IBM Tivoli Storage Manager for Databases
Version 7.1

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

IBM

IBM Tivoli Storage Manager for Databases
Version 7.1

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

IBM

Note

Before using this information and the product it supports, read the information in "Notices" on page 273.

First edition (December 2013)

This edition applies to version 7, release 1, modification 0 of IBM Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server (product number 5608-E04) and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

Chapter 5. Protecting data 83

Chapter 6. Automating 109

Chapter 7. Troubleshooting Data Protection for SQL with VSS backup-restore support 115

Chapter 8. Performance 129

Chapter 9. Reference information. . . 131

Tables

About this publication

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

Data Protection for Microsoft SQL Server is also known as Data Protection for SQL Server. This book explains how to install, configure, and administer the Data Protection for SQL Server component.

You can use the Data Protection for SQL Server software to perform online backups of Microsoft SQL Server databases to Tivoli Storage Manager storage.

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 users, Tivoli Storage Manager administrators, and system administrators.

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

- Microsoft SQL Server
- Tivoli Storage Manager server
- Tivoli Storage Manager backup-archive client
- Tivoli Storage Manager Application Programming Interface

It is also assumed that you have an understanding of one of the following operating systems:

- Windows Server 2008
- Windows Server 2008 R2
- Windows Server 2012

It is also assumed that you have an understanding of the following IBM storage system 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™ (DS8100, DS8300, or DS8700)

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.

Conventions used in this publication

This guide uses several conventions for special terms and actions, operating system-dependent commands and paths.

This guide uses the following typeface conventions:

Bold

- Commands, keywords, authorization roles, or other information that you must use.
- Example: Log on to the server as **root** user.

italics

- Values or variables that you must provide.
- Emphasized words and phrases.
- Example: The node name of the *production node* and *backup node* must not be the same.

bold italics

- Options and parameters.
- Example: Specify the value for the ***compression*** option.

`monospace`

- Directories, parameters, URLs, and output examples.
- Example: The product is installed in the `C:\program files\tivoli\tsm\client\ba` directory.

UPPER CASE

- Environment variables associated with Tivoli Storage Manager, operating systems, or SQL Server.
- Example: Make sure the DSM_DIR environment variable is set correctly.

Reading syntax diagrams

This section describes how to read the syntax diagrams used in this book. 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 appears 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.	**Abbreviations** `►►—KEYWOrd——————————————►◄`
Symbols: Enter these symbols exactly as they appear in the syntax diagram.	* Asterisk { } Braces : Colon , Comma = Equal Sign - Hyphen () Parentheses . Period Space
Variables: Italicized lowercase items (*var_name*) denote variables. In this example, you can specify a *var_name* when you enter the KEYWORD command.	**Variables** `►►—KEYWOrd—`*`var_name`*`—————————►◄`

Syntax Diagram Description	Example
Repetition: An arrow returning to the left means you can repeat the item. A character or space within the arrow means you must separate repeated items with that character or space. A footnote by the arrow references the number of times you can repeat the item.	**Repetition** 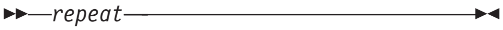 **Repetition** 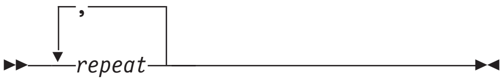 **Repetition** 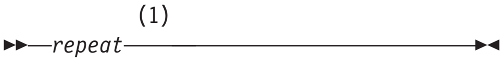 **Notes:** 1 Specify *repeat* as many as 5 times.
Required Choices: When two or more items are in a stack and one of them is on the line, you *must* specify one item. In this example, you *must* choose A, B, or C.	**Required choices** 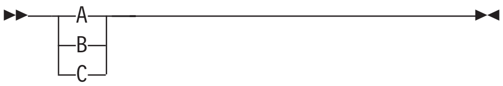
Optional Choice: When an item is below the line, that item is optional. In the first example, you can choose A or nothing at all. When two or more items are in a stack below the line, all of them are optional. In the second example, you can choose A, B, C, or nothing at all.	**Optional choice** 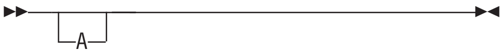 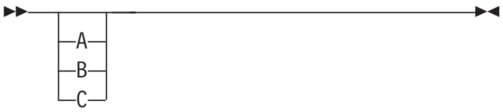
Defaults: Defaults are above the line. The default is selected unless you override it. You can override the default by including an option from the stack below the line. In this example, A is the default. You can override A by choosing B or C. You can also specify the default explicitly.	**Defaults** 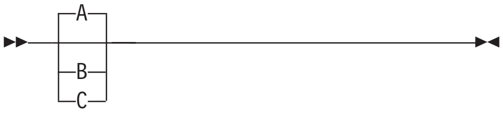
Repeatable Choices: A stack of items followed by an arrow returning to the left means you can select more than one item or, in some cases, repeat a single item. In this example, you can choose any combination of A, B, or C.	**Repeatable choices** 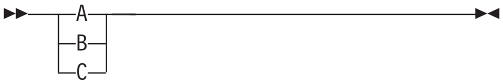

Syntax Diagram Description	Example
Syntax Fragments: Some diagrams, because of their length, must fragment the syntax. The fragment name appears between vertical bars in the diagram. The expanded fragment appears between vertical bars in the diagram after a heading with the same fragment name.	**Syntax fragments** 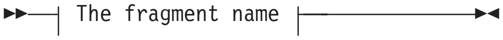 **The fragment name:** 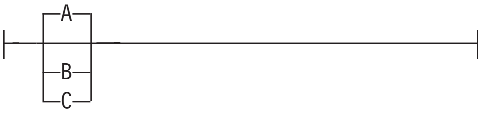

New for Version 7.1

Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server Version 7.1 contains several new features and changes.

“Restoring from virtual machine snapshots” on page 98
You can recover Microsoft SQL databases from a VM backup. To complete this task, use both Tivoli Storage Manager for Virtual Environments and Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server.

“Windows PowerShell and Data Protection for SQL Server” on page 111
The Data Protection for SQL Server software includes Windows PowerShell cmdlets to complement the command-line interface functions.

“Automated failover for data recovery” on page 10
When the primary Tivoli Storage Manager server is unavailable, Data Protection for Microsoft SQL Server can automatically fail over to the secondary server for data recovery.

Chapter 1. Getting started

With Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server software, you can continue to run primary applications on your database servers while data is backed up and restored.

Legacy-style and VSS backups are supported. The snapshots can be stored on a Tivoli Storage Manager server.

Backup overview

Data Protection for SQL Server provides several methods of backing up SQL Server data.

Legacy backup overview

A Legacy backup creates a copy of all or part of a SQL database or logs on Tivoli Storage Manager storage media.

Data Protection for SQL Server provides selection mechanisms and the logic that are required to back up and restore SQL data. When you initiate a legacy backup operation, Data Protection for SQL Server completes the following actions:

1. Begins a session with a Tivoli Storage Manager server by using the Tivoli Storage Manager API and information that is contained in a client options file.
2. Starts a session with the SQL Server by using the SQL-SMO interface.
3. Instructs the SQL Server by using the SQL VDI interface to begin a backup of the selected database objects.
4. Receives data from the SQL Server and sends it to the Tivoli Storage Manager server.
5. Informs the SQL Server that the backup is complete.
6. Ends the Tivoli Storage Manager server and SQL Server sessions.

Depending on the instructions that are provided, Data Protection for SQL Server software can either compress or instruct the SQL Server to compress the SQL data before it sends the data to the Tivoli Storage Manager server.

When a backup is run, Tivoli Storage Manager server retains information about the SQL Server and database. This information is available for query and restore operations after the backup is completed. The information about the names and sizes of the database filegroups and files is stored along with the database data as metadata.

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 SQL Server is the SqlServerWriter. The SqlServerWriter is provided by the SQL Server VSS Writer 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 SQL 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, respectively.

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 for each database or group of databases that are to be backed up and restored together on a separate, dedicated logical volume.
- Place logs for each database on a separate logical volume.
- Do not place non-SQL data on storage volumes that are dedicated to SQL data.
- When you use hardware snapshot providers, do not share storage group LUNs with other databases or applications.
- Read and follow specific installation and configuration instructions in the documentation that is provided by your VSS provider vendor.
- If you use XIV storage devices, install and configure IBM XIV Provider for Microsoft Windows Volume Shadow Copy Service (xProv) Version 2.3.0 and later.
- When a hardware provider is used, configure the disks that store SQL data and log files 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 SQL 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 SQL 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.

VSS backup management:

Some VSS backup characteristics are different from legacy backup characteristics. Examples of these differences are the backup characteristics for types that are supported, the backup granularity, and the backup storage location options.

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 run in a MicrosoftWindows Failover Clustering or Veritas Cluster Server (VCS) environments.

