IBM Tivoli Storage Manager for Mail
Version 7.1

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

IBM Tivoli Storage Manager for Mail
Version 7.1

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

Note
Before you use this information and the product it supports, read the information in "Notices" on page 201.

First edition (December 2013)

This edition applies to version 7, release 1 of IBM Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server (product number 5608-E06) and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

About this publication

The subject of this publication is Data Protection for Microsoft Exchange Server, a component of the IBM® Tivoli® Storage Manager for Mail product.

Data Protection for Microsoft Exchange Server is also known as Data Protection for Microsoft Exchange. You can use the Data Protection for Microsoft Exchange software to perform online backups of Microsoft Exchange Server databases to Tivoli Storage Manager storage. This integration with the Microsoft Exchange Server application program interface (API) maximizes the protection of data, thus providing a comprehensive storage management solution.

Tivoli Storage Manager is a client-server licensed product that provides storage management services in a multi-platform computer environment.

Who should read this publication

This publication is intended for system installers, system users, Tivoli Storage Manager administrators, and system administrators.

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

- Microsoft Exchange Server
- Tivoli Storage Manager server
- Tivoli Storage Manager Backup-Archive Client
- Tivoli Storage Manager 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
- 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

Product documentation formats

As applicable to your environment, Data Protection for Microsoft Exchange provides product documentation in the following formats.

Installation and User's Guide

The *IBM Tivoli Storage Manager for Mail Data Protection for Microsoft Exchange Server Installation and User's Guide 7.1* provides detailed information about how to install, configure, and use Data Protection for Microsoft Exchange 7.1 in a Windows server platform. This publication is provided online, in PDF and XHTML format, at the Tivoli Information Center: `http://pic.dhe.ibm.com/infocenter/tsminfo/v7r1/index.jsp`

GUI online help

GUI online help is provided for specific information that is related to tasks that are performed in the Data Protection for Microsoft Exchange GUI. After you start the GUI, click the **Help** menu item on the menu bar, and select **Help Topics**. The online help opens in a separate window.

Command-line help

Command-line help is also provided for specific information that is related to tasks that are performed on the Data Protection for Microsoft Exchange command line. Enter `tdpexcc help` on the Data Protection for Microsoft Exchange command-line interface for a list of available help topics. See "Help command" on page 128 for more information.

Publications

Publications for the Tivoli Storage Manager family of products are available online. The Tivoli Storage Manager product family includes IBM Tivoli Storage FlashCopy® Manager, IBM Tivoli Storage Manager for Space Management, IBM Tivoli Storage Manager for Databases, and several other storage management products from IBM Tivoli.

To search across all publications or to download PDF versions of individual publications, go to the Tivoli Storage Manager information center at http://pic.dhe.ibm.com/infocenter/tsminfo/v7r1.

You also can find the Tivoli Storage Manager product family information centers and other information centers that contain official product documentation for current and previous versions of Tivoli products at Tivoli Documentation Central. Tivoli Documentation Central is available at http://www.ibm.com/developerworks/community/wikis/home/wiki/Tivoli Documentation Central.

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. You can type the item in any combination of uppercase or lowercase letters. In this example, you can enter KEYWO, KEYWORD, or KEYWOrd.	►►—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
Variables: Italicized lowercase items (*var_name*) denote variables. In this example, you can specify a *var_name* when you enter the KEYWORD command.	►►—KEYWOrd—*var_name*—►◄

Syntax diagram description	Example
Repetition: An arrow that points to the left means you can repeat the item. A character or space within an arrow means you must separate the repeated items with that character or space.	
Required Choices: When two or more items are in a stack and one of them is on the line, specify one item. In this example, you *must* choose A, B, or C.	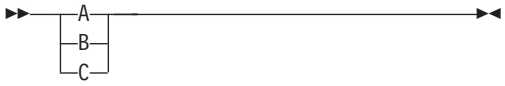
Optional Choice: When an item is below the line, that item is optional. In the first example, you can choose A or nothing at all. When two or more items are in a stack below the line, all of them are optional. In the second example, you can choose A, B, C, or nothing at all.	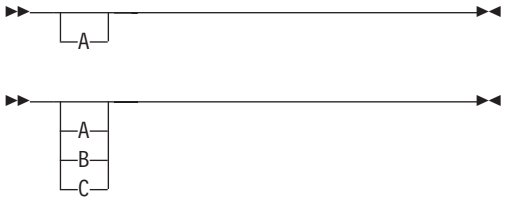
Defaults: Defaults are above the line. The default is selected unless you override it. You can override the default by including an option from the stack below the line. In this example, A is the default. You can override A by choosing B or C. You can also specify the default explicitly.	
Repeatable Choices: 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. In this example, you can choose any combination of A, B, or C.	
Syntax Fragments: 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.	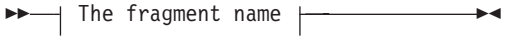 **The fragment name:** 

Syntax diagram description	Example
Footnote: A footnote in the diagram references specific details about the syntax that contains the footnote. In this example, the footnote by the arrow references the number of times you can repeat the item.	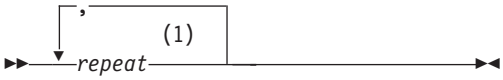 **Notes:** 1 Specify *repeat* as many as 5 times.

What's new

The following features are new for Data Protection for Microsoft Exchange:

Additional options for restoring personal storage folders
If you restore personal storage folders (`.pst` files), there are two options: **Restore Mail to Unicode PST file** and **Restore Mail to non-Unicode PST file**. Unicode `.pst` files can store messages in multiple languages, and are not limited to 2 GB of data. For non-Unicode `.pst` files, the file size must be less than 2 GB. For more information about restoring mailboxes and mailbox data, including `.pst` files, see "Restoring individual mailbox and mailbox item-level data" on page 78.

"Microsoft Exchange Server 2013 support" on page 26

When planning to restore Exchange Server 2013 mailboxes, there is a requirement for MAPI clients to use RPC over HTTPS (also known as Outlook Anywhere). RPC over HTTP is no longer supported. More planning information, and a link to the troubleshooting topic, is provided.

"Improving mailbox history handling" on page 57
To improve performance, mailbox history includes only the mailboxes from databases than are backed up. If you back up mailbox history with a version of Data Protection for Microsoft Exchange earlier than version 7.1, you can manually delete the old mailbox history. Details, including instructions, are provided.

"Restoring mailboxes directly from Exchange database files" on page 83
With all Data Protection for Microsoft Exchange configurations, administrators can complete an individual mailbox restore for an `.edb` file stored on a disk. For administrators who only want to complete individual mailbox restores from an `.edb` file on disk, the Mailbox Restore Only configuration option is available. The `.edb` file can come from a backup that mounted read-write using Tivoli Storage Manager for Virtual Environments, Tivoli Storage Manager restore files, or an offline file system copy.

"Windows PowerShell and Data Protection for Exchange" on page 91
The Data Protection for Microsoft Exchange software includes Windows PowerShell cmdlets to complement the command-line interface functions.

"Automated failover for data recovery" on page 14
If you use Data Protection for Exchange with the Tivoli Storage Manager configuration, when the Tivoli Storage Manager server is unavailable, Data Protection for Exchange can automatically fail over to the secondary server for data recovery.

Chapter 1. Getting started

IBM Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server provides online backups and restores of Microsoft Exchange Server components to Tivoli Storage Manager storage.

Data Protection for Microsoft Exchange Server provides a connection between an Exchange Server and a Tivoli Storage Manager server which allows Exchange data to be protected and managed by Tivoli Storage Manager. Data Protection for Microsoft Exchange Server protects Exchange Server data and improves the availability of Exchange databases.

Data Protection for Microsoft Exchange software can complete online backups and restores of Microsoft Exchange Server databases to Tivoli Storage Manager storage or local shadow volumes. You can run backups and restores by using a command-line or graphical user interface (GUI). Refer to your Exchange Server documentation for complete information about the back up and restore of Microsoft Exchange Servers.

Beginning with Exchange Server 2010, Microsoft no longer supports the Microsoft Legacy API (streaming) for backup and restore operations. It supports the use of VSS for the backup and restore.

Data Protection for Microsoft Exchange operations use the Tivoli Storage Manager application programming interface (API) to communicate with the Tivoli Storage Manager server, and use the Exchange API to communicate with Exchange Server. In addition to using these APIs, Data Protection for Microsoft Exchange VSS operations use the Tivoli Storage Manager backup-archive client (VSS Requestor) and Microsoft Volume Shadow Copy Service to produce an online snapshot (point-in-time consistent copy) of Exchange data that can be stored on local shadow volumes or on Tivoli Storage Manager server storage.

You must install Data Protection for Microsoft Exchange on the same system as the Exchange Server. For the required level of Tivoli Storage Manager server, see “Software and operating system requirements” on page 44. Data Protection for Microsoft Exchange also supports operations in a Database Availability Group (DAG) environment.

Data Protection for Microsoft Exchange product features

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

For Exchange 2013 users, you can back up and restore public folder mailboxes located on mailbox databases in standalone and TSM configurations. Item-level restore for public folders is not supported.

For Exchange 2010, public folder databases can be backed up and restored. For both Exchange 2010 and Exchange 2013, item-level restore for public folders is not supported.

Table 1. Data Protection for Microsoft Exchange key features

Feature	Referred to as:	More information:
Perform individual mailbox recovery and item-level recovery from Data Protection for Microsoft Exchange backups	Mailbox restore	"Restoring individual mailbox and mailbox item-level data" on page 78
Back up Exchange Server databases by using Microsoft Volume Shadow Copy Service (VSS)	VSS backup	"VSS backup" on page 4
Back up Exchange Server Database Availability Group databases to a common node to manage all DAG members by using a single policy	Back up to DAG node	"Managing Exchange Database Availability Group members by using a single policy" on page 33
Perform a backup to the Tivoli Storage Manager server by using an alternate system instead of a production system	Offloaded backup	"Offloaded VSS backups" on page 7
Restore backups that are on Tivoli Storage Manager server storage to their original location	VSS restore	"VSS restore" on page 11
Restore backups that are on local shadow volumes by using file-level copy mechanisms	VSS Fast Restore	"VSS fast restore" on page 11
Restore backups that are on local shadow volumes by using hardware-assisted volume-level copy mechanisms	VSS Instant Restore	"VSS instant restore" on page 12
Restore a backup to a recovery database, alternate database, or relocated database	Restore into	"Restoring VSS backups into alternate locations" on page 13
Query the managed capacity for backups that are on local shadow volumes	**query managedcapacity** command	"Query Managedcapacity command" on page 134
Delete a backup of an Exchange server database	**delete backup** command	"Delete backup command" on page 123
Manage policy for backups that are on local shadow volumes	**policy** commands	"Query policy command" on page 134
Integrate with Tivoli Storage FlashCopy Manager	Advanced VSS support	"Transitioning Exchange Server backups from Tivoli Storage FlashCopy Manager to Tivoli Storage Manager" on page 182
Tivoli Storage Manager policy-based management of backups	Server policy	"How Tivoli Storage Manager server policy affects Data Protection for Exchange" on page 28
Use the **restorefiles** command to restore backups to flat files without involving the Exchange Server.	**restorefiles** command	"Restorefiles command" on page 153

The term *local shadow volumes* is used throughout this document to describe data that is stored on shadow volumes that are localized to a disk storage subsystem.

Backup methods

Data Protection for Exchange uses the VSS method for backing up data.

Data Protection for Exchange tracks and stores mailbox location history, which is used to automate mailbox restore operations. This action causes a delay 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.

VSS framework

VSS provides software and hardware vendors with a common interface model for generating and managing snapshots.

The Microsoft VSS service manages and directs three VSS software components that are used during VSS operations: the VSS requestor, the VSS writer, and the VSS provider. The VSS requestor is the backup software. The VSS writer is the application software. Examples of application software include Microsoft Exchange Server and Microsoft SQL Server. The VSS provider is the specific combination of hardware and software that generates the snapshot volume.

VSS writer

The VSS writer for the Microsoft Exchanger Server is the Microsoft Exchange Writer. The Microsoft Exchange Writer is provided by the Microsoft Exchange Information Store service or the Microsoft Exchange Replication service.

VSS requestor

The Tivoli Storage Manager backup-archive client serves as the VSS requestor component and communicates with Microsoft VSS services to access data and create volume shadow copies. Because the Tivoli Storage Manager backup-archive client acts as the VSS interface, features such as LAN-free backup, client-side deduplication, data encryption, and data compression, are available. These feature are enabled by setting certain options defined in the backup-archive client options file.

This application initiates a snapshot operation. The application sends a command to the VSS service to create a shadow copy of a specified volume. The VSS requestor is the Tivoli Storage Manager backup-archive client.

VSS provider

This application produces the shadow copy and also manages the volumes where the Exchange data is located. A provider can be a system provider (such as the one included with the Microsoft Windows operating system). It can also be a software provider or a hardware provider (such as one that is included with a storage system).

VSS hardware providers require installation and configuration, including the installation of all required fix packages. For instructions, see the documentation for the VSS hardware provider.

For more information about VSS technology, see the Microsoft Technical Reference document *How Volume Shadow Copy Service Works*.

VSS system provider overview

A VSS system provider assists with creating and maintaining copies on local shadow volumes.

The VSS system provider refers to the default VSS provider that is available with Windows Server. If you are using the Windows VSS system provider, no configuration is required. However, you can make some configuration changes by using the VSSADMIN commands. See Microsoft documentation on the VSSADMIN commands for details.

VSS software or hardware provider overview

A software or hardware provider acts as an interface during VSS processing at the software or hardware level.

If you use a software or hardware provider, review the following operational requirements that are provided to help you plan for VSS backups:

- Place database files on a separate, dedicated logical volume.
- Place logs for each database on a separate logical volume.
- Do not place non-Exchange data on storage volumes that are dedicated to Exchange data.
- When you use hardware snapshot provider, do not share LUNs with other applications.
- Read and follow specific installation and configuration instructions in the documentation that is provided by your VSS provider vendor.
- When a hardware provider is used, configure the disks that store Exchange data as basic disks.

VSS backup

A VSS backup uses Microsoft Volume Shadow Copy Service technology to produce an online snapshot (point-in-time consistent copy) of Exchange data.

VSS backups eliminate the need for the server or file system to be in backup mode for an extended period of time. The length of time to perform the snapshot is usually 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.

VSS backups can be stored on local VSS shadow volumes, or, when integrated with Tivoli Storage Manager, in Tivoli Storage Manager server storage. Both of these storage destinations require that sufficient space be available for the snapshot.

When sufficient space is available for the snapshot, VSS backups stored locally on VSS shadow volumes are directly accessible by the system.

Restoring locally managed VSS backups is fast because the Exchange data is not transferred from Tivoli Storage Manager server storage over the network.

When you run VSS backups and store data on Tivoli Storage Manager server storage, sufficient space is temporarily required on local snapshot volumes. This space is used to hold the snapshot until transfer to the Tivoli Storage Manager server is complete. After the data transfer to the server is complete, the snapshot volume is released. The space can be reused.

If you also store VSS backup locally, in addition to Tivoli Storage Manager server storage, 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 Tivoli Storage Manager server storage. The maximum number of local backup version to be maintained is set in the Tivoli Storage Manager policy.

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

For data backed up to both VSS shadow volumes and Tivoli Storage Manager server storage, a local snapshot backup is run and the data on the local snapshot volume is sent to the Tivoli Storage Manager server. The local snapshot volume is retained as a local backup after the transfer to the Tivoli Storage Manager server is complete.

Data Protection for Exchange completes the following actions when a VSS backup operation is initiated:

1. Data Protection for Exchange validates the state of Exchange server objects.
2. Data Protection for Exchange begins a session with a Tivoli Storage Manager server.
3. Data Protection for Exchange verifies that the VSS service is running and that the Exchange writer is available.
4. The Tivoli Storage Manager VSS requestor lists the backup components through the VSS writer.
5. The Tivoli Storage Manager VSS requestor runs the VSS snapshot backup preparation stage.
6. The Tivoli Storage Manager VSS requestor runs the actual VSS backup.
7. The Tivoli Storage Manager VSS requestor runs an integrity check on the VSS backup.
8. Optionally, the integrity check can be offloaded to an alternative system that has the Tivoli Storage Manager VSS requestor installed and configured.
9. The Tivoli Storage Manager VSS requestor backs up the data, including metadata, to a Tivoli Storage Manager server. Optionally, the movement of data to a Tivoli Storage Manager server can be offloaded to an alternate system that has the Tivoli Storage Manager VSS requestor installed and configured.
10. The Tivoli Storage Manager VSS requestor marks the backup as complete in VSS.
11. Data Protection for Exchange ends the Tivoli Storage Manager server session.

VSS backup characteristics

VSS backups have characteristics that affect backup management tasks.

Backups can be stored on local shadow volumes, Tivoli Storage Manager server storage, or both locations. Backups to Tivoli Storage Manager server storage can be offloaded to another system as resource relief for production servers. In addition, backups can be restored to flat files without the involvement of the Exchange Server.

Backups provide an Exchange Server database integrity check function, but do not provide an Exchange Server database zeroing function. Full, copy, incremental, and differential backup types are supported. Different policy settings can be defined for each backup location and backup type (`FULL` or `COPY`).

For databases in an Exchange Server DAG that have two or more healthy copies, the database integrity check can be skipped. Exchange Server Database Availability Group (DAG) databases can be backed up under a common DAG node name, regardless of which DAG member runs the backup. The backup can be from an active or passive copy. When you back up data to a common node, the backups are managed by a common policy, and the user can restore to any Exchange Server.

VSS backup planning requirements

Plan a VSS backup strategy to optimize your backup operations performance and avoid potential problems.

Consider the following requirements when you plan for VSS backups:

- When you run VSS operations, ensure that you have at least 200 MB of free disk space on your Windows System Drive. This space is used to hold the metadata files for Data Protection for Exchange.
- Make sure to review optimal practice recommendations by Microsoft for your level of Exchange Server.
- Ensure that you have a well-defined and tested recovery plan that meets your service level objectives.
- Use single hardware LUNs for log and system files.
- Use single hardware LUNs for the database files.
- Use basic disks.
- If you plan to keep some VSS snapshot backups on local shadow volumes only, make sure to consider the VSS provider-specific implementation and configuration options when you set up your strategy. For example, if your VSS hardware provider supports a full-copy snapshot versus a copy-on-write (COW) snapshot mechanism, full-copy type implementations have greater disk storage requirements. However, full-copy type implementations are less risky because they do not rely on the original volume to restore the data. COW implementations require much less disk storage but rely completely on the original volume to process a restore. Since these implementations are entirely controlled by the VSS provider and not Data Protection for Exchange, make sure to consult your VSS provider documentation for a complete understanding of your VSS implementation.
- If you run parallel VSS backups, stagger the start of the backups by at least ten minutes. This interval ensures that the snapshot operations do not overlap. If you do not stagger the snapshots, errors can occur. In addition, configure the parallel instance backups so they do not take snapshots of the same volumes. Ensure that parallel backups do not make a snapshot of the same LUN.
- Do not place multiple volumes on the same LUN. Microsoft advises that you configure a single volume, single partition, and single LUN as 1 to 1 to 1.Do not set the `ASNODENAME` option in the `dsm.opt` file when you use Data Protection for Exchange. Setting `ASNODENAME` can cause VSS backups and VSS restores to fail.

IBM System Storage requirements

Specific database, log, file, and LUN settings are required for IBM System Storage.

The DS8000®, SAN Volume Controller, Storwize V7000, and XIV storage subsystems require these settings when you plan for VSS backups:

- 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, make sure the database LUNs are dedicated to only one database or application.
- If you delete a LOCAL snapshot that is stored on a SAN Volume Controller or Storwize V7000 Space Efficient volume (SEV) that has multiple dependent targets, you must delete them in the same order in which you created them. You must delete the oldest one first, followed by the second oldest, and so on. Failure to delete them in this order can cause removal of other snapshots of the same source.
- (SAN Volume Controller and Storwize V7000 only) If you use multiple target FlashCopy mappings, a mapping can stay in the copying state after all the source data is copied to the target. This situation can occur if mappings that were started earlier and use the same source disk are not yet fully copied. Because of this situation, initiate local backups for SAN Volume Controller and Storwize V7000 storage subsystems at intervals greater than the time required for the background copy process to complete.

Offloaded VSS backups

An offloaded backup uses another system to run the integrity check and to move the data to the Tivoli Storage Manager server.

This type of backup shifts the backup load from the production system to another system. An offloaded VSS backup requires a VSS hardware provider that supports transportable shadow copy volumes is installed on the production and secondary systems.

Offloaded VSS backups require a Tivoli Storage FlashCopy Manager license. Tivoli Storage FlashCopy Manager is a separately purchasable program.

Database Availability Group backups

Database Availability Group backups are a way to use Exchange Server 2010 and later features.

Database Availability Group (DAG) technology helps to protect your Exchange Server and possibly reduce the frequency of backup operations. A DAG is an Exchange Server high availability feature. Database Availability Groups replace LCR, CCR, and SCR replication features. They provide for enhanced data and service availability and automatic recovery from failures. Database copies are mirrored on any DAG member within the DAG. The active copy can also be moved to other DAG members. You can create a backup from the active copy or from any passive copy within the DAG that contains a database copy.

Data Protection for Exchange includes the following functions for Exchange Server DAGs:

- Querying of DAG database copies and their status

- Creating full, copy, incremental, and differential backups of active and passive databases that are managed within a DAG
- Querying of all DAG database copy backups
- Restoring of all DAG database copy backups
- Restoring into an active database, from either active or passive database copy backups
- Restoring into a Recovery (or alternate) database
- Restoring the mailbox from either active or passive database copy backups
- Deleting DAG database copy backups

When you use Data Protection for Exchange with Exchange Server DAGs, refer to the following information:

- Use a DAG member to store DAG database backups.
- All DAG members are to use the same VSS policy.
- When migrating backups to DAG member backups, the first backup is to be a `FULL` backup.
- When migrating backups to DAG member backups, the previous backups that are stored under the Data Protection member name must be manually deleted when they are no longer needed.
- Backups to `LOCAL` can be restored only to and expired from the Exchange server on which the backup was made.
- Database restores must be run on the active database.

Review your Microsoft documentation for important details that regard this new replication technology.

Database Availability Group backup optimal practices

When you use Database Availability Group backup, use the following optimal practices:

- To decrease the load on the production Exchange server, specify that the backups are taken from a healthy passive database copy. If no passive copy is available, the backup is made from the active copy of the database. To specify that backups are taken from a healthy passive database copy, add `/PREFERDAGPASSIVE` to a backup command.
- Use the command-line backup options `/EXCLUDEDAGPASSIVE`, `/EXCLUDEDAGACTIVE`, or `/EXCLUDENONDAGDBS` to exclude certain databases from backup processing.
- For databases in a DAG that have two or more healthy copies, you can skip the database integrity check by using the `/SKIPINTEGRITYCHECK` option.
- To maximize availability, schedule multiple servers to back up copies of each database. Use the `/MINIMUMBACKUPINTERVAL` option to ensure that only one copy is backed up for each backup cycle.

Database Availability Group restore optimal practices

When you use Database Availability Group restore, use the following optimal practices:

- Run restores to the active database copy.
- Backups to Tivoli Storage Manager can be restored to any Exchange server in the domain. However, backups to LOCAL can be restored only on the server where the backup was created.

- To restore to a server that is hosting a passive database copy, make the copy active before you run the restore.
- When the restore is complete, you can move the active database copy back to the passive state.

Backup types

Data Protection for Exchange supports different types of backups. The full backup, copy backup, incremental backup, and differential backup types can be performed with VSS operations.

Data Protection for Exchange backup types have the following characteristics:

Full backup (VSS)
A full backup backs up the specified database, and associated transaction logs. The Exchange Server deletes the committed log files after the database, and logs are successfully checked for integrity and backed up. If the database is not mounted, the backup fails and the transaction logs are not truncated.

Copy Backup (VSS)
A copy backup is similar to a full backup except that transaction log files are not deleted after the backup. A copy backup is used to make a full backup of the Exchange Server database without disrupting any backup procedures that use an incremental or differential backup.

Incremental Backup (VSS)
An incremental backup backs up only transaction logs. The Exchange Server deletes the committed log files after they are successfully backed up. These log files are not deleted if the backup fails. Restoration of an Exchange Server database from an incremental backup requires the following tasks:

- Restore of the last full backup
- Restore of any other incremental backups that are performed between the full backup and this incremental backup
- Restore of this incremental backup

The log files are not deleted if databases are not mounted.

Differential Backup (VSS)
A differential backup backs up transaction logs. The log files are not deleted. When you run a full backup followed by only differential backups, the last full backup and the latest differential backup contain all data that is required to bring the database back to the most recent state. This type of backup is also called a *cumulative incremental* backup.

Restoring an Exchange Server database from a differential backup requires the following tasks:

- Restore of the last full backup
- Restore of this differential backup, but no other differential backups

When circular logging is enabled, you cannot use differential or incremental backups. Data loss might occur if the log wrapped before an incremental or differential backup is finished. If you choose a backup strategy that involves incremental or differential backups, you must disable circular logging for the Exchange database from the Exchange program. See the Microsoft Exchange Server documentation for more information about circular logging.

Restore processing

A restore obtains backup copies of Exchange databases and transaction logs and returns them to the Exchange Server.

Restore processing: Actions

To perform restore processing, Data Protection for Microsoft Exchange requires that the Exchange Information Store service must be running. However, the databases that are being restored must be dismounted.

You can use Data Protection for Microsoft Exchange to restore mailbox databases to a recovery database. With Microsoft Exchange Server, you can also use the item recovery feature of the Exchange Client to recover messages and folders that are accidentally deleted. Exchange Server also provides a deleted mailbox feature to recover deleted mailboxes. For more information, see your Microsoft Exchange Server documentation. Data Protection for Microsoft Exchange also has a Mailbox Restore feature that enables mailbox and item level restore operations by using a batch or drag-and-drop selection method.

The **restorefiles** command restores the `.edb` and `.log` files from specified Data Protection for Microsoft Exchange VSS backups into a specified directory. This action enables individual mailbox recovery products from other vendors, including mailbox recovery tools to be used. You can run the **restorefiles** command from a system without an Exchange server that is installed to a specified directory on the same system as the Exchange server. For more information about this command, see "Restorefiles command" on page 153.

When a VSS restore operation is initiated, the Data Protection for Microsoft Exchange software completes the following tasks:

1. Validates the state of Exchange server objects.
2. When you use the Data Protection for Microsoft Exchange graphical user interface, you are prompted whether to dismount the databases you are restoring into.
3. Begins a session with a Tivoli Storage Manager server.
4. Verifies that the VSS service is running and that the Exchange writer is available.
5. The Tivoli Storage Manager VSS requestor performs the VSS snapshot restore preparation stage.
6. The Tivoli Storage Manager VSS requestor restores the backup data.
7. The Tivoli Storage Manager VSS requestor marks the restore as complete in VSS.
8. Optionally, mounts databases to run recovery.

Mailbox restore processing: Actions

When a mailbox restore operation is initiated, the Data Protection for Microsoft Exchange software completes the following tasks:

1. Starts a session with the Tivoli Storage Manager server.
2. Queries the Tivoli Storage Manager server for a list of available backups.
3. Selects an appropriate backup that is based on user input.
4. Creates an Exchange recovery database on a specified directory.
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 original mailbox or other location.
7. Removes the Exchange recovery database and all associated files.

Mailbox Restore Browser

When the Mailbox Restore Browser is started, the Data Protection for Microsoft Exchange software completes the following tasks:

- Detects if there is an existing recovery database. If a recovery database is found, the browser automatically connects to that database and displays the database contents. If there is not an existing recovery database, start the mailbox selection dialog. In this dialog, the user selects a mailbox or database to browse.
- Starts a session with the Tivoli Storage Manager server.
- Queries the Tivoli Storage Manager server for a list of available backups.
- Selects an appropriate backup that is based on user input.
- Creates a recovery database, and restores the selected backup into the recovery database.
- Connects to the recovery database and displays the contents.

After you select the items and they are restored, you can, optionally, remove the recovery database.

VSS restore

A VSS restore restores VSS backups (Exchange database files and log files) that are on Tivoli Storage Manager server storage to their original location.

The following characteristics are true of a VSS restore:

- VSS restore granularity is at the database level.
- Supports restoring one or more databases from a VSS snapshot backup that are located on Tivoli Storage Manager server storage.
- Restores can be run in a Database Availability Group (DAG) environment.
- Supports restoring a VSS backup (directly from Tivoli Storage Manager server storage) to an alternate system.
- Supports restoring a VSS backup to an alternate database.
- Full, copy, incremental, and differential backup types can be restored. Database copy backup types are not supported by VSS, and therefore cannot be restored.
- Supports restoring an Exchange Server 2010 and later backup that is taken from a DAG replica into the production server.
- Supports restoring a backup that is taken from a relocated database into the production server.

VSS fast restore

A VSS fast restore operation restores data from a local snapshot. The snapshot is the VSS backup that is on a local shadow volume. The restore operation retrieves the data by using a file-level copy method.

The following characteristics are true of VSS fast restores:

- Full, copy, incremental, and differential backup types can be restored. Database copy backup types are not supported by VSS and therefore, cannot be restored.
- Restore granularity is at the database level.
- The key component of producing a VSS fast restore is the speed with which the application can become operational with the data that are on local shadow

volume. Even though the data is restored relatively quickly, the transaction logs must still be replayed after the restore and therefore, the time of recovery for the application can increase.

- Supports restoring a VSS backup to an alternate database.
- Supports restoring a backup that is taken from a DAG replica into the production server.
- Supports restoring a backup that is taken from a relocated database into the production server.
- Restores can be run in a Database Availability Group (DAG) environment.
- Any backup to **LOCAL** can be restored only to the same system.

VSS instant restore

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.

The key component of producing a VSS instant restore is the speed with which the application can become operational with the data that is stored on local shadow volumes. Even though the data is restored relatively quickly, the transaction logs must be replayed after the restore. The time of recovery for the Exchange database increases as the number of logs to be replayed increases.

A VSS instant restore is only possible when all of the data from the database that is specified for restore is on storage subsystems that are supported by the VSS instant restore. If part of the data that is restored, including the log files, is on a local disk, a VSS fast restore is completed.

When you perform VSS instant restores, make sure that any previous background copies that include a copy of a volume that is being restored are completed before you initiate the VSS instant restore. However, this check is not necessary for XIV, SAN Volume Controller, or Storwize V7000 with space-efficient target volumes.

You cannot complete a partial restore by using the `/partial` option when a VSS instant restore is completed. Although Data Protection for Exchange allows this operation to begin, it either fails or completes incorrectly. If you restore only one database from a VSS backup that is stored on local VSS shadow volumes on DS8000, SAN Volume Controller, Storwize V7000, or XIV, set the `InstantRestore` option to *False* in the Restore Options pane in the **Recover** tab of the Data Protection for Exchange GUI, or specify `/instantrestore=`*no* on the command-line interface.

VSS instant restore capability is automatically disabled when a restore is completed to an Exchange recovery database.

VSS instant restore is the default restore method when all data specified for a restore is on storage subsystems that are supported by the VSS instant restore. A failover to VSS fast restore can still occur when an error is detected early enough in the VSS instant restore process to trigger the failover. In this situation, an error is logged in the `dsmerror.log` file. The `dsmerror.log` file is used by the DSMAGENT. However, a failover to VSS fast restore might not always be possible. For example, if an error occurs later in the restore process, VSS instant restore processing fails without a failover to VSS fast restore. An error can be a pending background copy on the storage subsystem, a failure to start the FlashCopy operation on the snapshot provider system, or other hardware error.

When you plan for VSS instant restore, backups can only be restored to the same DS8000, SAN Volume Controller, XIV, or Storwize V7000 storage subsystem from which they are backed up.

The list of devices that support instant restore is maintained online at http://www.ibm.com/support/docview.wss?uid=swg21455924.

Restoring VSS backups into alternate locations

An Exchange Server database backup, or DAG active or passive database copy backup can be restored into a recovery database or into an alternate (or relocated) database.

This restore capability is referred to as a *restore into* scenario and requires the following actions:

- If you are operating a VSS restore of a relocated database, you must use the *restore into* function. Also, specify the same database name as the one you are restoring. The restore fails if you do not specify the same name.
- Running any type of restore into function automatically disables VSS Instant Restore.

Backups to `LOCAL` can be restored only to the system where the backups were created.

Thin provisioning support

Thin provisioning or the ability to allocate less physical storage than the declared size of a logical storage volume is available with supported hardware. A thinly provisioned volume is referred to as a space-efficient (SE) volume.

The complete list of supported hardware for a space-efficient FlashCopy is available online at http://www.ibm.com/support/docview.wss?uid=swg21455924.

SAN Volume Controller and Storwize V7000 provide FlashCopy restore from SE target volumes and from fully allocated target volumes for which the background copy of the VSS backup is not yet completed. In addition, the hardware supports a restore from fully allocated target volumes for which the backgroud copy of the VSS backup has completed. You can retain multiple FlashCopy images of a source volume as backup generations at a much reduced storage cost. You do not have to allocate the full size of the source volume for each backup generation.

For SE target volumes, the SAN Volume Controller and Storwize V7000 hardware architectures minimize the space that is required to maintain multiple snapshots of the same source volume. Target volumes are placed into a cascade where each target is dependent on changes that are recorded in target volumes of subsequent snapshots. For example, assume that four VSS snapshots are created of a source volume. S is the source and T1 through T4 are the targets. T1 is the first, chronologically, and T4 is the last. The following cascade occurs:

```
S -> T4 -> T3 -> T2 -> T1
```

With this type of cascade relationship, a copy-on-write process is needed only between the source volume and the latest FlashCopy target. Any blocks that remain unchanged on the source volume are not copied at all. However, the cascaded relationship, where multiple SE target volumes have the same FlashCopy source, requires some special considerations when you use the target volumes as backup versions managed by Data Protection for Exchange.

Automated failover for data recovery

If you use Data Protection for Exchange with the Tivoli Storage Manager configuration, Data Protection for Exchange can automatically fail over to the secondary server for data recovery when there is an outage on the Tivoli Storage Manager server.

The Tivoli Storage Manager server that Data Protection for Exchange 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 Tivoli Storage Manager 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
- Backup-archive client node (also called the DSM agent node)
- Remote DSM agent node (for offloaded backups to the primary server)
- DAG node (for backups of databases in an Exchange Server Database Availability Group (DAG))

During normal operations, connection information for the secondary server is automatically sent to Data Protection for Exchange 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 was 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.

You can confirm that Data Protection for Exchange has failed over by looking for entries about the secondary server in the following log files:

- `Tivoli\tsm\TDPExchange\dsierror.log`
- `Tivoli\tsm\baclient\dsmerror.log`

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

- Data Protection for Exchange must be at the V7.1 level.
- The primary server, secondary server, and backup-archive client must be at the V7.1 level.
- 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
  - Backup-archive client 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, 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 during failover:

- Any operation that requires data to be stored on the Tivoli Storage Manager 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 attempt to 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.
- For more information about the failover capabilities of Tivoli Storage Manager components, see http://www.ibm.com/support/docview.wss?uid=swg21649484.

For more information about automated client failover with the Tivoli Storage Manager backup-archive client, see *Automated client failover configuration and use* in the Tivoli Storage Manager information center (http://pic.dhe.ibm.com/infocenter/tsminfo/v7r1/topic/com.ibm.itsm.client.doc/c_cfg_autoclientfailover.html).

Chapter 2. Planning

Planning information that regards backup strategies, Tivoli Storage Manager policy, available options, and preference settings is provided.

Security requirements

Data Protection for Exchange must be registered to the Tivoli Storage Manager server and use the appropriate node name and password when it connects to the Tivoli Storage Manager server. Standard Tivoli Storage Manager security requirements apply to Data Protection for Exchange.

Security requirements for Data Protection for Exchange backup and restore tasks

To run backup and restore tasks on Exchange Server, Data Protection for Exchange must be operating in an account with membership in the Organization Management group.

Security requirements for Data Protection for Exchange mailbox restore tasks on Exchange Server

To run mailbox restore tasks, Data Protection for Exchange must be operating in an account with membership in the Organization Management group. The account must have a live Exchange mailbox in the domain.

The account must have a live Exchange mailbox in the domain. By default, Windows adds the Exchange Organization Administrators group to other security groups, such as the local Administrators group and the Exchange Recipient Administrators group. If these default settings change, the account must be manually added to these other groups.

Backup and restore prerequisites

Before you run backup and restore tasks, review the following prerequisites.

Exchange Server 2010 and 2013

Tivoli Storage FlashCopy Manager software must operate in an account with membership in the Organization Management group. You must also have local Administrator privilege.

For mailbox restore and mailbox restore browser operations, membership in the Organization Management group is also required. The Exchange server must have the Client Access server role installed, or Tivoli Storage FlashCopy Manager must be configured to use a different client access server in the domain.

When you run backups, the Exchange database file size can increase because of increased database commitments that are triggered by backup operations.

Mailbox restore operations

There are several options to consider when restoring mailboxes and mailbox data. For example, you can choose where to restore the mail, and

how to restore the mail. You can restore mailbox data from the graphical or command-line interfaces of Tivoli Storage FlashCopy Manager. From these interfaces there are options for restoring interactively with the Mailbox Restore Browser or directly from Exchange database files. When restoring mailboxes and mailbox data, make sure your environment is set up to meet the following requirements:

- The administrator account that is being used to run the mailbox restore must have an active Exchange mailbox in the domain.
- Temporary space is required to accommodate the mailbox database during restore operations. Specify the temporary space on the General property page for the Exchange Server workload. On the General property page, set the following options:
  - **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 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 in the *TSM for Mail - All Requirement Documents* website at http://www.ibm.com/support/docview.wss?uid=swg21219345. When you are at the website, follow the link to the requirements technote for your specific release or update level.
- When restoring mailboxes directly from Exchange database files, verify that there is read and write access to the `EDB` file and verify that the Exchange transaction log files are available.

Microsoft Outlook cannot be installed on the server that is being used to run the mailbox restore.

When you submit a restore or mount request, all of the volumes that are contained in the original snapshot set are imported. If the number of volumes that are imported exceeds the maximum number of allowable mapped volumes for the environment, the restore or mount operation can fail.

Backup strategies

Depending on your specific requirements that regard network traffic, backup window, and acceptable restore times, you might choose to follow different backup strategies. It is important to understand all aspects of Exchange Server disaster recovery, and backup considerations that are recommended by Microsoft. Refer to your Exchange Server documentation for this information.

When you consider backup strategies, use the following guidelines:

- Do not use incremental and differential backups together.
- If you choose a strategy that involves incremental or differential backups, circular logging must be disabled on the databases of the Exchange Server.
- Consider applying DAG database replication technologies. Refer to your Microsoft documentation for details that regard this technology.

  For another Database Availability Group (DAG) backup strategy, set up all DAG members to back up all of the database copies. Use the **/MIN** and **/PREFERDAGPAS** flags.

Full backups only

Only full backups are done with this strategy.

This approach is best for Exchange Servers that are relatively small because each backup contains enough data to restore the entire database. Each backup takes longer to run. However, the restore process is the most efficient because only the most recent (or other appropriate) full backup is restored.

Full backup plus incremental backups

This strategy is commonly used when the normal backup window or network capacity cannot support a full backup each time.

