
- query tsm command
 - and /SHOWMAILBOXInfo parameter 183
 - example 184
 - overview 178
- quiet parameter
 - and delete backup command 166
 - and restorefiles command 198

R

- RBAC
 - permissions 87
- recovery database
 - procedure 106
- reference
 - Data Protection for Microsoft Exchange Server 151
- regional properties 60
- registration 52

- remote system configuration 68
- remotedsmagentnode parameter
 - and set command 223
- replay option, GUI 105
- replication copies 12, 91
- reports
 - viewing, printing, and saving 120
- requirements 29
 - IBM System Storage DS8000 series 29
 - SAN Volume Controller 29
 - Storwize family 29
 - XIV 29
- restore 25
 - command line 188
 - database 25
 - mailbox 26
 - restorefiles command 25
 - transaction log 25
 - types 25
- restore command
 - and /BACKUPDESTination parameter 189
 - and /CONFIGfile parameter 189
 - and /ERASEexistinglogs parameter 190
 - and /FROMEXCSErVer parameter 190
 - and /INSTANTREStore parameter 190
 - and /INTODB parameter 190
 - and /LOGFile parameter 191
 - and /LOGPrune parameter 191
 - and /MOUNTDAtabases parameter 192
 - and /MOUNTWait parameter 192
 - and /OBJect parameter 192
 - and /Quiet parameter 192
 - and /RECOVER parameter 192
 - and /TEMPLOGREStorepath parameter 193
 - and /tsmnode parameter 167
 - and /TSMNODE parameter 193
 - and /tsmoptfile parameter 167
 - and /TSMOPTFile parameter 193
 - and /tsmpassword parameter 167
 - and /TSMPassword parameter 194
 - overview 186
 - syntax diagram 187
- restore considerations
 - Exchange VSS 5
- restore guidelines
 - Exchange VSS 5
- restore operations
 - using the GUI
 - auto select option 104
 - from server option 104
 - instant restore 105
 - replay option 105
 - restore options 104
- restore options
 - GUI
 - mountdatabases 105
 - run recovery 105
- restorefiles 196
 - snapshot backup 231
- restorefiles command
 - and /configfile parameter 196
 - and /fromexcserver parameter 197
 - and /into parameter 197
 - and /logfile parameter 197
 - and /mountwait parameter 198
 - and /object parameter 198
 - and /quiet parameter 198

- restorefiles command (*continued*)
 - and /tsmnode parameter 198
 - and /tsmoptfile parameter 198
 - and /tspassword parameter 199
 - backups 194
 - syntax diagram 195
- restoremailbox
 - individual mailbox
 - command line 203
 - mailbox
 - command line 203
- restoremailbox command
 - and /configfile parameter 203
 - and /EXCLUDEDUMPster parameter 204
 - and /FROMArchive parameter 203
 - and /KEEPRDB parameter 204
 - and /logfile parameter 204
 - and /mailboxfilter parameter 205
 - and /mailboxoriglocation parameter 208
 - and /mailboxrestoredate parameter 209
 - and /mailboxrestoredestination parameter 210
 - and /mailboxrestoretime parameter 209
 - and /mountwait parameter 215
 - and /tempdbrestorepath parameter 216
 - and /templogrestorepath parameter 216
 - and /tsmnode parameter 217
 - and /tsmoptfile parameter 217
 - and /tspassword parameter 217
 - and /USEEXISTINGRDB parameter 217
 - example 218
 - overview 199
 - syntax diagram 201
- restoring data
 - Exchange Server 2010 114
 - Exchange Server 2013 114
 - Mailbox Restore Browser 114
- Role Based Access Control
 - permissions 87