The full and copy-only full backup types are supported. Log, differential, file, group, and set backup types are not supported. Legacy differential and legacy log backups can be applied after a full VSS backup is restored. Different policy settings can be defined for each storage location, backup method, and backup type (`FULL` or `COPY`).

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 SQL Server.
- Continue to schedule and run Legacy backups in your strategy.
- Ensure that you have a well-defined and tested recovery plan that meets your service level objectives.
- 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 SQL Server, 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 SQL Server. 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 for each database or group of databases that are going to be backed up and restored together as a unit on a separate and dedicated logical volume.
- Place logs for each database or group of databases that are going to be backed up and restored together as a unit on a separate logical volume.
- Do not place non-SQL data on storage volumes that are dedicated to SQL.
- 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 machine 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.

Backup types:

Data Protection for SQL Server offers an expanded range of backup types that allows flexibility for your environment and processing needs.

Data Protection for SQL Server backup types have the following characteristics:

Full database backup (Legacy and VSS)

Data Protection for SQL Server backs up an entire SQL Server database and the portion of the transaction log necessary to provide a consistent database state. With both full and differential backups, the copy includes enough information from any associated transaction logs to make a backup consistent with itself. The portion of the log included contains only the transactions that occur from the beginning of the backup until its completion.

Legacy backups are a stream of bytes that Data Protection for SQL Server stores on the Tivoli Storage Manager server. VSS backups differ since they are at the volume and file-level. When a SQL Server database is not fully allocated, a legacy backup might transfer a smaller amount of data for a Tivoli Storage Manager backup than for a VSS backup. This situation occurs because a VSS backup transfers the entire file, regardless of its allocation.

Copy-only full backup (Legacy and VSS)

A copy-only full backup is a type of backup that is independent of the sequence of conventional SQL Server backups. The copy-only full backup does not disturb the sequence for a differential backup. The differential backup is not associated with the copy-full backup, but is associated with the prior full backup that was completed. This type of backup can be used for special purpose backups that do not affect existing backup and restore procedures. In addition, when compared to conventional backups, this type of backup can be used for longer term retention. An example of a special purpose backup is a backup of a log before an online file restore. In this scenario, the copy-only full backup is used one time. After the backup is used, it is deleted.

Differential backup (Legacy only)

Data Protection for SQL Server backs up only the data pages in a SQL Server database instance that changed after the last full backup and a portion of the transaction log.

Differential backup is associated with the last full backup that was run. The last full backup might be completed by Data Protection for SQL Server or another tool or product. For example, if you run a full backup with SQL Server to disk backup, and run a differential backup with Data Protection for SQL Server, the differential backup is associated with the SQL Server disk backup.

(Microsoft SQL Server 2012 only) Differential backup is not supported for databases on the secondary replica.

Log backup (Legacy only)

Data Protection for SQL Server backs up only the contents of a SQL Server database transaction log since the last successful log backup. Before the first log backup, complete either a full backup or an equivalent type of backup. Log backups normally follow full backups. The portion of the log included in full and differential backups is not equivalent to a log backup. Additionally, in full and differential backups, the log is not truncated as it is during a log backup. However, a log backup that follow a full or differential backup includes the same transactions as a full or differential backup. Log backups are not cumulative as are differential; they must be applied against a base backup and in the correct order.

A log backup in SQL Server terms is not equivalent to an incremental backup in Tivoli Storage Manager terms.

File backup (Legacy only)

Data Protection for SQL Server backs up only the contents of a specified SQL Server logical file. This type of backup can ease the scheduling for backing up large databases. You can back up different sets of files during different scheduled backups. File, group, and set backups must be followed by a log backup, but a full is not required.

Group backup (Legacy only)

Data Protection for SQL Server backs up only the contents of a specified SQL Server filegroup. This backup enables you to back up the set of database tables and indexes within a specific group of files.

The "group" is specified as part of the setup within the SQL Server when you define the database files. If no group is specified and all the database files are part of the "primary" group, it is not possible to back up or restore just part of the database by using the group.

Set backup (Legacy only)

Data Protection for SQL Server backs up the contents of specified SQL Server file groups and files as a unit.

Restore methods

Data Protection for SQL Server provides several methods of restoring SQL Server data.

Legacy restore overview

A legacy restore obtains backup copies of SQL databases from Tivoli Storage Manager server storage and restores them to their original location.

Like a Legacy backup, it uses a specialized API restore that functions with the SQL Server.

A complete restore of a database involves restoring a full backup or the equivalent thereof (from group, file, or set backups) and restoring all transaction logs since the last full backup. For a Legacy Restore, Data Protection for SQL Server:

1. Starts a session with a Tivoli Storage Manager server by using the Tivoli Storage Manager API and information that is contained in a client options file.
2. Starts a session with the SQL Server by using the SQL-SMO interface.
3. Queries the Tivoli Storage Manager server for a list of database backups.
4. Instructs the SQL Server by using the SQL VDI interface to begin a restore of the selected database objects.
5. Receives data from the Tivoli Storage Manager server and forwards it to the SQL Server.
6. Ends the Tivoli Storage Manager and SQL Server sessions.

VSS restore

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

The following characteristics are true of a VSS restore:

- You can restore only SQL Server VSS backups to the same SQL Server instance.
- Full and copy-only full backup types can be restored. Differential, individual filegroups, individual files, and set backups are not supported by VSS and therefore, cannot be restored.
- 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 Microsoft Windows Failover Clustering or Veritas Cluster Server (VCS) environment.
- Supports restoring a VSS backup (directly from Tivoli Storage Manager server storage) to an alternate location by using the **/relocatedir** option.

VSS fast restore

A VSS fast restore 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.

The following characteristics are true of VSS fast restores:

- Full and copy-only full backup types can be restored. Differential, individual filegroups, individual files, and set backups are not supported by VSS and therefore, cannot be restored.
- You can restore only SQL Server VSS backups to the same SQL Server instance.
- VSS backups can be restored to an alternate location by using the **/relocatedir** option.
- Restore granularity is at the database level.
- Restores can be run in a Microsoft Windows Failover Clustering or Veritas Cluster Server environment.

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.

A VSS instant restore is only possible when all of the data from the storage group or 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 being restored, including the log files and full-text index 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 involve the volumes that are 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.

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.

SQL Server VSS backups can only be restored into the same SQL Server instance from which they were backed up. This limitation is a Microsoft SQL Server limitation. Full and copy-only full backup types can be restored. Legacy differential and legacy log backups can be applied after a full or copy-only full VSS backup is restored.

When you plan for VSS instant restore, use the following considerations:

- (DS8000, SAN Volume Controller, Storwize V7000) Requires IBM System Storage Support for Microsoft Volume Shadow Copy Service software. XIV has separate VSS Provider software.
- 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.

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 SQL Server.

Automated failover for data recovery

Data Protection for Microsoft SQL Server can automatically fail over to a 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 SQL Server connects to for backup services is called the *primary server*. If the primary server is set up for node replication, the client node data on the primary server can be replicated to another 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)
- AlwaysOn node (for backups of availability databases in an AlwaysOn Availability Group on SQL Server 2012)

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

Each time the backup-archive client logs on to the server for backup services, it attempts to contact the primary server. If the primary server is unavailable, the backup-archive client automatically fails over to the secondary server. In failover mode, you can restore data that 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 SQL Server has failed over by looking for entries about the secondary server in the following log files:

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

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

- Data Protection for SQL Server 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
  - AlwaysOn node, if applicable
- Before the connection information for the secondary server can be sent to Data Protection for SQL Server, the following processes must occur:
  - You must back up data at least one time to the primary server.

- The following nodes must be replicated at least one time to the secondary server:
  - Data Protection node
  - AlwaysOn node, if applicable

Restriction: The following restrictions apply to Data Protection for SQL Server 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

Guidelines about backup strategies, options, preferences, policy settings, and other planning information are provided to assist when planning for Data Protection for Microsoft SQL Server backup and restore operations.

Before you begin

For best results, review this information before performing any Data Protection for Microsoft SQL Server configuration tasks.

About this task

Consider your production environment capabilities and backup objectives when planning for SQL Server data. Install and configure Data Protection for Microsoft SQL Server before attempting a backup operation.

Security requirements

Data Protection for SQL Server requires certain settings in order to perform operations in a secure environment.

Windows administrator authority is required for installation. Data Protection for SQL Server must be registered to the Tivoli Storage Manager server and the appropriate node name and password must be used when it connects to the Tivoli Storage Manager server. In addition, standard Tivoli Storage Manager security requirements apply to Data Protection for SQL Server.

Three options are provided when you specify SQL Server logon information:

- Accept the default `sa` account and blank password.
- Use SQL user ID security and specify both the SQL user name and password. With SQL user ID security, the SQL Server administrator provides the logon ID and the password that provides access to the SQL Server.
- Use a trusted connection and allow Windows authenticate the logon.