In such cases, a periodic full backup followed by a series of incremental backups allows the backup window and network traffic to be minimized during peak usage times. For example, you can schedule full backups on the weekend and incremental backups during the week. The full backups can be done during low usage times when a larger backup window and increased network traffic can be tolerated. The restore process becomes more complex, however, because a full backup, and subsequent incremental backups, must be restored. In addition, transactions within the logs must be applied which increases process time. The more transactions are applied, the longer the recovery process.

If you use this backup strategy, modify the Tivoli Storage Manager storage management policies to ensure that all incremental backups are stored together on the Tivoli Storage Manager server (collocated). This action helps improve restore performance by reducing the number of media mounts necessary for restoring a series of incremental backups.

Full backup plus differentials

This strategy provides an easier restore than the full plus incremental backup strategy.

This approach might be useful if your backup window and network capacity can process the backup of all transaction logs that accumulate between full backups. This strategy requires the transfer of only one differential plus the last full backup to accomplish a restore. However, the same amount of data must be transferred in the differential image, as in the series of incremental backups.

Therefore, a full backup plus differential backup policy increases network traffic and Tivoli Storage Manager storage usage. This policy assumes that the differential backups are processed with the same frequency as the incremental backups.

Carefully consider whether there is sufficient advantage to justify the additional resource necessary to resend all prior transaction logs with each subsequent differential backup.

Back up to Tivoli Storage Manager storage versus back up to local shadow volumes

When you create policy for your backups, consider these differences between backing up data to Tivoli Storage Manager storage versus VSS disks.

Tivoli Storage Manager storage

A Tivoli Storage Manager backup operation stores the backed up data on Tivoli Storage Manager server storage. Although this type of backup typically takes longer to process than a backup to local shadow volumes, a Tivoli Storage Manager backup is necessary when long-term storage is needed. Saving Exchange data on tape for archival purposes is an example of needing long-term storage. Tivoli Storage Manager 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 Tivoli Storage Manager server storage, a point-in-time copy is available if backups on the local shadow volumes become corrupted or deleted.

Backups to Tivoli Storage Manager server storage are dictated by time, not by versions.

Local shadow volumes

Sufficient local storage space must be available on local shadow volumes for a VSS backup strategy to be successful. Ensure that there is enough available storage space that is assigned to the volumes to accommodate your Data Protection for Exchange backup operations. Environment and storage resources also impact how many backup versions are maintained on local shadow volumes (for VSS fast restore and VSS instant restore) and how many backup versions are maintained on Tivoli Storage Manager server (VSS restore and longer term storage). Create different sets of policies for backups to both local shadow volumes and to Tivoli Storage Manager server storage. If you are using a VSS provider other than the Windows VSS System Provider, make sure to review the documentation for that specific VSS provider.

Backups to local shadow volumes can be managed by both time and versions. However, because of a higher frequency of local snapshot creation, and VSS storage provisioning and space limitations, set up policy for local backups to be based on version limits. In addition, in DAG environments, all of the DAG members are to use the same local VSS policy.

Data Protection for Exchange with SAN Volume Controller and Storwize V7000

Data Protection for Exchange exploitation of SAN Volume Controller and Storwize V7000 FlashCopy capabilities on Windows is dependent on the Volume Shadow Copy Service (VSS) hardware provider for SAN Volume Controller and Storwize V7000.

Configuration of the VSS provider for SAN Volume Controller and Storwize V7000 controls what type of FlashCopy is run when a VSS snapshot is requested. It also controls the behavior that results when you use VSS snapshots.

The VSS provider that supports SAN Volume Controller and Storwize V7000 has the following characteristics:

- If the VSS provider is configured to use incremental FlashCopy, only one backup version is allowed. Only one backup version is the limit because each VSS snapshot request for a source volume causes an incremental refresh of the same target volume.

  In this scenario, deleting the VSS snapshot removes it from the VSS inventory, but the FlashCopy relationship remains with the SAN Volume Controller and Storwize V7000. A subsequent VSS snapshot of the same source volume results in an incremental refresh of the target volume.
- When the VSS provider is configured to use space-efficient target volumes - specifically, when the background copy rate is set to zero - the following is true:
  - The deletion of a VSS snapshot, that is represented by a target volume in a cascade, also causes all target volumes that are dependent on the volume that is being deleted (that is, the target volumes that were created earlier) to be deleted. For example, the deletion of a snapshot that is represented by target volume *T2* in the sample cascade *S -> T4 -> T3 -> T2 -> T1* causes *T2* and *T1* to be deleted. The cascade *S -> T4 -> T3* remains after the deletion.

    When you manually delete backups on the SAN Volume Controller and Storwize V7000 space-efficient target volumes, and multiple backup versions exist, the backup that is being deleted, and any older backups that contain the same volumes are deleted. The deletion might not occur until the next snapshot operation.
  - A FlashCopy restore of the source volume from a target volume in a cascade of multiple target volumes is destructive to the target volume that is being restored and to all newer targets in the cascade. For example, the restore of a snapshot that is represented by target volume *T3* in the sample cascade *S -> T4 -> T3 -> T2 -> T1* causes *T4* and *T3* to be deleted. The cascade *S -> T2 -> T1* remains after the restore.

    One exception to this pattern is that a FlashCopy restore from an space-efficient target that is the only target in the cascade is not destructive.
  - If an space-efficient target volume runs out of space to hold the data from changed blocks on the source volume, that target volume and all target volumes that are dependent on that target volume go offline and render those backup versions unusable.

A space-efficient backup version is defined by a FlashCopy to an space-efficient target volume that has a background copy rate of zero. The use of space-efficient target volumes with the `autoexpand` option that is enabled and a background copy rate set to greater than zero does not create space-efficient backup versions. The target volumes grow to the allocated size of the source volumes when the background copy completes.

Given these characteristics, the following requirements apply to Data Protection for Exchange support of SAN Volume Controller and Storwize V7000:

- Using a mix of space-efficient and fully allocated target volumes is not supported. You must choose to use either space-efficient or fully allocated volumes for FlashCopy targets, and set the VSS provider background copy rate parameter.

  A transition from fully allocated targets to space-efficient targets is accommodated by treating fully allocated targets as if they were 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. In addition, if you use space-efficient backup versions, refer to following item in this list.

- When you use space-efficient backup versions:
  - Do not mix persistent and nonpersistent VSS snapshots. Use of a nonpersistent VSS snapshot that follows one or more persistent snapshots causes the older persistent snapshots to be deleted when the nonpersistent snapshot is deleted.

    A VSS backup with `backupdestination` set to *TSM* creates a nonpersistent VSS snapshot. Therefore, do not follow a series of backups to local with `backupdestination` set to *TSM*. Instead, set `backupdestination` to *both* to send data to Tivoli Storage Manager while it preserves local snapshot backup versions. The settings `backupdestination=LOCAL` and `backupdestination=TSM` are mutually exclusive. Do not use both in a backup strategy.
  - Enable `autoexpand` for the space-efficient target volumes to avoid out-of-space conditions.
  - Allocate enough space for space-efficient target volumes to hold 120 percent 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 percent per day, VSS backups are done every six hours, and a steady rate of change throughout the day is assumed, the expected change rate between snapshots is 5 percent of the source volume (20/4). Therefore, the allocated space for the space-efficient target volumes is to be 1.2 times 5 percent equal to 6 percent 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.
  - Do not delete snapshots manually. Allow Data Protection for Exchange to delete backup versions that are based on the defined policy to ensure that deletion is done in the correct order. This process avoids deletion of more backup versions than expected.

Instant restore

Data Protection for Exchange supports VSS instant restore operations when multiple backup versions exist on SAN Volume Controller and Storwize V7000 space-efficient target volumes.

However, in this situation, VSS instant restore accesses snapshot volumes that contain dependent FlashCopy relationships. The snapshot volumes that create the dependency are typically backups that are created after the snapshot that is being restored. These snapshot volumes are removed for the VSS instant restore operation to complete successfully. The backups that included the deleted snapshots are deleted from storage. This destructive restore operation occurs only when VSS instant restore operations occur in an environment where Data Protection for Exchange manages multiple backup versions on SAN Volume Controller and Storwize V7000 space-efficient target volumes.

When multiple backup versions exist, all snapshots that are newer than the snapshot that is being restored are deleted during the VSS instant restore operation. The snapshot that is being restored is also deleted. When only one snapshot backup version exists, the snapshot that is being restored is not deleted.

More guidelines for SAN Volume Controller and Storwize V7000 environments

There are additional guidelines you can use when protecting data in SAN Volume Controller and Storwize V7000 environments. For example, you can change the background copy rate to have the background copies complete more quickly.

The default background copy rate is *50*. This value minimizes impact to response time for host system I/O, but it might not complete background copies as quickly as you want. Increasing the background copy rate that is used by the VSS provider to a value greater than *50* causes the background copies to complete more quickly. Do not set the background copy rate higher than *85* because this action can significantly lengthen response times to I/O from host systems.

You can review the following guidelines before you attempt backup operations:

- Determine whether to use space-efficient or fully allocated backup targets before you issue a backup operation. Provision enough target volumes in the SAN Volume Controller VSS_FREE volume group for as many of the backup versions you require. If you use fully allocated target volumes, their capacity size 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 issuing the `ibmvcfg set backgroundCopy 0` command. To activate the changes, restart the IBM VSS system service after you issue the command. For more details about configuring the IBM VSS Hardware Provider for space-efficient target volumes, make sure to read the appropriate VSS-related content in the SAN Volume Controller or Storwize V7000 documentation.
- Do not mix COPY and NOCOPY FlashCopy relationships from the same source volume or volumes.
- Do not mix fully allocated and space-efficient VDisks (used for backup targets) in the VSS_FREE pool.
- If the protected data is on SAN Volume Controller or Storwize V7000 volumes, and the VDisks in the VSS_FREE pool are space efficient, then VSS instant restore from multiple backups is possible. However, the VSS instant restore operation in this environment is destructive.
- Make sure that IBM VSS hardware provider is installed. This provider must be configured to accommodate multiple backup versions on SAN Volume Controller or Storwize V7000 space-efficient target volumes.

These guidelines apply specifically to NOCOPY FlashCopy backups on SAN Volume Controller or Storwize V7000:

- While NOCOPY FlashCopy backups can be mounted remotely, you must use either SAN Volume Controller or Storwize V7000 storage to restore a NOCOPY FlashCopy backup.
- You can create a NOCOPY FlashCopy to a space-efficient target. However, protection from physical failures to the source volume is not provided.

Make sure to review your IBM VSS hardware provider documentation for important information about these two issues:

- IBM VSS hardware provider prerequisites (for example, Microsoft VSS hotfixes).
- Configuration instructions for creating FlashCopy mappings of NOCOPY backups on SAN Volume Controller or Storwize V7000.

Space-efficient target volumes go offline when their capacity limit is exceeded. As a result, the current backup and all older backups (which are not reached

FULL_COPY status) are lost. To avoid this situation, use the AUTOEXPAND option when you create space-efficient targets. This option allocates more physical storage to prevent space-efficient target volumes that are going offline.

Restriction: When you use VSS instant restore operations with multiple backup versions that exist on SAN Volume Controller or Storwize V7000 space-efficient target volumes, use only full or copy type backups when the backup destination specifies local. A local backup (including any local backups that are created after the one being restored) is deleted by SAN Volume Controller or Storwize V7000 because of the destructive restore behavior. As a result, any full, copy, incremental, or differential local backup is removed and unavailable for restore operations. If you want to use incremental or differential local backups with SAN Volume Controller or Storwize V7000 space-efficient target volumes, disable VSS instant restore during any restore operations to avoid this situation.

VSS limitations for SAN Volume Controller and Storwize V7000

When you run a Data Protection for Exchange VSS backup (non-offloaded) with the backup destination of Tivoli Storage Manager server, in some cases the SAN Volume Controller or Storwize V7000 LUNs remain mapped to the Windows host even though the backup is complete. In this situation, the Exchange Server data is on SAN Volume Controller or Storwize V7000 disks and the IBM System Storage VSS Hardware Provider is used. To work around this issue, you can use a backup destination other than Tivoli Storage Manager server (BOTH or LOCAL). You can also manually unmap the volumes that are attached to the Windows host.

When you run two Data Protection for Exchange 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. In this situation, the Exchange Server data is on SAN Volume Controller or Storwize V7000 disks. However, the second backup is waiting for the SAN Volume Controller or Storwize V7000 background copy of the first backup to complete before proceeding. SAN Volume Controller or Storwize V7000 does not allow two background copies of the same volume to occur at the same time. There is no indication that the second backup is waiting for the first background copy to complete.

You might also see timeout errors if the previous SAN Volume Controller or Storwize V7000 background copy takes too long. To work around this issue, schedule your VSS backups far enough apart to accommodate this situation. You can also increase the copy rate of the SAN Volume Controller or Storwize V7000 background copy.

VSS operations in IBM N-series and NetApp environments

Be aware of these storage space guidelines when you perform VSS operations in IBM N-series and NetApp environments.

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. The following formula can be used to help 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 blocks equal to two times the specified size of the LUN to be used. This space reservation ensures writes for virtual disks. The following example demonstrates how to calculate the size of these 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
```

VSS limitations for NetApp FAS series or IBM N-series

NetApp FAS series and IBM N-series require certain limitations.

Because of the limitations in SnapDrive 4.2 and any supported prior versions, the VSS Provider for NetApp FAS series and IBM N-series, VSS-based operations that use Data Protection for Exchange with backup destination set to `LOCAL`, must be done in specific ways. Failure to comply with the following configuration and operational guidelines can lead to serious conditions such as premature deletion of snapshots that represent VSS backups to `LOCAL`, backup failure, and out of space conditions on the production volumes. When the limitations in the SnapDrive are addressed by NetApp, Data Protection for Exchange VSS operations can be fully used. However, this situation is not applicable when FlexVols are used.

Exchange Server storage configuration for NetApp FAS series or IBM N-series VSS operations

If you plan to run VSS backups with backup destination set to `LOCAL`, check your setup to ensure that following requirements are met.

- The NAS file server LUNs used by a database must be fully dedicated to the database. The Exchange Server databases cannot share LUNs.
- A NAS filer LUN used by the Exchange Server databases must be the only LUN on the filer volume. For example, if Exchange uses four LUNs, there must be four corresponding filer volumes, each volume that contains one LUN.

Guidelines for VSS backup operations for NetApp FAS series or IBM N-series

If you plan to run VSS backups with backup destination set to `LOCAL`, these backups must adhere to the following guidelines.

- If the NetApp volume type is Traditional, VSS backups with backup destination set to `LOCAL` must be bound to a management class that has `verExists=1`. This setting is not required if flexible volumes are used.
- VSS backups with backup destination set to `LOCAL` can either be a full or copy backup. You cannot mix local backups of type full and copy.
- VSS backups with backup destination set to `TSM` can be a full or copy backup. There are no restrictions on Tivoli Storage Manager backups.
- When running VSS backups, you must ensure that previous backup finishes completely before you start a new backup. Any overlap of backups can result in undesirable side-effects on the Exchange Server, the VSS service, and, the NAS filer.

Sample VSS backup procedure for NetApp FAS series or IBM N-series

Taking the prior considerations into account, the following section describes a sample backup procedure that can be used to run VSS backups by using both Tivoli Storage Manager and local backup destinations in an optimal manner. The following assumptions apply to this sample backup procedure:

- The configuration requirements that are stated are met.
- The VSS backup to Tivoli Storage Manager takes one hour to complete.
- The VSS backup to a local destination takes five minutes to complete.

Your backup procedure can consist of the following backups:

- Daily VSS full backups to a local destination at every four hours - 12 a.m., 4 a.m., 8 a.m., 12 p.m., 4 p.m., 8 p.m.
- Daily VSS full backups to Tivoli Storage Manager storage by one of the following two methods:
  - Set **backupdestination** to BOTH at 12 a.m. This specification creates a 12 a.m. backup to a local destination. Therefore, no separate 12 a.m. backup to local is required.
  - Full offloaded-backup at 1 a.m. No VSS local backup is available to restore from between 1 a.m. and 4 a.m. when next VSS backup to local takes place.
- run weekly VSS full backups to Tivoli Storage Manager (offloaded backup) 5 a.m.

Microsoft Exchange Server 2013 support

When planning to restore Exchange Server 2013 mailboxes, there is a requirement for MAPI clients to use RPC over HTTPS (also known as Outlook Anywhere). With Exchange Server 2013, Microsoft does not support RPC over HTTP.

When planning to protect Microsoft Exchange Server 2013 data, use the following list to verify your environment is set up correctly:

- Use Exchange 2013 CU2 or later and the correct MAPI download. These software requirements are documented in the Hardware and Software Requirements technote. This technote is available in the *TSM for Mail - All Requirement Documents* website at http://www.ibm.com/support/docview.wss?uid=swg21219345. When you are at the website, follow the link to the requirements technote for your specific release or update level.
- Verify that in the Exchange environment you can open the current administrator's mailbox in Outlook or Outlook Web App.
- Verify that in the Exchange environment you can open the target mailbox in Outlook or Outlook Web App.
- (Optional) Use MFCMapi to verify your MAPI configuration. To complete this verification, a link to a Microsoft blog post is provided http://blogs.msdn.com/b/dvespa/archive/2013/05/27/omniprof.aspx. If MFCMapi cannot connect to Exchange, Data Protection for Microsoft Exchange cannot connect to Exchange either.

If your environment is configured correctly, mailbox restore works as with previous versions of Microsoft Exchange Server. If you experience problems, for troubleshooting information, see Chapter 7, “Troubleshooting,” on page 95.

Preparing for VSS instant restore in DS8000, Storwize V7000, XIV, and SAN Volume Controller environments

When you prepare for a restore, consider the type of data that you want to restore and where the backups are located.

About this task

The information in the following list is specific to the DS8000, Storwize V7000, XIV, and SAN Volume Controller environments. When planning for a VSS instant restore, consider the following facts:

- Restore granularity is at the volume level.
- VSS requires that data must always be restored to the same drive letters and paths as existed during the original backup.
- VSS requires IBM System Storage Support for Microsoft Volume Shadow Copy Service software if you use a DS8000, Storwize V7000, or SAN Volume Controller disk subsystem.
- VSS requires IBM XIV Provider for Microsoft Windows Volume Shadow Copy Service if you are using an XIV disk subsystem.
- Backups must be located on the same XIV, DS8000, Storwize V7000, or SAN Volume Controller storage subsystem to which they are restored.
- In a DAG environment, stop the Microsoft Exchange Replication Service on the active node before you run the VSS instant restore operation.
- In an Exchange 2013 environment, stop the Exchange Search Host Controller Service on the active node before you run the VSS instant restore operation.

Policy management

With Data Protection for Microsoft Exchange software you can manage and configure storage management policy for backups.

Although Tivoli Storage Manager policy determines how Data Protection for Microsoft Exchange backups are managed on Tivoli Storage Manager 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 there is enough available storage space 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 is dependent on the VSS provider that you use.

How backups expire based on policy

Backups expire based on Data Protection for Exchange policy.

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

The data value 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 you can specify the number of snapshot backups to retain and the length of time to retain them.

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

For Exchange Database Availability Group backups that use the DAG node, only the system on which the backup was 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 Tivoli Storage Manager server. It 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.

A number of backup copies are retained. When the maximum number of backup copies is reached, the oldest backup expires and is deleted. The maximum number of backup copies is specified in the Data Protection for Exchange 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 policy.

How Tivoli Storage Manager server policy affects Data Protection for Exchange

Tivoli Storage Manager policy determines how Data Protection for Exchange backups are managed on Tivoli Storage Manager storage and on local shadow volumes when the environment is configured for VSS operations.

The Tivoli Storage Manager server recognizes Data Protection for Exchange as a node. Data that is backed up to Tivoli Storage Manager storage from this Data Protection for Exchange node is stored and managed according to settings specified for Tivoli Storage Manager server policy items.

Tivoli Storage Manager policy can manage the VSS backups that are placed in Tivoli Storage Manager server storage pools. The Tivoli Storage Manager server is responsible for managing VSS backups.

If you used IBM Tivoli Storage Manager for Copy Services and upgraded to Data Protection for Exchange, with the license for Tivoli Storage Manager for Copy Services, you can store VSS backups to local shadow volumes.

The number of local backup versions that are maintained by the Tivoli Storage Manager server is determined by the value that is specified by the Tivoli Storage Manager server **verexists** parameter that is defined in the copy group of the

management class to which the local backup belongs. Allocation of Target volume sets is not necessary when you use the system provider. When you are not using the system provider, the number of target volume sets allocated for local backups is to be equal to the **verexists** parameter. Target volume sets are not applicable to XIV.

For example, if `verexists=3`, then at least 3 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 Tivoli Storage Manager 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*.

When you use the configuration wizard, offered through the graphical user interface, the **VSSPOLICY** parameter that is to be configured is set in the `tdpexc.cfg` file.

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

The policy management of local backups is responsible for reconciling the local backup repository with the information stored on the Tivoli Storage Manager server. For example, if target volume LUNs that were used for a local backup are removed from the storage subsystem, the information that represents the backup on the Tivoli Storage Manager server must be reconciled. Likewise, if the Tivoli Storage Manager 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 subsystem. This release is necessary so that these LUNs can be used for future backup operations. Tivoli Storage Manager automatically detects these situations and does the reconciliation.

Consider the scenario where you use a two-member 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* are not expired until the next time you do a backup, query, or deletion operation on *MEMBER1*. This scenario can lead to using more storage than specified by **verexists**.

Storage space considerations for local shadow volumes

Tivoli Storage Manager requires that sufficient storage space is to be available to create shadow volumes that are required for VSS backup processing. Even when the VSS backup destination is the Tivoli Storage Manager server, storage space to create a shadow volume is still required, though on a temporary basis.

Because the value of the **verexists** parameter that is specified for your local backup policy, determine the number of backup versions to retain on local shadow volumes, a `verexists=1` setting causes the deletion of an existing backup on local shadow volumes (during a VSS backup to Tivoli Storage Manager server storage) to create enough temporary space for the new snapshot. Therefore, if you want to keep *N* backups on local shadow volumes and also do VSS backups to Tivoli Storage Manager server storage, provision enough storage space on local shadow volumes and specify `verexists=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, there are no backups.

If you keep multiple backups (snapshots), the oldest backup is removed before a new 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 there are to be five retained backups, but the latest backup fails, you might have only four backup versions.

Make sure to specify a **verexists** value that accommodates your VSS backup goals. If you have limited storage space for VSS operations and are restricted to a `verexists=1` setting, you can take advantage of the **Backup Destination**`BOTH` option. This stores the backup on local shadow volumes and sends a copy to Tivoli Storage Manager server storage.

It is possible for VSS backups (that Data Protection for Exchange creates and stores on local shadow volumes) to be changed and deleted from outside of Tivoli Storage Manager control. For example, the Microsoft VSSADMIN DELETE SHADOWS command can remove a VSS backup that are managed by Tivoli Storage Manager without Tivoli Storage Manager being able to prevent such a removal. In such a situation, Tivoli Storage Manager recognizes the backup removal and reconciles its index of available backups with what is on local shadow volumes. Because of this potential for removal, it is important to establish a strategy that protects VSS backup data that is stored on local shadow volumes from being compromised.

Policy considerations for VSS backups

The following issues impact your Tivoli Storage Manager policy for managing VSS backups:

- Overall backup strategy.
- Length of time that VSS backups are on Tivoli Storage Manager server storage.
- Number of VSS backup versions on Tivoli Storage Manager server storage.
- Types of VSS backups that are on Tivoli Storage Manager server storage.
- Number of VSS backup versions on local shadow volumes.
- Types of VSS backups on local shadow volumes.
- The amount of available target volume storage that is provisioned for VSS operations.

Specifying policy binding statements

Policy binding statements associate Exchange backups to a management policy.

About this task

Specify policy binding statements to use for binding snapshots to a policy. You can specify policy binding statements by using the GUI or by manually adding binding statements to the configuration file. A default policy binds any backups that are not explicitly bound to a named policy. Policy binding is available in environments with or without a Tivoli Storage Manager server.

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

A policy statement is defined in the respective configuration file. For example:

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

Binding backups to policies

A backup policy determines how backups on local shadow volumes are managed and retained. You can add, update, delete, or change the processing order of existing binding statements.

Procedure

1. From the Management Console, select the **Exchange Server** instance from the tree view.
2. In the **Protect** tab, click **Properties** in the **Action** pane.
3. Select **VSS Policy Binding** from the list of available property pages.
4. Add, update, delete, or change the processing order of existing binding statements.

   Tip: You can use a wildcard character (*) to mean "all". For example, specify a wildcard character in the **Server** field to bind the policy to all Exchange servers.
5. Optional: Use **Move Up** and **Move Down** to modify the processing order. Policies are processed from the bottom up and processing stops at the first match. To ensure that more specific statements are processed, the more general specification are to be listed before the more specific ones. This processes the general ones after the more specific specifications. Otherwise, the more general specifications match the target before the more specific specifications are seen.
6. Save any new or changed binding statement.
7. Optional: Verify new or updated policies and bindings.
   a. Run one or more test backups.
   b. In the **Recover** tab, verify the management classes that are bound to your test backups.

VSSPOLICY statements when changing backup types

When you change from legacy backups to VSS backups, pay attention to the VSSPOLICY statements set for the backup.

For VSS backups, VSSPOLICY statements are used to associate VSS backups with management classes. These VSSPOLICY statements are entered in the configuration file (for example, `tdpexc.cfg`).

A configuration files can include multiple VSSPOLICY statements. The configuration file is read from the bottom up. If you are familiar with the Tivoli Storage Manager backup-archive client configuration file, the VSSPOLICY statements in the `tdpexc.cfg` file are read like the INCLUDE statements configured in the `dsm.opt` file.

If there are no VSSPOLICY statements included in the configuration file, or if the VSSPOLICY statements do not match the type of backup that is set up, the default management class for the policy domain is used. Backup expiration parameters for

the default management class might differ from the settings 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. Check the parameters to verify the backups expire according to the business needs of your environment.

Any policy changes entered in the `tdpexc.cfg` files require that you restart the Tivoli Storage Manager Client Acceptor Daemon (CAD), Tivoli Storage Manager Remote Client Agent (DSMAgent), and the Tivoli Storage Manager Scheduler Service for Exchange. 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 code sample presents the syntax of a VSSPOLICY statement:

```
VSSPOLICY srv name "database-name" backup-type backup-dest mgmtclass
```

The Exchange server name is defined by the *srv name* variable. You can enter the wildcard character (*) to match all Microsoft Exchange servers.

The database name is defined by the *"database-name"* variable. You can enter the wildcard character (*) to match all Microsoft Exchange databases. Because the name can include a space, use the quotation marks to encapsulate the database name.

The *backup-type* variable specifies the backup type. For example, `FULL` or `COPY` or the wildcard character (*) to match all backup types.

The *backup-dest* variable specifies the backup destination. The options are `TSM` to backup to Tivoli Storage Manager or `LOCAL` for a backup to a local disk or the wildcard character (*) to match both backup types.

The *mgmtclass* variable specifies the Tivoli Storage Manager management class that is used to bind the types of specified backups.

The following code sample presents an example of a VSSPOLICY statement. This code sample is part of the sample configuration file that is included with the software that you installed. In this example, the VSSPOLICY statement is commented out. To make the VSSPOLICY statement effective, uncomment the VSSPOLICY statement by removing the initial asterisk character (*).

```
------------------------------------------------------------
* 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    "DB 1" *     TSM     IUG1
*********************************************************************
```

In this example, the following policy rules are specified:

- Any VSS backups of the *DB 1* database on the Exchange server *SERVER1* to Tivoli Storage Manager 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 Tivoli Storage Manager are bound to the management class *IUG_TSM_COPY*.
- Full VSS backups of any other database on any other Exchange server to Tivoli Storage Manager are bound to the management class *IUG_TSM*.
- This policy is complete. 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

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

About this task

For Microsoft Exchange Server databases in a 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 to back up database copies from different DAG members under a single DAG node.

All database copies can be managed as a single entity. This management is regardless of where they were backed up from, and whether they were active or passive at the time of the backup. You can then set up a minimum interval between database backups. The minimum interval ensures that the database copies are not backed up at the same time or backed up too frequently.

Procedure

To manage DAG members by using a single policy, complete the following steps:

1. Use the Tivoli Storage Manager Configuration Wizard to configure the DAG node.
   - For VSS backups to Tivoli Storage Manager, ensure that you specify a node name in the **DAG Node** field in the TSM Node Names page in the wizard. This node is used to back up all the databases in a Database Availability Group.
   - For best results, ensure that all the DAG members are configured with the same DAG node name.

2. Ensure that the Tivoli Storage Manager 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 Tivoli Storage Manager server, the proxies are to be defined. In addition, the backup archive client node and the Data Protection node need `proxynode` authority. The backup archive client node also needs `proxynode` authority to act on behalf of the Data Protection node. For example, the Tivoli Storage Manager administrator can issue the following commands on the Tivoli Storage Manager server:

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

3. Ensure that the DAG node and the Data Protection for Exchange node are in the same policy domain.
4. Create a backup schedule and specify the **/MINIMUMBACKUPINTERVAL** parameter in a backup command. You must use the Tivoli Storage Manager scheduler to run this schedule. For example, to use a single Tivoli Storage Manager schedule to back up exactly one copy of a database that contains multiple copies, do the following steps:
   a. Create a command script named `C:\BACKUP.CMD` with the 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 to this schedule.

   When the backup schedule is run, the minimum backup interval is observed and only one backup is created.
5. Optional: To decrease the load on the production Exchange server, you can specify that the backups are taken from a healthy passive database copy. If a healthy passive copy is not available, the backup is made from the active copy of the database. To do this specification, add **/PREFERDAGPASSIVE** to a backup command. For example:

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

Related concepts:

"Migrating backups to a DAG node" on page 56

Data Protection for Exchange DAG member name settings

Review the following settings when registering your Data Protection for Exchange DAG member name.

The system where Data Protection for Exchange is installed must be registered to the Tivoli Storage Manager server with a DAG member name. This DAG member name owns and manages all Data Protection for Exchange data that is backed up to the Tivoli Storage Manager server. Specify this DAG member name with the **nodename** option in the `dsm.opt` options file located (by default) in the Data Protection for Exchange installation directory. To run VSS operations, you might be required to register DAG member names for more systems.

Use the following Tivoli Storage Manager parameter conditions when you register your Data Protection for Exchange DAG member name (system) to the Tivoli Storage Manager server:

- **MAXNUMMP** This parameter determines the maximum number of mount points a client node is allowed to use on the Tivoli Storage Manager server during a backup operation.
- **TXNGroupmax** This parameter determines the number of files that are transferred as a group between Data Protection for Exchange and the Tivoli Storage Manager server between transaction commit points. This parameter must have a value of 12 or greater.
- **COMPRESSIon** This parameter determines whether the backup-archive client node compresses data before it sends the data to the Tivoli Storage Manager 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 Data Protection for Exchange options

Several Data Protection for Exchange parameters need to be configured.

The Tivoli Storage Manager administrator is to provide you with the node name, password, and the communications method with the appropriate parameters to connect to the Tivoli Storage Manager server. These values, with other parameters, are stored in an options file that are located (by default) in the Data Protection for Exchange installation directory. If needed, edit the `dsm.opt` file by using a text editor.

If you edit the `dsm.opt` file, make sure that the Data Protection for Exchange options file and the backup-archive client options file specify the same Tivoli Storage Manager server.

The options file includes the following parameters, which are required for initial configuration:

COMMMethod
This option specifies the communication protocol to use between the Data Protection for Exchange node with the Tivoli Storage Manager server. Data Protection for Exchange supports the same set of communication protocols that are supported by other Tivoli Storage Manager clients on Windows systems. Depending on the chosen `commmethod`, the connectivity parameters for that `commmethod` must be specified as well.

For all backups, specify the **commmethod** option in the Data Protection for Exchange options file. In addition, specify the **commmethod** 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 **commmethod** option in the backup-archive client options file that is used as the Remote DSMAGENT Node.

NODename
The Tivoli Storage Manager node name is the unique name by which Tivoli Storage Manager recognizes the system that runs Data Protection for Exchange.

The following options are not required for initial configuration. By default they are not specified, but you can modify the default settings:

CLUSTERnode
Leave this option blank. When the option is blank, the default value is used.

COMPRESSIon

This option instructs the Tivoli Storage Manager API to compress data before it is sent to the Tivoli Storage Manager server. This compression reduces traffic and storage requirements. If you enable compression, it affects performance in two ways:

- CPU usage increases on the system on which Data Protection for Exchange is running.
- Network bandwidth use is lower because fewer bytes are sent.
- Storage usage on the Tivoli Storage Manager server is reduced.

You might want to specify `compression yes` if any of the following conditions exist:

- The network adapter has a data overload.
- Communications between Data Protection for Exchange and Tivoli Storage Manager server are over a low bandwidth connection.
- There is heavy network traffic.
- You can also use the `compressalways yes` option (with the `compression yes` setting) to specify that file compression continues even if the file grows as a result of compression.

It might be better to specify `compression no` in the following cases:

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

The Tivoli Storage Manager administrator can override the compression option setting for the Data Protection for Exchange node when they register or update the node by specifying, on the Tivoli Storage Manager server side, that a particular node:

- Always uses compression.
- Never uses compression.
- Leaves the decision up to the client (default value).

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 available in the client documentation before you attempt to compress your data.

DEDUPLication

Client-side data deduplication is used by the Tivoli Storage Manager API to remove redundant data during backup and archive processing before the data is transferred to the Tivoli Storage Manager server. Specify whether the Tivoli Storage Manager API deduplicates data before it is sent to the Tivoli Storage Manager server. You can specify `Yes` or `No`. The default

value is No. The value of the deduplication option for Data Protection for Exchange applies only if the Tivoli Storage Manager administrator allows client-side data deduplication.

The deduplication and **enablelanfree** options are mutually exclusive. You can use either one option or the other, but not both options together.

You can turn on client-side data deduplication by adding DEDUPLICATION YES to the dsm.opt file and by making sure that the deduplication prerequisites are met.

ENABLECLIENTENCRYPTKEY

This option encrypts Exchange databases during backup and restore processing. One random encryption key is generated per session and is stored on the Tivoli Storage Manager server with the object in the server database. Although Tivoli Storage Manager manages the key, a valid database must be available to restore an encrypted object. Specify enableclientencryptkey yesin the options file. In addition, assign the type of encryption to use by specifying the encryptiontype option in this same options file. You can specify DES56 (56 bit) or AES128 (128 bit). The default is AES128. In this same file, you must also specify the databases that you want encrypted 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, 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 available in the client documentation before you attempt to encrypt your databases.

For example, encrypt your Exchange database backups by adding the following three options:

1. Add the enableclientencryptkey yes option.
2. Add the **encryptiontype** option with the type of encryption to use.

ENABLELANFree

This option allows Data Protection for Exchange to run in a LAN-free environment (if you are equipped to do so). To run a LAN-free VSS backup with Data Protection for Exchange, specify enablelanfree yes in the DSMAGENT (VSS Requestor) options file. See *Managed System for SAN Storage Agent User's Guide* for detailed information about LAN-free environments.

INCLUDE and EXCLUDE

Use the VSSPOLICY statement in the Data Protection for Exchange configuration file to set policy for VSS backups.

The general **include** and **exclude** syntax is displayed:

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

where objectNameSpecification is:

```
ExchangeServerName\ExchangeStorageGroupName\...\backupType
```

where backupType is one of the following:

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

The Tivoli Storage Manager API does not allow sending any of the three data types (meta, data, logs) that comprise an Exchange database backup to different storage destinations on the Tivoli Storage Manager server.

This example excludes `Database 1` from a backup:

```
EXCLUDE "SERVER1\Database 1\...\*"
```

This example binds all objects for database `DB2` to management class `CLASS1`:

```
INCLUDE "SERVER1\DB2\...\*" CLASS1
```

This example binds all `Directory` backups to management class `CLASS2`:

```
INCLUDE "SERVER2\Directory\...\*" CLASS2
```

This example binds all `incremental` objects to management class `CLASS3`:

```
INCLUDE "SERVER3\...\incr" CLASS3
```

This example binds mailbox history objects to management class `CLASS4:`

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

Consider the following behavior when you set **include** and **exclude** statements:

- The wildcard character (*) matches zero or more characters.
- The wildcard character (?) matches any one character.
- The wildcard character (*) within a qualifier replaces zero or more characters only within that qualifier. The qualifier itself must exist in the matching object name. To match zero or more qualifiers, use ellipses (\...\).
- Incremental object names are always unique. These names contain qualifiers whose values make them unique. Incremental object names are generated at the time of the backup and therefore are not predictable and cannot be specified.
- Include/exclude lists are processed from the bottom up and processing stops at the first match. To ensure that more specific specifications are processed at all, the more general specification are to be listed before the more specific ones, The general specifications are processed after the more specific ones. Otherwise, the more general specification matches the target before the more specific specifications are seen.
- When a match is found, processing of the list stops and the statement that matches is examined.
  - If it is an **exclude** statement, the matching object name is not backed up.
  - If it is an **include** statement, the matching object name is backed up.

  If the **include** statement contains a `ManagementClassName`, that management class is associated with the object name, for this backup and for all backups of the same name on the current node.
- If a match is not found, the object is backed up using the default management class for the current node.
- If a match is found for an **include** statement that specifies a management class but the specified management class is not valid for the current node, the default management class for the current node is used.

- Exchange database names must be of the correct case, as shown by the displayed results from the **query exchange** or **query tsm**. Data Protection for Exchange constants must be lower-case: `meta, data, logs`. However, now the Windows Tivoli Storage Manager API assumes the specifications are for a Windows file system and ignores case. Because they might be accepted in the future, the correct case must always be used.

PASSWORDAccess

This option instructs the Tivoli Storage Manager API to store the current password (encrypted) in the Windows registry and automatically generates a new one when the current one expires. This method of password management is recommended when you run scheduled, unattended backups since it ensures that the backup never fails because of an expired password. The default is `prompt`.

You can manage the password as stored in the registry by using a utility program named `dsmcutil.exe`. This utility program is distributed with the Tivoli Storage Manager backup-archive client package. For more information about using the dsmcutil program, see the `dsmcutil.hlp` file or the `dsmcutil.txt` file that are distributed with the Tivoli Storage Manager backup-archive client package.

You can create more Data Protection for Exchange options files to point to other Tivoli Storage Manager servers. You might also want to create more than one options file, each with different parameters to use with a single Tivoli Storage Manager server.

Specifying Data Protection for Exchange preferences

Data Protection for Exchange configuration parameters are defined in the Data Protection for Exchange configuration file (`tdpexc.cfg` and by default). These configuration parameters determine such preferences as the location of your log file, how date and time are displayed, and the performance tuning parameters.

You can set the values of the Data Protection for Exchange configuration parameters. Use the Management Console (MMC) GUI or the command-line interface:

- In the MMC GUI, set the value in Properties.
- Use the **tdpexcc set** command in the Data Protection for Exchange command-line interface. See "Set command" on page 173.

Bind VSS backups to Tivoli Storage Manager policy by selecting **Properties** > **VSS Policy Binding** in the MMC GUI, and then entering appropriate values in the fields.

Proxy node definitions (VSS backups)

Since Data Protection for Exchange VSS backup operations are implemented through the Tivoli Storage Manager backup-archive client, you must use node names specifically for VSS operations in addition to using a node name for where Data Protection for Exchange is installed.

As part of the configuration procedure, a proxy relationship is defined for these various node names. By default, this proxy relationship is defined when you run the configuration wizard. You can use this topic for information about manually completing the configuration.

This proxy relationship allows node names to process operations on behalf of another node name. When you register these nodes to the Tivoli Storage Manager server for VSS operations, do not specify the Tivoli Storage ManagerUSerid=NONE parameter. VSS operations fail when this parameter is specified.

There are two types of node names that are defined in proxy node relationships:

- *Target node*: A node name that controls backup and restore operations and that also owns the data on the Tivoli Storage Manager server. This node name is specified in the Data Protection for Exchange 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.

These nodes are defined by using the backup-archive client **`grant proxy`** command. For example:

```
GRANT PROXY TARGET=dpexc_node_name AGENT=dsmagent_node_name
```

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 2. Required node names for basic VSS operations

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

Target node

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

Agent node

This node name is where the backup-archive client and VSS provider are installed. This node is responsible for processing the VSS operations as Data Protection for Exchange does not process any direct 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 configuration file (`tdpexc.cfg` by default). You can use the Properties window of the Management Console (MMC) GUI by selecting **VSS backup**. From here, you can update the Local DSMAGENT Node name. Otherwise, use the **`tdpexcc set`** command to specify this parameter.

Note: The agent node and target node are on the same system for basic VSS operations.

Required node names for basic VSS offloaded backups

VSS offloaded backups require specific node name settings.

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

Table 3. Required node names for basic VSS offloaded backups

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

Target node

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

Agent node

This node name is where the backup-archive client and VSS provider are installed. This node is responsible for processing the VSS operations as Data Protection for Exchange itself does not process any direct 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 configuration file (tdpexc.cfg by default). You can use the Properties window of the Management Console (MMC) GUI by selecting **VSS backup**. From here, you can update the Local DSMAGENT Node name. Otherwise, use the **tdpexcc set** command to specify this parameter.

Agent 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 is responsible for moving VSS snapshot data from local shadow volumes to the Tivoli Storage Manager 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 configuration file (tdpexc.cfg by default). You can use the Properties window of the MMC GUI by selecting **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

transportable shadow copies. It is not supported if you are using the default VSS system provider. Refer to your VSS provider documentation for details.

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

Chapter 3. Installing and upgrading

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

Installation prerequisites

Before you install the software, ensure that your system meets the minimum hardware, software, and operating system requirements.

Details of the 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 most current requirements, review the Hardware and Software Requirements technote that is associated with the level of your Data Protection for Exchange program. This technote is available in the *TSM for Mail - All Requirement Documents* website at http://www.ibm.com/support/docview.wss?uid=swg21219345. When you are at the website, follow the link to the requirements technote for your specific release or update level.

Minimum hardware requirements

Before you install the software, ensure that your system meets the minimum hardware requirements.

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

Hardware for an x64 system
Compatible hardware that is supported by the Windows operating system and Exchange Server

Details of the hardware and software requirements change over time because of maintenance updates and the addition of operating system, application, and other software currency support.

For the most current requirements, see the Hardware and Software Requirements technote that is associated with your level of software. This technote is available from the following website: http://www.ibm.com/support/docview.wss?uid=swg21219345

When you go to this website, follow the link to the requirements technote for your specific release or update level.

Software and operating system requirements

Details of the software and operating system requirements for Data Protection for Microsoft Exchange can change over time.

For current requirements, see the *TSM for Mail - All Requirements Documents* website, http://www.ibm.com/support/docview.wss?uid=swg21219345.

Virtualization environment

Information about virtualization environments that can be used with Data Protection for Exchange is available.

For more information, see the *IBM Tivoli Storage Manager (TSM) guest support for Virtual Machines and Virtualization* website at http://www.ibm.com/support/docview.wss?uid=swg21239546.

Installing and configuring Data Protection for Microsoft Exchange

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

Before you begin

Before you install and configure, verify that the hardware and software requirements are met. Details of the hardware and software requirements change over time because of maintenance updates and the addition of operating system, application, and other software currency support.

For the most current requirements, see the Hardware and Software Requirements technote that is associated with your level of software. This technote is available from the following website: http://www.ibm.com/support/docview.wss?uid=swg21219345

When you go to this website, follow the link to the requirements technote for your specific release or update level.

Procedure

Follow these instructions to quickly install, configure, verify, and customize Data Protection for Exchange:

1. Install Data Protection for Exchange.
   a. Log on as an administrator.
   a. Insert the Data Protection for Exchange product DVD into your DVD drive. If autorun is enabled, the setup wizard starts automatically when the DVD loads. Otherwise, click **Start** > **Run**, and at the prompt, specify: `x:\setupfcm.exe`, where `x:` is your DVD drive. Click **OK**.
   b. Follow the installation instructions that are displayed on the screen.
   c. If prompted, restart your system before the installation completes.
   d. Click **Finish** to complete the installation of Data Protection for Exchange.
   e. If you plan to use VSS operations, you must install the most recent version of the Tivoli Storage Manager backup-archive client. The backup-archive client is also the VSS Requestor and is available separately.
2. Configure Data Protection for Exchange.

a. Start the Management Console (MMC GUI) by clicking **Start** > **All Programs** > **Tivoli Storage Manager** > **Data Protection for Microsoft Exchange Server** > **DP for Exchange Management Console**. If you did not previously configure Data Protection for Exchange, the Tivoli Storage Manager configuration wizard starts automatically.

b. If the Tivoli Storage Manager configuration wizard does not start automatically, click **Manage** > **Configuration** > **Wizards** in the tree view, select the wizard, and click **Start** in the Actions pane.

c. 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 for help on resolving potential issues.

TSM Node Names
Specify the Tivoli Storage Manager 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 Tivoli Storage Manager 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 application is installed. This node stores the Data Protection for Exchange 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 Tivoli Storage Manager 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 from making too many backups of the same database.

  Important: On the Tivoli Storage Manager 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 Tivoli Storage Manager server address, and choose whether to have the wizard configure the Tivoli Storage Manager server. Alternatively, you can view and change the commands that the configuration wizard uses to configure the Tivoli Storage Manager 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 Tivoli Storage Manager server, the Tivoli Storage Manager administrator must configure the Tivoli Storage Manager server before verification can be completed. If the wizard does not configure the Tivoli Storage Manager server, it provides a link to a macro that can be provided to the Tivoli Storage Manager administrator as an example of one way to configure the Tivoli Storage Manager server.

3. Verify the configuration.
   a. Verify that VSS is working correctly.

      If the **VSS Diagnostics** check box was 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.

      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**.

   b. Verify that Data Protection for Exchange is configured properly:
      1) 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.
      2) Change **PowerShell** to **Command Line**.
      3) Click the folder icon, and select the `verify_exc.txt` file. Then, click **Open**.
      4) 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 to verify your configuration. The configuration is verified when these commands run without warnings or errors.
      5) When verification is complete, you can use Data Protection for Exchange to back up and restore Exchange server data.
      6) Back up and restore a set of test data.
4. Customize Data Protection for Exchange.

After Data Protection for Exchange is configured and verified successfully, customize your settings by defining your policy settings and scheduled operations. This action ensures that your business requirements are satisfied.

What to do next

For detailed information about the installation and configuration procedures, or if you want to do these tasks manually, see Chapter 3, "Installing and upgrading," on page 43 and Chapter 4, "Configuring," on page 59.

Installing Data Protection for Microsoft Exchange on a local system

The setup wizard guides you through installing Data Protection for Exchange.

Before you begin

Before you install and configure, verify that the hardware and software requirements are met. Details of the hardware and software requirements change over time because of maintenance updates and the addition of operating system, application, and other software currency support.

For the most current requirements, see the Hardware and Software Requirements technote that is associated with your level of software. This technote is available from the following website: http://www.ibm.com/support/docview.wss?uid=swg21219345

When you go to this website, follow the link to the requirements technote for your specific release or update level.

About this task

Data Protection for Exchange 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.

See the `README.FTP` file for instructions about how to install a fix pack or interim fix package. The `README.FTP` file is available in the same directory where the maintenance package is downloaded.

Procedure

To install Data Protection for Exchange from a DVD, complete the following steps:

1. Install Data Protection for Exchange by using the setup wizard. The wizard installs the product and any prerequisites such as the .NET Framework and Report Viewer.
   a. Log on as an administrator.
   b. Insert the Data Protection for Exchange product DVD into your DVD drive.

If autorun is enabled, the installation dialog starts automatically when the DVD loads. Otherwise, select **Start** > **Run**, and at the prompt, specify: `x:\setupfcm.exe`, where `x:` is your DVD drive, and click **OK**.

c. Follow the installation instructions that are displayed on the screen.

d. If prompted, restart your system before the installation completes.

e. Click **Finish** to complete the installation of Data Protection for Exchange.

The Management Console (MMC) GUI is shared among Data Protection for Exchange, Data Protection for SQL Server, and Tivoli Storage FlashCopy Manager. If one of these products is installed in a location other than the default, 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`.

2. To install more language packs, see "Installing and activating the language packs."
3. If you plan to back up and restore local snapshots or run VSS offloaded backup operations, follow the tasks that are described in "Installing Tivoli Storage FlashCopy Manager." If not, for important configuration information see "Configuring Data Protection for Microsoft Exchange for TSM Configuration" on page 59.

Installing Tivoli Storage FlashCopy Manager

IBM Tivoli Storage FlashCopy Manager is a separately purchasable program that provides application-aware backups and restores by using the advanced snapshot technologies of storage system.

Before you begin

Before you begin, ensure that the Data Protection for Exchange product is installed.

About this task

For information about how to install Tivoli Storage FlashCopy Manager, see Installing Tivoli Storage FlashCopy Manager.

What to do next

After you install Data Protection for Exchange and Tivoli Storage FlashCopy Manager, see "Configuring Data Protection for Microsoft Exchange for TSM Configuration" on page 59 for important configuration information.

Installing and activating the language packs

Each language pack contains language-specific information for the Management Console (MMC) GUI, command-line output, and messages. The installation wizard automatically identifies the language of your geographical area, and loads the language pack for that language.

Installing more language packs

To view the Management Console (MMC) GUI, command-line output, and messages in a language other than English, install the language pack that you want. The language packs are executable program files that are in their respective language directory on the product DVD.

Before you begin

Make sure that Data Protection for Exchange is successfully installed before you attempt to install the language packs.

About this task

The `setupfcm.exe` program automatically starts the setup program for the MMC language pack if installation is done on a computer with a supported language other than English.

The configuration wizard automatically provisions a language pack for any components it provisions. The following instructions describe how to install a language pack manually.

Procedure

1. Insert the product DVD into the DVD drive and select **Run** from the **Start** menu.
2. Run the following commands:

   Data Protection for Exchange Management Console language packs

   `x:\fcm\aaa\mmc\4100\bbb\setup.exe`

   Data Protection for Exchange language packs

   `x:\fcm\aaa\languages\bbb\setup.exe`

   Where `x:` is your DVD drive, *aaa* is `x64`, and *bbb* is the three-letter country code that is associated with that language.
3. Follow the installation instructions that are contained in the prompt windows.
4. Click **Finish** to complete the installation.

What to do next

After you install the language pack, you must activate it.

Activating the language packs

After you install the language pack, you must activate the language by updating the Data Protection for Exchange configuration file (`tdpexc.cfg` by default).

Procedure

Activate the language by using either of the following methods:

- Use the **`set`** command with the **`language`** parameter to specify the language that you want. For example:

  ```
  tdpexcc set lang=fra
  ```

  See the description of the **`language`** parameter in "Set positional parameters" on page 174 for a list of available languages and their three-letter country codes.
- Use the property pages to set the language by doing the following steps:
  1. Select the Exchange server instance in the tree view.

2. Click **Properties** in the Actions pane.
3. Select the Regional property page.
4. Click **Regional and Language Options** to ensure that system settings match the language that you want to use. The Management Console (MMC) GUI uses system language settings.
5. Select the language from the list of installed language packs. The Data Protection components use language settings from a configuration file.
6. For the best results and correct operation, select the language that matches the system settings. Click **Match MMC language** to automatically update the language to match the system.

Installing Data Protection for Microsoft Exchange silently

Follow these instructions to create a silent installation package.

Before you start, 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.

The following example uses `c:\tdpdpkg` as a staging directory. Issue the following commands to create the package.

Table 4. Commands for creating a silent installation package

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

When you create the installation package, test the silent installation. When you complete the test, the package can be placed on a DVD or it can be made available from a shared directory.

Silently installing Data Protection for Microsoft Exchange with the setup program

Use the setup program to silently install Data Protection for Exchange.

Before you begin

You must install two components: Data Protection for Exchange 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 Management Console setup program
(64-bit) `x:\fcm\x64\mmc\4100\enu\setupfcm.exe`

Data Protection for Exchange setup program
(64-bit) `x:\fcm\x64\exc\7100\enu\setup.exe`

The Data Protection for Exchange Management Console and Data Protection for Exchange 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.

Run the following commands to silently install both components to the default installation directories:

```
x:\fcm\x64\mmc\4100\enu\setupfcm.exe /s /v/qn
x:\fcm\x64\exc\7100\enu\setup.exe /s /v/qn
```

where `x:` is your DVD drive.

You must substitute the appropriate feature when you install a language other than English. For more information, see "Silent installation features (Language Packages only)" in “Installing Data Protection for Microsoft Exchange silently” on page 50.

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

```
x:\fcm\x64\mmc\4100\enu\setupfcm.exe /s /v"INSTALLDIR=\"C:\Program Files\Tivoli\"
ADDLOCAL=\"Client\" TRANSFORM=1033.mst REBOOT=ReallySuppress /qn /l*v
\"C:\Temp\DpExcMmcSetupLog.txt\""
```

```
x:\fcm\x64\exc\7100\enu\setup.exe /s /v"INSTALLDIR=\"C:\Program Files\Tivoli\tsm\"
ADDLOCAL=\"Client\" TRANSFORM=1033.mst REBOOT=ReallySuppress /qn /l*v
\"C:\Temp\DpExcSetupLog.txt\""
```

The following list identifies a few additional facts to remember when completing this installation process:

- You must place a backslash (\) 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 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 software. The installation package includes a number of prerequisites that is installed by Data Protection for Exchange. 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.

Creating batch files

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

The following sample script (`c:\setup.bat`) demonstrates an unattended installation:

```
@echo off
 rem ==================================
 rem sample silent install script
 rem
 call x:\fcm\x64\mmc\3200\enu\setupfcm.exe /s
/v"INSTALLDIR=\"C:\Program Files\Tivoli\" ADDLOCAL=\"Client\" TRANSFORM=1033.mst
 REBOOT=ReallySuppress /qn /l*v \"C:\Temp\DpExcMmcSetupLog.txt\""
 rem
 call x:\fcm\x64\exc\6400\enu\setup.exe /s
 /v"INSTALLDIR=\"C:\Program Files\Tivoli\tsm\" ADDLOCAL=\"Client\"
 TRANSFORM=1033.mst REBOOT=ReallySuppress /qn /l*v \"C:\Temp\DpExcSetupLog.txt\""
 rem ==================================
 rem code could be added after the
 rem installation completes to
 rem customize the dsm.opt files
 rem if desired
 rem ==================================
```

Installing with MSI

Use the Microsoft Installer program, `msiexec.exe`, to silently install Data Protection for Exchange.

Before you begin

Data Protection for Exchange 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 `setup.exe` and `setupfcm.exe` programs, the `msiexec.exe` program does not install any prerequisites. When you use `msiexec.exe`, you must install all prerequisites manually.

For the most current requirements, review the *Hardware and Software Requirements* technote that is associated with the level of your Data Protection for Exchange program. This technote is available in the *TSM for Mail - All Requirements Documents* website at http://www.ibm.com/support/docview.wss?uid=swg21219345. When you are at the website, follow the link to the requirements technote for your specific release or update level.

Procedure

The following examples show how to use **`msiexec`** to install the Data Protection for Exchange Management Console and Data Protection for Exchange. Enter each **`msiexec`** command on a single line.

1. Install the Data Protection for Exchange Management Console.

   ```
   msiexec /i"x:\fcm\x64\mmc\3200\enu\IBM Tivoli Storage Manager for Mail
   - MS Exchange - Management Console.msi" RebootYesNo="No"
   Reboot="Suppress" ALLUSERS=1 INSTALLDIR="c:\program files\tivoli"
   ADDLOCAL="Client" TRANSFORM=1033.mst /qn /l*v "c:\temp\DpExcMmcLog.txt"
   ```

   Where *x:* is your DVD drive.

2. Install Data Protection for Exchange:

```
msiexec /i"x:\fcm\x64\exc\6400\enu\IBM Tivoli Storage Manager for Mail
- MS Exchange.msi" RebootYesNo="No" Reboot="Suppress" ALLUSERS=1
INSTALLDIR="c:\program files\tivoli\tsm" ADDLOCAL="Client"
TRANSFORM=1033.mst /qn /l*v "c:\temp\DpExcLog.txt"
```

Where *x:* is your DVD drive.

What to do next

You can install language packs in a similar way. MSI files for the language packs are in the language folders that are associated with each component. For language packs, use `ADDLOCAL="LanguageFiles"` instead of `ADDLOCAL="Client"`.

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.

Installation problems: Capturing a log of the installation

If a silent installation fails, record the symptoms and environment information for the failing installation and contact customer support with that information. You can create a detailed log file of the failed installation that can facilitate analysis of your situation.

The following environmental information can be helpful:

- Operating system level
- Service pack
- Hardware description
- Installation package (DVD or electronic download) and level
- Any Windows event log that is relevant to the failed installation
- Other Windows services active at the time of the installation (for example, antivirus software)

Before you contact support, check for the following items:

- You are logged on to the local system console, not through a terminal server.
- You are logged on as a local administrator, not a domain administrator. Cross-domain installations are not supported.

Assuming that all looks correct, gather a detailed log of the failing installation in to a file called `setup.log`. To generate a log file, ensure that `/l*v \"filename\"` is used on the command-line interface.

For example, issue the following command on a single line to generate a log file named `C:\Temp\DpExcSetupLog.txt`:

```
x:\fcm\x64\exc\6400\enu\setup.exe /s /v"INSTALLDIR=\"C:\Program
Files\Tivoli\tsm\" ADDLOCAL=\"Client\" TRANSFORM=1033.mst
REBOOT=ReallySuppress /qn /l*v \"C:\Temp\DpExcSetupLog.txt\""
```

Creating the package on a DVD or a file server

Use these instructions to create a silent installation package on a DVD or a file server.

The administrator has a choice of making the package available in different ways. These ways include burning a DVD or placing the package in a shared directory on a file server. Typically, the package contains the Data Protection for Exchange code distribution files and a batch file for a silent installation.

Creating a silent installation package

Follow these instructions to create a silent installation package.

Before you start, 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.

The following example uses `c:\tdpdpkg` as a staging directory. Issue the following commands to create the package.

Table 5. Commands for creating a silent installation package

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

When you create the installation package, test the silent installation. When you complete the test, the package can be placed on a DVD or it can be made available from a shared directory.

Playing back the silent installation

When the package is available on a DVD or from a shared directory, it can be played back (run) on another computer.

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

From a silent installation package on DVD:

If autostart is enabled, the silent installation begins as soon as the DVD is inserted into the drive. If autostart is not enabled, the silent installation can be run by starting the `setup.bat` file from the root of the DVD.

```
cd /d g:\
setup.bat
```

From a distribution directory:

If the package was placed in a shared directory that is called `tdpdpkg` at `\\machine1\d$`, another computer can run the command: `net use x: \\machine1\d$` to share the drive as drive `x`. You can issue the following command:

```
cd /d x:\tdpdpkg
setup.bat
```

In either case, the silent installation begins.

Setup error messages

The **setup.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

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

About this task

Upgrading the software is a three-step process:

Procedure

1. Download the updates.
2. To install the updates, run `setupfcm.exe`.
3. Start the Management Console by clicking **Start** > **All Programs** > **Tivoli Storage Manager** > **Data Protection for Microsoft Exchange Server** > **DP for Exchange Management Console**. When you start the Management Console 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 found on the system, the configuration process varies. The wizard provides instructions to guide you through the process.

   If the configuration wizard does not start automatically, click **IBM Tivoli Storage Manager** in the tree view, and click **Configuration**. Then, double-click **Wizards**.

Migration considerations

Migration from earlier versions of Data Protection for Exchange is supported.

After you upgrade and configure from the older version of Data Protection for Exchange to the newer version, use VSS restore for local VSS backups that were originally created with the older version of the software. .

If you used a previous version of Data Protection for Exchange in a Microsoft clustering environment and you upgrade to a newer version of Data Protection for Exchange, any existing backups that are completed on cluster disks do not count toward the maximum number of versions. New backups for clustered disks that are completed with the newer version of Data Protection for Exchange are managed logically for the cluster. Except for the active backup, older backups eventually expire. When you no longer must retain the active backup, the active backup must be deleted by using the **delete backup** command. The existing backups are restorable.

Migrating backups to a DAG node

If you are migrating from a version earlier than Data Protection for Exchange V6.4, when you configure Data Protection for Exchange to back up databases in an Exchange Server Database Availability Group (DAG) to a common DAG node, all DAG databases are backed up with the new DAG node name, and no longer with the previous Data Protection node name.

After migration, the first backup must be a full backup. To be successful, understand how to manage the backups from previous versions of Data Protection for Exchange.

Do not mix backups that were created with previous versions of Data Protection for Exchange 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\TDPExchange` directory, and use a new DAG node name to store the new backups. If you want to view or restore a backup that is stored under the previous node name, change the Data Protection for Exchange configuration. Over time, when the old backups are no longer useful, you must manually delete them.

To view and restore backups that are stored under the previous Data Protection node name:

1. Remove the **DAG Node** by using the General properties page, configuration wizard, or the **set** command on the command-line interface.
2. Restart or refresh the Management Console (MMC) GUI or command-line interface.
3. Click the **Recover** tab in the MMC GUI, or run a `tdpexcc query tsm *` command. Since the **DAG Node** parameter is not set, Data Protection for Exchange lists the backups that are stored under the Data Protection for Exchange node.
4. Proceed to restore one or more of the listed backups.

To manually delete old backups that are stored under the previous Data Protection node name:

1. Remove the **DAG Node** by using the General properties page, configuration wizard, or the **set** command on the command-line interface.
2. Restart or refresh the GUI or command-line interface.
3. Click the **Recover** tab in the MMC GUI, or run a `tdpexcc query tsm *` command. Since the **DAG Node** parameter is not set, Data Protection for Exchange lists the backups that are stored under the Data Protection for Exchange node.
4. Delete the backups that are expired.

Improving mailbox history handling

To improve performance, mailbox history includes only the mailboxes from databases than are backed up.

About this task

If you back up mailbox history with a version of Data Protection for Exchange earlier than version 7.1, you can manually delete the old mailbox history. Data Protection for Exchange backs up a new set of mailbox history data.

With the new mailbox history data, you can experience better performance when backing up mailbox history. It is also easier to find the mailbox when you restore a mailbox. And, 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 created by earlier versions of Data Protection for Exchange. When you restore deleted or moved mailboxes from the old backup copies, you have to specify the **/MAILBOXORIGLOCATION** parameter. After the old backup copies expire, mailbox history works without needing to specify the **/MAILBOXORIGLOCATION** parameter.

If you delete old mailbox history data, save a backup copy before completing the deletion task. You can use the backup copy when you need location information for the deleted and moved mailboxes.

If you delete the old mailbox history, the command you use is a delete filespace command. If you do not enter the command correctly, all previous backups, including backups of Exchange 2010 data, might be deleted.

Procedure

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

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

   Keep this file for reference. If you need to restore a mailbox from the old backup copies, and the mailbox location changes before the time of deleting mailbox history, use the **/MAILBOXORIGLOCATION** parameter to restore the mailbox.
2. Complete the following steps to delete the old mailbox history information from the Tivoli Storage Manager server.
   a. Launch Tivoli Storage Manager command-line administrative interface, `dsmadmc.exe`.
   b. Log on to the Tivoli Storage Manager server.
   c. Enter 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 node name when the **DAGNODE** is not being used.

      ```
      tsm: FCM>QUERY FILESPACE CXCLAB_EXC *MAILBOXINFO
      ```

The following results are displayed:

```
Node Name    Filespace    FSID   Platform   Filespace   Is Files-  Capacity   Pct
             Name                           Type        pace                  Util
                                            Unicode?
-----------  -----------  ----   --------   ---------   ---------  ---------  -----
CXCLAB_EXC   cxcserver.-    52   TDP MSE-   API:ExcD-   No            0 KB      0.0
             com\MAILB-          xchg       ata
             OXINFO
```

3. Enter the following command to delete the filespace for the old mailbox history:

   ```
   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
   ```

Chapter 4. Configuring

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

About this task

TSM Configuration
When you select the TSM Configuration configuration option, you configure Data Protection for Microsoft Exchange to work with Tivoli Storage Manager server. Data Protection for Microsoft Exchange must be installed on your system. A Tivoli Storage Manager server must be available to communicate with Data Protection for Microsoft Exchange.

Mailbox Restore Only
When you select the Mailbox Restore Only configuration option, you configure Data Protection for Microsoft Exchange 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 software functionality. The functionality offered with this configuration option is included in the other configuration options.

Configuring Data Protection for Microsoft Exchange for TSM Configuration

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

About this task

To configure Data Protection for Exchange for TSM Configuration, complete the following steps:

Procedure

1. Start the Management Console by clicking **Start** > **All Programs** > **Tivoli Storage Manager** > **Data Protection for Microsoft Exchange Server** > **DP for Exchange Management Console**.
2. From the start page, click **Configuration**. Alternatively, from the tree view, navigate to the **Configuration** node. Then, double-click **Wizards**.
3. In the results pane, double-click **TSM Configuration** to open the Tivoli Storage Manager 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 are configuring an application for which you do not have the necessary

license, 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.

c. In the TSM Node Names page, specify the Tivoli Storage Manager node names that 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 data.
   - In the **Data Protection for Exchange** field, enter the node name where the Data Protection application is installed. This is the node name that is used to store the Data Protection for Exchange 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 node. They 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 Tivoli Storage Manager 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 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 to all use the same DAG node. On the Tivoli Storage Manager 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 Tivoli Storage Manager server.

   Create a node name that can help you distinguish the type of backup that is run. 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 have to be related to the VSS Requestor or the Data Protection for Exchange node name. For example, you can name it `TSMDAG`.

d. Enter information for the Tivoli Storage Manager server that you are connecting to and click **Next** to continue.
   - In the **Tivoli Storage Manager Server Address** field, enter the TCP/IP domain name or a numeric IP address for the Tivoli Storage Manager server that contains the backups. Obtain this information from your Tivoli Storage Manager server administrator.
   - In the **Tivoli Storage Manager Server Port** field, enter the port number for the Tivoli Storage Manager server that contains the backups. Obtain this information from yourTivoli Storage Manager administrator.
   - Specify whether to have the wizard to configure the Tivoli Storage Manager 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 so that it can be provided to the Tivoli Storage Manager administrator as an example of one way to configure the Tivoli Storage Manager 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 field:

- The name of the Tivoli Storage Manager administrator account.
- The password for the Tivoli Storage Manager administrator.
- Click **Test Communications** if you want to test your connection with the Tivoli Storage Manager server. This button is not available until the VSS requestor is installed.
- Click **Review/Edit** to review or update the Tivoli Storage Manager macro file. Alternatively, you can review the macro file and directly run the commands on the Tivoli Storage Manager server.

e. Select the **Default** configuration setting. When you select the **Default** configuration setting, in addition to configuring the applications that you selected, the VSS Requestor is configured. 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 to complete the wizard.
6. Optional: For a VSS configuration, verify that the **Run VSS diagnostics when this wizard exits** option is selected. When this option is selected, after you complete the wizard, a diagnostic process tests the VSS snapshots on your system.

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

What to do next

The configuration wizard automatically installs the Tivoli Storage Manager backup-archive client.

After you configure Data Protection for Exchange, complete the following steps to verify the configuration:

1. In the Management Console, 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
   ```

4. Click **Enter** to run the commands to verify your configuration.

Manually configuring Data Protection for Microsoft Exchange for TSM Configuration

If you manually configure Data Protection for Exchange, complete the following steps.

Perform these tasks on the computer that runs the Exchange Server

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

Before you begin

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

Procedure

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

1. Specify your Data Protection for Exchange node name and communication method in the `dsm.opt` file located (by default) in the Data Protection for Exchange installation directory. More options are also available.
2. Using the **set** command, specify your Data Protection for Exchange preferences (language, date format, log file) in the `tdpexc.cfg` file in the Data Protection for Exchange installation directory.
3. (VSS only) If you are configuring Data Protection for Exchange in an Exchange Server Database Availability Group (DAG) environment, use 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 Tivoli Storage Manager 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, to 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 Tivoli Storage Manager 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 made from the active database copy. This parameter must be specified as part of a **backup** command script that is run by the Tivoli Storage Manager 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 configuration file. Exchange servers that use the same DAG node name are to share the VSS Policy.
7. (VSS only) Configure the Tivoli Storage Manager 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 configuration file (`tdpexc.cfg`).

   For more information, see *Tivoli Storage Manager for Windows Backup-Archive Clients Installation and User's Guide* and "Proxy node definitions".
8. (VSS only) Install and configure the Tivoli Storage Manager 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 Tivoli Storage Manager 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) If you want to manage local persistent VSS snapshots, which are created for VSS backups to LOCAL, VSS Instant Restores, and want to run offloaded backups, you must install IBM Tivoli Storage FlashCopy Manager.
11. (VSS only) Add the Microsoft Exchange Server binary path to the PATH statement in the system environment variables. For example:

    ```
    "C:\Program files\Exchsrvr\bin"
    ```

    Verify that **ESEUTIL.EXE** tool exists in this directory. This tool is used by Data Protection for Exchange 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. There is no installation or configuration that is required if you are using 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 hold all copies of the VSS backups as designated by your policies. Provisioning storage space to manage VSS snapshots is dependent on the VSS provider that you use. Consult the VSS Provider documentation for more details.

Perform these tasks on the Tivoli Storage Manager server

Ensure sure that the Tivoli Storage Manager server is available before you process this task.

Procedure

Follow these steps on the Tivoli Storage Manager server:

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 backup and restore requirements. For VSS operations, Tivoli Storage Manager server authentication must be on.

2. Register your Data Protection for Exchange node name and password with the Tivoli Storage Manager **register node** command. For example, for VSS operations, this node is the target node. When you register nodes to the Tivoli Storage Manager server specifically for VSS operations, do not specify the Tivoli Storage Manager **USerid**=NONE parameter. VSS operations fail when this parameter is specified.
3. If not already defined, register your Tivoli Storage Manager 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 Tivoli Storage Manager 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) with the Tivoli Storage Manager **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 Tivoli Storage Manager backup-archive client and DAG node names and passwords with the Tivoli Storage Manager **register node** command.
   b. Ensure that the Tivoli Storage Manager 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 Tivoli Storage Manager server, the proxies are to be defined. In addition, the backup archive client node and the Data Protection node need proxynode authority. The backup archive client node also needs proxynode authority to act on behalf of the Data Protection node. For example, the Tivoli Storage Manager administrator can issue the following commands on the Tivoli Storage Manager server:

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

What to do next

For any warning messages that 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 Tivoli Storage Manager. Other warnings have links to web sites where you can download the packages that you are to successfully complete the configuration process.

Perform these tasks on the system that runs the offloaded backups

This task is for VSS operations only.

Procedure

Perform the following steps on the computer that is running the offloaded backups:

1. Configure the Tivoli Storage Manager 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 configuration file (`tdpexc.cfg`) on the local, not offload, system.
2. Install and configure the Tivoli Storage Manager Client Acceptor (CAD) Service and the Remote Client Agent Service (DSMAGENT) if they 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\Exchsrvr\bin`). Verify that the `ESEUTIL.EXE` tool is stored in this directory. Data Protection for Exchange 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 are required to 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\Exchsrvr\bin"
   ```

5. Install and configure a VSS provider if you are not using the default system VSS provider. Consult the VSS provider documentation for information about the configuration of that software.

Perform these tasks to configure your system for mailbox-level and item-level restore operations

To use the Data Protection for Microsoft Exchange mailbox restore feature, there are more configuration steps that the configuration wizard addresses.

About this task

Because of an Exchange Server requirement, the Data Protection for Microsoft Exchange configuration wizard checks the Microsoft Exchange Server MAPI Client and Collaboration Data Objects versions on the Exchange Server from which you are running the mailbox restore. If the incorrect version is used, a warning is displayed with a link to a site where you can download the correct version.

The Client Access Server Role must also be configured to run Mailbox Restore operations on Exchange Server 2010 and later. For more information about specifying the Client Access Server with the **set** command, see "Set syntax" on page 173.

The Microsoft Exchange Server MAPI Client and Collaboration Data Objects download is required to process mailbox restore operations. Microsoft does not support installing the Exchange MAPI and Outlook MAPI on the same system. For more information, see the Microsoft documentation.

Perform these tasks to test your configuration

Before you start a backup or restore operation, verify that Data Protection for Microsoft Exchange is installed and configured correctly.

Verifying the configuration from the integrated command line

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
   ```

4. With the cursor in the command-line panel press Enter to run the commands to verify your configuration.

The Data Protection for Microsoft Exchange server configuration is verified as correct when these commands complete without errors or warnings.

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. For best results, complete these tests before you install Tivoli Storage Manager.

When these tests complete without errors, you can install Tivoli Storage Manager. 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 are to be fixed before you finish the configuration and start a Data Protection for Microsoft Exchange operation.

Using the DISKSHADOW command

Before you install Data Protection for Exchange, test the core VSS function first. VSS function can be validated with the Windows Server-embedded command DISKSHADOW. DISKSHADOW is available for Windows Server 2008, Windows Server 2008 R2, and later operating systems. The following list identifies the DISKSHADOW tests to complete before any Tivoli Storage Manager components are installed.

1. Test non-persistent shadow copy creation and deletion.
   - 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 the Windows Event Log contains no errors.

2. Test Persistent shadow copy creation and deletion.
   - Run DISKSHADOW on 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 Database and log volumes. Repeat this test four times, verify the Windows Event Log contains no errors.
3. Test Non-persistent transportable shadow copy creation and deletion.
   - 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 Database and log volumes. Repeat this test four times, and verify the Windows Event Log contains no errors.

After the tests complete satisfactorily, you can install Tivoli Storage Manager components.

Diagnose the cause of common errors returned from VSS operations

The following two errors are commonly returned when a VSS operation is running. Information is provided to help locate the cause of the error.

ANS1017E (RC-50) Session rejected: TCP/IP connection failure
This message is displayed when the Tivoli Storage Manager 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 Tivoli Storage Manager server is not configured correctly for the proxy nodes.

Configuring Data Protection for Microsoft Exchange for Mailbox Restore Only

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

About this task

To configure Data Protection for Exchange for Mailbox Restore Only, complete the following steps:

Procedure

1. Start the Management Console by clicking **Start** > **All Programs** > **Tivoli Storage Manager** > **Data Protection for Microsoft Exchange Server** > **DP for Exchange Management Console**.
2. From the start page, click **Configuration**. Alternatively, from the tree view, navigate 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 and click **Next** to move to the next page.
5. Click **Finish** in the Completion page to complete the wizard.

SAN Volume Controller and Storwize V7000 configuration examples

The following sections provide examples of configurations. These examples can be used when planning to set up backup and recovery solutions.

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

SVC and Storwize V7000 settings
Create 14 SE target volumes for each source volume to be

protected. Enable autoexpand for the SE target volumes. Add the SE target volumes to the VSS free pool.

VSS Provider settings
Set background copy rate equal to *0*.

Data Protection for Exchange settings
Set the policy to retain 14 local backup versions. Schedule snapshot backups as required by using backup destination equal 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 are 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 also send the backup to Tivoli Storage Manager.

SVC and Storwize V7000 settings
Create two SE target volumes for each source volume to be protected. Enable autoexpand for the SE target volumes. Add the SE target volumes to the VSS free pool.

VSS Provider settings
Set background copy rate equal to *0*.

Data Protection for Exchange settings
Set the policy to retain two local backup versions. Schedule snapshot backups as required by using backup destination equal 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 are 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 Tivoli Storage Manager.

SVC 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 (*50*). Configure to use incremental FlashCopy.

Data Protection for Exchange settings
Set the policy to retain one local backup version. Schedule snapshot backups as required by using backup destination equal to both.

When you use incremental FlashCopy, 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 one.

Production application data are on standard volumes. Keep two snapshot backup versions. Full physical copies are required for local backup versions. Perform VSS backups every 12 hours with one backup daily sent to Tivoli Storage

Manager.

SVC 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 default background copy rate (*50*).

Data Protection for Exchange settings
Set the policy to retain three local backup versions. Schedule VSS backups as follows: backup destination equal to local at *11:00*, backup destination equal to both at *23:00*.

Set the policy for local snapshot backups to retain *N+1* backup versions so that *N* snapshot backups are available for restore.

Production application data are 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 Tivoli Storage Manager.

SVC and Storwize V7000 settings
Create five SE target volumes for each source volume to be protected. Enable autoexpand for the SE target volumes. Add SE target volumes to the VSS free pool.

VSS Provider settings
Set background copy rate equal to *0*.

Data Protection for Exchange settings
Set the policy for local snapshot backups to retain five local backup versions. Schedule VSS backups as follows: backup destination equal to local at *06:00*, *12:00*, and *18:00*, backup destination equal 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 are on SE 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 Tivoli Storage Manager.

SVC and Storwize V7000 settings
Create three SE target volumes for each source volume to be protected. Allocate the same percentage of real storage as for source volumes. Add SE target volumes to the VSS free pool.

VSS Provider settings
Use default background copy rate *50*.

Data Protection for Exchange settings
Set the policy to retain three local backup versions. Schedule VSS backups as follows: backup destination equal to local at *06:00*, *12:00*, and *18:00*, backup destination equal 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. This setting allows thin provisioning for both source and target volumes and allows them to grow together.

Chapter 5. Protecting data

After you complete the configuration process, start theManagement Console to protect your Exchange Server data.

About this task

To start the Management Console, click **Start** > **All Programs** > **Data Protection for Microsoft Exchange Server** > **DP for Exchange Management Console**. If you try to use Data Protection for Microsoft Exchange Server before you complete the configuration process, the software does not function correctly.

The Management Console that is displayed is the Microsoft Management Console (MMC) with Data Protection for Microsoft Exchange Server software displayed as a plug-in. The console uses a navigation tree to organize the computer data that you have registered. Each computer icon that is followed by the word *Dashboard* represents a physical computer.

When you register a computer, information about this particular system 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 assigned to each system and is used when backing up and restoring 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 additional computers in the tree to include in an operation.

Managing remotely

From one Data Protection for Microsoft Exchange installation you can manage all of the Data Protection for Microsoft Exchange installations in your enterprise.

Before you begin

To remotely manage systems, Microsoft Windows PowerShell Version 3.0 is required. The Windows PowerShell software needs to be installed and enabled on all Data Protection for Microsoft Exchange installations that you want to manage. For information about downloading, installing, and enabling Windows PowerShell, see the following Microsoft documentation. At the time of publication, the following web site is available: http://www.microsoft.com/en-us/download/details.aspx?id=34595 .