S

- saving reports 120
- security requirements 87
- sending support files by using email 146
- Service Management Console 147
- set command
 - and /backupdestination parameter 220
 - and /configfile parameter 225
 - and /dateformat parameter 220
 - and /IMPORTVSSSNAPSHOTONLYWhenneeded parameter 221
 - and /localdsmagentnode parameter 221
 - and /logfile parameter 221
 - and /mountwait parameter 222
 - and /numberformat parameter 223
 - and /remotedsmagentnode parameter 223
 - and /tempdbrestorepath parameter 223
 - and /templogrestorepath parameter 224
 - and /timeformat parameter 224
 - and /USESNAPOFASNAPTmount parameter 225
 - example 225
 - overview 218
 - syntax diagram 219
- SHOWMAILBOXInfo parameter
 - and query tsm command 183
- silent installation of Data Protection for Microsoft Exchange Server 39

- silent installation
 - with spinstall.exe 40
- spinstall.exe
 - used for silent installation 40
- starting
 - Data Protection for Exchange Server GUI 87, 94
 - MMC GUI 87, 94
- storage
 - determining managed capacity 98
- storage group
 - backup
 - command line 164
 - VSS backup
 - GUI 98, 100
 - VSS restore
 - GUI 104
- storage management, policy 14
- Storwize family
 - requirements 29
- syntax diagrams
 - changetsmppassword command 160
 - delete backup command 163
 - help command 168
 - mount backup command 169
 - query exchange command 173
 - query tdp command 176
 - restore command 187
 - restorefiles command 195
 - restoremailbox command 201
 - set command 219
 - unmount backup command 226

T

- tdpexc.cfg file
 - and backup command 155
 - and changetsmppassword command 161
 - and delete backup command 165
 - and mount backup command 170
 - and query tdp command 177
 - and restore command 189, 193
 - and restorefiles command 196
 - and restoremailbox command 203, 216
 - and set command 225
 - and unmount backup command 227
 - query exchange 173
- tdpexc.log file
 - and backup command 156
 - and changetsmppassword command 161
 - and delete backup command 165
 - and mount backup command 170, 171, 215, 222
 - and query exchange command 173
 - and query tdp command 177
 - and restore command 191
 - and restorefiles command 197
 - and restoremailbox command 204
 - and set command 221
 - and unmount backup command 227
- tdpexcc.exe
 - overview 152
- tempdbrestorepath parameter
 - and restoremailbox command 216
 - and set command 223
- templogrestorepath parameter
 - and restoremailbox command 216
 - and set command 224

- timeformat parameter
 - and set command 224
- transaction log
 - restore 25, 186
- TSM Configuration
 - wizard 68
- tsmnode parameter
 - and changetsmpassword command 162
 - and mount backup command 172
 - and restore command 167
 - and restorefiles command 198
 - and restoremailbox command 217
 - and unmount backup command 228
- tsmoptfile parameter
 - and changetsmpassword command 162
 - and mount backup command 172
 - and restore command 167
 - and restorefiles command 198
 - and restoremailbox command 217
 - and unmount backup command 228
- tsmpassword parameter
 - and mount backup command 172
 - and restore command 167
 - and restorefiles command 199
 - and restoremailbox command 217
 - and unmount backup command 229

- VSS Requestor 3
- VSS writer 3
- VSSPOLICY, statements 22

X

- XIV
 - requirements 29

U

- unmount backup command
 - and /configfile parameter 227
 - and /logfile parameter 227
 - and /tsmnode parameter 228
 - and /tsmoptfile parameter 228
 - and /tsmpassword parameter 229
 - syntax diagram 226
- USEEXISTINGRDB parameter
 - and restoremailbox command 217
- using Data Protection for Exchange Server
 - with IBM SAN Volume Controller and IBM Storwize V7000 28, 83

V

- viewing reports 120
- viewing system information for Data Protection for Microsoft Exchange Server 145
- VSS
 - cluster 133
 - N-series and NetApp
 - storage 6
 - node names 48
 - overview 3
 - proxy nodes 47
 - restore into alternate locations 26, 106
- VSS backup
 - characteristics 3
 - overview 11
 - planning requirements 4
- VSS fast restore
 - method 25
- VSS instant restore
 - method 25
- VSS planning 3
- VSS properties 61
- VSS provider 3



Product Number: 5725-X02

Printed in USA