The SQL logon user or Windows user name must be added to the SQL Server SYSADMIN fixed server role before it can be used by Data Protection for SQL Server.

Backup strategies

Different backup strategies are available depending on specific requirements regarding network traffic, backup window and acceptable restore times.

Strategies defined by backup type

Some commonly used strategies (based upon backup type) are described as follows:

Full backup only (Legacy and VSS)

This approach is best for SQL databases that are relatively small because it implies that the entire database is backed up each time. Each full backup takes longer to perform, but the restore process is most efficient because only the most recent (or other appropriate) full backup need be restored. This is the appropriate strategy for system databases such as *master*, *model*, and *msdb* due to their normally small size.

Full plus log backup (Legacy and VSS)

A full plus transaction log backup 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 log backups allows the backup window and network traffic to be minimized. For example, you can perform full backups on the weekend and log 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, as well as subsequent log backups, must be restored. It is also possible to do a point-in-time restore to restore a transaction log to a specified date and time.

You can apply legacy log backups after a full VSS backup has been restored. to do this, you must leave the database in a recovering state by specifying **`/recovery`**=`no` on the command-line interface or by making sure that the **Recovery** option in the GUI Restore Databases or Restore Groups/Files is not selected when restoring the VSS backup.

Full plus differential backup (Legacy and VSS)

This strategy can be used *between* full backups. A differential database backup can save both time and space. Space is saved because the backup consists of only the changed portions of a database since the last full backup (it is cumulative). Time is saved because you can avoid applying all individual log backups within that time to the operation. This applies to restore operations as well; only the last differential backup (latest version) need be restored.

Although VSS supports full backups only, legacy differential backups can be applied to the VSS full backup. To do this, you must leave the database in a recovering state by specifying **`/recovery`**=*no* on the command-line interface or by making sure that the Recovery option is not selected when restoring the VSS backup.

Full plus differential plus log backup (Legacy and VSS)

This strategy allows for a faster restore scenario by reducing the number of transactions that may need to be restored and applied. If, for example, a full legacy or VSS backup is done weekly, a differential nightly, and a log backup every four hours, the restore would involve the full backup, a differential, and at most five log backups. However, simply a full plus log backup scheme on the same cycle could require a full plus up to forty-one log backups to be restored (six days times six log backups per day plus up to five backups on the day the full backup was done). Although VSS supports full backups only, legacy log backups and legacy differential backups can be applied to the VSS full backup.

File or group backups (Legacy only)

When a group is created on the SQL Server, database files are identified with that group. The group used for the group backup is dependant on the group to which the database files are defined.

Use a file backup strategy when it is impractical to backup an entire database because of size and accompanying time and performance issues. When performing restore operations for a file or filegroup, provide a separate backup of the transaction log.

File or group options can also save both backup and restore time in cases when certain tables or indexes have more updates than others and need to be backed up more often. It is time-effective to place such data in their own filegroup or files and then back up only those items.

The PRIMARY filegroup must be restored prior to restoring a user-defined filegroup. Ensure that you are able to restore the PRIMARY filegroup backup by taking a full backup or a group backup of the PRIMARY filegroup before taking the user-defined backup.

Consult your Microsoft SQL Server documentation for more details on SQL Server backup strategy and planning.

Strategies defined by other considerations

Some commonly used strategies (based upon various considerations) are described as follows:

Saving time:

- If a SQL Server volume fails, restoring only the files that are on that volume can save restore time.
- Using multiple data stripes can speed up both backup and restore time. If backing up directly to sequential storage media such as tape pool, use as many stripes as there are tape drives that can be allocated to the SQL backup; otherwise, the separate sessions will queue up waiting for a tape. Striping is available with legacy operations only.
- Using data compression will reduce network traffic and storage requirements. However, whether it increases or decreases total backup time depends on several factors including the speed of the processors doing the compression and available network bandwidth. For fast networks, compression can increase the backup and restore times.

Data striping (Legacy only):

- If you use data striping, also use Tivoli Storage Manager server file space collocation to try to keep each stripe on a different storage volume. Use the Tivoli Storage Manager command **update stgpool** to set this parameter. Metadata (counted as a separate file space) is not to go to tape media.
- The maximum number of data stripes you can use must be smaller than the maximum supported by the SQL Server and less than the value of the Tivoli Storage Manager server `txngroupmax` option in the `dsmserv.opt` file.

Clustering:

If you use Microsoft Failover Clustering or Veritas Cluster Server clustering for fail-over support, you must install Data Protection for SQL Server on each cluster node and configure it identically. Additional setup is required to complete the fail-over installation. You must identify a clustered SQL Server by its virtual server name and use that name in Data Protection for SQL Server to access that SQL Server.

Truncate log on checkpoint option:

When you choose to perform only *full* backups in SQL, you can also indicate that you want to truncate the log after checkpoints. This will prevent the log from growing without bounds.

Truncate log option:

When you choose to perform a transaction log backup, you can indicate that you do not want to truncate the log. In general, you do not want to truncate the log when rebuilding a corrupt database. This option enables the server to back up the transaction log but does not try to touch the data in any way. It writes all transaction log entries from the time of the last log backup to the point of database corruption.

Collocation:

If you use the *full plus log* backup strategy, you must decide whether to modify Tivoli Storage Manager storage management policies to ensure that all log backups are stored together on the Tivoli Storage Manager server (collocated). This helps improve restore performance by reducing the number of media mounts necessary for restoring a series of log backups. Consult your Tivoli Storage Manager administrator for details on collocation.

Multiple SQL Servers:

If multiple instances of SQL Server are running, the additional instances are identified by name. You must use that name in Data Protection for SQL Server to access that SQL Server.

Miscellaneous:

- VSS backups cannot be restored to an alternate SQL Server. This is a Microsoft SQL Server limitation.
- You must use the **`maxnummp`** parameter on a Tivoli Storage Manager **`register node`** or **`update node`** command to allow a node to use multiple sessions to store data on removable media (which requires multiple mount points to be allocated to that node).
- Set backups are intended for special circumstances. If you plan to back up a set of filegroups and files regularly, back up each separately in order to use version limits within the management class.
- You cannot back up the *tempdb* database. It is a temporary database that is re-created each time the SQL Server is started.
- SQL databases with the `truncate log on checkpoint` option (*master* or *msdb*) or that use the *Simple* recovery model do not have transaction logs that can be backed up.
- Regardless of the frequency of database backups, run **dbcc checkdb** and **dbcc checkcatalog** on a database just before backing it up to check the logical and physical consistency of the database. See your SQL Server documentation for more information on using the SQL Server database consistency checker.
- Data Protection for SQL Server provides backup and restore functions for SQL databases and associated transaction logs. However, Data Protection for SQL Server does not provide a complete disaster recovery solution for a SQL Server by itself. There are many other files that are part of the SQL Server installation. These files would need to be recovered in a disaster recovery situation. Examples of these files are executable and configuration files. A comprehensive disaster recovery plan can be obtained by using the normal Tivoli Storage Manager backup-archive client for Windows, together with Data Protection for SQL Server.

Using VSS and Legacy Backups together

Using VSS backups and Legacy Backups together can implement a highly-effective backup solution for Data Protection for SQL Server data.

Although VSS supports only full backups, legacy differential and legacy log backups can be applied after a full VSS backup is restored.

Use the following good application practices:

- Legacy and VSS backups to Tivoli Storage Manager server storage are dictated by time, not by versions.
- Backups to local shadow volumes are dictated by versions because of space limitations and provisioning of VSS storage.
- When you run VSS operations, have at least 200 megabytes of free disk space on your Windows system folder. You can set this location with the **vssaltstagingdir** parameter.

Table 1. Backup strategy characteristics

Strategy characteristics	Legacy backup only	Legacy backup plus VSS backup
Available backup types	• FULL 1+ per week • DIFF 1+ per day • LOG 1+ per day	• Legacy FULL 1+ per week • VSS FULL 1+ per day • Legacy DIFF 1+ per day • Legacy LOG 1+ per day
Available restore types	Restore to production SQL Server or alternate SQL server	VSS, including VSS restore [1] Legacy: Restore to the production SQL Server or alternate SQL server
Restore attributes	• FULL, LOG, DIFF, GROUP, FILE • Server, database, filegroup, and individual file-level restore granularity • Point-in-time recovery • Roll-forward recovery • Restore to alternate system	VSS: • FULL • Database level restore granularity • Point-in-time recovery[2] • Roll-forward recovery[2] Legacy: • FULL, LOG, DIFF, GROUP, FILE • Server, database, filegroup, and individual file-level restore granularity • Point-in-time recovery • Roll-forward recovery • Restore to alternate system

Note:

1. Files are copied from the Tivoli Storage Manager server directly to the production source volumes.
2. To acquire these restore attributes, these backups must be Legacy log backups or Legacy differential backups that are applied to a full VSS backup.