About this task

For reference, the following commands are provided:

- In most environments, to enable remote management for Data Protection for Microsoft Exchange installations, use the following command:

  ```
  Enable-PSRemoting -force
  ```

- If you are using Microsoft Exchange, complete the following steps:

  1. On the primary system, enter the following command:

```
enable-wsmancredssp -role client -delegatecomputer remote_computername
```

2. On each remote system that runs Microsoft Exchange, enter the following command:

```
enable-wsmancredssp -role server
```

- Add the Data Protection for Microsoft Exchange servers to the trusted hosts list by entering the following command:

```
Set-Item WSMan:\localhost\Client\TrustedHosts -Value remote_server_name -Force
```

- After making configuration changes, restart the winrm service by entering the following commands:

```
Restart-Service winrm
```

To remotely manage Data Protection for Microsoft Exchange installations, complete the following steps:

Procedure

1. From the Management Console, in the Actions pane, click **Manage Computers**.
2. Navigate the interface to add, remove, or configure remote systems. For systems that are not in the domain, provide the fully-qualified address.

Determining managed storage capacity

You can track the capacity of managed storage assets. This feature can be useful for storage planning during license renewal.

About this task

Typically there is a difference between the capacity that is used by server data and the capacity of the volume that contains that data. For example, a set of databases might require a capacity of 1 GB and are on a 10 GB volume. When a snapshot of the volume is performed, the Data Protection for Microsoft Exchange managed capacity measurement is 10 GB.

Procedure

To determine managed storage capacity:

1. Select an Exchange instance from the Management Console.
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. View a list of the volumes that contain backups and their respective managed capacities by clicking **Show Details**.
5. Close this dialog window.

Backing up Exchange data

You can back up Exchange Server data by using Microsoft Volume Shadow Copy Service (VSS).

Before you begin

To create VSS backups, you must have a VSS provider that is configured for your environment.

If you are backing 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). (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 completed the backup.) You can set up the DAG node name in the **DAG Node** field in the TSM Node Names page of the Tivoli Storage Manager configuration wizard, or in the **Back up DAG databases to common node** field in the General properties page for your Exchange workload.

Procedure

To back up Exchange Server data, complete the following steps:

1. Start the Management Console and click **Exchange Server** in the tree view.
2. On the **Protect** tab, select one or more databases to back up. Alternatively, click the **Protect Data** shortcut in the start page of the Management Console.

   Filter the list of available databases in the results pane by entering a keyword in the **Search** field.

   If you are running 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 in the **Protect** tab.
3. Verify the backup options. If the backup options are not displayed, click **Show Backup Options**.
   - Set the **Offload** option to **True** to use offloaded backups. An offloaded backup uses another system (specified with the **RemoteDSMAGENTNode** parameter) to move Exchange data to Tivoli Storage Manager server storage and runs the Exchange Integrity Check. An offloaded backup can reduce the load on network, I/O, and processor resources during backup processing.

     If you are going to use offloaded backups, make sure a **RemoteDSMAGENTNode** is specified. This option applies only to VSS backups.
   - Click **Skip Exchange Integrity Check** if you do not want to run the Exchange Integrity check to verify that the databases and log files to be backed up do not contain integrity issues. According to Microsoft advised practices, you can skip the integrity check if the database you are backing up has two or more healthy copies within a DAG environment.

     Attention: If you do not verify that the backups are valid by using the Exchange integrity check, and there is an integrity error when you restore the database, you must run repairs on the database, which can result in data loss. If you skip the integrity check and the database is not recoverable because of integrity errors, contact Microsoft support for help in recovering your data.
   - If you are scheduling the backup of databases in an Exchange Server DAG, use **Minimum Backup Interval** to set 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 backup for the database is determined from the Exchange Server and not the Tivoli Storage Manager server.

This option specifies that only one database copy can be backed up within a timeframe. This option prevents all of the members in a DAG from backing up the database, which would be redundant and invalidate the Tivoli Storage Manager storage management policy.

This setting is intended to be used with tasks that are scheduled to be run with the **Run Scheduled** button, or in a script to be run with the Tivoli Storage Manager Scheduler.

- If you are scheduling the backup of databases in an Exchange Server DAG, set **PreferDAGPassive** to **True** to skip the backup for an active database copy unless no healthy passive copy is available. If no healthy passive copy is available, the backup is made from the healthy active database copy. There is no default value for **PreferDAGPassive** because **PreferDAGPassive** is a switch option.

  This setting is intended to be used with tasks that are scheduled to be run with the **Run Scheduled** button, or in a script to be run with the Tivoli Storage Manager Scheduler.

4. In the Action pane, click **Backup Destination** to specify whether you want the data to be backed up to your local server, a Tivoli Storage Manager 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 starts, complete with the appropriate command that is required to complete the task.
6. Create the backup by selecting the backup action from the **Action** pane. You can run a full, copy, incremental, or differential backup with the VSS backup method.

Restore options

Descriptions of the options available in the Management Console Restore tab are provided.

From the Recover tab, select **Database Restore** and click the **Show Restore Options** to modify the default restore options.

AutoSelect

Set this option to `True` (default) to quickly select the backup objects to restore. With auto-selection, when you select the most recent backup to restore, all other necessary backups are automatically selected for you, up to the previous full backup. **AutoSelect** provides these characteristics:

- Operates when you click a full, differential, or incremental backup.
- Ignores copy backups.
- When you click a full backup, the latest associated differential or all associated incremental backups are selected.

- 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. This statement is not applicable when you back up to Tivoli Storage Manager.

AutoSelect does not make more selections when a differential or incremental backup is selected and no associated full backup can be found.

FromServer
Enter the name of the server where the original backup was done. The default value is a wildcard character (*).

Instant Restore
Set this option to **True** to use volume level snapshot restore (instant restore) for local VSS backups if the backup exists on SAN-attached volumes. Set this option to **False** 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 **True**, which uses **volume level snapshot restore** if it is supported.

This option is available for VSS operations only. When you use instant restore for SAN Volume Controller earlier than version 5.1 or DS8000, a best practice is to make sure that any previous background copies (that involve the volumes that are being restored) are completed before you initiate the instant restore.

This option is automatically set to **False** during **restoreinto** operations.

Instant restore overwrites all files on the destination file system. Also, instant restore of incremental and differential backups are automatically converted to file-level restores. Instant restore requires that the drive or volume where the mailbox database is located must be available. There must be no access to the drive or volume by any other process or application.

Mount Databases After Restore
Select the **MountDatabasesAfterRestore** option to automatically mount databases after the recovery completes.

Replay Restored and Current Logs
Use the **ReplayRestoredANDCurrentLogs** option to replay any transaction log entries that display in the current active-transaction log. This log includes both current and restored logs. This option is the default value. This option is not supported for instant restore.

Replay Restored Logs Only
Use the **ReplayRestoredLogsONLY** option to replay any transactions that display in the restored-transaction logs. After you run this type of restore, do a new full backup.

RunRecovery
Select this option to specify whether to replay just the restored logs or to replay both the restored and current logs. When recovery is not run, the databases are not online.

VSS restore considerations

Review the following list before completing a VSS restore.

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

- Install any Microsoft VSS-related urgent fixes.

VSS instant restore considerations:

- In a DAG environment, stop the Microsoft Exchange Replication Service on the active node before you run the VSS instant restore operation.
- In an Exchange 2013 environment, stop the Exchange Search Host Controller Service on the active node before you run the VSS instant restore operation.
- Performing any type of **Restore Into** function automatically disables VSS instant restore.
- When you perform VSS instant restores, you must make sure that any previous background copies (that involve the volumes that are being restored) are completed before you initiate the VSS instant restore. This situation applies to DS8000, Storwize V7000, XIV, and SAN Volume Controller (non-space-efficient) volumes only.
- A VSS instant restore operation overwrites the entire contents of the source volumes. However, you can avoid overwriting the source volumes by setting the **Instant Restore** option to **False**. This option bypasses volume-level copy and uses file-level copy instead to restore the files from a VSS backup that are on local shadow volumes.
- Any backup to **LOCAL** can be restored only to the same system.

If you are performing a VSS restore of a database that was relocated (system file path, log file path, or database file path), you must use the **Instant Restore** function and specify the same database name as the one you are restoring. The restore fails if you do not specify the same database name.

When a VSS restore from local shadow volumes is performed, the bytes transferred displays *0*. This display occurs because no data (*0*) is restored from the Tivoli Storage Manager server.

Restoring VSS backups into alternate locations

An Exchange Server database backup, or DAG active or passive database copy backup can be restored into a recovery database or into an alternate (or relocated) database.

This restore capability is referred to as a *restore into* scenario and requires the following actions:

- If you are operating a VSS restore of a relocated database, you must use the *restore into* function. Also, specify the same database name as the one you are restoring. The restore fails if you do not specify the same name.
- Running any type of restore into function automatically disables VSS Instant Restore.

Backups to LOCAL can be restored only to the system where the backups were created.

Preparing for VSS instant restore in DS8000, Storwize V7000, XIV, and SAN Volume Controller environments

When you prepare for a restore, consider the type of data that you want to restore and where the backups are located.

About this task

The information in the following list is specific to the DS8000, Storwize V7000, XIV, and SAN Volume Controller environments. When planning for a VSS instant restore, consider the following facts:

- Restore granularity is at the volume level.
- VSS requires that data must always be restored to the same drive letters and paths as existed during the original backup.
- VSS requires IBM System Storage Support for Microsoft Volume Shadow Copy Service software if you use a DS8000, Storwize V7000, or SAN Volume Controller disk subsystem.
- VSS requires IBM XIV Provider for Microsoft Windows Volume Shadow Copy Service if you are using an XIV disk subsystem.
- Backups must be located on the same XIV, DS8000, Storwize V7000, or SAN Volume Controller storage subsystem to which they are restored.
- In a DAG environment, stop the Microsoft Exchange Replication Service on the active node before you run the VSS instant restore operation.
- In an Exchange 2013 environment, stop the Exchange Search Host Controller Service on the active node before you run the VSS instant restore operation.

Complete restore or replacement

Information regarding a complete restore or replacement is provided.

For information about how to recover an Exchange Server 2010, see the article "Understanding Backup, Restore, and Disaster Recovery" at the following URL:
`http://support.microsoft.com/default.aspx?scid=kb;en-us;326052`

Individual mailbox recovery

Data Protection for Microsoft Exchange backs up at the database level, but can restore individual items from the database backup.

Backing up Exchange servers at the 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.

With the Data Protection for Microsoft Exchange mailbox restore feature, you can run individual mailbox and item-level recovery operations in Microsoft Exchange Server 2007 or Microsoft Exchange Server 2010 environments that utilize Data Protection for Microsoft Exchange backups. See "Restoring individual mailbox and mailbox item-level data" and "Restoring mailbox messages interactively with the Mailbox Restore Browser" on page 81.

Restoring individual mailbox and mailbox item-level data

Use Data Protection for Exchange to restore mailboxes and mailbox data.

Procedure

Follow these instructions to restore an Exchange Server mailbox or mailbox items.

1. Start the Management Console and select **Exchange Server** in the tree view.
2. In the **Recover** tab for the Exchange instance, and change the selected view to **Mailbox Restore**.
3. Select one or more mailboxes to restore. A list of mailboxes that are backed up is displayed. If you are restoring a mailbox that was deleted or re-created since the time of the backup, enter a mailbox with sufficient space to temporarily store the messages during the restore. This mailbox is set by using the `Alias of temporary mailobx` option from the Properties page, under the General tab.
4. Optional: By default, Data Protection for Exchange restores the most current backup available for the specified mailbox. If you want to restore data to a different point in time, use the **Backup Date** option to select an earlier date and time. When you specify a backup date, Data Protection for Exchange looks for a backup with that exact date. If a backup with that exact date is not found, Data Protection for Exchange looks for and selects the first backup **after** that date. For example, if you have a backup at *9:51* and a backup at *10:09*, and you specify *10:00*, Data Protection for Exchange selects the backup at *10:09*. This backup is selected so the software does not miss any messages. If the backup at *9:51* was selected, the software would miss messages from *9:51* to *10:00*.

   By default, the entire mailbox is restored. Use the **Item-Level Mailbox Filters** to identify individual messages to restore.

   a. Click **Show Filter Options** and **Add Row**.
   b. Click the down arrow in the **Column Name** field and select an item to filter. You can filter by **Backup Date**, **Folder Name**, **Subject Text**, **Sender Name**, **Message Body Text**, **All Content**, **Attachment Name**, and **Received Date**. When restoring to a Unicode `.pst` file, except for the **Folder Name** and **All Content** filters, the filters are ignored.

      When you click **All Content**, the mailbox items are filtered by attachment name, sender, subject, and message body.

      To filter by **Backup Date**, click the default date and time to edit the table cell. To change the date, click the drop-down icon that is displayed 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 using the 12-hour clock time convention.
   c. Select an operator in the **Operator** field.
   d. Specify a value to filter on in the **Value** field.

e. In you want to filter on more items, click **Add Row**.

5. Verify restore options. If the restore options are not currently displayed, click **Show Restore Options**.

   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.

   Mark restored messages as unread
   Use this option to automatically mark the mailbox messages as unread after restore operation completes. The default value is **True**.

   Mailbox Original Location
   Use this option only if the mailbox is deleted or recreated since the time of the selected backup, and mailbox history is disabled. Specify the Exchange Server and the database where the mailbox is at the time of the backup. Use the following format: `server-name,db-name`

6. Click one of the **Restore** actions in the **Action** pane to begin the restore operation.

   Restore Mail to Original Location
   Select this action to restore the mail back to where the mail items existed at the time of backup.

   Restore Mail to Alternate Location
   Select this action to restore the mail items to a different mailbox. A dialog is displayed for you to specify the mailbox

   Restore Mail to non-Unicode PST file
   Select this action to restore the mail items to a non-Unicode personal folders (`.pst`) file. When you restore to a `.pst` file with one selected mailbox, you are prompted for a file name. When you restore 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 bears the name of the mailbox at the specified directory.

   If the `.pst` file exists, the file is used. If it does not exist, the file is created.

   Restore Mail to Unicode PST file
   Select this action to restore the mail items to a Unicode personal folders (`.pst`) file. When you restore to a `.pst` file with one selected mailbox, you are prompted for a file name. When you restore to a `.pst` file with more than one selected mailbox, you are prompted for a directory location.

   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.

   Each mailbox is restored to a separate `.pst` file that bears the name of the mailbox at the specified directory. If the `.pst` file exists, the file is used. If it does not exist, the file is created.

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.

Restoring a deleted mailbox or items from a deleted mailbox

You can use the Data Protection for Microsoft Exchange mailbox restore operation to restore a mailbox (or items from a mailbox) that was deleted from an Exchange Server.

About this task

To restore a deleted mailbox or items from a deleted mailbox before you run the mailbox restore operation, run a Data Protection for Microsoft Exchange mailbox restore to restore the deleted mailbox.

Data Protection for Microsoft Exchange requires a temporary mailbox to run mailbox restore operations on mailboxes that were deleted or recreated since the time of the backup you are restoring from. Use the **/TEMPMAILBOXAlias** parameter to specify the temporary mailbox. If the **/TEMPMAILBOXAlias** parameter is not set, the default is the logon user mailbox. Ensure that the temporary mailbox is active and has enough storage capacity to accommodate all items of the mailboxes that are being restored. See "Restoremailbox optional parameters" on page 161 for details about the **/TEMPMAILBOXAlias** parameter.

With the mailbox restore operation there are three options for choosing where to direct the restore of mailbox data from a deleted mailbox:

- Restore the deleted mailbox data to the original location of the original mailbox.
- Restore the deleted mailbox data into an active alternative mailbox in an online Exchange Server.
- Restore the deleted mailbox data into an Exchange Server personal folders (`.pst`) file.

If you are restoring the deleted mailbox data to the original location, before you run the mailbox restore, recreate the mailbox using Exchange.

If the backup that contains the deleted mailbox was taken with a version of Data Protection for Microsoft Exchange before version 6.1, or if the mailbox history is disabled, and the mailbox was relocated since the time it was backed up, you must specify the Exchange Server and the database where the mailbox was at the time of backup. Use the **Mailbox Original Location** option in the GUI to specify this information. You can also use the **restoremailbox** command parameter, **/MAILBOXORIGLOCATION**, for this task.

See "Restoremailbox command" on page 158 for details about this command.

The Exchange Server and the database where the mailbox is located are to be specified.

See "Restoremailbox optional parameters" on page 161 for details about the **/MAILBOXORIGLOCATION** parameter.

Restoring mailbox messages interactively with the Mailbox Restore Browser

You can interactively restore a mailbox or items from a mailbox on Exchange Server by using the Mailbox Restore Browser.

About this task

In addition, review the mailbox restore characteristics before you attempt a restore operation:

- When the Management Console is started, it detects whether there exists a recovery database that was previously created by Data Protection for Exchange. If one exists, then the Management Console automatically connects to the existing recovery database and displays its contents. Otherwise, you are prompted for the mailbox or database to restore into the recovery database.
- The Management Console also detects a recovery database that was created outside of Data Protection for Exchange and 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** action item.
- When a mailbox is selected, it is first restored to the recovery database. It is from this location that the mailbox becomes available for browsing. When the restore operation to this location completes, the restored mailbox and folders are shown in the results pane.
- If you plan to restore mail or folders by using a Simple Mail Transfer Protocol (SMTP) Server, make sure to configure the SMTP Server before you attempt a restore operation. Set the configuration in the Management Console by right-clicking **Dashboard** in the tree view and selecting **Properties**. Then, go to the E-mail property page. Enter the SMTP server and port in this property page.
- If you select a mailbox to restore, you can click **Restore Mailbox to Original Mailbox**. If you select a folder, you can click **Restore Folder to Original Mailbox**, or **Restore Folder to SMTP Server**. If messages are selected, you can click **Restore Messages to the Original Mailbox**, **Restore Messages to SMTP Server** or **Save Mail Message Content**.
- Data Protection for Exchange restores the mailbox backup to its original mailbox location. However, you can also restore a mailbox to either of the following locations:
  - To restore a mailbox item to a different mailbox, use the **Open Exchange Mailbox** task in the Action pane. Enter the alias of the mailbox to identify it as the restore destination. This mailbox restore destination is shown in the lower results pane. Drag the source mailbox from the upper results pane to the destination mailbox in the lower results pane.
  - To restore a mailbox to an Outlook personal folders (`.pst`) file, use the **Open PST File** task in the Action pane. A Windows File dialog opens so that you can select an existing `.pst` file or create a `.pst` file. This specified destination `.pst` file is shown in the lower results pane. Drag the source mailbox from the upper results pane to the destination `.pst` file in the lower results pane.

    The mailbox restore browser only supports non-Unicode `.pst` files.

  In either case, a merge operation is done during the restore. If the object exists, Data Protection for Exchange does not create a duplicate. Data Protection for Exchange restores only items that do not exist in the restore destination.

  When a mailbox is restored to its original mailbox location, the items are merged. When a mailbox is restored to a different mailbox or to a `.pst` file, the items are restored to a folder that bears the original mailbox name.

- The **Close Exchange Mailbox** and **Close PST File** tasks in the Action pane are only shown when a destination mailbox or `.pst` file is opened.

Restriction: Only mailboxes within the same database can be restored in a single mailbox restore action.

Procedure

To restore mailbox messages with the mailbox restore browser, complete the following steps:

1. Start the Management Console.
2. Under the **Protect and Recover Data** node in the tree, select **Exchange Server**.
3. In the Recover panel, click **View** > **Mailbox Restore Browser**. The Select Source Mailbox to Restore dialog opens.
4. Specify the mailbox to restore in the Select Source dialog:
   a. To browse mailboxes, select **Browse Mailboxes**. You can also switch to the databases view by selecting **Browse Databases** in the drop-down list.

      Enter the name of the mailbox in the field in the **Mailbox Name** field, or scroll down through the list and select a mailbox. The list is populated by using mailbox history that is taken at the time of the backup. If mailbox history is disabled, you can type a mailbox name. Otherwise, use the **Search** field to filter the mailboxes. You can also sort the mailboxes by columns. Click **OK**.

      You can also specify a date and time in the **Backup Date/Time** field when you want to restore a backup that was created at a specific point in time. To filter by date and time, click the default date and time to edit the table cell. To change the date, click the drop-down icon that is displayed 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 using the 12-hour clock time convention.

   b. To browse all mailboxes in a particular backup, specify **Browse Databases**. A list of available backups is displayed. Scroll down the list and select a database. Use the **Search** field to filter the databases. You can also sort the databases by columns. Click **OK**.

   c. To restore a mailbox that was deleted or recreated after the time of the backup, go to the **Properties** > **General** tab. Enter the temporary mailbox alias. If this alias is not entered, the mailbox restore operation uses the current administrator user mailbox as a temporary storage location.

   After the specified mailbox is restored to the recovery database, the restored mailbox and folders are shown in the results pane.
5. Use the results pane to browse the folders and messages that are contained within your mailbox. The following features are available:
   - Preview: When a mailbox item is selected, its content is shown in the preview panel. When an item contains an attachment, click the attachment icon to preview its contents (click **Open**) or save it (click **Save**).
   - Filter: Use the filter options to narrow the list of folders and messages in the result pane.
     a. Click **Show Filter Options** and **Add Row**.
     b. Click the down arrow in the **Column Name** field and select an item to filter. You can filter by Folder Name, Subject Text, Sender Name, Message Body Text, All Content, Attachment Name, Size (in KB), Created Date, Modified Date, Sent Date, and Received Date.

When you select **All Content**, the mailbox items are filtered by attachment name, sender, subject, and message body.

c. Select an operator in the **Operator** field.

d. Specify a value to filter on in the **Value** field.

e. In you want to filter on more items, click **Add Row**.

f. Click **Apply Filter** to filter your messages and folders.

Select the mailbox, folder, or message to restore before proceeding.

6. Click the restore task in the Action pane. Depending on the item that you selected, the following restore actions are available:
   - **Restore Folder to Original Mailbox**
   - **Restore Messages to Original Mailbox**
   - **Restore Folder to SMTP Server**
   - **Restore Mail to SMTP Server**

     If the SMTP Server was not configured, you must configure it before you run the restore action. Right-click the dashboard and select **Properties**, then click **E-mail** to complete the configuration.
   - **Save Mail Message Content**: A Windows Save File dialog is displayed. Specify the location and message name and click **Save**. The Save Mail Message Content action becomes available when a message is selected in the preview pane.

     When you restore an email with an attachment larger than 3 MB, a Microsoft fix is required. The fix resolves the following issue: *QFD: System.Net.Mail - SmtpClient class throws exceptions if file attachment is over 3MB*. The fix is available online at https://connect.microsoft.com/VisualStudio/Downloads/DownloadDetails.aspx?DownloadID=30226.

   The Restore Progress dialog opens and shows operation details.

   The **Close Mailbox to Restore** button is displayed after a recovery database is created. When you click this button, Data Protection for Exchange removes the recovery database that was created 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 the Management Console.

Restoring mailboxes directly from Exchange database files

When 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.

About this task

If you use Data Protection for Microsoft Exchange Server to back up the Exchange server, the database files can be restore to a local disk with the following command:

```
tdpexcc RESTOREFILES
```

If you are using Tivoli Storage Manager for Virtual Environments software, review the following statements before restoring the mailbox:

- You can use Tivoli Storage Manager for Virtual Environments to back up an Exchange server in a virtual machine. For more information see the *Tivoli Storage Manager for Virtual Environments User's Guide*.
- To restore mailboxes from the backups created by Tivoli Storage Manager for Virtual Environments, mount the virtual volumes that contain the EDB file and

log files with read-write access. Read-write access can be gained by clearing the **Mount virtual volume as read only** checkbox.

- If the log files reside on a different volume than the EDB file, mount the volume containing the log files on an unused drive letter. With this action you can apply the transaction logs to the EDB file.

To complete the restore from the graphical-user interface, complete the following steps:

Procedure

1. From the Exchange server, launch Data Protection for Microsoft Exchange Server.
2. After you log on to Data Protection for Microsoft Exchange Server, in the navigation area, select the Exchange Server node and Recover tab. The Mailbox Restore Browser view opens.
3. From the Actions pane, click **Open EDB File on Disk**.
4. In the dialog, enter or browse to the location of the backup database (EDB) file.
5. In the dialog, enter or browse to the location of the log file directory.
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 is complete, click **Close Mailbox to Restore**. During the closure process, a message is displayed. The message asks you to decide if you want to delete the recovery database folder.

Restore by using the Recovery Database

In some rare cases, you might want to manually restore mailboxes using the recovery database, instead of using mailbox restore or the mailbox restore browser.

To manually restore mailboxes using the recovery database, use VSS restore.

Requirements for using the recovery database

These requirements must be met for this procedure to be successful.

- For VSS restores, the mailbox database to be restored must be in the same Administrative Group. The backup must be taken on the same version of Exchange Server for the restore to complete.
- You must run the restore from an account that has Receive As and Send As permissions on all mailboxes to be restored.
- You cannot use multiple instances of Data Protection for Microsoft Exchange to restore databases into the recovery database simultaneously.

Restoring data to a recovery database

Restore your data to a recovery database. You must already back up your database before you attempt this task.

About this task

Information regarding recovery database processing is written to the Data Protection for Microsoft Exchange activity log file (`tdpexc.log` by default).

Note: When you restore to a recovery database, you must specify the option to replay restored logs only, otherwise the restore can fail. Select **Replay Restored Logs ONLY** in the GUI Restore tab or specify **/recover**=*applyrestoredlogs* on the command line.

Procedure

To restore data to a recovery database, complete the following steps:

1. Use the Exchange Management Console to 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 Microsoft Exchange to restore the mailbox database.

   For VSS Restores, select the name of a database into which a VSS backup is restored (use the INTO command to select the database to restore into). To restore into a recovery database, a recovery database must exist.

   Tip: From the Management Console, you can 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.

   Only transaction logs that are contained in the backup are applied to the mailbox database when a recovery database restore is processing.

Restoring a Database Availability Group database copy

Perform these steps to restore a replicated database copy in a Database Availability Group (DAG). This procedure assumes that you already backed up your database.

About this task

You can do some of these steps by using either the Exchange Management Console or the Exchange Management Shell commands, which are provided in parentheses.

Procedure

To restore a Database Availability Group database copy:

1. Make the database that you want to restore active (**Move-ActiveMailboxDatabase**).
2. Suspend replication of the all passive copies of the database (**Suspend-MailboxCopy**).
3. Unmount the active mailbox database (**Dismount-Database**).
4. Stop the replication service on all copies of the database. Do this step only for a VSS Instant Restore operation.
5. Restore the database and logs by using the Data Protection for Microsoft Exchange command line or GUI.

Restriction: The database must not be mounted automatically after the restore. If you use the GUI, ensure that the **MountDatabasesAfterRestore** option is set to **False** in the Restore panel, you must clear it. If you use the command line, the **/mountdatabases** restore option must be set to NO.

6. If the service was stopped, start the replication service first 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**). This step avoids potential transaction log synchronization problems that might arise if replication were resumed directly.
9. Move the active database to the server that you want. (**Move-ActiveMailboxDatabase**)

Mounting backups

From the Recover tab, you can mount a backup.

About this task

To mount backups, complete the following steps:

Procedure

1. Start the Management Console.
2. Click **Recover Data** in the welcome page of the Management Console.
3. In the Recover tab, go to the Action pane. 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**.

Results

In the Recover tab, the backup that you mounted is displayed. You can use the **Explore** and **Unmount Backup** options in the Action pane to complete tasks with the backup that you mounted.

Deleting Exchange Server Backups

Use this procedure to remove an Exchange Server backup object that was created with the VSS backup method.

Before you begin

Attention: Do not use this procedure for typical delete tasks as backups are deleted automatically, based on user-defined policy management settings. This procedure is necessary for those deletions that are outside the scope of standard policy management deletions. Perform this task with caution and only as a last resort.

For backups of Exchange Server Database Availability Group (DAG) databases to Tivoli Storage Manager, backups of a database from a DAG member to LOCAL can be deleted only from the Exchange Server on which the backup was created.

Procedure

To delete Exchange Server backups:

1. Start the Management Console.
2. Click **Recover Data** in the welcome page of the Management Console.
3. In the **Recover** tab for the Exchange instance, select **View: Database Restore**. Use the results pane to browse and select one or more database backups to delete.
4. Click **Delete Backup** in the **Action** pane to delete the backups of the selected databases.

   Note: When a delete backup is in progress, two tasks display in the task window to show the deletion is in progress, and that the view is being refreshed. The view content is updated when both tasks are finished.

What to do next

For special considerations about multiple backups on space-efficient target volumes with SAN Volume Controller and Storwize V7000, see "More guidelines for SAN Volume Controller and Storwize V7000 environments" on page 23.

Viewing, printing, and saving reports

Access reports on recent activity, historical managed capacity, and which licenses and software are installed.

About this task

Follow these steps to view, save, or print reports.

Procedure

1. Select **Reporting** in the tree view's **Manage** section. A list of available reports is displayed. Each report has a description of what data the report contains.
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 6. Automating

The term *automation* applies to Data Protection for Exchange. You can run commands from the command line, create scripts, schedule tasks, and use the graphical user interface to start tasks. The tasks are based on scripts and schedules that you create.

The software supports running tasks from both the command-line interface or Microsoft Windows PowerShell command prompt (Version 3.0 and later). You can also use the **Automate** tab in the Management Console.

Automating tasks

You can use the Automate view to work with commands. You can save a command and run the command at a scheduled time.

About this task

You can use the Automate view to create, save, store, and schedule commands. Open the Automate view by selecting a workload that you want to work with and clicking **Automate**. 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.

Procedure

1. Change **PowerShell** to **Command Line**.
2. Type a command in the details pane and click the **Execute** icon to run the command. You can also run a saved task by clicking the **Open** icon, selecting the command file, and clicking the **Execute** icon.

   The commands can be entered 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
   ```
3. Click the **Save** icon and follow the prompts to save a command for future use.
4. 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.
5. The output of the command is displayed in the results pane. The output can be saved or sent to an email address.

What to do next

You can automate commands from the Protect, Recover, Schedule, and Task List views in the Management Console:

1. Start the Management Console and select the **Exchange Server** instance in the tree view.
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 schedule wizard starts, enter the information for each prompt to create a scheduled task.

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.

Scheduling

Review these characteristics when you are defining a Tivoli Storage Manager schedule.

- If you use the Tivoli Storage Manager server-prompted scheduling mode, ensure that the Data Protection for Exchange option file has the `tcpclientaddress` and `tcpclientport` options specified. 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). An example of running more than one scheduler service is when you schedule Data Protection for Exchange and the regular Windows backup client.

  Server-prompted scheduling is supported only when TCP/IP communication is being used. By default, Data Protection for Exchange uses the client polling schedule mode.

- If any changes that affect the scheduler are made to the Data Protection for Exchange options file, restart the scheduler to activate the changes. Examples of what can affect the scheduler are the Tivoli Storage Manager server address, the schedule mode, or the client TCP address or port. To restart the scheduler, issue the following commands:

  ```
  net stop "Data Protection for Exchange Scheduler"
  net start "Data Protection for Exchange Scheduler"
  ```

- The default Tivoli Storage Manager scheduler log file (`dsmsched.log`) contains status information for the Tivoli Storage Manager scheduler. In this example, the file is in this path:

  ```
  d:\Program Files\Tivoli\TSM\TDPExchange\dsmsched.log
  ```

  You can override this file name by specifying the `schedlogname` option in the Data Protection for Exchange options file.

- Data Protection for Exchange 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 is an indication that Tivoli Storage Manager successfully issued the scheduled command that is associated with the *eventname*. No attempt is made to determine the success or failure of the command. Assess the success or failure

of the command by evaluating the return code from the scheduled command in the scheduler log. The scheduler log entry for the command's 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 Tivoli Storage Manager password is to be specified on the **tdpexcc** command. To specify the password, use the **/tsmpassword** parameter in the command file that is being run by the scheduler (`excfull.cmd`). You can also specify the password on the Data Protection for Exchange command line. For example:

  ```
  tdpexcc query tsm /tsmnode=mars1 /tsmpassword=newpassword
  ```

Windows PowerShell and Data Protection for Exchange

Data Protection for Exchange includes a set of Windows PowerShell cmdlets to help manage Data Protection for Exchange components in your environment. Because cmdlets can be chained together to form commands and because there is a large body of existing cmdlets from other vendors the Data Protection for Exchange cmdlets help support a seamless management environment. Remote management and automation capabilities are greatly improved when using the Data Protection for Exchange cmdlets.

Getting started

The cmdlets can be used in supported Windows environments.

About this task

Before you use the cmdlets provided with Data Protection for Exchange, complete the following steps:

Procedure

1. Log on to the system as an administrator.
2. From a Windows PowerShell command prompt, enter the following command:

   ```
   set-executionpolicy remotesigned
   ```

3. Import the Windows PowerShell modules from the `TDPExchange` folder:
   - FmModuleExc.dll
   - FmModuleMMC.dll

   To import modules, with the administrator credentials, from a Windows PowerShell command prompt, complete the following steps:

   a. Navigate to the `TDPExchange` folder.

   b. Enter the following commands:

   ```
   import-module .\FmModuleExc.dll
   import-module .\FmModuleMMC.dll
   ```

   c. (Optional) To use the cmdlets in these modules any time you start Windows PowerShell, add the following lines to your profile:

   ```
   $path = (get-itemproperty -path "HKLM:\SOFTWARE\IBM\TDPExchange\
    currentversion\mmc" -ea SilentlyContinue).path
   if ($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 documentation. Information about Windows PowerShell cmdlets consistent naming patterns, parameters, arguments, and syntax is also provided in the Windows PowerShell documentation. The following web site is a starting point for this type of documentation: http://technet.microsoft.com/en-us/library/hh857337.aspx.

Cmdlets for protecting Microsoft Exchange server data

The following table identifies the cmdlets that are available for use when protecting Microsoft Exchange server data.

Table 6. 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.
Backup-DpExcComponent	**tdpexcc backup**	Back up a Microsoft Exchange database.
Copy-DpExcPolicy	**tdpexcc copy policy**	Copy an existing policy.
Dismount-DpExcBackup	**tdpexcc unmount backup**	Dismounts 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**	Displays configuration information.
Get-DpExcConnection	**tdpexcc query tsm**	Query a list of the current values set in the configuration file for Tivoli Storage Manager.
Get-DpExcInformation	**tdpexcc query exchange**	Query general local Exchange Server information.
Get-DpExcMailboxLocationHistory		Query the mailbox location history.
Get-DpExcManagedCapacity	**tdpexcc query managedcapacity**	Query managed capacity for Microsoft Exchange Server.
Get-DpExcPolicy	**tdpexcc query policy**	Displays policy information.
Mount-DpExcBackup	**tdpexcc mount backup**	Mounts a backup to provide access to the files that the backup contains.
Remove-DpExcBackup	**tdpexcc delete backup**	Removes the backup.
Remove-DpExcPolicy	**tdpexcc delete policy**	Deletes the policy.
Reset-DpExcTsmPassword	**tdpexcc changetsmpassword**	Change the Tivoli Storage Manager password used by Data Protection for Exchange.
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.

Table 6. Cmdlets to protect Microsoft Exchange Server data (continued). 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
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
```

To go to the information center, 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.

Cmdlets for the Management Console

The following list identifies the cmdlets that are available for use when interacting with the Management Console.

- **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
```

To go to the information center, 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.

Chapter 7. Troubleshooting

Data Protection for Exchange provides support for protecting Microsoft Exchange databases.

If you encounter a problem during Data Protection for Exchange processing when using VSS for backup and restore, complete the following steps:

1. Retry the operation that failed.
2. Restart the Tivoli Storage Manager 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 (antivirus applications, for example). 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. There might be some log entries that help you identify a VSS event that triggers the issue.

If you do not find a resolution to the problem in the log files, you can complete the following procedure:

1. Shut down the Exchange server.
2. Restart the Exchange server.
3. Run the operation that failed.

Alternatively, you can also complete this procedure:

1. Shut down the entire computer.
2. Restart the computer.
3. Run the operation that failed.

Another troubleshooting procedure to consider: Determine if this problem is reproducible on other Exchange servers.

When troubleshooting a mailbox restore error, you can use the **TDPMAPI TESTMAPI** command. The command helps diagnose MAPI connection issues when connecting to the mailbox. The following parameters can be used:

/MAILBOXALIAS
This parameter is the alias name for the mailbox that was specified when the mailbox was originally created. The parameter refers to the email alias for the user and is the portion of the email address on the left side of the @ symbol. You should run this command against both the mailbox to be restored and the mailbox of the administrator you are currently logged in as.

/EXCSERVER
(Exchange 2010 environments) This parameter is the name of the Exchange Server that has the Client Access Server (CAS) role. The default is the local server. The **get-ExchangeServer | fl** Exchange PowerShell command can be used to determine which Exchange Server has the CAS role defined for the mailbox database. It is mandatory to specify this parameter when there is a CAS Load Balancer within the environment.

(Exchange 2013 environments) Depending on the information that you need, use one of the following formats:

Retrieve the ExchangeGUID for the user that is logged on

Enter the following command:

```
whoami | Get-Mailbox | fl ExchangeGUID
```

Retreive diagnostic information

Enter the following command:

```
tdpmapi.exe /excserver={ExchangeGUID@domain}
```

/TRACEFILE

This parameter is the filename used to hold the output from the tracing. By default, tracing is turned off. The filename can be qualified with a drive and full path location, and must have write permissions for the user running the command.

For example:

```
TDPMAPI TESTMAPI /MAILBOXALIAS=alias /EXCSERVER=cas /TRACEFILE=filename
```

Exchange 2013

The following list identifies troubleshooting information when protecting data for an Exchange 2013 server:

- For the target mailbox, grant `full access` permission. When using the Administrator mailbox, Exchange 2013 usually, by default, blocks full access permission for this administrator.
- Log on as an Exchange administrator with a mailbox on an Exchange 2013 database.
- Make sure both the administrator mailbox and the target mailbox are accessible in either Microsoft Outlook or Outlook Web Access.
- Use an Exchange 2013 CAS. Specifically, set the **CLIENTACCESSServer** parameter to an Exchange 2013 CAS.
- If you are using a load balancer, for troubleshooting, set the **CLIENTACCESSServer** parameter to an actual server instead of the load balancer.
- To open the administrator mailbox and the target mailbox, try using the Mailbox Restore Browser with Open Exchange Mailbox.
- Check the MAPI registry key. The key is located at `HKEY_CURRENT_USER\Software\Microsoft\Windows NT\Current Version\Windows Messaging Subsystem`. The name of the key is `RpcHttpProxyMap_TSM`. Depending on your environment, you might have to change the HTTP to HTTPS, or change the authentication method, or change the domain name to `*`. The MAPI download contains the Microsoft documentation for setting this registry key.

  If the default registry key is something like the following example:

  ```
  contoso.com=http://mail.contoso.com,ntlm,ntlm,false
  ```

  You might make one or more of the following updates:

  ```
  contoso.com=https://mail.contoso.com,ntlm,ntlm,false
  *=http://mail.contoso.com,ntlm,ntlm,false
  contoso.com=http://mail.contoso.com,negotiate,negotiate,false
  ```

  For all of the information about registry keys, see the Microsoft documentation.

In addition, when completing a backup and restore tasks for an Exchange 2013 server, if a Client Access Server (CAS) is used and later removed as the CAS, or, if the CAS is enabled for SSL authentication, an Enter Password window might be

displayed. This window prompts you to enter the domain, user name, and password. To verify and workaround this problem is occurring, complete the following steps:

1. From a command prompt, enter **regedit.exe**.
2. Use the Microsoft documentation to backup the registry.
3. Locate the following registry key: `RpcHttpProxyMap_TSM`. This key is located at `HKCU\Software\Microsoft\WindowsNT\CurrentVersion\Windows Messaging Subsystem`
4. To go from HTTPS to HTTP, change the `subdomain.domain.com=https://server.subdomain.domain.com,ntlm,ntlm,false` value to `subdomain.domain.com=http://server.subdomain.domain.com,ntlm,ntlm,false`.

Debugging installation problems with an installation-log file

If a problem occurs during the installation process, gather details about the installation process. This information can assist IBM Software Support when you evaluate your situation. You can create a detailed log file of the failed installation that can facilitate analysis of your situation.

The installation wizard collects log files for the installation process. To view the log files, on the navigation pane go to **Manage** > **Diagnostics** > **Trace And Log Files**. 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.

To quickly resolve problems, the following information is needed:

- 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)

You can create a detailed log file (`setup.log`) of the failed installation. To create this log file, enter the following command to run the setup program (`setup.exe`):

```
setup /v"l*v setup.log"
```

Troubleshooting VSS and SAN Volume Controller, Storwize V7000, or DS8000

The troubleshooting tips included here are designed to help you accelerate your problem determination task. Check these items first to disqualify some common configuration issues.

The following areas are where you can troubleshoot when you are having VSS and SAN Volume Controller, Storwize V7000, or DS8000 problems:

- CIMOM (Common Information Model Object Manager) connectivity issues.

  To verify connectivity to the CIMOM, complete the following steps:

1. Refer to your SAN Volume Controller, Storwize V7000, or DS8000 documentation.
2. Run the **IBMVCFG LIST** command. The default location is `%Program Files%\IBM\Hardware Provider for VSS-VDS`.
3. Issue the **IBMVCFG SHOWCFG** command to view the provider configuration information.

- CIMOM operational issues.

  If your backup or restore fails, check the `IBMVSS.log` file. If the 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 131 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. To determine why it 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.

- Host configuration issues.

  If the failure seems to be for a different reason than a CIMOM failure, verify your configuration. Run the latest support levels of the software for SAN Volume Controller, Storwize V7000, or DS8000.

- Collecting logs in this environment.

  If you are unable to resolve these problems, provide the following information to IBM Support:

  – Information that is listed in the Tivoli Storage Manager 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

Determine the source of the problem

You can help determine whether a problem is a Data Protection for Microsoft Exchange issue or an Exchange server issue.

For VSS operations, recreate the problem with the Microsoft diskshadow tool. These applications can run backups that use the Microsoft Exchange VSS APIs. If the problem is recreated with the diskshadow tool, the problem probably exists within the VSS provider or the Exchange server.

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. The first step is to determine whether the problem is with the Microsoft VSS service or with the Tivoli Storage Manager.

About this task

To isolate the source of the error, complete the following steps:

Procedure

1. Test the connectivity between the Data Protection client and the Tivoli Storage Manager dsmagent. Select the Exchange workload that you want to work with and click the **Automate** tab to open the **Automate** view. Issue the **Query Exchange** command in the lower details pane and click **Execute** (or **Enter**). The results are displayed in the pane. As an alternative, issue the **TDPEXCC QUERY EXCHANGE** command on the computer where the Exchange server is installed to verify that your installation and configuration is correct. 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 Tivoli Storage Manager 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, use the vssadmin and diskshadow tools to re-create the VSS issue. On failure, use these programs to re-create the error to determine whether it is a general VSS problem or a problem within the Tivoli Storage Manager code.

   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. The following commands are examples of possible **VSSADMIN** commands:

```
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.

diskshadow

The diskshadow tool is available on Windows 2008 server and 2008 R2. Before you install Tivoli Storage Manager for Mail, test the core VSS function. The following diskshadow testing can be done before any Tivoli Storage Manager components are installed:

a. Test non-persistent shadow copy creation and deletion by running 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 the**DISKSHADOW** commands four times and verify that the Windows event log file contains no errors.

b. Test persistent shadow copy creation and deletion by running 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 the diskshadow commands four times and verify that the Windows event log file contains no errors.

c. Test persistent transportable shadow copy creation and deletion by running 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 the **diskshadow** commands four times and verify that the Windows event log file contains no errors.

- The following items can be determined by using the vssadmin or diskshadow tool:
  - Verify VSS provider configurations
  - Rule out any possible VSS problems before you run the Tivoli Storage Manager 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 a Tivoli Storage Manager problem if an operation works with diskshadow or vssadmin but not with the Tivoli Storage Manager
- Perform the following tests to ensure that VSS is working correctly:

  Test nonpersistent shadow copy creation and deletion

  a. Run "`DISKSHADOW k: l:`" where *k*: and *l*: are the Exchange Server database and log volumes.
  b. Repeat the previous step 4 times.
  c. Inspect the Windows Event Log to ensure that things look appropriate.

  Test persistent shadow copy creation and deletion

  a. 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.
  b. Repeat the previous step 4 times.
  c. Inspect the Windows Event Log to ensure that things look appropriate.

  Test nonpersistent transportable shadow copy creation and deletion (VSS Hardware Provider environments only)

  a. Run "`DISKSHADOW -p -t=export.xml k: l:`" where *k*: and *l*: are the Exchange Server database and log volumes.
  b. Copy the resultant "`export.xml`" file from computer 1 to computer 2 before you continue to the next step.
  c. On the computer you have set aside for offload, run "`DISKSHADOW -i=export.xml`"
  d. Inspect the Windows Event Log to ensure that things look appropriate.

  If any of these tests fail repeatedly, there is 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 3.

3. 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 Tivoli Storage Manager are simulated by running a diskshadow nonpersistent snapshot.
- Exchange VSS backups to Local and to the Tivoli Storage Manager are simulated by running a diskshadow persistent snapshot.
- Offloaded Exchange VSS backups to the Tivoli Storage Manager 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 a Tivoli Storage Manager or Data Protection for Exchange client problem.

What to do next

For more information, see the Verifying VSS functionality for the Data Protection Exchange backup Technote: Verifying VSS functionality for the Data Protection Exchange backup.

Diagnosing VSS issues

Test VSS snapshots on your system.

Before you begin

The wizard performs persistent and non-persistent snapshot testing.

Attention: 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.

Procedure

Follow these steps to test persistent and non-persistent VSS snapshots:

1. Start the Management Console.
2. Click **Diagnostics** in the results pane of the welcome page. Click the **VSS Diagnostics** icon in the action pane. The diagnostics wizard opens, a list of volumes are displayed, and the status of each test is displayed when it is completed.
3. Select the volumes or mount points to test and click **Next**. Click **Show VSS Information** to view details about the VSS providers, writers, and snapshots available on your system. The results of the persistent and non-persistent snapshot testing displays as Passed or Failed.
4. 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.
   - 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.

What to do next

Return to the Management window and begin backup operations.

For more information about troubleshooting of VSS operations, refer to the *Tivoli Storage Manager Problem Determination Guide*.

Viewing trace and log files

View files that are used during troubleshooting tasks.

Before you begin

You can collect trace and log files in the Diagnostics property page for a workload.

About this task

When you encounter a problem in the Management Console, you can create trace files by using the Diagnostics property page. Click **Properties** > **Diagnostics**, and click **Begin**. Then, close the property page and reproduce the problem. Finally, open the Diagnostics property page and click **Stop**. The log files are displayed in the Trace and Log Files view, and you can click a file to view it.

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.

Data Protection for Exchange uses several components. Each component is in its own directory along with its respective troubleshooting files. The Trace and Log Files view brings these files into a central location for easy viewing. Examples including default log and trace files are provided:

- Examples of Data Protection for Exchange default log and trace files:
  - Installation directory: `C:\Program Files\Tivoli\TSM\TDPExchange`
  - `dsierror.log`
  - `tdpexc.log`
  - *`TraceFileExc.trc`*

  If the `tdpexc.log` is defined in a path other than the default `C:\Program Files\Tivoli\TSM\TDPExchange\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`

Click the trace or log file you want to view. The contents of the file are displayed in the results pane. Use the toolbar icons to create, save, edit, or email a file.

Tracing the Data Protection client when using VSS technology

You must gather traces for Data Protection for Microsoft Exchange, the Tivoli Storage Manager application programming interface (API), and the DSMAGENT processes to ensure a good diagnosis of the Volume Shadow Copy Service (VSS) operation.

The following traces are the different traces to gather when you diagnose Data Protection for Exchange VSS operational problems:

Data Protection for Exchange trace

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
```

Enable tracing for FlashCopy Manager. See the *IBM Tivoli Storage FlashCopy Manager Installation and User's Guide* for information about how to enable tracing.

Tivoli Storage Manager API trace

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
```

DSMAGENT trace

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).

Exchange VSS Writer tracing

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:

1. Open the Exchange Management Console.
2. Find the server object.
3. 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.
4. Click the **Diagnostics Logging** tab.
5. Expand the **MSExchangeIS** node in the **Services** pane and click **System**.
6. Click **Exchange writer** in the **Categories** pane and select the logging level.
7. Click **Apply** and then **OK** to close the Properties dialog box.

Enable the Volume ShadowCopy service debug trace features in Windows

See the following websites for information about enabling debug tracing:

- http://support.microsoft.com/kb/887013
- 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 is dependent upon the operating system and the Exchange application. Collecting all the necessary information about the environment can significantly assist with determining the problem.

The Management Console (MMC) is able to collect information in a package file. The package file can be sent to IBM Software Support.

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 Tivoli Storage Manager API.
- The exact level of the Tivoli Storage Manager server.
- The exact level of the Tivoli Storage Manager backup-archive client.
- The exact level of the Tivoli Storage Manager storage agent (if LAN-free environment).
- The Tivoli Storage Manager server platform and operating system level.
- The output from the Tivoli Storage Manager 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 files from Exchange with VSS before calling IBM

Several log files and other data can be collected for Data Protection for Microsoft Exchange server diagnosis.

Gather as many of the following files as possible before you contact IBM Support. To collect information manually, refer to the following list. The MMC automatically collects this information.

- The contents of the `C:\adsm.sys\vss_staging` directory and subdirectories. Or gather the appropriate directories if you are using the `VSSALTSTAGINGDIR` option.
- The Data Protection for Microsoft Exchange configuration file. The default configuration file is `tdpexc.cfg`.
- The Data Protection for Microsoft Exchange Tivoli Storage Manager API options file. The default options file is `dsm.opt`.
- The Tivoli Storage Manager registry hive export.
- The Exchange Server registry hive export.
- The Tivoli Storage Manager Server activity log. The Data Protection client logs information to the server activity log. A Tivoli Storage Manager administrator can view this log for you if you do not have a Tivoli Storage Manager administrator user ID and password.
- If the Data Protection client is configured for LAN-free data movement, also collect the options file for the Tivoli Storage Manager storage agent. The default name for this file is `dsmsta.opt`.
- Any screen captures or command-line output of failures or problems.

Log files 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 following files are the Tivoli Storage Manager log files that you can gather:

- The Data Protection for Microsoft Exchange log file. The default location of this file is `C:\Program Files\Tivoli\TSM\TDPExchange\tdpexc.log`
- The Tivoli Storage Manager API Error log file. The default location of this file is `C:\Program Files\Tivoli\TSM\TDPExchange\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.

You can 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 and 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. The following are the different methods of specifying 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 Tivoli Storage Manager 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 /?
```

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 DS8000 - `%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

Emailing support files

Send diagnostic information to IBM support personnel.

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

Follow these steps to send diagnostic information to IBM support personnel:

1. Start the Management Console.
2. Click **Diagnostics** in the results pane of the welcome page. Click the **E-Mail Support files** icon in the action pane.
3. 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 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 them to another site by using another method such as FTP.

Online IBM support

There are multiple resources for support.

About this task

The following list identifies the various ways that you can find information online:

- Tivoli Storage Manager wiki at http://www.ibm.com/developerworks/mydeveloperworks/wikis/home/wiki/Tivoli Storage Manager.
- Service Management Connect site at https://www.ibm.com/developerworks/servicemanagement/sm/index.html.
- Data Protection for Microsoft Exchange Server product support at http://www.ibm.com/software/tivoli/products/storage-mgr-mail/. Enter the search term to narrow the search criteria for your support need. Examples of search terms are an authorized program analysis report (APAR) number, release level, or operating system

Viewing system information

View or edit scripts that provide information about system components. Examples of some system components are Data Protection for Microsoft Exchange related Windows Services, 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 customize scripts.

Procedure

To work with scripts, follow these steps:

1. Open the System Information view by doing the following steps:
   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 Tivoli Storage Manager scripts.

2. Add, update, or delete your scripts.
   - To add your own scripts, click **New** in the Actions pane. You can also copy your scripts directly to the `ProgramFiles\Tivoli\FlashCopyManager\Scripts` directory.

     Tivoli Storage FlashCopy Manager uses the file type extension to determine how to run the script. As a result, make sure that your scripts follow these extension requirements:
     - PowerShell scripts: *`filename`*`.ps1`
     - Windows Management Instrumentation (WMI) scripts: *`filename`*`.wmi`
     - Tivoli Storage Manager scripts: *`filename`*`.tsm`
   - To view or edit an existing script:
     a. 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.

     b. Click **Command Editor** in the Actions pane to open the script file for viewing or editing.

     c. View or edit the script. Click **OK** to save your changes, or click **Cancel** to exit the System Information Command Editor without saving any changes.
   - To delete a script:
     a. 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.

     b. Click **Delete** in the Actions pane.

Chapter 8. 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 Microsoft Exchange. Some options that are related to Data Protection for Microsoft Exchange can be tuned for optimum performance. See "Specifying Data Protection for Exchange options" on page 35 for details that regard these options. In addition, the following issues affect performance:

- Backups to local shadow volumes eliminates the transfer of data to the Tivoli Storage Manager server.
- During VSS backup processing, integrated Exchange integrity checking reads every page in the files to be backed up. Backup processing time can be significant. Performance can improve if you specify the **/SKIPINTEGRITYCHECK** parameter to bypass integrity checking. There is a risk that is associated with using this parameter. For information about skipping Exchange integrity checking, refer to the Microsoft documentation.
- The time that is required to complete a snapshot, ranges from seconds to minutes, depending on the type of VSS provider used. If an integrity check is run, it can delay the completion of the backup depending on the size of the database and log files.
- Backup-archive client settings can affect performance when you back up data to the Tivoli Storage Manager server. Review the information that is provided in the "VSS backup" on page 4 and "How Tivoli Storage Manager server policy affects Data Protection for Exchange" on page 28 sections.
- Performing Data Protection for Microsoft Exchange VSS backups from an Exchange Server DAG passive copy can offload I/O and possibly CPU resources from the production server.

If you have not applied the update for Tivoli Storage Manager server APAR IC86558, apply the update.

For VSS backups, the **RESOURCEUTILIZATION** client option is also important. This option increases or decreases the ability 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*. For more information about **RESOURCEUTILIZATION**, see the *IBM Tivoli Storage Manager Performance Tuning Guide*.

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.

LAN-free data movement

Running Data Protection for Microsoft Exchange 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. The *Tivoli Storage Manager Managed System for SAN Storage Agent User's Guide* provides detailed information about setting up a LAN-free environment.

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.

lanfreetcpport
Specifies the TCP/IP port number where the Tivoli Storage Manager Storage Agent is listening.

lanfreetcpserveraddress
Specifies the TCP/IP address for a Tivoli Storage Manager Storage Agent.

For more information about these options, see *IBM Tivoli Storage Manager for Windows Backup-Archive Clients Installation and User's Guide*.

Chapter 9. Reference information

Data Protection for Exchange reference information is provided.

Command overview

The name of the Data Protection for Microsoft Exchange command-line interface is **tdpexcc.exe**. This program is in the directory where Data Protection for Microsoft Exchange is installed.

Using the Data Protection for Microsoft Exchange command-line interface from the GUI

You can start a Windows command prompt with administrative privileges. Or, you can follow these steps to start the Data Protection for Microsoft Exchange command-line interface:

1. Start the Management Console (MMC) GUI.
2. In the tree view, select the computer node where you want to run the commands.
3. Expand the **Protect and Recover Data** node.
4. In the tree view, 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. Change **PowerShell** to **Command Line**.

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.

Backup command

Use the **backup** command to run Exchange Server database backups from the Exchange Server to Tivoli Storage Manager 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 Tivoli Storage Manager 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 Tivoli Storage Manager server (full)
- VSS Tivoli Storage Manager server (copy)

The Exchange database file size might increase resulting from increase database commitments that are triggered by backup operations. This behavior is standard for the Microsoft Exchange server.

For SAN Volume Controller and Storwize V7000 storage subsystems, 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 completes. As a result, local backups for SAN Volume Controller and Storwize V7000 storage subsystems are to be initiated at a frequency greater than the time required for the background copy process to complete.

See “Backup strategies” on page 18 for more information that is related to the **backup** command.

Data Protection for Microsoft Exchange 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

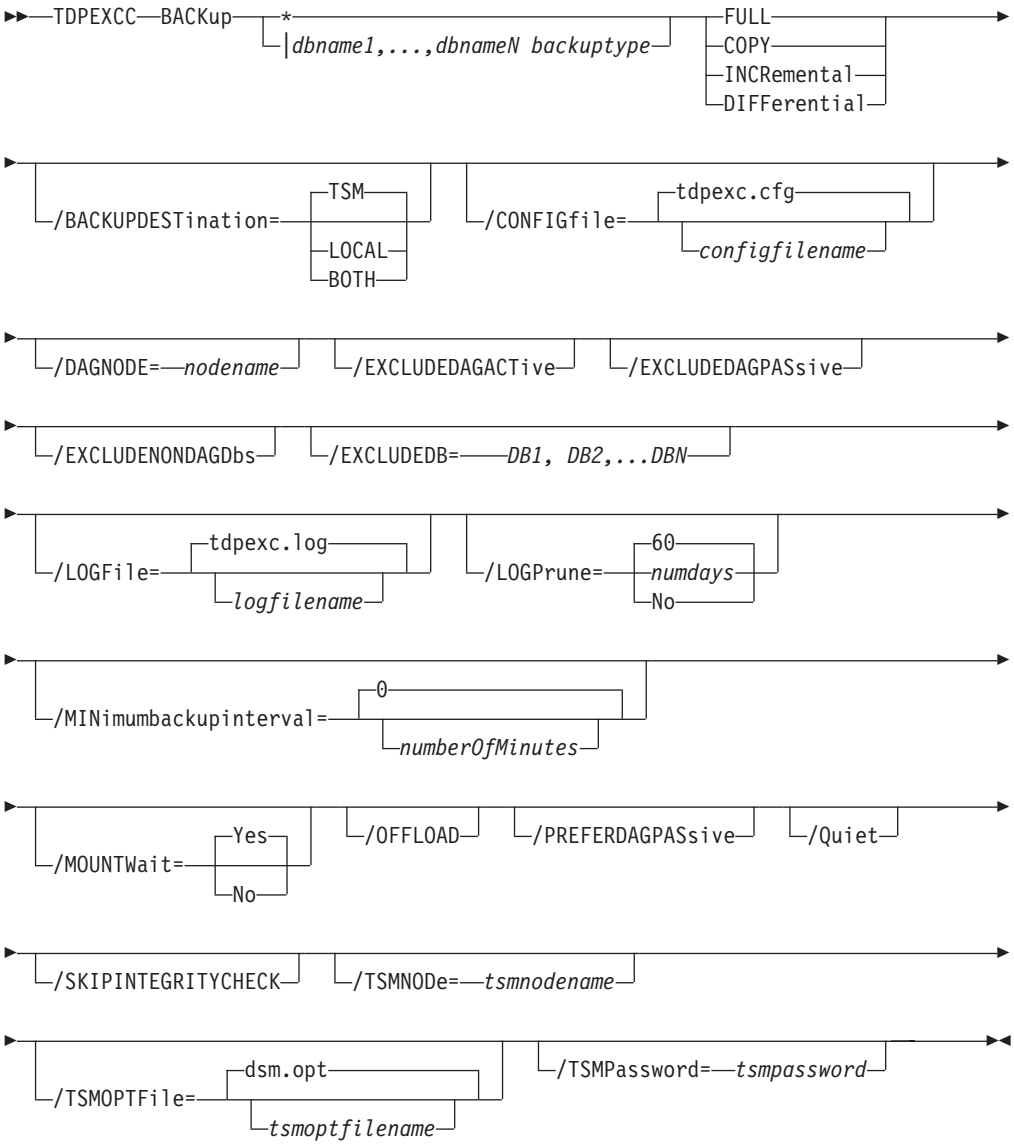

Backup positional parameters

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 contains commas or 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 Tivoli Storage Manager 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 Tivoli Storage Manager server storage and local shadow volumes.

/CONFIGfile=*configfilename*
Use the **/CONFIGfile** parameter to specify the name (*configfilename*) of the Data Protection for Microsoft Exchange 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 Microsoft Exchange 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 Tivoli Storage Manager 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 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 Microsoft Exchange.

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 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 Microsoft Exchange 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**
The **/LOGPrune** parameter prunes the Data Protection for Microsoft Exchange activity log and specifies how many days of entries are saved. By default, log pruning is enabled and done once each day that Data Protection for Microsoft Exchange is run; however, you can use this option to disable log pruning or explicitly request a prune of the log for one command run even if the log file is already pruned for the day. The *numdays* variable represents the number of days to save log entries. By default, *60* days of log entries are saved in the prune process.

- If you specify *numdays*, it can range from *0* to *9999*. A value of *0* deletes all entries in the Data Protection for Microsoft Exchange activity log file except for the current command entries.

- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the Data Protection for Microsoft Exchange configuration file. Specifying this parameter does not change the value in the configuration file.
- Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the Data Protection for Microsoft Exchange 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, do one of the following to prevent undesired 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. If you use the parameter, but do not specify a value, you can back up the database again immediately after a backup operation of that database completes.

Setting this parameter specifies that only one database copy can be backed up within a timeframe. This option prevents all of the members in a Database Availability Group from backing up the database, which would be redundant and invalidate the Tivoli Storage Manager storage management policy.

/MOUNTWait=Yes|No

Use the **/MOUNTWait** parameter to specify whether Data Protection for Microsoft Exchange is to wait for removable media to mount (such as tapes or CDs) or to stop the current operation. This situation occurs when the Tivoli Storage Manager 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 perform the integrity check and backup of files to Tivoli Storage Manager on the system 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. It is not supported 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 healthy passive copy is available. If no healthy passive copy is available, the backup is made 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

Specify this parameter to bypass the Exchange integrity check typically

performed during a backup. During VSS backup processing, integrated Exchange integrity checking reads every page in the files to be backed up. Backup processing time can be significant. You can specify the**/SKIPINTEGRITYCHECK** parameter to bypass integrity checking. This parameter is valid for all VSS backups, only skip these checks in accordance with recommendations from Microsoft.

When using this parameter, it is possible that the stored backup is not valid because it is not being verified with the Exchange integrity check utility. Make sure that you have a valid backup stored on Tivoli Storage Manager server storage. If you are using a Database Availability Group (DAG) configuration, and there are at least two viable copies of the database, you can skip the integrity check in order to speed up the backup.

/TSMNODe=*tsmnodename*

Use the *tsmnodename* variable to refer to the Tivoli Storage Manager node name that Data Protection for Microsoft Exchange uses to log on to the Tivoli Storage Manager server. You can store the node name in the Tivoli Storage Manager options file (`dsm.opt`). This parameter overrides the value in the Tivoli Storage Manager 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 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 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 Tivoli Storage Manager password that Data Protection for Microsoft Exchange uses to log on to the Tivoli Storage Manager server. If you specified **PASSWORDACCESS** `GENERATE` in the Data Protection for Microsoft Exchange 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 Tivoli Storage Manager password the first time Data Protection for Microsoft Exchange connects to the Tivoli Storage Manager 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 Tivoli Storage Manager password that Data Protection for Microsoft Exchange uses to log on to the Tivoli Storage Manager server can be up to 63 characters in length.

Examples: backup command

Examples of how to use the **backup** command are given.

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 to back up only the database KEENVM1_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 a Tivoli Storage Manager schedule that starts this command script, and associate all DAG nodes to this schedule.

```
tdpexcc backup KEENVM1_M_DB1 full /minimumbackupinterval=60
```

This example shows how to run a full VSS backup of one healthy passive copy of a database that contains multiple copies in an Exchange Server Database Availability Group (DAG). If a healthy passive copy is not available, the active database copy is backed up. The command instructs Data Protection for Exchange to back up only the passive copy of database KEENVM1_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 a Tivoli Storage Manager schedule that starts this command script, and associate all DAG nodes to this schedule.

```
tdpexcc backup KEENVM1_M_DB1 full /minimumbackupinterval=60 /preferdagpassive
```

Changetsmpassword command

To change the Tivoli Storage Manager password that is used by Data Protection for Microsoft Exchange, use the **changetsmpassword** command. The password is used to log on to the Tivoli Storage Manager server.

If you do not enter the old and new passwords, Data Protection for Microsoft Exchange prompts you for the old and new passwords. Data Protection for Microsoft Exchange does not display the password on the screen. The Tivoli Storage Manager password Data Protection for Microsoft Exchange uses to log on to the Tivoli Storage Manager 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

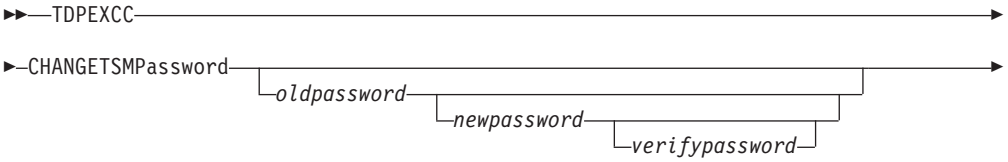

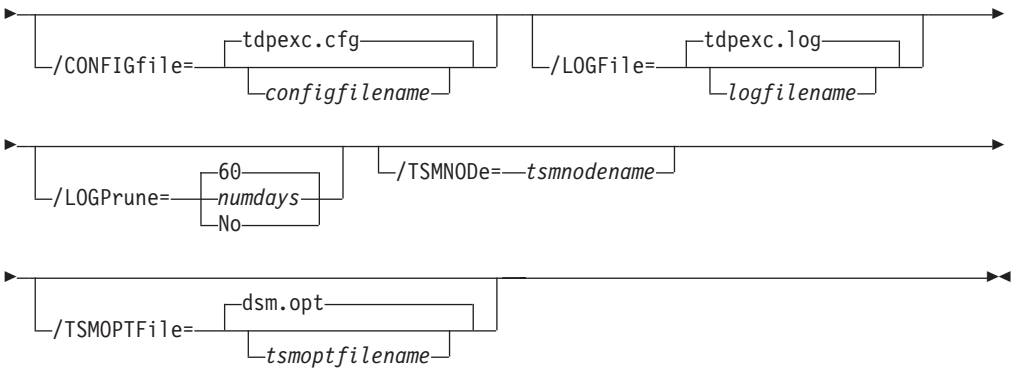

Changetsmpassword positional parameters

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 Microsoft Exchange.

newpassword
Specifies the new password that is used by Data Protection for Microsoft Exchange.

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 Microsoft Exchange configuration file that contains the values for the Data Protection for Microsoft Exchange configuration options. See "Set command" on page 173 for details about the contents of the file.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the Data Protection for Microsoft Exchange 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=*logfilename*
Use the **/logfile** parameter to specify the name of the activity log file that is generated by Data Protection for Microsoft Exchange.

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 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 Microsoft Exchange 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**

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 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 Microsoft Exchange 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 is already done for the day.

Changes to the value of the **timeformat** or **dateformat** parameter can result in prunning the log file unintentionally. If the value of the **timeformat** or **dateformat** parameter is changed, before issuing a Data Protection for Microsoft Exchange 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.

/TSMNODe=*tsmnodename*

Use the *tsmnodename* variable to refer to the Tivoli Storage Manager node name that Data Protection for Microsoft Exchange uses to log on to the Tivoli Storage Manager server. You can store the node name in the Tivoli Storage Manager options file (`dsm.opt`). This parameter overrides the value in the Tivoli Storage Manager 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 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 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 Tivoli Storage Manager password that is used by Data Protection for Microsoft Exchange:

```
tdpexcc changetsmpassword oldpw newpw newpw
```

Delete backup command

To delete a VSS backup of an Exchange Server database, use the **delete backup** command.

You must have local registry rights for all versions of Exchange Server to complete a Data Protection for Microsoft Exchange delete backup. When a full VSS snapshot backup is completed, the backup remains active until the backup version is deleted with the delete backup command, or expired by VSS according to the defined policy. Two different active backups can exist at the same time:

- Full backup, along with any associated incremental backups and differential backups.
- Copy backup, along with 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 erroneously seems to be 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

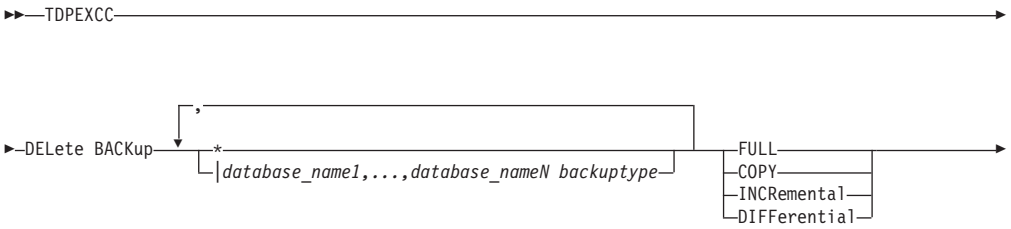

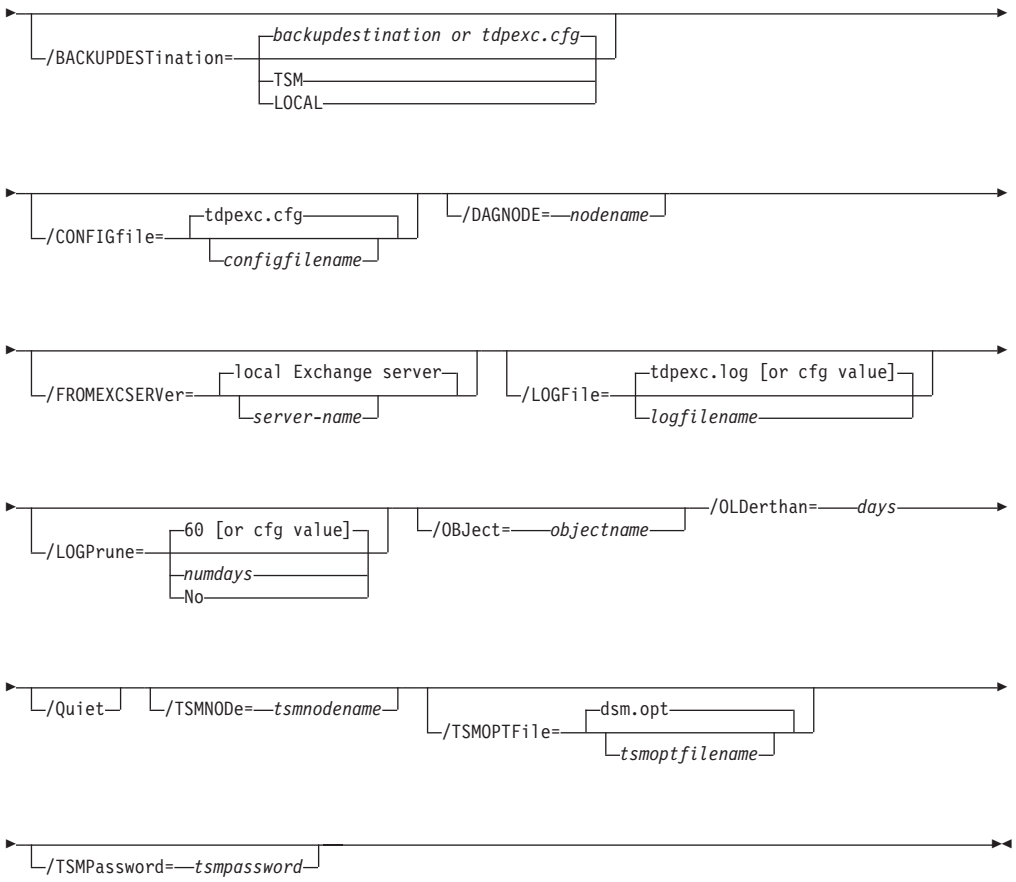

Delete Backup positional parameters

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 contains commas or 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 Microsoft Exchange preferences file (`tdpexc.cfg`). If no value is present, the backup is deleted from Tivoli Storage Manager server storage.

You can specify:

TSM The backup is deleted from Tivoli Storage Manager 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 Microsoft Exchange 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 Microsoft Exchange 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"
```

See "Set positional parameters" on page 174 for descriptions of available configuration parameters.

/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 Tivoli Storage

Manager 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 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 Microsoft Exchange.

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 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 using multiple simultaneous instances of Data Protection for Microsoft Exchange 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 Microsoft Exchange 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 Microsoft Exchange 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 you want to delete. The object name uniquely identifies each backup object and is created by Data Protection for Microsoft Exchange.

Use the Data Protection for Microsoft Exchange `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=*tsmnodename*

Use the *tsmnodename* variable to refer to the Tivoli Storage Manager node name that Data Protection for Microsoft Exchange uses to log on to the Tivoli Storage Manager server. You can store the node name in the Tivoli Storage Manager options file (`dsm.opt`). This parameter overrides the value in the Tivoli Storage Manager 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 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 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 Tivoli Storage Manager password that Data Protection for Microsoft Exchange uses to log on to the Tivoli Storage Manager server. If you specified **PASSWORDACCESS** `GENERATE` in the Data Protection for Microsoft Exchange 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 Tivoli Storage Manager password the first time Data Protection for Microsoft Exchange connects to the Tivoli Storage Manager 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 Tivoli Storage Manager password that Data Protection for Microsoft Exchange uses to log on to the Tivoli Storage Manager server can be up to 63 characters in length.

Help command

To display help for Data Protection for Microsoft Exchange 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

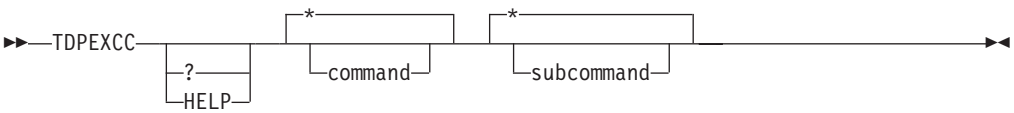

Help optional parameters

Optional parameters follow the Data Protection for Microsoft Exchange **help** command.

The following optional parameters specify the help to be displayed:

***|** *command*

Identifies the specific Data Protection for Microsoft Exchange command that is to be displayed. If the wildcard character (*) is used, help for all Data Protection for Microsoft Exchange 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 the wildcard character (*), help for all Data Protection for Microsoft Exchange **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

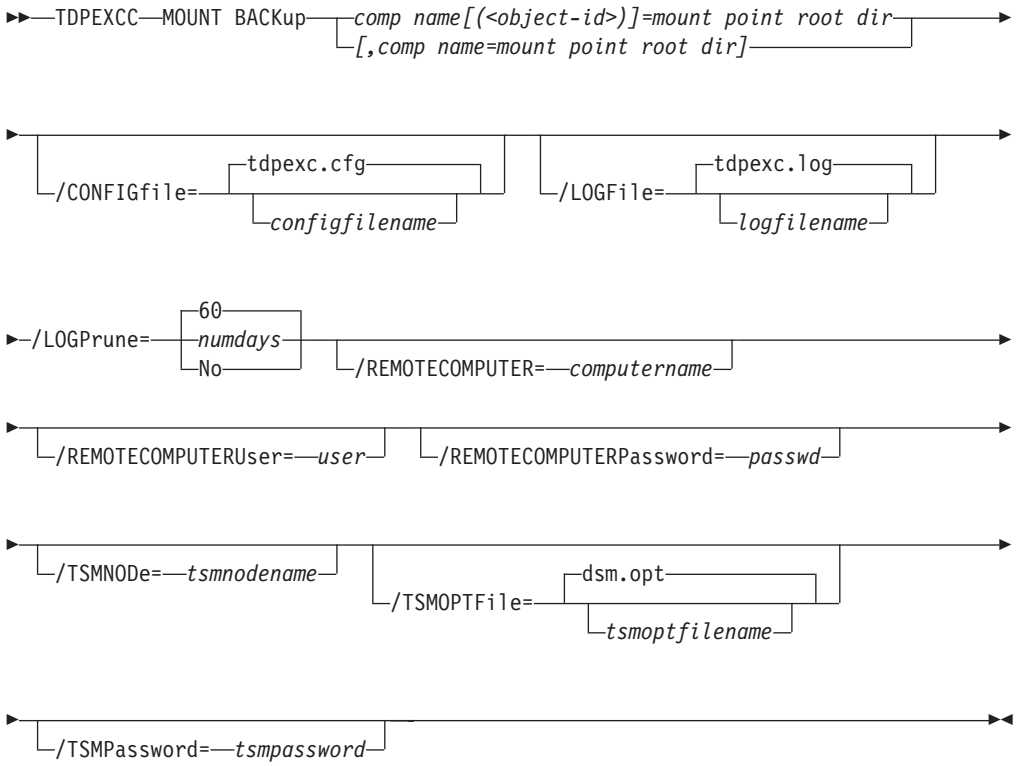

Mount backup positional parameter

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:

```
tdpexcc 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 Tivoli Storage FlashCopy Manager for Exchange 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 Tivoli Storage FlashCopy Manager for Exchange 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=*logfilename*

Use the **/logfile** parameter to specify the name of the activity log file that is generated by Tivoli Storage FlashCopy Manager for Exchange. 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 Tivoli Storage FlashCopy Manager for Exchange installation directory.

If the *logfilename* 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**

Use the **/logprune** parameter to disable log pruning or to explicitly request that the log is to be pruned 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 GUI or the **update config** 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 is already done for the day.

Changes to the value of the **timeformat** or **dateformat** parameter can result in pruning the log file unintentionally. If the value of the **timeformat** or **dateformat** parameter is changed, before issuing a Tivoli Storage

FlashCopy Manager for Exchange 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.

/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 Tivoli Storage Manager node name that Tivoli Storage FlashCopy Manager uses to log on to the Tivoli Storage Manager server. You can store the node name in the Tivoli Storage Manager options file (`dsm.opt`). This parameter overrides the value in the Tivoli Storage Manager 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 Tivoli Storage Manager options file.

The file name can include a fully qualified path name. If no path is specified, the directory where Tivoli Storage FlashCopy Manager 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 Tivoli Storage Manager password that Tivoli Storage FlashCopy Manager uses to log on to the Tivoli Storage Manager server. If you specified **PASSWORDACCESS** `GENERATE` in the Tivoli Storage FlashCopy Manager 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 Tivoli Storage Manager password the first time Tivoli Storage FlashCopy Manager connects to the Tivoli Storage Manager 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 Tivoli Storage Manager password that Tivoli Storage FlashCopy Manager uses to log on to the Tivoli Storage Manager 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

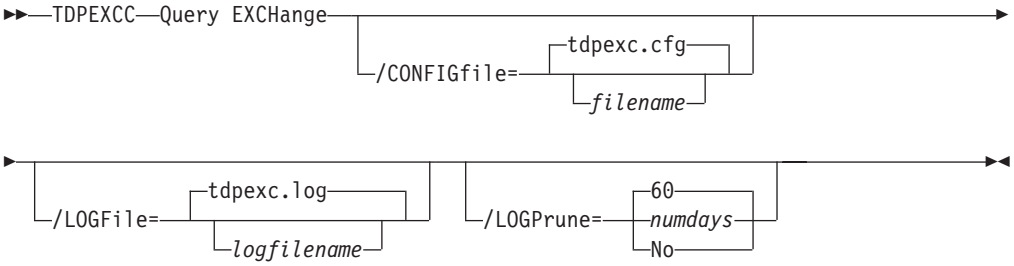

Query Exchange optional parameters

Optional parameters follow the **query exchange** command.