Using VSS operations in a SQL Server Failover Cluster environment

Data Protection for SQL Server supports VSS operations in a SQL Server Failover Cluster environment.

These requirements and limitations must be understood in order for Data Protection for SQL Server to successfully perform VSS operations in a clustered SQL Server environment.

SQL Server Failover Cluster requirements for VSS

The following requirements must be met for VSS operations to perform successfully in a SQL Server Failover Cluster environment. These requirements are met when you use the configuration wizard.

- The ***vssaltstagingdir*** option must be specified when the following circumstances are true of your cluster environment:
  - Tivoli Storage Manager performs the VSS operations.
  - VSS backups are stored on local shadow volumes.
  - Make sure you have at least 200 megabytes of free disk space on the drive that the ***vssaltstagingdir*** option specifies. This space is used to hold the meta data files for Data Protection for SQL Server.

  This option must be specified in the dsm.opt file for all potential ***localdsmagentnode*** nodes that could be running the Tivoli Storage Manager Remote Client Agent Service (DSMAGENT):

  ```
  vssaltstagingdir d:\dir
  ```

  `d:` represents a shared drive that is accessible by all nodes in the cluster. It can also be a disk that follows the Virtual SQL Server. `\dir` represents a directory located on the shared drive. This option must be specified on all nodes that are used in the cluster. For example:

  ```
  vssaltstagingdir Q:\TSMVSS
  ```

- Specify the following options in each of the dsm.opt files that are used for the LOCALDSMAGENT and REMOTEDSMAGENT machines:

  ```
  CLUSTERNODE NO
  CLUSTERDISKSONLY NO
  ```

- When changing the ***vssaltstagingdir*** option, change the value to the same value in both the `dsm.opt` file for the DSMAGENT and the `dsm.opt` file for Data Protection for SQL Server.

SQL Server Failover Cluster limitations for VSS

Be aware of these limitations when performing VSS operations in a SQL Server Failover Cluster environment:

- All servers within the cluster must use the same levels of Tivoli Storage Manager, Windows, and other applicable software.
- The Tivoli Storage Manager Client Acceptor Daemon (CAD) must be installed on each cluster node so that it can continue operations in the event of a failover. Make sure the CAD service name is the same on all cluster nodes so that it can be started by a generic cluster service.
- The Local DSMAgent client node should be a separate node from your normal backup-archive client, as this CAD service needs to be made a non-cluster option.

- When using the Remote DSMAgent client node, you do not need to register a separate node for each server within the cluster.
- Use the Microsoft **vssadmin** and **diskshadow** commands to verify the environment. For more information about using these commands, see the VSS diagnostic wizard.
- A Data Protection for SQL Server configuration file should be configured for each node in the cluster. These files are almost identical, except that the ***localdsmagentnode*** parameter points to the corresponding local DSMAgent on each node.

SQL Server Failover Cluster VSS limitations for scheduled operations

If you plan to perform scheduled VSS operations in a SQL Server Failover Cluster environment, be aware of these considerations:

- Install the Tivoli Storage Manager scheduler as a Windows service on all cluster nodes.
- If the command file resides on a local drive, you must make sure that it remains consistent on all cluster nodes. Optionally, you can create the command file on a shared drive. Make sure the ***objects*** parameter (specified with the **define schedule** command on the Tivoli Storage Manager server) points to this command file.

Considerations for using Data Protection for SQL Server in a Windows Failover Cluster environment

Data Protection for SQL Server supports SQL Server 2012 running in a Windows Failover Cluster environment.

Consider the following information before you use Data Protection for SQL Server to protect SQL Server 2012 databases in a Windows Failover Cluster environment. References to the SQL Server pertain to the virtual SQL Server name in a Windows Failover Cluster environment.

Cluster setup considerations:

- A Windows Failover Cluster environment is required for AlwaysOn Availability Groups.
- A SQL Server instance must be installed on a node in a Windows Failover Cluster environment. The cluster node must be online.
- Each availability replica of an availability group must be on a different node in the same Windows Failover Cluster environment.

Installation and configuration considerations:

- Install Data Protection for SQL Server on all nodes from where you intend to run backup and restore operations.
- When you use shared disk clusters, install Data Protection for SQL Server on all nodes on a disk that is local to each node and not on a shared cluster disk.
- Use the configuration wizard to register an AlwaysOn Node on the Tivoli Storage Manager server. The AlwaysOn Node manages backups of availability databases. This node is a shared node that allows backups and restores of availability databases from any replica.

- Databases that are not in an availability group are backed up under the standard Data Protection for SQL Server node name. To migrate your database backups to the AlwaysOn node, an option is available for you to back up all databases to the AlwaysOn node.
- Use identical configurations in the Data Protection for SQL Server options file when you configure Data Protection for SQL Server on each node of the cluster.
- If you are using the Tivoli Storage Manager scheduler for automating backups, install the scheduler service on each node of the cluster to enable failover support.

Operational considerations:

- The Tivoli Storage Manager server treats backups as coming from a single SQL Server (the virtual server) regardless of which node of the cluster was backed up.
- When you run a VSS full backup on a secondary replica, the copy-only type of backup is used to take the snapshot.
- Legacy full backups of availability databases on secondary replicas are copy-only. The copy-only option is not automatically used with log backups because log backups that truncate logs are supported on secondary replicas.
- Because of a limitation with the SQL Server, you cannot restore a VSS backup to an alternate instance. VSS backups must be restored on the same SQL Server instance where the snapshot was taken.

Using Data Protection for SQL Server in a Windows Failover Cluster environment

Data Protection for SQL Server supports SQL Server running in a Windows Failover Cluster environment.

The list below provides information to consider when running Data Protection for SQL Server in a Windows Failover Cluster environment.

- References to the SQL Server made in this section pertain to the virtual SQL Server name in a Windows Failover Cluster environment.
- You must install Data Protection for SQL Server on all nodes of the cluster. In addition, when installing Data Protection for SQL Server, you must install it on a disk local to each node (not on a shared cluster disk).
- Use identical configurations in the Data Protection for SQL Server options file when configuring Data Protection for SQL Server on each node of the cluster.
- If you are using the Tivoli Storage Manager scheduler for automating backups, you must install the scheduler service on each node of the cluster to enable failover support.
- The Tivoli Storage Manager treats backups as coming from a single server (the virtual server) regardless of which node of the cluster a backup was performed on.

Using Data Protection for SQL Server in a Veritas Cluster Server environment

Data Protection for SQL Server supports SQL Server running in a VCS environment.

The following list provides information to consider when running Data Protection for SQL Server in a Veritas Cluster Server Environment.

- References to the SQL Server made in this section pertain to the virtual SQL Server name in an VCS environment.
- You must install Data Protection for SQL Server on all nodes of the cluster. In addition, when installing Data Protection for SQL Server, you must install it on a disk local to each node (not on a shared cluster disk).
- Use identical configurations in the Data Protection for SQL Server options file when configuring Data Protection for SQL Server on each node of the cluster.
- If you are using the Tivoli Storage Manager scheduler for automating backups, you must install the scheduler service on each node of the cluster to enable failover support.
- The Tivoli Storage Manager treats backups as coming from a single server (the virtual server) regardless of which node of the cluster a backup was performed on.

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

When creating 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 SQL 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 SQL Server 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.

Data Protection for SQL Server with SAN Volume Controller and Storwize V7000

Data Protection for SQL Server 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 SQL Server 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 SQL Server 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 SQL Server 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 SQL Server 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.

- The Windows host must be attached to a SAN Volume Controller or Storwize V7000 cluster. The volumes that are assigned to the Windows host must be participating in the SAN Volume Controller or Storwize V7000 cluster that are attached to the SAN Volume Controller.
- 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 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.

VSS limitations for SAN Volume Controller and Storwize V7000

When you run a Data Protection for SQL Server 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 SQL 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 SQL Server VSS backups and if the volumes are large, or the background copy rate is set to a low number, or both conditions occur, the second VSS backup might be presented to be in a hang state. In this situation, the SQL 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.