/CONFIGfile=*filename*

Use the **/CONFIGfile** parameter to specify the name (*filename*) of the Data Protection for Microsoft Exchange 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 Microsoft Exchange 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 Microsoft Exchange. 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 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 Microsoft Exchange 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 that the log be pruned 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 Microsoft Exchange graphical user interface 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 prunning the log file unintentionally. If the value of the **timeformat** or **dateformat** parameter is changed, before issuing a Data Protection for Microsoft Exchange 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.

Query Managedcapacity command

Use the **query managedcapacity** command to assist with storage planning by determining the amount of managed capacity in use.

Purpose

The **query managedcapacity** command displays capacity that is related information about the volumes that are represented in local inventory that is managed by Data Protection for Microsoft Exchange. This command is valid for all Windows operating systems that are supported by Data Protection for Microsoft Exchange.

TDPEXCC command: Query MANAGEDCAPacity

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 : 100.01 GB (107,381,026,816 bytes)

Volume           : D:
Managed Capacity : 100.01 GB (107,381,026,816 bytes)

Completed
```

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:

```
Total Managed Capacity : 1,019.72 MB (1,069,253,632 bytes)

Volume           : I:
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

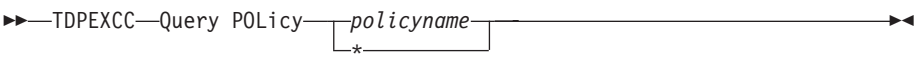

When you specify *, all policies are queried. The results of the query are displayed:

```
                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 Microsoft Exchange, 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

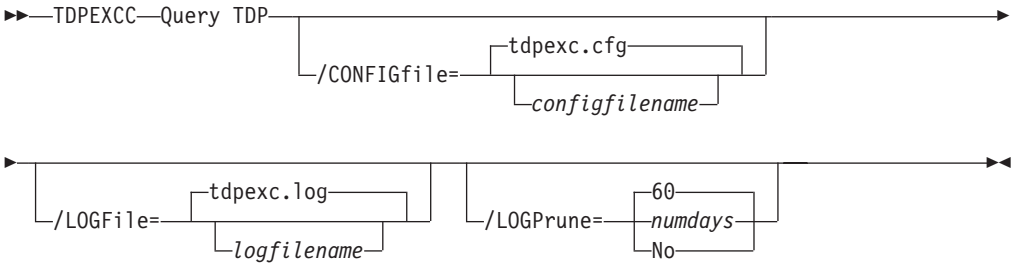

Query TDP optional parameters

Optional parameters follow the **query TDP** command.

/CONFIGfile=*configfilename*

Use the **/CONFIGfile** parameter to specify the name (*configfilename*) of the Data Protection for Microsoft Exchange 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 Microsoft Exchange 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=*logfilename*

Use the **/LOGFile** parameter to specify the name of the activity log file that is generated by Data Protection for Microsoft Exchange.

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 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 Microsoft Exchange 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 that the log be pruned for one command run. By default, log pruning is enabled and done once each 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 Microsoft Exchange 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 is already done for the day.

Changes to the value of the **timeformat** or **dateformat** parameter can result in prunning the log file unintentionally. If the value of the **timeformat** or **dateformat** parameter is changed, before issuing a Data Protection for Microsoft Exchange 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.

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 0.0
(C) Copyright IBM Corporation 1998, 2013. 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
REMOTEDSMAgentnode..................
TEMPDBRestorepath...................
TEMPLOGRestorepath..................
TIMEformat ......................... 1
IMPORTVSSSNAPSHOTSONLYWhenneeded.... No

Completed
```

Query TSM command

Use the **query tsm** command to display Tivoli Storage Manager server information.

This command displays the following information:

- Tivoli Storage Manager node name
- Network host name of the server
- Tivoli Storage Manager 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 Tivoli Storage Manager 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

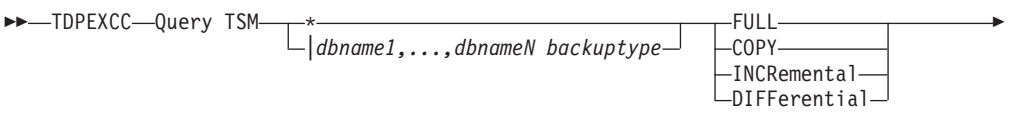

Query TSM positional parameters

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 Tivoli Storage Manager API and Tivoli Storage Manager 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 Microsoft Exchange configuration file that contains the values for the Data Protection for Microsoft Exchange configuration options.

The *filename* variable can include a fully qualified path. If the *filename* variable does not include a path, the Data Protection for Microsoft Exchange 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 Tivoli Storage Manager 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 Microsoft Exchange from making too many backups of the same database.

/DEtail
Use the **/DEtail** parameter to display comprehensive information about the status of the Tivoli Storage Manager server, including the following details:

- Backup compressed

Table 7. Backup compressed values

Value	Status
Yes	One or more objects are compressed.
No	No objects are compressed.
Unknown	It is not known if the backup is compressed. This applies to backups before Data Protection for Microsoft Exchange version 6.3.

- Backup encryption type

Table 8. Backup encryption type values

Value	Status
None	None of the objects are encrypted.
AES-128	The objects are encrypted with AES-128 encryption.

Table 8. Backup encryption type values (continued)

Value	Status
DES-56	The objects are encrypted with DES-56 encryption.
Unknown	It is not known whether the objects in the database are encrypted. This applies to backups before Data Protection for Microsoft Exchange version 6.3.

- Backup client-deduplicated

Table 9. 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 if the backup is client-side deduplicated. This applies to backups before Data Protection for Microsoft Exchange version 6.3.

- Backup supports instant restore

Table 10. Backup supports instant restore values

Value	Status
Yes	One or more objects support instant restore.
No	No objects support instant restore.
Unknown	The backup might not support instant restore. This setting applies to backups before Data Protection for Microsoft Exchange version 6.3.

/FROMEXCSERVer=*servername*

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.

If a DAG node is specified by using the **dagnode** parameter, Data Protection for Microsoft Exchange uses this node name instead of the Data Protection for Microsoft Exchange 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=*logfilename*

Use the **/LOGFile** parameter to specify the name of the activity log file that is generated by Data Protection for Microsoft Exchange.

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 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 Microsoft Exchange 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 that the log be pruned 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 Microsoft Exchange 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 interface, 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 prunning the log file unintentionally. If the value of the **timeformat** or **dateformat** parameter is changed before issuing a Data Protection for Microsoft Exchange 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.

/REMOTECOMPUTER=*computername*

Enter the IP address or host name for the remote system where you want to mount the data.

/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.

/SHOWMAILBOXINFO

The mailbox history processing task gathers information from the Active Directory service for all mail users within a Microsoft Exchange Server environment. The information that is gathered includes references for the location of the mailbox for each user. This information can be used to

automate the mailbox restore processing when mailboxes are moved from one database or server to another database or server between backup operations. To show the information that is saved during the Mailbox History backup, enter the following command from the Data Protection for Exchange command-line interface:

```
tdpexcc q tsm /showmailboxinfo
```

/TSMNODe=*nodename*

Use the *tsmnodename* variable to refer to the Tivoli Storage Manager node name that Data Protection for Microsoft Exchange uses to log on to the Tivoli Storage Manager server. You can store the node name in the Tivoli Storage Manager options file (`dsm.opt`). This parameter overrides the value in the Tivoli Storage Manager 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 Microsoft Exchange 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 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 Tivoli Storage Manager password that Data Protection for Microsoft Exchange uses to log on to the Tivoli Storage Manager server. If you specified `PASSWORDACCESS GENERATE` in the Data Protection for Microsoft Exchange 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 Tivoli Storage Manager password the first time Data Protection for Microsoft Exchange connects to the Tivoli Storage Manager 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 Tivoli Storage Manager password that Data Protection for Microsoft Exchange uses to log on to the Tivoli Storage Manager server can be up to 63 characters in length.

Examples: `query tsm` command

The **query tsm** command displays information about the Tivoli Storage Manager API and Tivoli Storage Manager 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 0.0
(C) Copyright IBM Corporation 1998, 2013. 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 0.0

Server Name ............................ GIJOE_SERVER1_230
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

Completed
```

For backups that are made 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 Tivoli Storage Manager 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 0.0
(C) Copyright IBM Corporation 1998, 2013. 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
-------------------  ----------- - ---- ----   --- ------------------------
03/27/2013 16:11:14     149.07MB A VSS  full   Srv 20120327161114

                         13.01MB                  Logs
                        136.06MB                  File
03/27/2013 18:02:01      14.00MB A VSS  incr   Srv 20120327180201
                         14.00MB                  Logs
```

The following example queries the Tivoli Storage Manager 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 0.0
(C) Copyright IBM Corporation 1998, 2013. 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 ..... 03/27/2013 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 / 20120327161114 (From DBC
opy)
Backup Object Size / Name .....    13.01MB / Logs
Backup Object Size / Name .....   136.06MB / File
```

Restore command

To restore a backup from Tivoli Storage Manager storage to an Exchange server, use the **restore** command.

Before completing a VSS restore, review the following topics:

- "VSS restore considerations" on page 146
- "Restoring VSS backups into alternate locations" on page 13

When you use the restore command, remember the following facts:

- 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 Tivoli Storage Manager storage. You can issue a **tdpexcc query tsm *** command to obtain a list of the object names.

  If the tdpexcc restore sgname 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 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 completes.

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.

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 Microsoft Exchange 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 *Enn*`restore.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. See the following Microsoft Exchange documentation for further details: http://msdn.microsoft.com/en-us/library/bb204044.aspx

Data Protection for Microsoft Exchange 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.

VSS restore considerations

Refer to the following considerations when you run VSS restores.

Unless otherwise specified, *VSS restores* refers to all restore types that use VSS (VSS restore, VSS fast restore, VSS instant restore):

- A VSS restore ignores the recovery database and is placed directly into the production database unless the **/intodb** parameter is specified.
- A VSS instant restore overwrites the entire contents of the source volumes. However, you can avoid overwriting the source volumes by specifying **/INSTANTRESTORE**=*NO*. This parameter bypasses volume-level copy and uses file-level copy instead to restore the files from a VSS backup that are on local shadow volumes.
- A VSS restore requires the restored database to be dismounted.
- If a VSS hardware provider is used, the disks that contain Exchange data are configured as basic disks.
- When a VSS restore from local shadow volumes is complete, the bytes transferred display *0*. This display occurs because no data (*0*) is restored from the Tivoli Storage Manager server.
- Do not set the ASNODENAME option in the dsm.opt file when you use Data Protection for Microsoft Exchange Server. Setting ASNODENAME can cause VSS backups and VSS restores to fail.

Restore syntax

Use the **restore** command syntax diagrams as a reference to view available options and truncation requirements.

TDPEXCC command

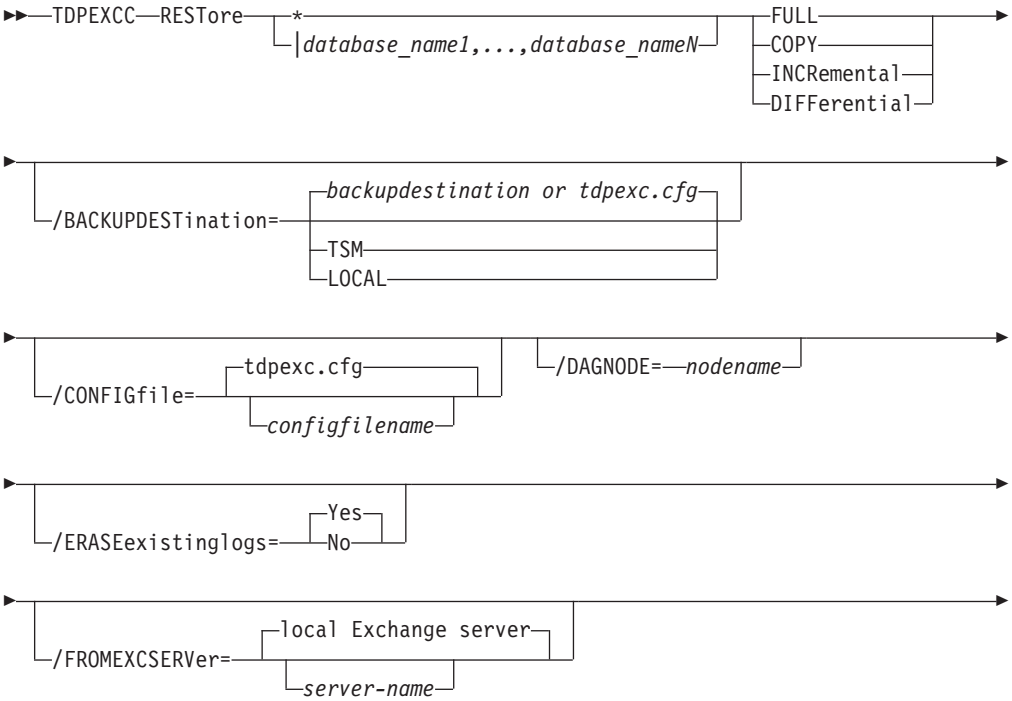

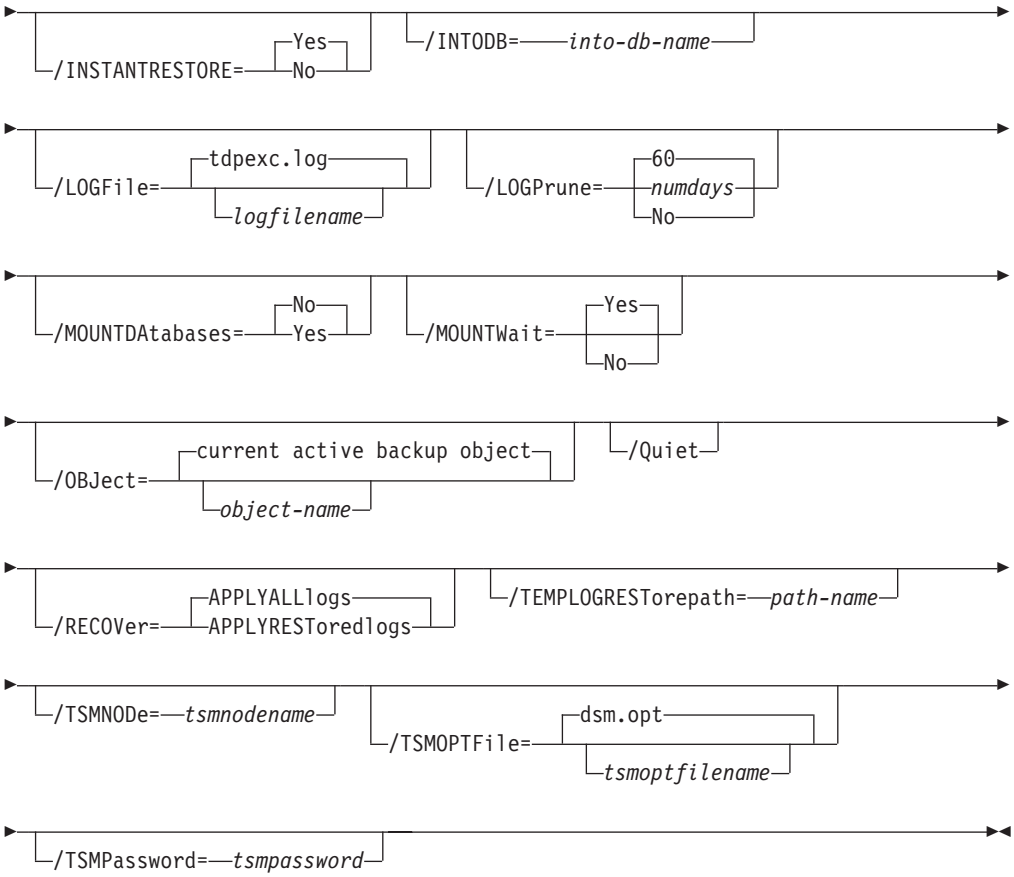

Restore positional parameters

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 contains commas or 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 Microsoft Exchange preferences file (`tdpexc.cfg`). If no value is present, the backup is restored from Tivoli Storage Manager server storage.

You can specify:

TSM The backup is restored from Tivoli Storage Manager 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 Microsoft Exchange configuration file that contains the values for the Data Protection for Microsoft Exchange configuration options. See "Set command" on page 173 for details about the contents of the file.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the Data Protection for Microsoft Exchange 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 Tivoli Storage Manager 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 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 Microsoft Exchange uses this node name instead of the Data Protection for Microsoft Exchange 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 V7000, or DS8000 storage subsystem is required to perform VSS Instant Restores.

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 Restores, 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 Microsoft Exchange.

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 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 Microsoft Exchange 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 Microsoft Exchange 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 prunning the log unintentionally. If the value of the **TIMEformat** or **DATEformat** parameter is changed, before you issue a Data Protection for Microsoft Exchange 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 completes. Specify one of the following values:

Yes Mount the databases after the restore operation completes.

No Do not mount the databases after the restore operation completes. `No` is the default.

/MOUNTWait=Yes|No

Use the **/MOUNTWait** parameter to specify whether Data Protection for

Microsoft Exchange is to wait for removable media to mount (such as tapes or CDs) or to stop the current operation. This situation occurs when the Tivoli Storage Manager 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 Microsoft Exchange.

Use the Data Protection for Microsoft Exchange **query tsm** command to view the names of the backup objects.

If the `tdpexcc restore sgname 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 Microsoft Exchange configuration file. The default Data Protection for Microsoft Exchange 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 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 Tivoli Storage Manager node name that Data Protection for Microsoft Exchange uses to log on to the Tivoli Storage Manager server. You can store the node name in the Tivoli Storage Manager options file (`dsm.opt`). This parameter overrides the value in the Tivoli Storage Manager 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 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 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 Tivoli Storage Manager password that Data Protection for Microsoft Exchange uses to log on to the Tivoli Storage Manager server. If you specified **PASSWORDACCESS** `GENERATE` in the Data Protection for Microsoft Exchange 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 Tivoli Storage Manager password the first time Data Protection for Microsoft Exchange connects to the Tivoli Storage Manager 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 Tivoli Storage Manager password that Data Protection for Microsoft Exchange uses to log on to the Tivoli Storage Manager server can be up to *63* characters in length.

Restorefiles command

To restore the flat files from a specified Data Protection for Microsoft Exchange backup into a specified directory, use the **restorefiles** command.

The following information provides details about this command:

- 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 Tivoli Storage Manager 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 Microsoft Exchange graphical user interface.

The following provides details about the **restorefiles** command:

- 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 STG1 FULL /into=d:\temprestore
  ```

  For an incremental restore with backed up log files, enter the following command:

  ```
  tdpexcc restorefiles STG1 INCR /into=d:\temprestore
  ```

- 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.

Restorefiles syntax

Use the **restorefiles** command syntax diagrams as a reference to view available options and truncation requirements.

TDPEXCC command

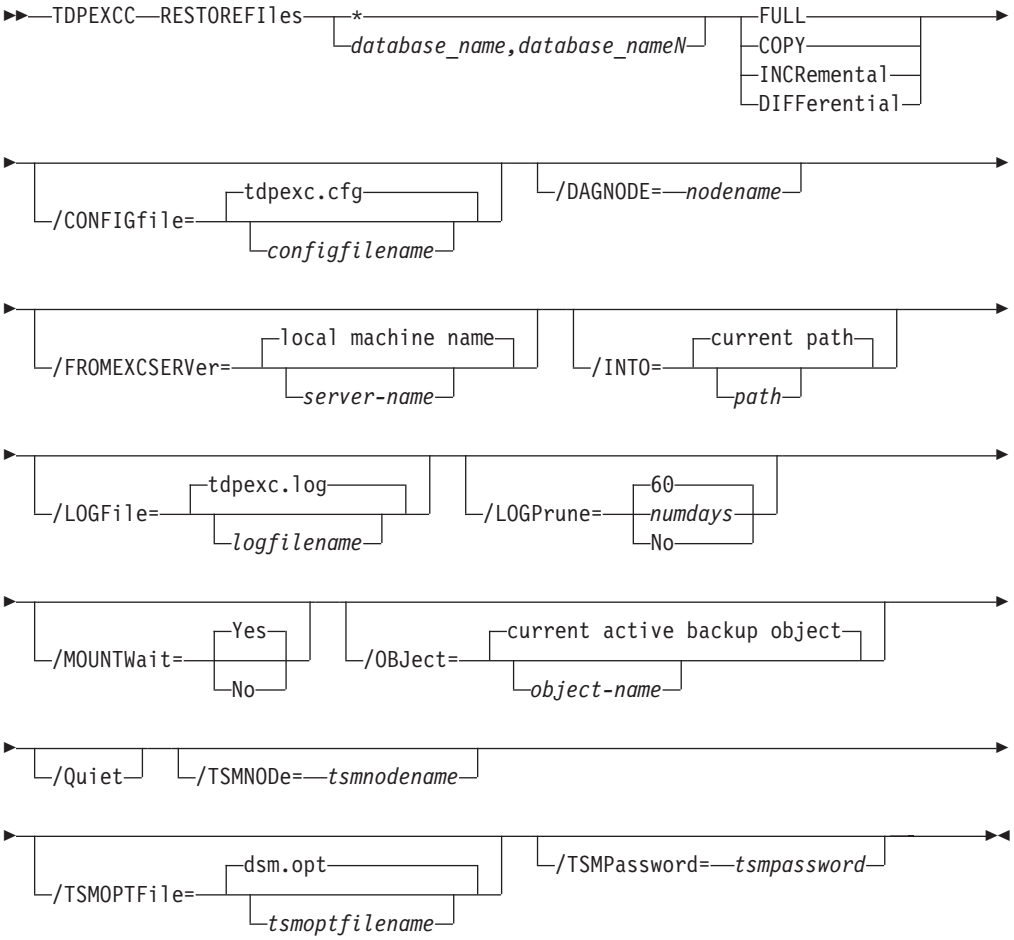

Restorefiles positional parameters

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 Microsoft Exchange configuration file that contains the values for the Data Protection for Microsoft Exchange configuration options. See "Set command" on page 173 for details about the contents of the file.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the Data Protection for Microsoft Exchange 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 Tivoli Storage Manager 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 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 Microsoft Exchange uses this node name instead of the Data Protection for Microsoft Exchange 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 are to 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.

/LOGFile=*logfilename*

Use the **/logfile** parameter to specify the name of the activity log file that is generated by Data Protection for Microsoft Exchange.

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 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 Microsoft Exchange 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**

Use the **/logprune** parameter to disable log pruning or to explicitly request that pruning of 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 Microsoft Exchange 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 is done for the day.

Changes to the value of the **timeformat** or **dateformat** parameter can result in pruning the log file unintentionally. If the value of the **timeformat** or **dateformat** parameter is changed, before you issue a Data Protection for Microsoft Exchange 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.

/MOUNTWait=Yes|**No**

Use the **/mountwait** parameter to specify whether Data Protection for Microsoft Exchange waits for removable media to mount (such as tapes or CDs) or to stop the current operation. This situation occurs when the Tivoli Storage Manager 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 Microsoft Exchange.

Use the Data Protection for Microsoft Exchange **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 Tivoli Storage Manager node name that Data Protection for Microsoft Exchange uses to log on to the Tivoli Storage Manager server. You can store the node name in the Tivoli Storage Manager options file (`dsm.opt`). This parameter overrides the value in the Tivoli Storage Manager 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 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 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 Tivoli Storage Manager password that Data Protection for Microsoft Exchange uses to log on to the Tivoli Storage Manager server. If you specified **PASSWORDACCESS** `GENERATE` in the Data Protection for Microsoft Exchange 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 Tivoli Storage Manager password the first time Data Protection for Microsoft Exchange connects to the Tivoli Storage Manager 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 Tivoli Storage Manager password that Data Protection for Microsoft Exchange uses to log on to the Tivoli Storage Manager server can be up to 63 characters in length.

Restoremailbox command

Use the **restoremailbox** command to restore mailbox-level data or mailbox-item-level data from Data Protection for Microsoft Exchange backups.

The following information provides details about this command:

- Before you start, ensure that you have the required privileges to run Mailbox restores.
- The **restoremailbox** command applies to any Data Protection for Microsoft Exchange VSS backups:
  - VSS backups that are stored on Tivoli Storage Manager server
  - VSS backups that are stored on local shadow volumes
- You can use the mailbox restore operation in the GUI to restore mailbox-level data or mailbox-item-level data. The GUI also provides the Mailbox Restore Browser, an interactive action panel that lists all available mailbox actions. Some features of the **restoremailbox** command are only available on the command-line interface:
  - The **/mailboxoriglocation** parameter is necessary when the mailbox history is disabled and if the mailbox being restored has either moved or been deleted after the time of backup.

    The **/mailboxoriglocation** parameter is available when you select **Show Restore Options** from the Recover tab. Use this setting to specify the server and the database where the mailbox was at the time of backup.
  - Alternatively use the command-line interface from the Automate tab to specify the **/mailboxoriglocation**.
  - Set the **/tempmailboxalias** optional parameter by selecting **Properties** from the Actions pane. In the **Data Protection Properties** dialog, click the **General** tab to specify the alias of the temporary mailbox. Use this selection for mailbox restore operations on mailboxes that were deleted or recreated after the time of the backup you are restoring from.

    Select **Properties** from the Actions pane to open the Data Protection for Exchange Server Properties form. 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.
- With Data Protection for Microsoft Exchange you can restore multiple mailboxes with the same mailbox restore operation.
- You can use the **restoremailbox** command to restore data into a mailbox in an online Exchange Server or to restore data as an Exchange Server personal folders (`.pst`) file.
- You can use the **restoremailbox** command on the primary Exchange Server or on an alternate Exchange Server that is in the same domain.
- 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

When restoring 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.

Prerequisites for Data Protection for Microsoft Exchange mailbox restore tasks

Review these prerequisites before you run mailbox restore tasks on Exchange Server 2010 and later:

- See *Security requirements for Data Protection for Microsoft Exchange mailbox tasks* in "Security requirements" on page 17
- See "Perform these tasks to configure your system for mailbox-level and item-level restore operations" on page 65
- Temporary space is needed to accommodate the mailbox database during mailbox restore operations. Specify the location of this temporary space by setting these two optional parameters in the Data Protection for Exchange configuration file with the **tdpexcc set** command:
  - **TEMPDBRESTorepath**

    If you choose to not enter a path, the default value of **TEMPDBRESTorepath** is the value of the TEMP environment variable.
  - **TEMPLOGRESTorepath**

    If you choose to not enter a path, the default value of **TEMPLOGRESTorepath** is the value of the TEMP environment variable.

  The temporary restore locations must have enough space to restore the entire restored databases and log files.
- Ensure that 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 Data Protection for Exchange program. This technote is available in the *TSM for Mail - All Requirement Documents* website at http://www.ibm.com/support/docview.wss?uid=swg21219345. When you are at the website, 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.

Restoremailbox syntax

Use the **restoremailbox** command syntax diagrams as a reference to view available options and truncation requirements.

TDPEXCC command

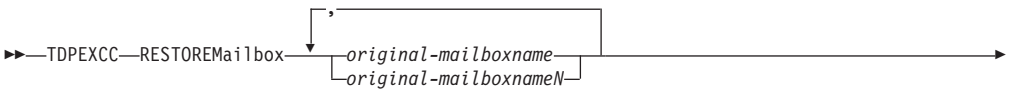

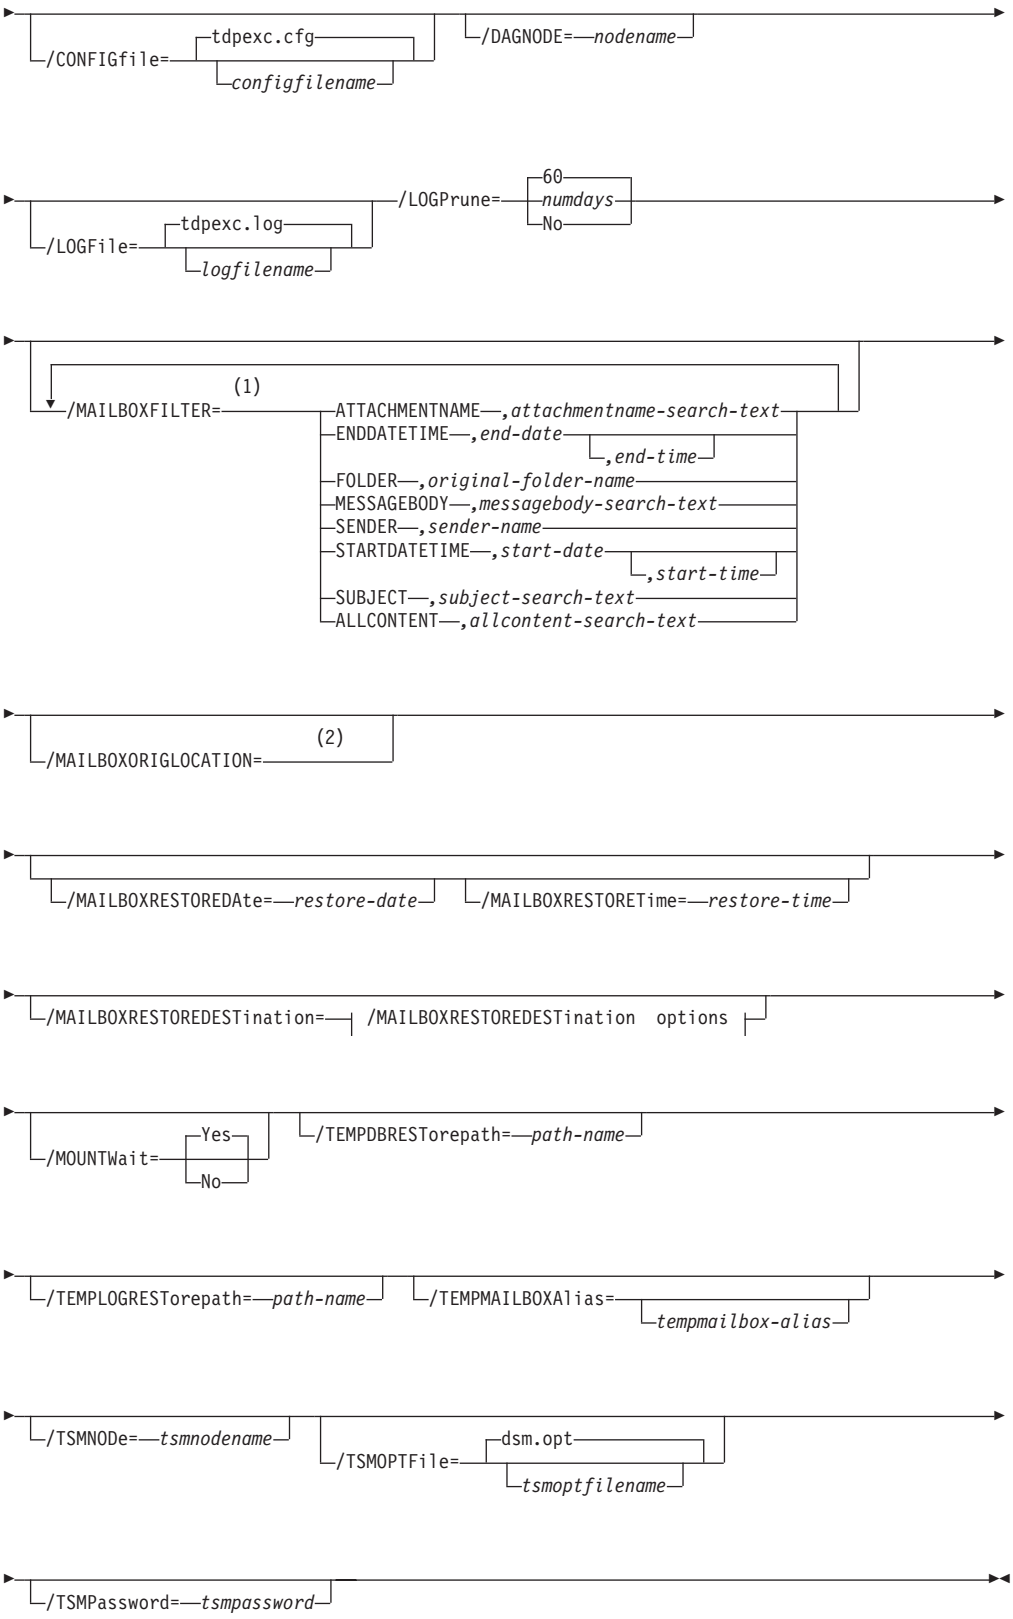
/CONFIGfile=
tdpexc.cfg
configfilename
/DAGNODE=
nodename
/LOGFile=
tdpexc.log
logfilename
/LOGPrune=
60
numdays
No
/MAILBOXFILTER=
(1)
ATTACHMENTNAME,attachmentname-search-text
ENDDATETIME,end-date
,end-time
FOLDER,original-folder-name
MESSAGEBODY,messagebody-search-text
SENDER,sender-name
STARTDATETIME,start-date
,start-time
SUBJECT,subject-search-text
ALLCONTENT,allcontent-search-text
/MAILBOXORIGLOCATION=
(2)
/MAILBOXRESTOREDATe=
restore-date
/MAILBOXRESTORETime=
restore-time
/MAILBOXRESTOREDESTination=
/MAILBOXRESTOREDESTination options
/MOUNTWait=
Yes
No
/TEMPDBRESTorepath=
path-name
/TEMPLOGRESTorepath=
path-name
/TEMPMAILBOXAlias=
tempmailbox-alias
/TSMNODe=
tsmnodename
/TSMOPTFile=
dsm.opt
tsmoptfilename
/TSMPassword=
tsmpassword

/MAILBOXRESTOREDESTination options:

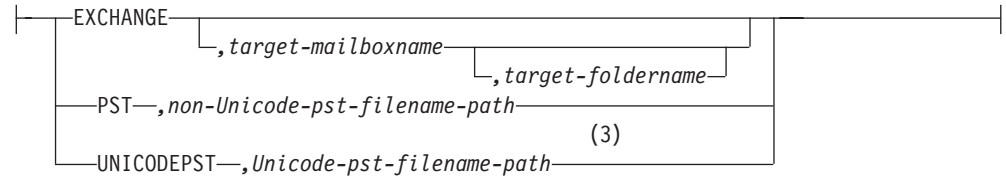

Notes:

1 You can specify the **/MAILBOXFILTER** parameter multiple times however, you must specify each **/MAILBOXFILTER** subparameter only once.

2 server-name,db-name

3 If you restore personal storage folders (`.pst` files), there are two options: **Restore Mail to Unicode PST file** and **Restore Mail to non-Unicode PST** file. Unicode `.pst` files can store messages in multiple languages, and are not limited to 2 GB of data. For non-Unicode `.pst` files, the file size must be less than 2 GB.

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 or the mailbox-display name. The *original-mailboxname* parameter is required.

To specify more than one name, separate them by commas.

If any mailbox name contains commas or blank spaces, 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 configuration file that contains the values for the Data Protection for Microsoft Exchange configuration options. See "Set command" on page 173 for details about the contents of the file.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the Data Protection for Microsoft Exchange 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 Tivoli Storage Manager 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 from making too many backups of the same database.

/LOGFile=*logfilename*

Use the **/logfile** parameter to specify the name of the activity log file that is generated by Data Protection for Microsoft Exchange.

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 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 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.

Attention: Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=*numdays* **| No | 60**

Use the **/logprune** parameter to disable log pruning or to explicitly request that the log is pruned 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 Microsoft Exchange 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 is already done for the day.

Changes to the value of the **timeformat** or **dateformat** parameter can result in pruning the log file unintentionally. If the value of the **timeformat** or **dateformat** parameter changes before you issue a Data Protection for Microsoft Exchange command that might prune the log file, select 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 and *logfilename* setting.

/MAILBOXFILTER=ATTACHMENTNAME | ENDDATETIME | FOLDER | MESSAGEBODY | SENDER | STARTDATETIME | SUBJECT | ALLCONTENT

Use the **/MAILBOXFILTER** parameter to specify filters to restrict what mailbox data is restored. When restoring to a Unicode `.pst` file, except for the `FOLDER` and `ALLCONTENT` filters, the filters are ignored.

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:

```
tdpexcc.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.

Attention: 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.

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 prObe.

Enclose the *messagebody-search-text* variable in double quotation marks.

Attention: 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.

Attention: 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 was 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 was at the time of backup.

db-name

The name of the database where the mailbox was at the time of backup.

The **/mailboxoriglocation** parameter is only necessary if the mailbox to be restored from is moved or deleted since the time of the backup, and no mailbox location history is available. This parameter is case-sensitive. Data Protection for Microsoft Exchange 6.1 (and later) maintains mailbox location history.

A **restoremailbox** operation from a backup that is processed by Data Protection for Microsoft Exchange before version 6.1 fails if the **/mailboxoriglocation** parameter is not specified for mailboxes that meet one or both of the following the conditions:

- The mailbox to be restored is moved. The mailbox is not in the same server and the same database where the mailbox was at the time of the backup.
- The mailbox to be restored was deleted and the restore destination is to an alternate mailbox or to a `.pst` file.

For example:

```
TDPEXCC RESTOREMAILBOX annjones /MAILBOXORIGLOCATION=serv1,sg1,mbdb1
/MAILBOXRESTOREDAte=02/21/2013
/MAILBOXRESTOREDESTination=PST,c:\team99\rcvr.pst
```

```
TDPEXCC RESTOREMAILBOX johngrimshawe /MAILBOXORIGLOCATION=serv1,mbdb1
/MAILBOXRESTOREDAte=03/06/2013
/MAILBOXRESTOREDESTination=PST,c:\team54\rcvr.pst
```

The deleted mailbox is to be re-created.

/MAILBOXRESTOREDAte=*restore-date*

Use the **/mailboxrestoredate** parameter with or without the **/mailboxrestoretime** parameter to establish a date and time to restore mailbox data from. A mailbox is restored from the earliest backup that was done after the date and time was established by the **/mailboxrestoredate** and the **/mailboxrestoretime** parameters.

The backup after the specified time is selected because, if an earlier backup is selected, Data Protection for Microsoft Exchange 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 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 options file.

If neither *restore-date* nor *restore-time* is specified, then no date and time are 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 for the time and date, you use a format other than the default, it is ignored.

/MAILBOXRESTORETime=*restore-time*

Use the **/mailboxrestoretime** 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 was taken after the date and time was established by the **/mailboxrestoredate** and the **/mailboxrestoretime** parameters.

The backup after the specified time is selected because, if an earlier backup is selected, Data Protection for Microsoft Exchange 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 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 options file.

If neither *restore-date* nor *restore-time* is specified, then no date and time are 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 was done 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

Use the **/mailboxrestoredestination** parameter to specify the destination to restore the mailbox data to.

If you do not specify the **/mailboxrestoredestination** parameter, the default 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 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.

If you specify the **/mailboxrestoredestination** EXCHANGEoption without specifying any variables, /mailboxrestoredestination=EXCHANGE, 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 choice of options restores mailbox data into a subfolder (designated by each original mailbox-alias) of the specified target folder in an active mailbox. In each subfolder are the folders (from the corresponding original mailbox) that contain the restored mailbox messages.

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 target mailbox can be the original mailbox or an alternate mailbox.

In the target mailbox, the specified folder (in the target mailbox) contains a subfolder (designated by the original-mailbox alias name). In the subfolder are sub-subfolders that contain the restored mailbox messages. These sub-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.

target-foldername

The *target-foldername* variable specifies the mailbox folder in the target mailbox to restore mailbox messages to. If you specify the *target-mailboxname* variable and the target mailbox is not the original mailbox, you must specify a folder name.

If the mailbox folder specified by the *target-foldername* variable does not exist in the target mailbox, a folder with the *target-foldername* is created in the target mailbox.

The target folder contains one subfolder for each original-mailbox that is restored (designated by each original-mailbox alias). In each subfolder are the folders from the original mailbox that contain the restored mailbox messages. 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). In each subfolder are the folders from the original mailbox that contain the restored mailbox messages.

For example, this **restoremailbox** operation restores mailboxes "andrew baker" and "sally wood" to the folder "previous_acctmngr" in the target mailbox "mary brown":

```
restoremailbox "andrew baker","sally wood"
/mailboxrestoredest=exchange,"mary brown",previous_acctmngr
```

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_acctmngr     (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 Accnts (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
  ```

  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 top-level folder name is the mailbox-alias-name, with the rest of the mailbox folders below it.

- You can specify *non-Unicode-pst-filename-path* as a directory path to have Tivoli Storage FlashCopy Manager for Exchange 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 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
  ```

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. 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), Tivoli Storage FlashCopy Manager for Exchange tries to convert it to a UNC path for you, but it may 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\vpo.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 top-level folder name is the mailbox-alias-name, with the rest of the mailbox folders below it.

- You can specify *Unicode-pst-filename-path* as a directory path to have Tivoli Storage FlashCopy Manager for Exchange 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 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
```

If the *Unicode-pst-filename-path* variable includes spaces, enclose the entire *Unicode-pst-filename-path* variable entry in double quotation marks. For example:

```
TDPEXCC RESTOREMAILBOX "george clark"
  /mailboxrestoredestination=UNICODEPST,"c:\mb\dept54\access group\"
```

/MOUNTWait=Yes | No

Use the **/mountwait** parameter to specify whether Data Protection for Microsoft Exchange waits for removable media to mount (such as tapes or CDs) or to stop the current operation. This situation occurs when the Tivoli Storage Manager 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.

/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 configuration file is tdpexc.cfg. If the TEMPDBRESTOREPATH value does not exist in the Data Protection for Microsoft Exchange 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:

```
TDPEXCC RESTOREMAILBOX richgreene
  /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, place the current active-transaction logger on a different physical device from the paths that are specified by the values of the **/tempdbrestorepath** parameter and the **/tempdbrestorepath** parameter. The paths that are specified by the values of the **/tempdbrestorepath** parameter and the **/tempdbrestorepath** parameter 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=*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 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 configuration file, the TEMP environment variable value is used.

Attention:

- 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.

Tip: For better performance, put the current active-transaction logger 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.

Restriction: Do not specify double-byte characters (DBCS) within the temporary log-restore path.

/TSMNODe=*tsmnodename*

Use the *tsmnodename* variable to refer to the Tivoli Storage Manager node name that Data Protection for Microsoft Exchange uses to log on to the Tivoli Storage Manager server. You can store the node name in the Tivoli Storage Manager options file (`dsm.opt`). This parameter overrides the value in the Tivoli Storage Manager 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 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 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 Tivoli Storage Manager password that Data Protection for Microsoft Exchange uses to log on to the Tivoli Storage Manager server. If you specified **PASSWORDACCESS** `GENERATE` in the Data Protection for Microsoft Exchange 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 Tivoli Storage Manager password the first time Data Protection for Microsoft Exchange connects to the Tivoli Storage Manager 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 Tivoli Storage Manager password that Data Protection for Microsoft Exchange uses to log on to the Tivoli Storage Manager server can be up to 63 characters in length.

Set command

To set the Data Protection for Microsoft Exchange configuration parameters in a configuration file, use the **set** command.

The values that you set are saved in a Data Protection for Microsoft Exchange configuration file. The default file is `tdpexc.cfg`. Configuration values can also be set in the Data Protection Properties window in the Management Console.

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 Microsoft Exchange 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

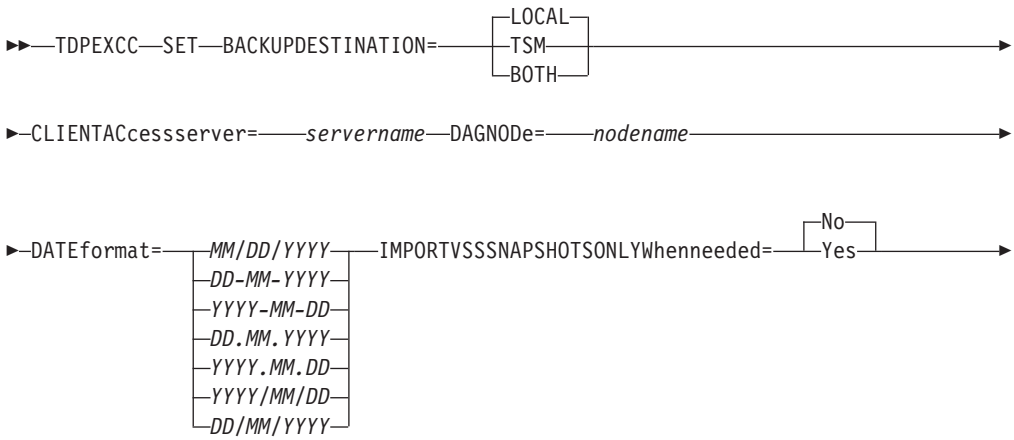

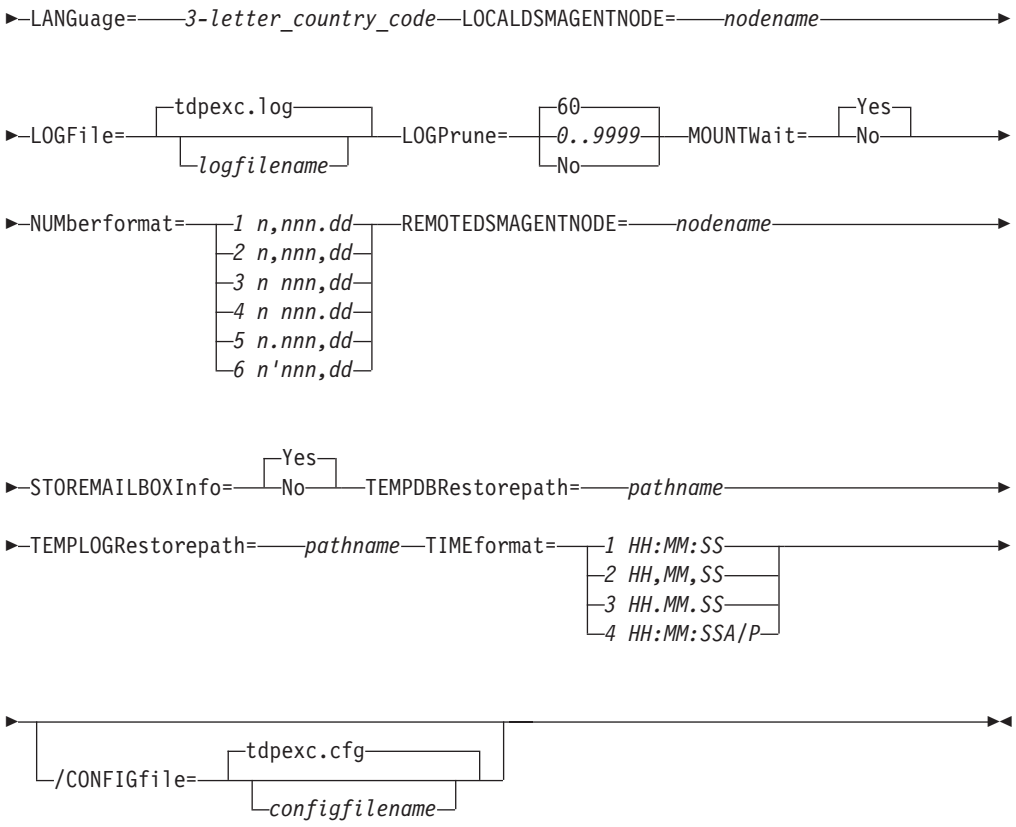

Set positional parameters

Positional parameters immediately follow the **set** command and precede the optional parameters.

The following positional parameters specify the values in the Data Protection for Microsoft Exchange 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 Tivoli Storage Manager 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 Tivoli Storage Manager 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 Tivoli Storage

Manager 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 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 Microsoft Exchange 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 Microsoft Exchange.
- Specify a new log file with the **/logfile** parameter.

IMPORTVSSSNAPSHOTSONLYWhenneeded=Yes|No
By default, the parameter is set to `No`. This default setting means that local persistent VSS snapshots are automatically imported to the Windows system where the snapshots are created. By importing the VSS snapshots only when needed, the snapshots are imported to a host for FlashCopy Manager operations. To automatically import local persistent snapshots to the Windows system where the snapshots are created, set the parameter to `Yes`.

LANGuage=*language*
Specify the three-character code of the language that you want to use to display messages:

CHS Simplified Chinese

CHT Traditional Chinese

DEU Standard German

ENU (Default) American English

ESP Standard Spanish

FRA Standard French

ITA Standard Italian

JPN Japanese

KOR Korean

PTB Brazilian Portuguese

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 Microsoft Exchange. The Data Protection for Microsoft Exchange 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 Microsoft Exchange installation directory.

LOGPrune=*numdays*|**No**
Use the **LOGPrune** positional parameter to disable log pruning or to set log pruning parameters. By default, log pruning is enabled and done once per day. The *numdays* variable represents the number of days to save log entries. You can specify a value of `No` or `0 - 9999`. By default, `60` days of log entries are saved in the pruning process.

MOUNTWait=Yes|No
Use the **MOUNTWait** positional parameter to specify whether Data Protection for Microsoft Exchange wait for removable media to mount (such as tapes or CDs) or to stop the current operation. This situation occurs when the Tivoli Storage Manager 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 Microsoft Exchange to wait until all initial volumes of any required removable media are made available to the Tivoli Storage Manager server before the command completes.

Specify `No` for Data Protection for Microsoft Exchange to end the command (if removable media are required). An error message displays.

NUMberformat=*fmtnum*
Use the **NUMberformat** 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 Tivoli Storage Manager server storage during offloaded backups.

STOREMAILBOXInfo=Yes|No

The **STOREMAILBOXInfo** 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 Microsoft Exchange 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 **STOREMAILBOXInfo** to No. When you set the option for **STOREMAILBOXInfo** 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 **Restoremailbox** 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 **TEMPDBRESTorepath** 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:

```
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, the current active-transaction logger should 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 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

4 HH:MM:SSA/P

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 Microsoft Exchange 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 Microsoft Exchange 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 Microsoft Exchange log file instead of the default Data Protection for Microsoft Exchange log file, `tdpexc.log`, in the directory where Data Protection for Microsoft Exchange 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 0.0
(C) Copyright IBM Corporation 1998, 2013. All rights reserved.

ACO5054I 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 0.0
(C) Copyright IBM Corporation 1998, 2013. 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 Tivoli Storage FlashCopy Manager for Exchange, 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

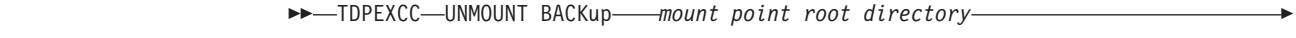

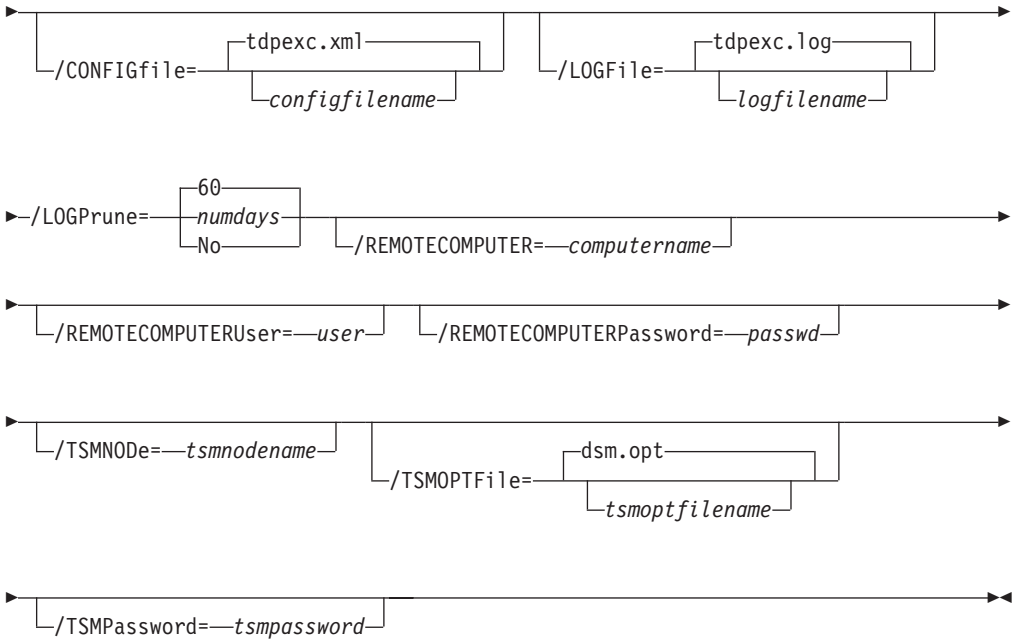

Unmount Backup positional parameter

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=*logfilename*
Use the **/logfile** parameter to specify the name of the activity log file that is generated by Tivoli Storage FlashCopy Manager for Exchange. 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 Tivoli Storage FlashCopy Manager for Exchange installation directory.

If the *logfilename* 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**

Use the **/logprune** parameter to disable log pruning or to explicitly request that the log is to be pruned 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 GUI or the **update config** 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 is done for the day.

Changes to the value of the **timeformat** or **dateformat** parameter can result in pruning the log file unintentionally. If the value of the **timeformat** or **dateformat** parameter is changed, before you issue a Tivoli Storage FlashCopy Manager for Exchange 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.

/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 Tivoli Storage Manager node name that Tivoli Storage FlashCopy Manager uses to log on to the Tivoli Storage Manager server. You can store the node name in the Tivoli Storage Manager options file (`dsm.opt`). This parameter overrides the value in the Tivoli Storage Manager 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 Tivoli Storage Manager options file.

The file name can include a fully qualified path name. If no path is specified, the directory where Tivoli Storage FlashCopy Manager 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 Tivoli Storage Manager password that Tivoli Storage FlashCopy Manager uses to log on to the Tivoli Storage Manager server. If you specified **PASSWORDACCESS** `GENERATE` in the Tivoli Storage FlashCopy Manager 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 Tivoli Storage Manager password the first time Tivoli Storage FlashCopy Manager connects to the Tivoli Storage Manager 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 Tivoli Storage Manager password that Tivoli Storage FlashCopy Manager uses to log on to the Tivoli Storage Manager server can be up to 63 characters in length.

Transitioning Exchange Server backups from Tivoli Storage FlashCopy Manager to Tivoli Storage Manager

Configure Tivoli Storage FlashCopy Manager so that you can access both a local and Tivoli Storage Manager server. Use this approach if you move to a Tivoli Storage Manager environment and want to continue interacting with locally managed snapshots until policy marks them for expiration.

About this task

To configure the Tivoli Storage FlashCopy Manager, use the **Standalone** and Tivoli Storage Manager server configuration wizards from the Tivoli Storage FlashCopy Manager. To interact with a Tivoli Storage Manager server, run the **TSM** configuration wizard. To interact with a Tivoli Storage FlashCopy Manager 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.

Some command examples that are provided here are formatted on multiple lines. Issue each command on a single line.

Completing these tasks on the Tivoli Storage Manager server

About this task

Coordinate efforts with your Tivoli Storage Manager server administrator to get these tasks completed:

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 Tivoli Storage FlashCopy Manager stand-alone environment. Use the GUI or the command-line interface of Data Protection for Microsoft Exchange to retrieve this information.
2. Register your Data Protection for Microsoft Exchange node name and password with the Tivoli Storage Manager **register node** command. For example:

   ```
   register node DPnodename DPpassword
   ```
3. If not already defined in the Tivoli Storage Manager server, register the Tivoli Storage Manager backup-archive client node name and password for the workstation where the Exchange server is installed. For example:

   ```
   register node BAnodename BApassword
   ```
4. Define the proxy node relationship for the Target Node and agent nodes with the Tivoli Storage Manager **grant proxynode** command. For example:

   ```
   grant proxynode target=DP agent=BAnodename
   ```

Completing these tasks on the workstation that running the Exchange Server

Procedure

1. In the directory where the Data Protection for Microsoft Exchange is installed, make a copy of the options file named `dsm.opt`. After you begin by using the Tivoli Storage Manager server, the copy is used for access to the Tivoli Storage FlashCopy Manager stand-alone environment. One method of making the copy is to start the Exchange command-line prompt from the Tivoli Storage FlashCopy Manager Snapin: In the Tivoli Storage FlashCopy Manager 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 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 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 Tivoli Storage FlashCopy Manager.
   - You are updating the policy bindings to reflect changes in your policy specifications for your Tivoli Storage Manager server usage.

3. In the Tivoli Storage Manager 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 `TCPSERVERADDRESS` line. Replace `FLASHCOPYMANAGER` with the IP address of the Tivoli Storage Manager server. For example:

   ```
   TCPServeraddress 9.52.170.67
   ```

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

5. To access the Tivoli Storage FlashCopy Manager stand-alone environment during the transition period, open a Windows command prompt and change the directory to the Tivoli Storage Manager backup-archive client installation directory. This path is the default:

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

   Create an alternative Windows service for the Tivoli Storage Manager 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, refer to the information about the client service configuration utility in the *Tivoli Storage Manager Windows Backup-Archive Clients Installation and User's Guide*.

6. Create an alternate Windows service for the Tivoli Storage Manager 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 installation directory. Add this line:

   ```
   HTTPPORT 1583
   ```

8. Start the alternate Tivoli Storage Manager Client Acceptor service:

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

9. Stop and restart the original Tivoli Storage Manager 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 Tivoli Storage Manager server environment, phase out the remaining backups that are created in the Tivoli Storage FlashCopy Manager stand-alone environment. You can choose between two ways of achieving the phase-out:

    a. In the Tivoli Storage FlashCopy Manager 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 Backup/Restore GUI. Information about how to start the GUI is in the section that describes how to access the Tivoli Storage FlashCopy Manager 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 **delete backup** command, or the **Delete Backup** (right mouse-click menu option) in the GUI **Restore** tab.

11. To access the Tivoli Storage FlashCopy Manager stand-alone environment:
    a. Open the Automate tab to access the integrated command-line prompt.
    b. Start Tivoli Storage FlashCopy Manager 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 GUI (from the Command Line prompt) by issuing the GUI invocation command, for example:

    ```
    tdpexc /tsmoptfile=dsm_local.opt
    /configfile=tdpexc_local.cfg
    ```