Related concepts:

"More guidelines for SAN Volume Controller and Storwize V7000 environments" on page 24

VSS operations in IBM N-series and NetApp environments

Be aware of these storage space guidelines when performing 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:

```
SQL Database size:  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 SQL Server 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 SQL Server VSS operations can be fully used. However, this situation is not applicable when FlexVols are used.

SQL 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 SQL Server databases cannot share LUNs.
- A NAS filer LUN used by the SQL Server databases must be the only LUN on the filer volume. For example, if SQL 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.
- 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 SQL 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.

AlwaysOn Availability Groups

Microsoft SQL Server 2012 uses AlwaysOn Availability Groups to provide high availability and disaster recovery capabilities.

Data Protection for SQL Server protects availability databases in an AlwaysOn Availability Group and AlwaysOn Failover Cluster Instances. It provides high availability and disaster recovery at the SQL Server database level and SQL Server instance level.

Backups of availability databases

Data Protection for SQL Server backs up and restores each availability database as a single object, regardless of which availability replica is used for backup and restore operations.

An AlwaysOn Availability Group can contain a set of primary databases and one to four copies of the set of primary databases, called secondary databases. Databases in an availability group are called availability databases, and they fail over together as a group.

An AlwaysOn Availability Group requires SQL Server instances on Windows Failover Cluster nodes. An availability group can have a number of replicas. For example, availability group *1* might have replicas *node1*, *node2*, and *node3*.

A cluster node might be an availability replica for one or more availability groups. For example, *node1* might be a replica for availability group *1* and another availability group. For the secondary replica, *read-only* is an option that can be set at the availability group level.

The AlwaysOn Node is used to manage backups of availability databases. When you work in a Tivoli Storage Manager environment, the AlwaysOn Node is to be common in a Windows Failover Cluster. This presence enables the management of backups of an availability database in a single location, regardless of the replica that is used to perform the backup.

The following types of VSS backup operations are supported:

- Full VSS backups of the primary availability replica
- VSS copy-only full backups of availability replicas

The following types of legacy backup operations are supported:

- On the primary replica, legacy full, differential, file, set, group, and log backups are supported.
- On the secondary replica, legacy full, file, set, group, and log backups are supported.
- VSS and legacy copy-only full backups, legacy copy-only file, set, or group backups, and legacy copy-only and normal log backups are supported.

For all backup operations of secondary availability replicas, the secondary replicas must be in the synchronized or synchronizing state.

To assist with scheduling and load balancing, scheduled backup preference settings of availability groups are also available.

Restores of availability databases

Depending on how availability databases are backed up, legacy restore and VSS restore operations are available to restore the availability databases on primary or secondary availability replicas.

There are some restrictions with restoring availability databases:

Legacy restore
You can restore an availability database on either a primary or secondary replica.

During the restore process, the restored database is removed from the availability group. When a database is removed from the availability group, the database becomes a local database on that replica. The database is restored as a local database. After this restore is complete, manually add the database back to the availability group. However, before you add the database to the availability group, verify that the data on all replicas is transactionally consistent.

To verify data is transactionally consisent, verify that the backup copy contains data and transaction log records. Full backups and differential backups contain data and transaction log records so that the restored database is transactionally consistent.

After you verify that the data is transactionally consistent, the database can be added to the availability group.

VSS restore

Because of a SQL Server limitation, you cannot restore a VSS backup to an alternative SQL server instance. Therefore, VSS backups must be restored to the same SQL server instance where the snapshot was taken.

Deployment of Data Protection for SQL Server on Windows Server Core

If you are protecting Microsoft SQL Server 2012 and later databases, you can install and use Data Protection for SQL Server on Windows Server 2008 R2 Server Core SP1 and later.

Server Core is a minimal and low-maintenance server environment where you can run the minimum services necessary to maintain Windows Server 2008 and later. You can install and operate Data Protection for SQL Server in this minimal server environment.

Because of the minimal environment, only the command-line interface is available for Data Protection for SQL Server on Windows Server Core. Also, if you use Windows Installer (MSI) to install Data Protection for SQL Server, only the unattended mode is supported.

How Tivoli Storage Manager server policy affects Data Protection for SQL Server

Tivoli Storage Manager policy determines how Data Protection for SQL Server 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 SQL Server as a node. Data that is backed up to Tivoli Storage Manager storage from this Data Protection for SQL Server 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 SQL Server, 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 tdpsql.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.

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 SQL Server 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.

Policy binding statements

Policy binding statements associate SQL backups to a management policy.

About this task

Although Tivoli Storage Manager policy determines how Data Protection for SQL Server 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. Ensure 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 required is dependent on the VSS provider that is used.

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.

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

VSSPOLICY	**<server name>**	**<object name>**	**<backup type>**	**<backup dest>**	**<mgmt class>**
VSSPOLICY	*	acctdb1	FULL	LOCAL	MC_1
VSSPOLICY	SERVER_3	hrdb	INCR	LOCAL	MC_6

How backups expire based on policy

Backups expire based on Data Protection for SQL Server policy.

Expiration is the process by which SQL 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 SQL Server 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 SQL Server session.

For AlwaysOn Availability Groups on SQL Server 2012, only the system on which the backup was created can cause a local backup to expire. As an example, 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 and 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 SQL Server policy.

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

Binding backups to a policy

A backup policy determines how backup versions are retained. You can add, update, delete, or change the processing order of existing binding statements.

Procedure

1. From the Management Console, select the **SQL Server** instance from the tree view.
2. On 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 SQL 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 at all, the more general specification should be listed before the more specific ones, so as to be processed after the more specific specifications. Otherwise, the more general specification will 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.

Data Protection for SQL Server node name: recommended settings

Review these recommended settings when registering your Data Protection for SQL Server node name.

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

Use the following Tivoli Storage Manager parameter conditions when you register your Data Protection for SQL Server node name (system) to the Tivoli Storage Manager server:

- **BACKDELete** This parameter determines whether the Data Protection for SQL Server node can delete its own backup files from the Tivoli Storage Manager server. This parameter must have a value of `yes`.
- **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. If you are to use SQL data-striping with data that is sent directly to a tape pool, this parameter must be set to a number greater than the default value of *1*. For example, set this value to be at least the maximum number of stripes to be used for backup or restore when removable media such as tapes are used or if migration occurs during the backup or restore operation. If other backups or restores occur at the same time, the value of this parameter must be large enough to allow for all of the needed mount points. If the storage pool for the backup has Active Data or Backup Stgpools that are going to be written to simultaneously, this **MAXNUMMP** parameter must also include these mount points.
- **TXNGroupmax** This parameter determines the number of files that are transferred as a group between Data Protection for SQL Server and the Tivoli Storage Manager server between transaction commit points. This parameter must have a value of one more than the maximum number of stripes to be used for backup or restore operations regardless of media.
- **COMPRESSIon** (Legacy only) This parameter determines whether the Data Protection for SQL Server node compresses data before it sends the data to the Tivoli Storage Manager server during a backup operation. Specify `COMPression=Client` to allow the Data Protection for SQL Server node to decide whether to compress data by using the value of the client **COMPRESSIon** option that is specified in the options file (`dsm.opt`) in the Data Protection for SQL Server directory.

If you are running Data Protection for SQL Server on a Microsoft Failover Clustering or Veritas Cluster Server, the node name cannot be the name of the local system. Instead, the node name is to match the SQL virtual server name. See the *IBM Tivoli Storage Manager Window's Administrator's Reference* for complete information about these parameters.

Proxy node definitions (VSS backups)

Since Data Protection for SQL Server 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 SQL Server 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 Manager`USerid=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 SQL Server dsm.opt file.

- *Agent node*: A node name that processes operations on behalf of a target node. This node name is specified in the Backup-Archive Client dsm.opt file.

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

```
GRANT PROXY TARGET=dpsql_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 SQL Server node name.	Use the nodename option in the Data Protection for SQL Server 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 SQL Server configuration file (tdpsql.cfg)

Target node

This node name is where Data Protection for SQL Server 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 SQL Server 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 SQL Server 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 SQL Server configuration file (tdpsql.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 **tdpsqlc 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 SQL Server node name	Use the **nodename** option in the Data Protection for SQL Server options file (dsm.opt)

Table 3. Required node names for basic VSS offloaded backups (continued)

Proxy node type	Node name	Where to specify
Agent node	Local DSMAGENT Node	Use the **localdsmagentnode** parameter in the Data Protection for SQL Server configuration file (tdpsql.cfg)
Agent node	Remote DSMAGENT Node	Use the **remotedsmagentnode** parameter in the Data Protection for SQL Server configuration file (tdpsql.cfg)

Target node

This node name is where Data Protection for SQL Server is installed. This node name (specified with the **nodename** option in the dsm.opt file) is referred to as the Data Protection for SQL Server 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 SQL Server 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 SQL Server configuration file (tdpsql.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 **tdpsqlc set** command to specify this parameter.

Agent node

This node name is a separate system that must also have the backup-archive client and VSS provider installed. This node is responsible for moving VSS snapshot data from local shadow volumes to the Tivoli Storage Manager server. This node name is referred to as the Remote DSMAGENT Node and is specified with the **remotedsmagentnode** parameter in the Data Protection for SQL Server configuration file (tdpsql.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 **tdpsqlc 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 SQL Server options file (dsm.opt) and the backup-archive client options file (also dsm.opt).