12. If necessary, start the Tivoli Storage FlashCopy Manager 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 Tivoli Storage FlashCopy Manager 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
```

Appendix A. Frequently asked questions

Answers related to frequently asked questions about Data Protection for Microsoft Exchange are provided.

How do I compress my Data Protection for Microsoft Exchange backups?

Use the **compression** option to instruct the Tivoli Storage Manager API to compress data before the data is sent to the Tivoli Storage Manager 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 available in the client documentation before you compress your data.

For more information about the **compression** option, see "Specifying Data Protection for Exchange options" on page 35.

How do I encrypt my Data Protection for Microsoft Exchange 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 available in the client documentation before you encrypt your databases.

For more information about the **enableclientencryptkey** and **encryptiontype** options, see "Specifying Data Protection for Exchange options" on page 35.

How do I deduplicate my Data Protection for Microsoft Exchange backups?

Use the **deduplication** option to enable client-side data deduplication. Client-side data deduplication is used by the Tivoli Storage Manager API to remove redundant data during backup processing before the data is transferred to the Tivoli Storage Manager 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 available in the client documentation before you encrypt your databases.

For more information about the **deduplication** option, see "Specifying Data Protection for Exchange options" on page 35.

What must I do before I complete Data Protection for Microsoft Exchange mailbox-level and mailbox item-level restores?

Review these prerequisites before you complete Data Protection for Microsoft Exchange mailbox restore tasks:

"Security requirements" on page 17

"Prerequisites for Data Protection for Microsoft Exchange mailbox restore tasks" on page 159

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 mailbox restore operations on my Exchange Server?

When you use the configuration wizard in the Management Console (MMC) GUI to configure Data Protection for Microsoft Exchange, 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 mailbox restore operation really do mailbox-level and mailbox item-level restores?

When a mailbox restore operation is initiated, Data Protection for Microsoft Exchange completes the following actions:

1. Starts a session with the Tivoli Storage Manager server.
2. Queries the Tivoli Storage Manager 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.

How do I use Data Protection for Microsoft Exchange to restore a deleted mailbox or items from a deleted mailbox?

Review "Restoring a deleted mailbox or items from a deleted mailbox" on page 80

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. For more information, see "Restoring a Database Availability Group database copy" on page 85.

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. For more information, see "Restoring VSS backups into alternate locations" on page 13.

Are VSS restores restored into the recovery database?

Yes, VSS restores can be restored into the recovery database or into an alternate database. For more information, see "VSS restore considerations" on page 76 and "Restoring VSS backups into alternate locations" on page 13.

Why is the VSS instant restore failing over to a VSS fast restore?

A failover can occur if the Exchange data is on storage subsystems that are not supported for VSS instant restore. For more information, see "VSS instant restore" on page 12.

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 V7000, XIV, or DS8000 storage subsystem is required to complete VSS instant restores. For more information, see "VSS instant restore" on page 12.

Now that I am completing VSS operations, why are there so many active backups?

Tivoli Storage Manager policy manages VSS backups on local shadow volumes and on Tivoli Storage Manager server storage. This management allows for different policies, which can lead to an increase in the number of active backups. For more information, see "How Tivoli Storage Manager server policy affects Data Protection for Exchange" on page 28 and "Back up to Tivoli Storage Manager storage versus back up to local shadow volumes" on page 20.

Can I use UNC drive letters with VSS offloaded backups?

No, Data Protection for Microsoft Exchange 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 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 Tivoli Storage Manager 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 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 backups. For more information, see "Restoring individual mailbox and mailbox item-level data" on page 78.

How do I set up my policy settings for Data Protection for Microsoft Exchange?

For more information, see the following topics for more information about Data Protection for Microsoft Exchange policy settings:

- "How Tivoli Storage Manager server policy affects Data Protection for Exchange" on page 28

- "Specifying Data Protection for Exchange options" on page 35

Can I restore a Data Protection for Microsoft Exchange database backup to flat files without using an Exchange Server? Can I restore a Data Protection for Microsoft Exchange database backup to a flat file without interrupting the Data Protection for Microsoft ExchangeServer?

Yes, use the **restorefiles** command. For more information, see "Restorefiles command" on page 153.

How do I schedule Data Protection for Microsoft Exchange backups?

You can schedule Data Protection for Microsoft Exchange backups by using the Tivoli Storage Manager backup-archive client scheduler or the Management Console scheduler.

What do I do if I get an "unknown Exchange API error" when I run Data Protection for Microsoft Exchange?

For more information about what to do when you encounter a problem, see Chapter 7, "Troubleshooting," on page 95.

How do I know whether my backup ran successfully?

A message displays that states the backup completed successfully. In addition, messages from the TDPExchange service for backup start and backup finish are displayed in the Event Viewer. The Task Manager in the Management Console provides centralized information about the status of your tasks. Processing information is also available in the following files:

- Data Protection for Microsoft Exchange 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.
- Tivoli Storage Manager server activity log

  Data Protection for Microsoft Exchange logs information about backup and restore commands to the Tivoli Storage Manager server activity log. A Tivoli Storage Manager administrator can view this log for you if you do not have a Tivoli Storage Manager administrator user ID and password.
- Tivoli Storage Manager 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 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, Tivoli Storage Manager, 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 Tivoli Storage Manager server error message is displayed: "ANR9999D snmode.c(xxxx): Error validating inserts, and so on"

You do not have to do anything as this message can be ignored. Installing a later version of Tivoli Storage Manager server prevents this message from being displayed.

What authority is needed to complete a Data Protection for Microsoft Exchange backup and restore?

For more information about the required authority to complete Data Protection for Microsoft Exchange backup and restore tasks, see "Security requirements" on page 17.

Do I use the same `nodename` as used by my backup-archive client?

No, you must use different node names. For more information, see "Specifying Data Protection for Exchange options" on page 35.

How do I set up LAN-free to back up Data Protection for Microsoft Exchange over my SAN?

See "LAN-free data movement" on page 112.

Can I run Data Protection for Microsoft Exchange with multiple backup sessions?

For more information, see "Backup strategies" on page 18.

Appendix B. Tivoli support information

You can find support information for Tivoli and other IBM products from various sources.

From the IBM Support Portal at http://www.ibm.com/support/entry/portal/, you can select the products that you are interested in and search for a wide variety of relevant information.

Communities and other learning resources

In addition to product documentation, many forms of assistance are available to help you get started as you deploy and use the Tivoli Storage Manager family of products. These resources can also help you to solve problems that you might have.

You can use forums, wikis, and other social media tools to ask questions, talk to experts, and learn from others.

User groups

Tivoli Global Storage Virtual User Group

Access this user group at http://www.tivoli-ug.org/storage.

This group makes it possible for individuals from many different industries and types of organizations to share information and work directly with the IBM product experts. Local chapters also exist where members meet in person to share experiences and hear from guest speakers.

ADSM.ORG

Access this mailing list at http://adsm.org.

This independently managed Storage Management discussion forum started when Tivoli Storage Manager was known as ADSTAR Distributed Storage Manager (ADSM). The members of this forum have many years of experience with Tivoli Storage Manager in almost every type of IT environment.

To subscribe to the forum, send an email to listserv@vm.marist.edu. The body of the message must contain the following text: `SUBSCRIBE ADSM-L` *your_first_name your_family_name*.

Tivoli Storage Manager community on Service Management Connect

Access Service Management Connect at http://www.ibm.com/developerworks/servicemanagement. In the Storage Management community of Service Management Connect, you can connect with IBM in the following ways:

- Become involved with transparent development, an ongoing, open engagement between users and IBM developers of Tivoli products. You can access early designs, sprint demonstrations, product roadmaps, and prerelease code.
- Connect one-on-one with the experts to collaborate and network about Tivoli and the Tivoli Storage Manager community.
- Read blogs to benefit from the expertise and experience of others.

- Use wikis and forums to collaborate with the broader user community.

Tivoli Storage Manager wiki on developerWorks®

Access this wiki at https://www.ibm.com/developerworks/servicemanagement/sm/index.html.

Find the latest best practices, white papers, and links to videos and other resources. When you log on, you can comment on content, or contribute your own content.

Tivoli Support Technical Exchange

Find information about upcoming Tivoli Support Technical Exchange webcasts at http://www.ibm.com/software/sysmgmt/products/support/supp_tech_exch.html. Replays of previous webcasts are also available.

Learn from technical experts who share their knowledge and then answer your questions. The sessions are designed to address specific technical issues and provide in-depth but narrowly focused training.

Other social media sites

LinkedIn

You can join groups on LinkedIn, a social media site for professionals. For example:

- **Tivoli Storage Manager Professionals**: http://www.linkedin.com/groups/Tivoli-Storage-Manager-Professionals-54572
- **TSM**: http://www.linkedin.com/groups?gid=64540

Twitter

Follow `@IBMStorage` on Twitter to see the latest news about storage and storage software from IBM.

Tivoli education resources

Use these education resources to help you increase your Tivoli Storage Manager skills:

Tivoli Education and Certification website

View available education at http://www.ibm.com/software/tivoli/education.

Use the Search for Training link to find local and online offerings of instructor-led courses for Tivoli Storage Manager.

Education Assistant

Access resources at http://publib.boulder.ibm.com/infocenter/ieduasst/tivv1r0/index.jsp.

Scroll to view the list of available training videos. Recorded product demonstrations are also available on a YouTube channel.

Searching knowledge bases

If a problem occurs while you are using one of the Tivoli Storage Manager family of products, you can search several knowledge bases.

Begin by searching the Tivoli Storage Manager Information Center at http://pic.dhe.ibm.com/infocenter/tsminfo/v7r1. Within the information center, you can enter words, phrases, or message numbers in the **Search** field to find relevant topics.

Searching the Internet

If you cannot find an answer to your question in the Tivoli Storage Manager information center, search the Internet for the information that might help you resolve the problem.

To search multiple Internet resources, go to the IBM support website at http://www.ibm.com/support/entry/portal/. You can search for information without signing in.

Sign in using your IBM ID and password if you want to customize the site based on your product usage and information needs. If you do not already have an IBM ID and password, click **Sign in** at the top of the page and follow the instructions to register.

From the support website, you can search various resources:

- IBM technotes.
- IBM downloads.
- IBM Redbooks® publications.
- IBM Authorized Program Analysis Reports (APARs). Select the product and click **Downloads** to search the APAR list.

Using IBM Support Assistant

IBM Support Assistant is a complimentary software product that can help you with problem determination. It is available for some Tivoli Storage Manager and Tivoli Storage FlashCopy Manager products.

IBM Support Assistant helps you gather support information when you must open a problem management record (PMR), which you can then use to track the problem. The product-specific plug-in modules provide you with the following resources:

- Support links
- Education links
- Ability to submit problem management reports

You can find more information and download the IBM Support Assistant web page at http://www.ibm.com/software/support/isa.

You can also install the stand-alone IBM Support Assistant application on any workstation. You can then enhance the application by installing product-specific plug-in modules for the IBM products that you use. Find add-ons for specific products at http://www.ibm.com/support/docview.wss?uid=swg27012689.

Finding product fixes

A product fix to resolve a software problem might be available from the IBM software support website.

Procedure

Determine what fixes are available by checking the IBM software support website at http://www.ibm.com/support/entry/portal/.

If you previously customized the site based on your product usage:

1. Click the link for the product, or a component for which you want to find a fix.
2. Click **Downloads**, and then click **Search for recommended fixes**.

If you have not previously customized the site:

Click **Downloads** and search for the product.

Receiving notification of product fixes

You can receive notifications about fixes, flashes, upgrades, and other news about IBM products.

Procedure

1. From the support page at http://www.ibm.com/support/entry/portal/, click **Sign in** and sign in using your IBM ID and password. If you do not have an ID and password, click **register now** and complete the registration process.
2. Click **Manage all my subscriptions** in the Notifications pane.
3. Click the **Subscribe** tab, and then click **Tivoli**.
4. Select the products for which you want to receive notifications and click **Continue**.
5. Specify your notification preferences and click **Submit**.

Contacting IBM Software Support

You can contact IBM Software Support if you have an active IBM subscription and support contract, and if you are authorized to submit problems to IBM.

Procedure

1. Ensure that you have completed the following prerequisites:
   a. Set up a subscription and support contract.
   b. Determine the business impact of the problem.
   c. Describe the problem and gather background information.
2. Follow the instructions in "Submitting the problem to IBM Software Support" on page 198.

Setting up and managing support contracts

You can set up and manage your Tivoli support contracts by enrolling in IBM Passport Advantage. The type of support contract that you need depends on the type of product you have.

Procedure

Enroll in IBM Passport Advantage in one of the following ways:

- **Online:** Go to the Passport Advantage website at http://www.ibm.com/software/lotus/passportadvantage/, click **How to enroll**, and follow the instructions.
- **By telephone:** For critical, system-down, or high-severity issues, you can call 1-800-IBMSERV (1-800-426-7378) in the United States. For the telephone number to call in your country, go to the IBM Software Support Handbook web page at http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html and click **Contacts**.

Determining the business impact

When you report a problem to IBM, you are asked to supply a severity level. Therefore, you must understand and assess the business impact of the problem you are reporting.

Severity level	Description
Severity 1	**Critical** business impact: You are unable to use the program, resulting in a critical impact on operations. This condition requires an immediate solution.
Severity 2	**Significant** business impact: The program is usable but is severely limited.
Severity 3	**Some** business impact: The program is usable with less significant features (not critical to operations) unavailable.
Severity 4	**Minimal** business impact: The problem causes little impact on operations, or a reasonable circumvention to the problem has been implemented.

Describing the problem and gathering background information

When explaining a problem to IBM, it is helpful to be as specific as possible. Include all relevant background information so that IBM Software Support specialists can help you solve the problem efficiently.

To save time, know the answers to these questions:

- What software versions were you running when the problem occurred?
- Do you have logs, traces, and messages that are related to the problem symptoms? IBM Software Support is likely to ask for this information.
- Can the problem be re-created? If so, what steps led to the failure?
- Have any changes been made to the system? For example, hardware, operating system, networking software, and so on.
- Are you using a workaround for this problem? If so, be prepared to explain it when you report the problem.

Submitting the problem to IBM Software Support

You can submit the problem to IBM Software Support online or by telephone.

Online

Go to the IBM Software Support website at http://www.ibm.com/support/entry/portal/Open_service_request/Software/Software_support_(general). Sign in to access IBM Service Requests and enter your information into the problem submission tool.

By telephone

For critical, system-down, or severity 1 issues, you can call 1-800-IBMSERV (1-800-426-7378) in the United States. For the telephone number to call in your country, go to the IBM Software Support Handbook web page at http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html and click **Contacts**.

Appendix C. Accessibility features for the Tivoli Storage Manager product family

Accessibility features help users who have a disability, such as restricted mobility or limited vision to use information technology products successfully.

Accessibility features

The IBM Tivoli Storage Manager family of products includes the following accessibility features:

- Keyboard-only operation using standard operating-system conventions
- Interfaces that support assistive technology such as screen readers

The command-line interfaces of all products in the product family are accessible.

Tivoli Storage Manager Operations Center provides the following additional accessibility features when you use it with a Mozilla Firefox browser on a Microsoft Windows system:

- Screen magnifiers and content zooming
- High contrast mode

The Operations Center and the Tivoli Storage Manager Server can be installed in console mode, which is accessible.

The Tivoli Storage Manager Information Center is enabled for accessibility. For information center accessibility information, see "Accessibility features in the information center" (http://pic.dhe.ibm.com/infocenter/tsminfo/v7r1/topic/com.ibm.help.ic.doc/iehs36_accessibility.html).

Vendor software

The Tivoli Storage Manager 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 the accessibility information about its products.

IBM and accessibility

See the IBM Human Ability and Accessibility Center (http://www.ibm.com/able) for information about the commitment that IBM has to accessibility.

Notices

This information was developed for products and services offered in the U.S.A.

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
Armonk, NY 10504-1785
U.S.A.

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

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:
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 states 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 want 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 Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758
U.S.A.

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 information 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.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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 is for planning purposes only. The information herein is subject to change before the products described become available.

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_.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

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 http://www.ibm.com/legal/copytrade.shtml.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

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 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

This glossary provides terms and definitions for Tivoli Storage Manager, Tivoli Storage FlashCopy Manager, and associated products.

The following cross-references are used in this glossary:

- *See* refers you from a nonpreferred term to the preferred term or from an abbreviation to the spelled-out form.
- *See also* refers you to a related or contrasting term.

For other terms and definitions, see the IBM Terminology website at www.ibm.com/software/globalization/terminology.

A

absolute mode
In storage management, a backup copy-group mode that specifies that a file is considered for incremental backup even if the file has not changed since the last backup. See also mode, modified mode.

access control list (ACL)
In computer security, a list associated with an object that identifies all the subjects that can access the object and their access rights.

access mode
An attribute of a storage pool or a storage volume that specifies whether the server can write to or read from the storage pool or storage volume.

ACK See acknowledgment.

acknowledgment (ACK)
The transmission of acknowledgment characters as a positive response to a data transmission.

ACL See access control list.

activate
To validate the contents of a policy set and then make it the active policy set.

active-data pool
A named set of storage pool volumes that contain only active versions of client backup data. See also server storage, storage pool, storage pool volume.

active file system
A file system to which space management has been added. With space management, tasks for an active file system include automatic migration, reconciliation, selective migration, and recall. See also inactive file system.

active policy set
The activated policy set that contains the policy rules currently in use by all client nodes assigned to the policy domain. See also policy domain, policy set.

active version
The most recent backup copy of a file stored. The active version of a file cannot be deleted until a backup process detects that the user has either replaced the file with a newer version or has deleted the file from the file server or workstation. See also backup version, inactive version.

activity log
A log that records normal activity messages that are generated by the server. These messages include information about server and client operations, such as the start time of sessions or device I/O errors.

adaptive subfile backup
A type of backup that sends only changed portions of a file to the server, instead of sending the entire file. Adaptive subfile backup reduces network traffic and increases the speed of the backup.

administrative client
A program that runs on a file server, workstation, or mainframe that administrators use to control and monitor the server. See also backup-archive client.

administrative command schedule
A database record that describes the planned processing of an administrative command during a specific time period. See also central scheduler, client schedule, schedule.

administrative privilege class
See privilege class.

administrative session
A period of time during which an administrator user ID communicates with a server to perform administrative tasks. See also client node session, session.

administrator
A person responsible for administrative tasks such as access authorization and content management. Administrators can also grant levels of authority to users.

agent node
A client node that has been granted proxy authority to perform operations on behalf of another client node, which is the target node.

aggregate
An object, stored in one or more storage pools, consisting of a group of logical files that are packaged together. See also logical file, physical file.

aggregate data transfer rate
A performance statistic that indicates the average number of bytes that were transferred per second while processing a given operation.

application client
A program that is installed on a system to protect an application. The server provides backup services to an application client.

archive
To copy programs, data, or files to another storage media, usually for long-term storage or security. See also retrieve.

archive copy
A file or group of files that was archived to server storage

archive copy group
A policy object containing attributes that control the generation, destination, and expiration of archived files. See also copy group.

archive-retention grace period
The number of days that the storage manager retains an archived file when the server is unable to rebind the file to an appropriate management class. See also bind.

association
The defined relationship between a client node and a client schedule. An association identifies the name of a schedule, the name of the policy domain to which the schedule belongs, and the name of a client node that performs scheduled operations.

audit To check for logical inconsistencies between information that the server has and the actual condition of the system. The storage manager can audit information about items such as volumes, libraries, and licenses. For example, when a storage manager audits a volume, the server checks for inconsistencies between information about backed-up or archived files that are stored in the database and the actual data that are associated with each backup version or archive copy in server storage.

authentication rule
A specification that another user can use to either restore or retrieve files from storage.

authority
The right to access objects, resources, or functions. See also privilege class.

authorization rule
A specification that permits another user to either restore or retrieve a user's files from storage.

authorized user
A user who has administrative authority for the client on a workstation. This user changes passwords, performs open registrations, and deletes file spaces.

AutoFS
See automounted file system.

automatic detection
A feature that detects, reports, and updates the serial number of a drive or library in the database when the path from the local server is defined.

automatic migration
The process that is used to automatically move files from a local file system to storage, based on options and settings that are chosen by a root user on a workstation. See also demand migration, threshold migration.

automounted file system (AutoFS)
A file system that is managed by an

automounter daemon. The automounter daemon monitors a specified directory path, and automatically mounts the file system to access data.

B

backup-archive client
A program that runs on a workstation or file server and provides a means for users to back up, archive, restore, and retrieve files. See also administrative client.

backup copy group
A policy object containing attributes that control the generation, destination, and expiration of backup versions of files. A backup copy group belongs to a management class. See also copy group.

backup retention grace period
The number of days the storage manager retains a backup version after the server is unable to rebind the file to an appropriate management class.

backup set
A portable, consolidated group of active versions of backup files that are generated for a backup-archive client.

backup set collection
A group of backup sets that are created at the same time and which have the same backup set name, volume names, description, and device classes. The server identifies each backup set in the collection by its node name, backup set name, and file type.

backup version
A file or directory that a client node backed up to storage. More than one backup version can exist in storage, but only one backup version is the active version. See also active version, copy group, inactive version.

bind To associate a file with a management class name. See also archive-retention grace period, management class, rebind.

C

cache To place a duplicate copy of a file on random access media when the server migrates a file to another storage pool in the hierarchy.

cache file
A snapshot of a logical volume created by Logical Volume Snapshot Agent. Blocks are saved immediately before they are modified during the image backup and their logical extents are saved in the cache files.

CAD See client acceptor daemon.

central scheduler
A function that permits an administrator to schedule client operations and administrative commands. The operations can be scheduled to occur periodically or on a specific date. See also administrative command schedule, client schedule.

client A software program or computer that requests services from a server. See also server.

client acceptor
A service that serves the Java applet for the web client to web browsers. On Windows systems, the client acceptor is installed and run as a service. On AIX®, UNIX, and Linux systems, the client acceptor is run as a daemon.

client acceptor daemon (CAD)
See client acceptor.

client domain
The set of drives, file systems, or volumes that the user selects to back up or archive data, using the backup-archive client.

client node
A file server or workstation on which the backup-archive client program has been installed, and which has been registered to the server.

client node session
A session in which a client node communicates with a server to perform backup, restore, archive, retrieve, migrate, or recall requests. See also administrative session.

client option set
A group of options that are defined on the server and used on client nodes in conjunction with client options files.

client options file
An editable file that identifies the server and communication method, and provides the configuration for backup, archive, hierarchical storage management, and scheduling.

client-polling scheduling mode
A method of operation in which the client queries the server for work. See also server-prompted scheduling mode.

client schedule
A database record that describes the planned processing of a client operation during a specific time period. The client operation can be a backup, archive, restore, or retrieve operation, a client operating system command, or a macro. See also administrative command schedule, central scheduler, schedule.

client/server
Pertaining to the model of interaction in distributed data processing in which a program on one computer sends a request to a program on another computer and awaits a response. The requesting program is called a client; the answering program is called a server.

client system-options file
A file, used on AIX, UNIX, or Linux system clients, containing a set of processing options that identify the servers to be contacted for services. This file also specifies communication methods and options for backup, archive, hierarchical storage management, and scheduling. See also client user-options file, options file.

client user-options file
A file that contains the set of processing options that the clients on the system use. The set can include options that determine the server that the client contacts, and options that affect backup operations, archive operations, hierarchical storage management operations, and scheduled operations. This file is also called the dsm.opt file. For AIX, UNIX, or Linux systems, see also client system-options file. See also client system-options file, options file.

closed registration
A registration process in which only an administrator can register workstations as client nodes with the server. See also open registration.

collocation
The process of keeping all data belonging to a single-client file space, a single client node, or a group of client nodes on a minimal number of sequential-access volumes within a storage pool. Collocation can reduce the number of volumes that must be accessed when a large amount of data must be restored.

collocation group
A user-defined group of client nodes whose data is stored on a minimal number of volumes through the process of collocation.

commit point
A point in time when data is considered to be consistent.

communication method
The method by which a client and server exchange information. See also Transmission Control Protocol/Internet Protocol.

communication protocol
A set of defined interfaces that permit computers to communicate with each other.

compression
A function that removes repetitive characters, spaces, strings of characters, or binary data from the data being processed and replaces characters with control characters. Compression reduces the amount of storage space that is required for data.

configuration manager
A server that distributes configuration information, such as policies and schedules, to managed servers according to their profiles. Configuration information can include policy and schedules. See also enterprise configuration, managed server, profile.

conversation
A connection between two programs over a session that allows them to communicate with each other while processing a transaction. See also session.

copy backup
A full backup in which the transaction log files are not deleted so that backup procedures that use incremental or differential backups are not disrupted.

copy group
A policy object containing attributes that control how backup versions or archive copies are generated, where backup versions or archive copies are initially located, and when backup versions or archive copies expire. A copy group belongs to a management class. See also archive copy group, backup copy group, backup version,management class.

copy storage pool
A named set of volumes that contain copies of files that reside in primary storage pools. Copy storage pools are used only to back up the data that is stored in primary storage pools. A copy storage pool cannot be a destination for a backup copy group, an archive copy group, or a management class (for space-managed files). See also destination, primary storage pool, server storage, storage pool, storage pool volume.

D

daemon
A program that runs unattended to perform continuous or periodic functions, such as network control.

damaged file
A physical file in which read errors have been detected.

database backup series
One full backup of the database, plus up to 32 incremental backups made since that full backup. Each full backup that is run starts a new database backup series. A number identifies each backup series. See also database snapshot, full backup.

database snapshot
A complete backup of the entire database to media that can be taken off-site. When a database snapshot is created, the current database backup series is not interrupted. A database snapshot cannot have incremental database backups associated with it. See also database backup series, full backup.

data center
In a virtualized environment, a container that holds hosts, clusters, networks, and data stores.

data deduplication
A method of reducing storage needs by eliminating redundant data. Only one instance of the data is retained on storage media. Other instances of the same data are replaced with a pointer to the retained instance.

data manager server
A server that collects metadata information for client inventory and manages transactions for the storage agent over the local area network. The data manager server informs the storage agent with applicable library attributes and the target volume identifier.

data mover
A device that moves data on behalf of the server. A network-attached storage (NAS) file server is a data mover.

data storage-management application-programming interface (DSMAPI)
A set of functions and semantics that can monitor events on files, and manage and maintain the data in a file. In an HSM environment, a DSMAPI uses events to notify data management applications about operations on files, stores arbitrary attribute information with a file, supports managed regions in a file, and uses DSMAPI access rights to control access to a file object.

data store
In a virtualized environment, the location where virtual machine data is stored.

deduplication
The process of creating representative records from a set of records that have been identified as representing the same entities.

default management class
A management class that is assigned to a policy set. This class is used to govern backed up or archived files when a file is not explicitly associated with a specific management class through the include-exclude list.

demand migration
The process that is used to respond to an out-of-space condition on a file system for which hierarchical storage management (HSM) is active. Files are migrated to server storage until space usage drops to the low threshold that was set for the file system. If the high threshold and low threshold are the same, one file is migrated. See also automatic migration, selective migration, threshold migration.

desktop client
The group of backup-archive clients that includes clients on Microsoft Windows, Apple, and Novell NetWare operating systems.

destination
A copy group or management class attribute that specifies the primary storage pool to which a client file will be backed up, archived, or migrated. See also copy storage pool.

device class
A named set of characteristics that are applied to a group of storage devices. Each device class has a unique name and represents a device type of disk, file, optical disk, or tape.

device configuration file
1. For a storage agent, a file that contains the name and password of the storage agent, and information about the server that is managing the SAN-attached libraries and drives that the storage agent uses.
2. For a server, a file that contains information about defined device classes, and, on some servers, defined libraries and drives. The information is a copy of the device configuration information in the database.

disaster recovery manager (DRM)
A function that assists in preparing and using a disaster recovery plan file for the server.

disaster recovery plan
A file that is created by the disaster recover manager (DRM) that contains information about how to recover computer systems if a disaster occurs and scripts that can be run to perform some recovery tasks. The file includes information about the software and

hardware that is used by the server, and the location of recovery media.

domain
A grouping of client nodes with one or more policy sets, which manage data or storage resources for the client nodes. See also policy domain.

DRM See disaster recovery manager.

DSMAPI
See data storage-management application-programming interface.

dynamic serialization
Copy serialization in which a file or folder is backed up or archived on the first attempt regardless of whether it changes during a backup or archive. See also shared dynamic serialization, shared static serialization, static serialization.

E

EA See extended attribute.

EB See exabyte.

EFS See Encrypted File System.

Encrypted File System (EFS)
A file system that uses file system-level encryption.

enterprise configuration
A method of setting up servers so that the administrator can distribute the configuration of one of the servers to the other servers, using server-to-server communication. See also configuration manager, managed server, profile, subscription.

enterprise logging
The process of sending events from a server to a designated event server. The event server routes the events to designated receivers, such as to a user exit. See also event.

error log
A data set or file that is used to record error information about a product or system.

estimated capacity
The available space, in megabytes, of a storage pool.

event An occurrence of significance to a task or system. Events can include completion or failure of an operation, a user action, or the change in state of a process. See also enterprise logging, receiver.

event record
A database record that describes actual status and results for events.

event server
A server to which other servers can send events for logging. The event server routes the events to any receivers that are enabled for the sending server's events.

exabyte (EB)
For processor, real and virtual storage capacities and channel volume, 2 to the power of 60 or 1 152 921 504 606 846 976 bytes. For disk storage capacity and communications volume, 1 000 000 000 000 000 000 bytes.

exclude
The process of identifying files in an include-exclude list. This process prevents the files from being backed up or migrated whenever a user or schedule enters an incremental or selective backup operation. A file can be excluded from backup, from space management, or from both backup and space management.

exclude-include list
See include-exclude list.

expiration
The process by which files, data sets, or objects are identified for deletion because their expiration date or retention period has passed.

expiring file
A migrated or premigrated file that has been marked for expiration and removal from storage. If a stub file or an original copy of a premigrated file is deleted from a local file system, or if the original copy of a premigrated file is updated, the corresponding migrated or premigrated file is marked for expiration the next time reconciliation is run.

extend
To increase the portion of available space that can be used to store database or recovery log information.

extended attribute (EA)
Names or value pairs that are associated with files or directories. There are three

classes of extended attributes: user attributes, system attributes, and trusted attributes.

external library
A collection of drives that is managed by the media-management system other than the storage management server.

F

file access time
On AIX, UNIX, or Linux systems, the time when the file was last accessed.

file age
For migration prioritization purposes, the number of days since a file was last accessed.

file device type
A device type that specifies the use of sequential access files on disk storage as volumes.

file server
A dedicated computer and its peripheral storage devices that are connected to a local area network that stores programs and files that are shared by users on the network.

file space
A logical space in server storage that contains a group of files that have been backed up or archived by a client node, from a single logical partition, file system, or virtual mount point. Client nodes can restore, retrieve, or delete their file spaces from server storage. In server storage, files belonging to a single file space are not necessarily stored together.

file space ID (FSID)
A unique numeric identifier that the server assigns to a file space when it is stored in server storage.

file state
The space management mode of a file that resides in a file system to which space management has been added. A file can be in one of three states: resident, premigrated, or migrated. See also migrated file, premigrated file, resident file.

file system migrator (FSM)
A kernel extension that intercepts all file system operations and provides any space management support that is required. If no space management support is required, the operation is passed to the operating system, which performs its normal functions. The file system migrator is mounted over a file system when space management is added to the file system.

file system state
The storage management mode of a file system that resides on a workstation on which the hierarchical storage management (HSM) client is installed. A file system can be in one of these states: native, active, inactive, or global inactive.

frequency
A copy group attribute that specifies the minimum interval, in days, between incremental backups.

FSID See file space ID.

FSM See file system migrator.

full backup
The process of backing up the entire server database. A full backup begins a new database backup series. See also database backup series, database snapshot, incremental backup.

fuzzy backup
A backup version of a file that might not accurately reflect what is currently in the file because the file was backed up at the same time as it was being modified.

fuzzy copy
A backup version or archive copy of a file that might not accurately reflect the original contents of the file because it was backed up or archived the file while the file was being modified.

G

GB See gigabyte.

General Parallel File System (GPFS™)
A high-performance shared-disk file system that can provide data access from nodes in a clustered system environment. See also information lifecycle management.

gigabyte (GB)
For processor storage, real and virtual storage, and channel volume, 10 to the

power of nine or 1,073,741,824 bytes. For disk storage capacity and communications volume, 1,000,000,000 bytes.

global inactive state
The state of all file systems to which space management has been added when space management is globally deactivated for a client node.

Globally Unique Identifier (GUID)
An algorithmically determined number that uniquely identifies an entity within a system. See also Universally Unique Identifier.

GPFS See General Parallel File System.

GPFS node set
A mounted, defined group of GPFS file systems.

group backup
The backup of a group containing a list of files from one or more file space origins.

GUID See Globally Unique Identifier.

H

hierarchical storage management (HSM)
A function that automatically distributes and manages data on disk, tape, or both by regarding devices of these types and potentially others as levels in a storage hierarchy that range from fast, expensive devices to slower, cheaper, and possibly removable devices. The objectives are to minimize access time to data and maximize available media capacity. See also hierarchical storage management client, recall, storage hierarchy.

hierarchical storage management client (HSM client) A client program that works with the server to provide hierarchical storage management (HSM) for a system. See also hierarchical storage management, management class.

HSM See hierarchical storage management.

HSM client
See hierarchical storage management client.

I

ILM See information lifecycle management.

image A file system or raw logical volume that is backed up as a single object.

image backup
A backup of a full file system or raw logical volume as a single object.

inactive file system
A file system for which space management has been deactivated. See also active file system.

inactive version
A backup version of a file that is either not the most recent backup version, or that is a backup version of a file that no longer exists on the client system. Inactive backup versions are eligible for expiration processing according to the management class assigned to the file. See also active version, backup version.

include-exclude file
A file containing statements to determine the files to back up and the associated management classes to use for backup or archive. See also include-exclude list.

include-exclude list
A list of options that include or exclude selected files for backup. An exclude option identifies files that should not be backed up. An include option identifies files that are exempt from the exclusion rules or assigns a management class to a file or a group of files for backup or archive services. See also include-exclude file.

incremental backup
The process of backing up files or directories, or copying pages in the database, that are new or changed since the last full or incremental backup. See also selective backup.

individual mailbox restore
See mailbox restore.

information lifecycle management (ILM)
A policy-based file-management system for storage pools and file sets. See also General Parallel File System.

inode The internal structure that describes the individual files on AIX, UNIX, or Linux systems. An inode contains the node, type, owner, and location of a file.

inode number
A number specifying a particular inode file in the file system.

IP address
A unique address for a device or logical unit on a network that uses the Internet Protocol standard.

J

job file
A generated file that contains configuration information for a migration job. The file is XML format and can be created and edited in the hierarchical storage management (HSM) client for Windows client graphical user interface. See also migration job.

journal-based backup
A method for backing up Windows clients and AIX clients that exploits the change notification mechanism in a file to improve incremental backup performance by reducing the need to fully scan the file system.

journal daemon
On AIX, UNIX, or Linux systems, a program that tracks change activity for files residing in file systems.

journal service
In Microsoft Windows, a program that tracks change activity for files residing in file systems.

K

KB See kilobyte.

kilobyte (KB)
For processor storage, real and virtual storage, and channel volume, 2 to the power of 10 or 1,024 bytes. For disk storage capacity and communications volume, 1,000 bytes.

L

LAN See local area network.

LAN-free data movement
The movement of client data between a client system and a storage device on a storage area network (SAN), bypassing the local area network.

LAN-free data transfer
See LAN-free data movement.

leader data
Bytes of data, from the beginning of a migrated file, that are stored in the file's corresponding stub file on the local file system. The amount of leader data that is stored in a stub file depends on the stub size that is specified.

library
1. A repository for demountable recorded media, such as magnetic disks and magnetic tapes.
2. A collection of one or more drives, and possibly robotic devices (depending on the library type), which can be used to access storage volumes.

library client
A server that uses server-to-server communication to access a library that is managed by another storage management server. See also library manager.

library manager
A server that controls device operations when multiple storage management servers share a storage device. See also library client.

local
1. Pertaining to a device, file, or system that is accessed directly from a user system, without the use of a communication line. See also remote.
2. For hierarchical storage management products, pertaining to the destination of migrated files that are being moved. See also remote.

local area network (LAN)
A network that connects several devices in a limited area (such as a single building or campus) and that can be connected to a larger network.

local shadow volume
Data that is stored on shadow volumes localized to a disk storage subsystem.

LOFS See loopback virtual file system.

logical file
A file that is stored in one or more server storage pools, either by itself or as part of an aggregate. See also aggregate, physical file, physical occupancy.

logical occupancy
The space that is used by logical files in a storage pool. This space does not include the unused space created when logical files are deleted from aggregate files, so it might be less than the physical occupancy. See also physical occupancy.

logical unit number (LUN)
In the Small Computer System Interface (SCSI) standard, a unique identifier used to differentiate devices, each of which is a logical unit (LU).

logical volume
A portion of a physical volume that contains a file system.

logical volume backup
A back up of a file system or logical volume as a single object.

Logical Volume Snapshot Agent (LVSA)
Software that can act as the snapshot provider for creating a snapshot of a logical volume during an online image backup.

loopback virtual file system (LOFS)
A file system that is created by mounting a directory over another local directory, also known as mount-over-mount. A LOFS can also be generated using an automounter.

LUN See logical unit number.

LVSA See Logical Volume Snapshot Agent.

M

macro file
A file that contains one or more storage manager administrative commands, which can be run only from an administrative client using the MACRO command. See also Tivoli Storage Manager command script.

mailbox restore
A function that restores Microsoft Exchange Server data (from IBM Data Protection for Microsoft Exchange backups) at the mailbox level or mailbox-item level.

managed object
A definition in the database of a managed server that was distributed to the managed server by a configuration manager. When a managed server subscribes to a profile, all objects that are associated with that profile become managed objects in the database of the managed server.

managed server
A server that receives configuration information from a configuration manager using a subscription to one or more profiles. Configuration information can include definitions of objects such as policy and schedules. See also configuration manager, enterprise configuration, profile, subscription.

management class
A policy object that users can bind to each file to specify how the server manages the file. The management class can contain a backup copy group, an archive copy group, and space management attributes. See also bind, copy group, hierarchical storage management client, policy set, rebind.

maximum transmission unit (MTU)
The largest possible unit of data that can be sent on a given physical medium in a single frame. For example, the maximum transmission unit for Ethernet is 1500 bytes.

MB See megabyte.

media server
In a z/OS® environment, a program that provides access to z/OS disk and tape storage for Tivoli Storage Manager servers that run on operating systems other than z/OS.

megabyte (MB)
For processor storage, real and virtual storage, and channel volume, 2 to the 20th power or 1,048,576 bytes. For disk storage capacity and communications volume, 1,000,000 bytes.

metadata
Data that describes the characteristics of data; descriptive data.

migrate
To move data to another location, or an application to another computer system.

migrated file
A file that has been copied from a local file system to storage. For HSM clients on UNIX or Linux systems, the file is replaced with a stub file on the local file system. On Windows systems, creation of the stub file is optional. See also file state, premigrated file, resident file, stub file.

migration
The process of moving data from one computer system to another, or an application to another computer system.

migration job
A specification of files to migrate, and actions to perform on the original files after migration. See also job file, threshold migration.

migration threshold
High and low capacities for storage pools or file systems, expressed as percentages, at which migration is set to start and stop.

mirroring
The process of writing the same data to multiple disks at the same time. The mirroring of data protects it against data loss within the database or within the recovery log.

mode A copy group attribute that specifies whether to back up a file that has not been modified since the last time the file was backed up. See also absolute mode, modified mode.

modified mode
In storage management, a backup copy-group mode that specifies that a file

is considered for incremental backup only if it has changed since the last backup. A file is considered a changed file if the date, size, owner, or permissions of the file have changed. See also absolute mode, mode.

mount limit
The maximum number of volumes that can be simultaneously accessed from the same device class. The mount limit determines the maximum number of mount points. See also mount point.

mount point
A logical drive through which volumes are accessed in a sequential access device class. For removable media device types, such as tape, a mount point is a logical drive associated with a physical drive. For the file device type, a mount point is a logical drive associated with an I/O stream. See also mount limit.

mount retention period
The maximum number of minutes that the server retains a mounted sequential-access media volume that is not being used before it dismounts the sequential-access media volume.

mount wait period
The maximum number of minutes that the server waits for a sequential-access volume mount request to be satisfied before canceling the request.

MTU See maximum transmission unit.

N

Nagle algorithm
An algorithm that reduces congestion of TCP/IP networks by combining smaller packets and sending them together.

named pipe
A type of interprocess communication that permits message data streams to pass between peer processes, such as between a client and a server.

NAS file server
See network-attached storage file server.

NAS file server node
See NAS node.

NAS node
A client node that is a network-attached storage (NAS) file server. Data for the NAS node is transferred by a NAS file server that is controlled by the network data management protocol (NDMP). A NAS node is also called a NAS file server node.

native file system
A file system that is locally added to the file server and is not added for space management. The hierarchical storage manager (HSM) client does not provide space management services to the file system.

native format
A format of data that is written to a storage pool directly by the server. See also non-native data format.

NDMP
See Network Data Management Protocol.

NetBIOS (Network Basic Input/Output System)
A standard interface to networks and personal computers that is used on local area networks to provide message, print-server, and file-server functions. Application programs that use NetBIOS do not have to handle the details of LAN data link control (DLC) protocols.

network-attached storage file server (NAS file server)
A dedicated storage device with an operating system that is optimized for file-serving functions. A NAS file server can have the characteristics of both a node and a data mover.

Network Basic Input/Output System
See NetBIOS.

Network Data Management Protocol (NDMP)
A protocol that allows a network storage-management application to control the backup and recovery of an NDMP-compliant file server, without installing vendor-acquired software on that file server.

network data-transfer rate
A rate that is calculated by dividing the total number of bytes that are transferred by the data transfer time. For example, this rate can be the time that is spent transferring data over a network.

node A file server or workstation on which the

backup-archive client program has been installed, and which has been registered to the server.

node name
A unique name that is used to identify a workstation, file server, or PC to the server.

node privilege class
A privilege class that gives an administrator the authority to remotely access backup-archive clients for a specific client node or for all clients in a policy domain. See also privilege class.

non-native data format
A format of data that is written to a storage pool that differs from the format that the server uses for operations. See also native format.

O

offline volume backup
A backup in which the volume is locked so that no other system applications can access it during the backup operation.

online volume backup
A backup in which the volume is available to other system applications during the backup operation.

open registration
A registration process in which users can register their workstations as client nodes with the server. See also closed registration.

operator privilege class
A privilege class that gives an administrator the authority to disable or halt the server, enable the server, cancel server processes, and manage removable media. See also privilege class.

options file
A file that contains processing options. See also client system-options file, client user-options file.

originating file system
The file system from which a file was migrated. When a file is recalled, it is returned to its originating file system.

orphaned stub file
A file for which no migrated file can be found on the server that the client node is contacting for space management services. For example, a stub file can be orphaned when the client system-options file is modified to contact a server that is different than the one to which the file was migrated.

P

packet In data communication, a sequence of binary digits, including data and control signals, that are transmitted and switched as a composite whole.

page A defined unit of space on a storage medium or within a database volume.

partial-file recall mode
A recall mode that causes the hierarchical storage management (HSM) function to read just a portion of a migrated file from storage, as requested by the application accessing the file.

password generation
A process that creates and stores a new password in an encrypted password file when the old password expires. Automatic generation of a password prevents password prompting.

path An object that defines a one-to-one relationship between a source and a destination. Using the path, the source accesses the destination. Data can flow from the source to the destination, and back. An example of a source is a data mover (such as a network-attached storage [NAS] file server), and an example of a destination is a tape drive.

pattern-matching character
See wildcard character.

physical file
A file that is stored in one or more storage pools, consisting of either a single logical file, or a group of logical files that are packaged together as an aggregate. See also aggregate, logical file, physical occupancy.

physical occupancy
The amount of space that is used by physical files in a storage pool. This space includes the unused space that is created when logical files are deleted from aggregates. See also logical file, logical occupancy, physical file.

plug-in
A separately installable software module that adds function to an existing program, application, or interface.

policy domain
A grouping of policy users with one or more policy sets, which manage data or storage resources for the users. The users are client nodes that are associated with the policy domain. See also active policy set, domain.

policy privilege class
A privilege class that gives an administrator the authority to manage policy objects, register client nodes, and schedule client operations for client nodes. Authority can be restricted to certain policy domains. See also privilege class.

policy set
A group of rules in a policy domain. The rules specify how data or storage resources are automatically managed for client nodes in the policy domain. Rules can be contained in management classes. See also active policy set, management class.

premigrated file
A file that has been copied to server storage, but has not been replaced with a stub file on the local file system. An identical copy of the file resides both on the local file system and in server storage. Premigrated files occur on UNIX and Linux file systems to which space management has been added. See also file state, migrated file, resident file.

premigrated files database
A database that contains information about each file that has been premigrated to server storage.

premigration
The process of copying files that are eligible for migration to server storage, but leaving the original file intact on the local file system.

premigration percentage
A space management setting that controls whether the next eligible candidates in a file system are premigrated following threshold or demand migration.

primary storage pool
A named set of volumes that the server uses to store backup versions of files, archive copies of files, and files migrated from client nodes. See also copy storage pool, server storage, storage pool, storage pool volume.

privilege class
A level of authority that is granted to an administrator. The privilege class determines which administrative tasks the administrator can perform. See also authority, node privilege class, operator privilege class, policy privilege class, storage privilege class, system privilege class.

profile
A named group of configuration information that can be distributed from a configuration manager when a managed server subscribes. Configuration information can include registered administrator IDs, policies, client schedules, client option sets, administrative schedules, storage manager command scripts, server definitions, and server group definitions. See also configuration manager, enterprise configuration, managed server.

profile association
On a configuration manager, the defined relationship between a profile and an object such as a policy domain. Profile associations define the configuration information that is distributed to a managed server when it subscribes to the profile.

Q

quota

1. For HSM on AIX, UNIX, or Linux systems, the limit (in megabytes) on the amount of data that can be migrated and premigrated from a file system to server storage.
2. For HSM on Windows systems, a user-defined limit to the space that is occupied by recalled files.

R

randomization
The process of distributing schedule start times for different clients within a specified percentage of the schedule's startup window.

raw logical volume
A portion of a physical volume that is comprised of unallocated blocks and has no journaled file system (JFS) definition. A logical volume is read/write accessible only through low-level I/O functions.

rebind
To associate all backed-up versions of a file with a new management class name. For example, a file that has an active backup version is rebound when a later version of the file is backed up with a different management class association. See also bind, management class.

recall To copy a migrated file from server storage back to its originating file system using the hierarchical storage management client. See also selective recall.

receiver
A server repository that contains a log of server and client messages as events. For example, a receiver can be a file exit, a user exit, or the server console and activity log. See also event.

reclamation
The process of consolidating the remaining data from many sequential-access volumes onto fewer, new sequential-access volumes.

reclamation threshold
The percentage of space that a sequential-access media volume must have before the server can reclaim the volume. Space becomes reclaimable when files are expired or are deleted.

reconciliation
The process of ensuring consistency between the original data repository and the larger system where the data is stored for backup. Examples of larger systems where the data is stored for backup are storage servers or other storage systems. During the reconciliation process, data that is identified as no longer needed is removed.

recovery log
A log of updates that are about to be written to the database. The log can be used to recover from system and media failures. The recovery log consists of the active log (including the log mirror) and archive logs.

register
To define a client node or administrator ID that can access the server.

registry
A repository that contains access and configuration information for users, systems, and software.

remote
For hierarchical storage management products, pertaining to the origin of migrated files that are being moved. See also local.

resident file
On a Windows system, a complete file on a local file system that might also be a migrated file because a migrated copy can exist in server storage. On a UNIX or Linux system, a complete file on a local file system that has not been migrated or premigrated, or that has been recalled from server storage and modified. See also file state.

restore
To copy information from its backup location to the active storage location for use. For example, to copy information from server storage to a client workstation.

retention
The amount of time, in days, that inactive backed-up or archived files are kept in the storage pool before they are deleted. Copy group attributes and default retention grace periods for the domain define retention.

retrieve
To copy archived information from the storage pool to the workstation for use. The retrieve operation does not affect the archive version in the storage pool. See also archive.

root user
A system user who operates without restrictions. A root user has the special rights and privileges needed to perform administrative tasks.

S

SAN See storage area network.

schedule
A database record that describes client operations or administrative commands to be processed. See also administrative command schedule, client schedule.

scheduling mode
The type of scheduling operation for the server and client node that supports two scheduling modes: client-polling and server-prompted.

scratch volume
A labeled volume that is either blank or contains no valid data, that is not defined, and that is available for use. See also volume.

script A series of commands, combined in a file, that carry out a particular function when the file is run. Scripts are interpreted as they are run. See also Tivoli Storage Manager command script.

Secure Sockets Layer (SSL)
A security protocol that provides communication privacy. With SSL, client/server applications can communicate in a way that is designed to prevent eavesdropping, tampering, and message forgery.

selective backup
The process of backing up certain files or directories from a client domain. The files that are backed up are those that are not excluded in the include-exclude list. The files must meet the requirement for serialization in the backup copy group of the management class that is assigned to each file. See also incremental backup.

selective migration
The process of copying user-selected files from a local file system to server storage and replacing the files with stub files on the local file system. See also demand migration, threshold migration.

selective recall
The process of copying user-selected files from server storage to a local file system. See also recall, transparent recall.

serialization
The process of handling files that are modified during backup or archive processing. See also shared dynamic serialization, shared static serialization, static serialization.

server A software program or a computer that provides services to other software programs or other computers. See also client.

server options file
A file that contains settings that control various server operations. These settings affect such things as communications, devices, and performance.

server-prompted scheduling mode
A client/server communication technique where the server contacts the client node when tasks must be done. See also client-polling scheduling mode.

server storage
The primary, copy, and active-data storage pools that are used by the server to store user files such as backup versions, archive copies, and files migrated from hierarchical storage management client nodes (space-managed files). See also active-data pool, copy storage pool, primary storage pool, storage pool volume, volume.

session
A logical or virtual connection between two stations, software programs, or devices on a network that allows the two elements to communicate and exchange data for the duration of the session. See also administrative session.

session resource usage
The amount of wait time, processor time, and space that is used or retrieved during a client session.

shadow copy
A snapshot of a volume. The snapshot can be taken while applications on the system continue to write data to the volumes.

shadow volume
The data stored from a snapshot of a volume. The snapshot can be taken while applications on the system continue to write data to the volumes.

shared dynamic serialization
A value for serialization that specifies that a file must not be backed up or archived if it is being modified during the operation. The backup-archive client retries the backup or archive operation a number of times; if the file is being modified during each attempt, the backup-archive client will back up or archive the file on its last try. See also dynamic serialization, serialization, shared static serialization, static serialization.

shared library
A library device that is used by multiple storage manager servers. See also library.

shared static serialization
A copy-group serialization value that specifies that a file must not be modified during a backup or archive operation. The client attempts to retry the operation a number of times. If the file is in use during each attempt, the file is not backed up or archived. See also dynamic serialization, serialization, shared dynamic serialization, static serialization.

snapshot
An image backup type that consists of a point-in-time view of a volume.

space-managed file
A file that is migrated from a client node by the hierarchical storage management (HSM) client. The HSM client recalls the file to the client node on demand.

space management
See hierarchical storage management.

space monitor daemon
A daemon that checks space usage on all file systems for which space management is active, and automatically starts threshold migration when space usage on a file system equals or exceeds its high threshold.

sparse file
A file that is created with a length greater than the data it contains, leaving empty spaces for the future addition of data.

special file
On AIX, UNIX, or Linux systems, a file that defines devices for the system, or temporary files that are created by processes. There are three basic types of special files: first-in, first-out (FIFO); block; and character.

SSL See Secure Sockets Layer.

stabilized file space
A file space that exists on the server but not on the client.

stanza A group of lines in a file that together have a common function or define a part of the system. Stanzas are usually separated by blank lines or colons, and each stanza has a name.

startup window
A time period during which a schedule must be initiated.

static serialization
A copy-group serialization value that specifies that a file must not be modified during a backup or archive operation. If the file is in use during the first attempt, the backup-archive client cannot back up or archive the file. See also dynamic serialization, serialization, shared dynamic serialization, shared static serialization.

storage agent
A program that enables the backup and restoration of client data directly to and from storage attached to a storage area network (SAN).

storage area network (SAN)
A dedicated storage network tailored to a specific environment, combining servers, systems, storage products, networking products, software, and services.

storage hierarchy
A logical order of primary storage pools, as defined by an administrator. The order is typically based on the speed and capacity of the devices that the storage pools use. The storage hierarchy is defined by identifying the next storage pool in a storage pool definition. See also storage pool.

storage pool
A named set of storage volumes that is the destination that is used to store client

data. See also active-data pool, copy storage pool, primary storage pool, storage hierarchy.

storage pool volume
A volume that has been assigned to a storage pool. See also active-data pool, copy storage pool, primary storage pool, server storage, volume.

storage privilege class
A privilege class that gives an administrator the authority to control how storage resources for the server are allocated and used, such as monitoring the database, the recovery log, and server storage. See also privilege class.

stub A shortcut on the Windows file system that is generated by the hierarchical storage management (HSM) client for a migrated file that allows transparent user access. A stub is the sparse file representation of a migrated file, with a reparse point attached.

stub file
A file that replaces the original file on a local file system when the file is migrated to storage. A stub file contains the information that is necessary to recall a migrated file from server storage. It also contains additional information that can be used to eliminate the need to recall a migrated file. See also migrated file, resident file.

stub file size
The size of a file that replaces the original file on a local file system when the file is migrated to server storage. The size that is specified for stub files determines how much leader data can be stored in the stub file. The default for stub file size is the block size defined for a file system minus 1 byte.

subscription
In a storage environment, the process of identifying the subscribers to which the profiles are distributed. See also enterprise configuration, managed server.

system privilege class
A privilege class that gives an administrator the authority to issue all server commands. See also privilege class.

T

tape library
A set of equipment and facilities that support an installation's tape environment. The tape library can include tape storage racks, mechanisms for automatic tape mounting, a set of tape drives, and a set of related tape volumes mounted on those drives.

tape volume prefix
The high-level-qualifier of the file name or the data set name in the standard tape label.

target node
A client node for which other client nodes (called agent nodes) have been granted proxy authority. The proxy authority allows the agent nodes to perform operations such as backup and restore on behalf of the target node, which owns the data.

TCA See trusted communications agent.

TCP/IP
See Transmission Control Protocol/Internet Protocol.

threshold migration
The process of moving files from a local file system to server storage based on the high and low thresholds that are defined for the file system. See also automatic migration, demand migration, migration job, selective migration.

throughput
In storage management, the total bytes in the workload, excluding overhead, that are backed up or restored, divided by elapsed time.

timeout
A time interval that is allotted for an event to occur or complete before operation is interrupted.

Tivoli Storage Manager command script
A sequence of Tivoli Storage Manager administrative commands that are stored in the database of the Tivoli Storage Manager server. The script can run from any interface to the server. The script can include substitution for command parameters and conditional logic. See also macro file, script.

tombstone object
A small subset of attributes of a deleted object. The tombstone object is retained for a specified period, and at the end of the specified period, the tombstone object is permanently deleted.

Transmission Control Protocol/Internet Protocol (TCP/IP)
An industry-standard, nonproprietary set of communication protocols that provides reliable end-to-end connections between applications over interconnected networks of different types. See also communication method.

transparent recall
The process that is used to automatically recall a migrated file to a workstation or file server when the file is accessed. See also selective recall.

trusted communications agent (TCA)
A program that handles the sign-on password protocol when clients use password generation.

U

UCS-2 A 2-byte (16-bit) encoding scheme based on ISO/IEC specification 10646-1. UCS-2 defines three levels of implementation: Level 1-No combining of encoded elements allowed; Level 2-Combining of encoded elements is allowed only for Thai, Indic, Hebrew, and Arabic; Level 3-Any combination of encoded elements are allowed.

UNC See Universal Naming Convention.

Unicode
A character encoding standard that supports the interchange, processing, and display of text that is written in the common languages around the world, plus many classical and historical texts.

Unicode-enabled file space
Unicode file space names provide support for multilingual workstations without regard for the current locale.

Universally Unique Identifier (UUID)
The 128-bit numeric identifier that is used to ensure that two components do not have the same identifier. See also Globally Unique Identifier.

Universal Naming Convention (UNC)
The server name and network name combined. These names together identify the resource on the domain.

UTF-8 Unicode Transformation Format, 8-bit encoding form, which is designed for ease of use with existing ASCII-based systems. The CCSID value for data in UTF-8 format is 1208. See also UCS-2.

UUID See Universally Unique Identifier.

V

validate
To check a policy set for conditions that can cause problems if that policy set becomes the active policy set. For example, the validation process checks whether the policy set contains a default management class.

version
A backup copy of a file stored in server storage. The most recent backup copy of a file is the active version. Earlier copies of the same file are inactive versions. The number of versions retained by the server is determined by the copy group attributes in the management class.

virtual file space
A representation of a directory on a network-attached storage (NAS) file system as a path to that directory.

virtual mount point
A directory branch of a file system that is defined as a virtual file system. The virtual file system is backed up to its own file space on the server. The server processes the virtual mount point as a separate file system, but the client operating system does not.

virtual volume
An archive file on a target server that represents a sequential media volume to a source server.

volume
A discrete unit of storage on disk, tape or other data recording medium that supports some form of identifier and parameter list, such as a volume label or input/output control. See also scratch volume, server storage, storage pool, storage pool volume.

volume history file
A file that contains information about volumes that have been used by the server for database backups and for export of administrator, node, policy, or server data. The file also has information about sequential-access storage pool volumes that have been added, reused, or deleted. The information is a copy of volume information that is recorded in the server database.

Volume Shadow Copy Service (VSS)
A set of Microsoft application-programming interfaces (APIs) that are used to create shadow copy backups of volumes, exact copies of files, including all open files, and so on.

VSS See Volume Shadow Copy Service.

VSS Backup
A backup operation that uses Microsoft Volume Shadow Copy Service (VSS) technology. The backup operation produces an online snapshot (point-in-time consistent copy) of Microsoft Exchange data. This copy can be stored on local shadow volumes or on Tivoli Storage Manager server storage.

VSS Fast Restore
An operation that restores data from a local snapshot. The snapshot is the VSS backup that resides on a local shadow volume. The restore operation retrieves the data by using a file-level copy method.

VSS Instant Restore
An operation that restores data from a local snapshot. The snapshot is the VSS backup that resides on a local shadow volume. The restore operation retrieves the data by using a hardware assisted restore method (for example, a FlashCopy operation).

VSS offloaded backup
A backup operation that uses a Microsoft Volume Shadow Copy Service (VSS) hardware provider (installed on an alternate system) to move IBM Data Protection for Microsoft Exchange data to the Tivoli Storage Manager server. This type of backup operation shifts the backup load from the production system to another system.

VSS Restore
A function that uses a Microsoft Volume Shadow Copy Service (VSS) software provider to restore VSS Backups (IBM Data Protection for Microsoft Exchange database files and log files) that reside on Tivoli Storage Manager server storage to their original location.

W

wildcard character
A special character such as an asterisk (*) or a question mark (?) that can be used to represent one or more characters. Any character or set of characters can replace the wildcard character.

workload partition (WPAR)
A partition within a single operating system instance.

workstation
A terminal or personal computer at which a user can run applications and that is usually connected to a mainframe or a network.

worldwide name (WWN)
A 64-bit, unsigned name identifier that is unique.

WPAR See workload partition.

WWN See worldwide name.

Index

Special characters

A

B

C

D

E

F

G

H

I

K

L

M

N

O

P

Q

R

S

T

U

V

X

Product Number: 5608-E06

Printed in USA