Specifying Data Protection for SQL Server options

Once Data Protection for SQL Server is registered to Tivoli Storage Manager, several Data Protection for SQL Server 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 SQL Server 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 SQL Server 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 SQL Server node with the Tivoli Storage Manager server. Data Protection for SQL Server 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 SQL Server options file. In addition, for VSS backups, 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 SQL Server.

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

`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 SQL Server 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 SQL Server and Tivoli Storage Manager server are over a low bandwidth connection.
- There is heavy network traffic.

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

- The computer that runs Data Protection for SQL Server has a CPU overload; the added CPU usage can impact other applications that include the SQL 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 SQL Server 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 legacy backups, specify the **compression** option in the Data Protection for SQL Server options file.
- 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.

See "How to enable SQL Server backup compression" on page 43.

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 SQL Server 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 SQL 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 yes`in 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 Legacy backups, specify these encryption options in the Data Protection for SQL Server options file.
- 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 SQL 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.
3. (Legacy backups only) Add your include statements with the `include.encrypt` option. For example, to encrypt all SQL data, specify the following statement:

   ```
   include.encrypt *\...\*
   ```

   To encrypt all objects with database name *Db1* beginning with *Db*, specify the following statement:

   ```
   include.encrypt \...\Db*\...\*
   ```

   To encrypt all full or differential objects with database name *Db1*, specify the following statements:

   ```
   include.encrypt \...\Db1\full*
   include.encrypt \...\Db1\diff*
   ```

ENABLELANFree

This option allows Data Protection for SQL Server to run in a LAN-free environment (if you are equipped to do so). To run a LAN-free Legacy backup with Data Protection for SQL Server, a Tivoli Storage Manager Storage Agent must be installed on the same system and `enablelanfree yes` must be specified in the Data Protection for SQL Server options file. To run a LAN-free VSS backup with Data Protection for SQL Server, 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.

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 SQL Server 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.

Related concepts:

"Specifying Data Protection for SQL Server preferences"

Specifying Data Protection for SQL Server preferences

Data Protection for SQL Server configuration parameters are defined in the Data Protection for SQL Server configuration file (tdpsql.cfg by default). These configuration parameters determine such preferences as the location of your log file, how date and time stamps display, and the number of buffers to use.

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

- In the MMC GUI, set the value in Properties.
- Use the **`tdpsqlc set`** command in the Data Protection for SQL Server command-line interface. See "Set command" on page 238.

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.

Note the following characteristics of Data Protection for SQL Server configuration parameters:

- The value of a configuration parameter that is specified on a command-line invocation overrides (but does not change) the value of the configuration parameter that is specified in the Data Protection for SQL Server configuration file.
- During a command-line invocation that does not specify an overriding value for a configuration file parameter, the values in the default Data Protection for SQL Server configuration file (`tdpsql.cfg`) are used.

See "Set command" on page 238 for descriptions of available configuration parameters.

Set policy for VSS backups by specifying the **`VSSPOLICY`** statement in your Data Protection for SQL Server configuration file, or by selecting **Properties** > **VSS Policy Binding** in the MMC GUI, and then configuring the policy. You must specify this statement manually. You cannot specify it using the **`tdpsqlc set`** command. See "Setting automatic expiration (VSS and legacy)" on page 41 for detailed information and examples.

Setting automatic expiration (VSS and legacy)

Data Protection for SQL allows you to utilize Tivoli Storage Manager automatic expiration and version control by policy. You set automatic policy for backup data by editing the Data Protection for SQL Server options file, or by specifying them in the Management Console (MMC) GUI (**Utilities** > **VSS Policy Binding**). If you edit the options file, use include and exclude statements to define which files are subject to automatic processing, and to assign specific management classes to files using object naming conventions.

Setting automatic expiration for VSS

Set policy for VSS backups by specifying the VSSPOLICY statement in your Data Protection for SQL Server configuration file, or by specifying them in the MMC GUI (**Properties** > **VSS Policy Binding**). Note that you cannot specify it using the **`tdpsqlc set`** command.

VSSPOLICY statements 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 should be listed before the more specific ones, so as to be processed after the more specific specifications. Otherwise, the more general specification will match the target before the more specific specifications are seen.

Specify the following information in the VSSPOLICY statement:

```
VSSPOLICY srvname dbname backuptype backupdest mgmtcls
```

The statement contains the following values:

srvname
Name of the SQL Server or wildcard character (*)

dbname
Name of database or wildcard character (*)

backuptype
Backup type: FULL or wildcard character (*). Because VSS supports only full backup types, when you specify a wildcard character for *backuptype*, a FULL backup type is used.

backupdest
Backup destination: TSM or LOCAL or wildcard character (*)

mgmtcls
Management Class name. This sets the management class for the specified class of backup.

Setting automatic expiration for legacy

Data Protection for SQL Server allows you to utilize Tivoli Storage Manager automatic expiration and version control by policy. Setting automatic policy for backup data is accomplished through the Data Protection for SQL options file. Use include and exclude statements in the options file to define which files are subject to automatic processing, and to assign specific management classes to files using object naming conventions.

Ensure metadata is available for query without causing a volume mount. The metadata is stored as a data object on the Tivoli Storage Manager server and is available for migration to removable media if its policy allows this to occur.

A Data Protection for SQL backup object name is composed of a series of qualifiers separated by \.

The general include/exclude syntax for object naming is:

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

where:

objectNameSpecification is:

```
SqlServerName[\InstanceName]\dataType\...\DatabaseName
[\typeInfo]\backupType*
```

dataType is:
`meta | data`

typeInfo is:
`LogicalFileName` (for **file** backup type)

`GroupName` (for **group** backup type)

`...` (for **log** and **set** backup types)

not used for **full** and **diff** backup types

backupType is:
`full | diff | log | group | file | set`

Considerations:

- 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: **\...**
- All specifications must end with the wildcard character *. This is required because the specification must match both object names and temporary names. Temporary names are used to enable rolling back a backup transaction if an error occurs. Temporary names are object names with a unique string appended to the *backupType* qualifier.
- An *objectNameSpecification* should be placed within double quotation marks. If the specification includes spaces or special characters, the double quotes are required.
- For exclude statements, **meta** should be a match in the specification, either explicitly, or by wildcard or ellipses. Excluding **meta** excludes the entire object.
- Include statements can specify either **meta** or **data** separately and explicitly, or both by wildcard or ellipses.
- You may specify both data and meta objects in options file include lists in order to assign management classes. However, when you use exclude statements, you should specify only the meta objects. If a data object is not backed up, its meta object will not be created.
- **Log** and **set** object names are always unique. The *typeInfo* contains the qualifiers whose values make them unique. Because they are generated at the time of the backup, they are not predictable and you cannot specify them.
- 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, you should list the more general specifications before the more specific ones so

that they will be processed after the specific. Otherwise, the more general will match the target before the more specific 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 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.
- Include/exclude processing does not produce error messages for invalid specifications. Therefore, you should thoroughly test all include/exclude lists. Specifying an invalid management class name will generate an error message in the dsierror.log.
- For case-sensitivity, the Windows Tivoli Storage Manager API currently assumes the specifications are for a Windows file system and ignores case. However, because case may be honored in the future, you should always use the correct case. Specifically,
  - Use correct case for SQL names (server, database, group, or file names) as displayed by the **query sql** or **query tsm** commands.
  - Use lowercase for the Data Protection for SQL constants: **meta**, **data**, **full**, **diff**, **log**, **group**, **file**, and **set**.

The following are examples of individual *objectNameSpecifications* as they might appear in include/exclude statements:

SqlServerNames:
SQL2008, SQL2012

InstanceNames:
INST1, INST2

DatabaseNames:
Db1, Db2, Db3

GroupNames:
g1, g2, g3

LogicalFileNames:
f1, f2, f3

How to enable SQL Server backup compression

Support for SQL Server backup compression is available on Data Protection for SQL Server. You can use either the MMC GUI or the command line to enable this feature.

- From the MMC GUI, specify SQL native backup compression from the SQL Properties window. After you have set this option, the **SQL Workload** column on the Recover tab shows the SQL compression status for legacy backups.
- From the command line, add this statement to the SQL configuration file (`tdpsql.cfg`). Edit the file and enter the command as shown here:

```
SQLCOMPression Yes | No
```

The default value is `No`.

SQL Server backup compression is only available with Legacy backups on SQL Server. For SQL Server, backup compression is only supported on Enterprise Edition. SQL Server 2008 R2, backup compression is supported on Standard, Enterprise, and Datacenter editions. Starting with SQL Server 2008, any edition can restore a compressed backup.

SQL Server backup compression is generally faster and more effective than using it together with Tivoli Storage Manager compression.

Chapter 3. Installing

Review the appropriate prerequisite information, including hardware and software requirements, before installing Data Protection for SQL Server.

Quick installation and configuration

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

Before you begin

Before you begin, verify that hardware and software requirements are met. 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 SQL Server program. This technote is available in the *TSM for Databases - All Requirement Documents* website at http://www.ibm.com/support/docview.wss?uid=swg21218747. When you are at the website, click 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 SQL Server:

1. Install Data Protection for SQL Server.
   a. Log on as an administrator.
   a. Insert the Data Protection for SQL Server 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 SQL Server.
   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 SQL Server.
   a. Start the Management Console (MMC GUI) by clicking **Start** > **All Programs** > **Tivoli Storage Manager** > **Data Protection for Microsoft SQL Server** > **DP for SQL Management Console**. If you did not previously configure Data Protection for SQL Server, 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 **SQL 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 SQL 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 SQL** field, enter the node name where the Data Protection for SQL Server application is installed. This node stores the Data Protection for SQL Server backups. Do not use double-byte characters (DBCS).
- If you are configuring Data Protection for SQL Server in an SQL Server 2012 environment, enter a node name in the **AlwaysOn Node** field. This node name is used when the availability databases are backed up in an AlwaysOn Availability Group.

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 SQL Server 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_sql.txt` file. Then, click **Open**.
4) These commands are displayed in the command-line panel:

```
query tdp
query tsm
query sql
```

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 SQL Server to back up and restore SQL server data.
6) Back up and restore a set of test data.

4. Customize Data Protection for SQL Server.

After Data Protection for SQL Server 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

If you are installing Data Protection for SQL Server in a Windows Failover Cluster environment or Veritas Cluster server environment, repeat the installation procedure on the nodes of your cluster that you want to protect.

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 SQL Server program. This technote is available in the *TSM for Databases - All Requirement Documents* website at http://www.ibm.com/support/docview.wss?uid=swg21218747. 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 SQL Server:

Hardware for an x86 system
Compatible hardware that is supported by the Windows operating system and SQL Server in use.

Hardware for an x64 system
Compatible hardware that is supported by the Windows operating system and SQL Server in use.

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=swg21218747

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

Minimum software and operating system requirements

Before installing Data Protection for SQL Server, ensure that your system meets the minimum software and operating requirements.

Details of the software and operating system requirements for Data Protection for SQL Server can change over time. For current software requirements, see the *TSM for Databases - All Requirements Documents* website at http://www.ibm.com/support/docview.wss?uid=swg21218747.

Virtualization environment

Information about virtualization environments that can be used with Data Protection for SQL Server 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 Data Protection for SQL Server on a local system

The setup wizard guides you through installing Data Protection for SQL Server.

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=swg21218747

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 SQL Server is available in both licensed and maintenance packages. The installation process differs between these two package types.

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

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

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 SQL Server from a DVD, complete the following steps:

1. Install Data Protection for SQL Server 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 SQL Server 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.

      If you are configuring Data Protection for SQL Server in an SQL Server 2012 environment, enter a node name in the **AlwaysOn Node** field in the TSM Node Names page in the configuration wizard. This node name is used to back up the availability database backups.
   d. If prompted, restart your system before the installation completes.
   e. Click **Finish** to complete the installation of Data Protection for SQL Server.

   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. If you are installing Data Protection for SQL Server in a Microsoft Windows Failover Clustering environment or Veritas Cluster server environment, repeat the installation procedure on all nodes of your cluster.
3. To install more language packs, see “Installing and activating the language packs” on page 51.
4. 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” on page 51. If not, for important configuration information see “Configuring Data Protection for SQL Server” on page 65.

Installing Data Protection for SQL Server on Windows Server Core

If you are protecting Microsoft SQL Server 2012 and later service packs data in a Windows Server Core environment, you can use the setup wizard to install Data Protection for SQL Server.

Before you begin

Before you begin, verify that your environment meets the hardware and software prerequisites.

For the most current requirements, review the Hardware and Software Requirements technote that is associated with the level of your Data Protection for SQL Server. This technote is available in the *TSM for Databases - All Requirements Documents* website, at http://www.ibm.com/support/docview.wss?uid=swg21218747. When you are at the website, follow the link to the requirements technote for your specific release or update level.

Procedure

Follow these instructions to install Data Protection for SQL Server on Windows Server Core from a DVD:

1. Log on as an administrator.
2. Install Data Protection for SQL Server by using the setup wizard.
   a. Insert the Data Protection for SQL Server 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:\setup.exe`, where `x:` is your DVD drive, and click **`OK`**.

   b. Follow the installation instructions that are displayed on the screen.
   c. Click **Finish** to complete the installation. If prompted, restart your system.

What to do next

You are ready to configure Data Protection for SQL Server.

You can also do an unattended installation of Data Protection for SQL Server on Windows Server Core.

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 SQL Server 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 SQL Server and Tivoli Storage FlashCopy Manager, see “Configuring Data Protection for SQL Server” on page 65 for important configuration information.

Installing Data Protection for SQL Server in a cluster environment

You can install Data Protection for SQL Server on a Windows Failover Cluster environment to protect clustered SQL Server 2008 and SQL Server 2012 data.

About this task

Installing Data Protection for SQL Server in a Windows Failover Cluster environment requires the following:

Procedure

1. Install Data Protection for SQL Server on all nodes of your cluster from where you intend to run backups and restores.
2. If you are using a shared disk cluster, install Data Protection for SQL Server on all nodes on a disk that is local to each node and not on a shared cluster disk.
3. Follow the instructions in “Installing Data Protection for SQL Server on a local system” on page 48 for all nodes of your cluster.

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 SQL Server 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 SQL Server Management Console language packs
   `x:\fcm\`*`aaa`*`\mmc\4100\`*`bbb`*`\setup.exe`

   Data Protection for SQL Server language packs
   `x:\fcm\`*`aaa`*`\languages\`*`bbb`*`\setup.exe`

   Where `x:` is your DVD drive, *aaa* is either `x86` or `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 SQL Server configuration file (`tdpsql.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:

  ```
  tdpsqlc set lang=fra
  ```

  See the description of the **`language`** parameter in "Set positional parameters" on page 239 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 SQL 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 SQL Server silently

A silent installation runs on its own without any user interaction, and is considered unattended. Administrators can install Data Protection for SQL Server by using silent installation.

Silent installation is useful when Data Protection for SQL Server must be installed on a number of different computers with identical hardware. For example, a company might have 25 SQL servers that are installed across 25 different sites.

To ensure a consistent configuration and to avoid having 25 different people enter Data Protection for SQL Server parameters, an administrator can choose to produce an unattended installation package and make it available to the 25 sites. The installation package can be placed on a DVD and sent to each of the remote sites, or the package can be placed on a file server for distribution across the different sites.

You can run a silent installation by using one of the following methods:

Setup Program
Use the **`setup`** command with the command-line invocation and special silent installation options.

Microsoft Installer (MSI)
Use `msiexec.exe` to install the MSI package.

The following options can be used with both silent installation methods:

Table 4. Silent installation options

Option	Description
`/i`	Specifies the program is to install the product.
`/l*v`	Specifies verbose logging.
`/qn`	Runs the installation without running the external user interface sequence.
`/s`	Specifies silent mode.

Table 4. Silent installation options (continued)

Option	Description
/v	Specifies the Setup Program to pass the parameter string to the call it makes to the MSI executable program (`msiexec.exe`). Note the following syntax requirements when you use the **/v** option: • A backslash (\) must be placed in front of any quotation marks (" ") that are within existing quotation marks. • Do not include a space between the **/v** command-line option and its arguments. • Multiple parameters that are entered with the **/v** command-line option must be separated with a space. • You can create a log file by specifying the directory and file name at the end of the command. The directory must exist when a silent installation is done.
/x	Specifies the program is to uninstall the product.
addlocal	Specifies features to install.
allusers	Specifies which users can use the installation package.
installdir	Specifies the directory where Data Protection for SQL Server is to be installed.
reboot	Specifies whether to prompt the user to restart the system after silent installation. **Force** Always prompts user to restart after silent installation. **Suppress** Suppresses prompt to restart after silent installation. **ReallySuppress** Suppresses all restarts and prompts to restart after silent installation.
rebootyesno	Specifies whether to restart the system after silent installation. Specify *Yes* to restart the system after silent installation. Specify *No* not to restart the system after silent installation.
transform	Specifies language to install.

Note: Setting the **rebootyesno** option to *No* applies only to the installation of theData Protection for SQL Server software. The installation package includes a number of prerequisites that is installed by Data Protection for SQL Server if they are not installed as prerequisites onto the system. Ensure that all the prerequisites are installed before you start the silent installation, then set the **rebootyesno** option to *No* so that no system restart is required after the silent installation process finishes.

The following features are used in this procedure and are case-sensitive:

Table 5. Silent installation features (base client only)

Feature	Description
Client	Data Protection for SQL Server code

Table 6. Silent installation features (Language Packages only)

Feature	Description
LanguageFiles	Language-specific files

The following transforms are used in this procedure.

Table 7. Silent installation transforms

Transform	Language
1028.mst	CHT Chinese (Traditional)
1031.mst	DEU German
1033.mst	ENG English
1034.mst	ESP Spanish
1036.mst	FRA French
1040.mst	ITA Italian
1041.mst	JPN Japanese
1042.mst	KOR Korean
1046.mst	PTB Portuguese
2052.mst	CHS Chinese (Simplified)

Installing Tivoli Storage Manager client silently

Before you can install Data Protection for SQL Server on Windows Server Core, you must first install the Tivoli Storage Manager client on the same computer as Data Protection for SQL Server.

Before you begin

Before you begin, ensure that the most recent supported version of the Tivoli Storage Manager backup-archive client is available on your computer. You must be an administrator to do this installation.

About this task

Use Windows Installer program (**`msiexec.exe`**) to install the client.

To silently install the Tivoli Storage Manager client:

Run the command to silently install Tivoli Storage Manager client. For example, issue the following command on a single line from a "Run as Administrator" command prompt window:

```
msiexec /i "G:\tsm_images\TSM_BA_Client\IBM Tivoli Storage
Manager Client.msi" RebootYesNo="No" REBOOT="ReallySuppress" ALLUSERS=1
INSTALLDIR="c:\program files\tivoli\tsm" ADDLOCAL="Client,AdministrativeCmd"
TRANSFORMS="G:\tsm_images\TSM_BA_Client\1033.mst"
/qn /l*v "C:\downloads\logs\ba_logs.log"
```

Silent installation with the setup program

You can use the **setup** program (setup.exe) to silently install Data Protection for SQL Server. If you are protecting Microsoft SQL Server 2012 data, you can also use the setup program to silently install Data Protection for SQL Server on Windows Server Core.

Silently installing Data Protection for SQL Server with the setup program

Use the setup program to silently install Data Protection for SQL Server.

Before you begin

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

Data Protection for SQL Server Management Console setup program
(64-bit) `x:\fcm\x64\mmc\4100\enu\setupfcm.exe`

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

The Data Protection for SQL Server Management Console and Data Protection for SQL Server must be installed from an account that is a member of the local Administrators group for the system on which the SQL 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\sql\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 SQL Server silently” on page 53.

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\DpSqlMmcSetupLog.txt\""
```

```
x:\fcm\x64\sql\7100\enu\setup.exe /s /v"INSTALLDIR=\"C:\Program Files\Tivoli\tsm\"
ADDLOCAL=\"Client\" TRANSFORM=1033.mst REBOOT=ReallySuppress /qn /l*v
\"C:\Temp\DpSqlSetupLog.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 SQL Server software. The installation package includes a number of prerequisites that is installed by Data Protection for SQL Server. Ensure that all the prerequisites are installed before starting the silent installation, then set the **rebootyesno** option to *No* so that no system restart is required after the silent installation process finishes.

Silently installing Data Protection for SQL Server with the setup command on Windows Server core

If you are protecting Microsoft SQL Server 2012 data on Windows Server Core, you can use the **setup** program (`setup.exe`) to silently install Data Protection for SQL Server.

Before you begin

The Data Protection for SQL Server setup program is on the installation media (where `x:` is your DVD drive):

- (32-bit) `x:\fcm\x86\sql\7100\enu\setup.exe`
- (64-bit) `x:\fcm\x64\sql\7100\enu\setup.exe`

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

Run the following command to silently install Data Protection for SQL Server to the default installation directory:

```
x:\fcm\aaa\sql\7100\enu\setup.exe /s /v/qn
```

where `x:` is your DVD drive and `aaa` is either *x64* or *x86*.

You must substitute the appropriate feature when you install a language other than English. For more information, see “Installing Data Protection for SQL Server silently” on page 53.

The following command is an example that specifies the target directory, the features, language transform, start suppression, and logging. Specify the command on a single line from a Run as Administrator command prompt window.

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

```
x:\fcm\x64\sql\7100\enu\setup.exe /s /v"INSTALLDIR=\"C:\Program Files\Tivoli\tsm\"
ADDLOCAL=\"Client,License_Paid\" TRANSFORM=1033.mst /qn /l*v
\"C:\temp\logs\fcm.log\""
```

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 SQL Server software. The installation package includes a number of prerequisites that are installed by Data Protection for SQL Server. Ensure that all the prerequisites are installed before you start the silent installation, then set the **rebootyesno** option to *No* so that no system restart is required after the silent installation process finishes.

  Tip: For the most current requirements, review the Hardware and Software Requirements technote that is associated with the level of your Data Protection for SQL Server program. This technote is available in the *TSM for Databases - All Requirements Documents* website, at http://www.ibm.com/support/docview.wss?uid=swg21218747. When you are at the website, follow the link to the requirements technote for your specific release or update level.

What to do next

You are ready to configure Data Protection for SQL Server.

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\DpSqlMmcSetupLog.txt\""
 rem
 call x:\fcm\x64\sql\6400\enu\setup.exe /s
 /v"INSTALLDIR=\"C:\Program Files\Tivoli\tsm\" ADDLOCAL=\"Client\"
 TRANSFORM=1033.mst REBOOT=ReallySuppress /qn /l*v \"C:\Temp\DpSqlSetupLog.txt\""
 rem ====================================
 rem code could be added after the
 rem installation completes to
 rem customize the dsm.opt files
 rem if desired
 rem ====================================
```

Silent installation with MSI (msiexec.exe)

You can use the Microsoft Installer (MSI) program, (`msiexec.exe`) to silently install Data Protection for SQL Server. If you are protecting Microsoft SQL Server 2012 data, you can also use MSI to silently install Data Protection for SQL Server on Windows Server Core.

Silently installing Data Protection for SQL Server with MSI (msiexec.exe)

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

Before you begin

Data Protection for SQL Server must be installed from an account that is a member of the local Administrators group for the system on which the SQL 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 SQL Server program. This technote is available in the *TSM for Databases - All Requirements Documents* website at http://www.ibm.com/support/docview.wss?uid=swg21218747. 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 SQL Server Management Console and Data Protection for SQL Server. Enter each **`msiexec`** command on a single line.

1. Install the Data Protection for SQL Server Management Console:

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

   Where *x:* is your DVD drive, and *aaa* is either `x86` or `x64`.
2. Install Data Protection for SQL Server:

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

   Where *x:* is your DVD drive, and *aaa* is either `x86` or `x64`.

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.

Silently installing Data Protection for SQL Server with MSI (msiexec.exe) on Windows Server Core

If you are protecting Microsoft SQL Server 2012 data on Windows Server Core, use the Windows Installer (MSI) program, `msiexec.exe`, to silently install Data Protection for SQL Server.

Before you begin

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

Important: Unlike `setup.exe`, 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 SQL Server. This technote is available in the *TSM for Databases - All Requirements Documents* website, at http://www.ibm.com/support/docview.wss?uid=swg21218747. When you are at the website, follow the link to the requirements technote for your specific release or update level.

About this task

The following example shows how to use **msiexec** to install the Data Protection for SQL Server. Enter the **msiexec** command on a single line from a Run as Administrator command prompt window.

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

Where *x:* is your DVD drive, and *aaa* is either `x86` or `x64`.

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"`. For more information, see "Silent installation features (Language Packages only)" in “Installing Data Protection for SQL Server silently” on page 53.

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.

You are ready to configure Data Protection for SQL Server.

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\DpSqlSetupLog.txt`:

```
x:\fcm\x64\sql\6400\enu\setup.exe /s /v"INSTALLDIR=\"C:\Program
Files\Tivoli\tsm\" ADDLOCAL=\"Client\" TRANSFORM=1033.mst
REBOOT=ReallySuppress /qn /l*v \"C:\Temp\DpSqlSetupLog.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 SQL Server 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 8. 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 from a previous version of the software. Upgrading the software is a three-step process.

About this task

The first step is to download the updates. After you download the updated code, run `setupfcm.exe` to install the updates. The final step is to run the configuration wizard.

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 slightly varies. The wizard provides instructions to guide you through the process.

To start the configuration wizard, start the Management Console. To start the Management Console, click **Start** > **All Programs** > **Tivoli Storage Manager** > **Data Protection for Microsoft SQL Sever** > **DP for SQL Management Console**.

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

Chapter 4. Configuring

Configure Data Protection for SQL Server before you start protecting your applications.

Before you begin

Data Protection for SQL Server must be installed on your system. A Tivoli Storage Manager server must be available to communicate with Data Protection for SQL Server.

Configuring Data Protection for SQL Server

Configuration requirements for Data Protection for SQL Server, Tivoli Storage Manager, and other applications vary. The requirements depend on which Data Protection for SQL Server 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 SQL Server, complete the following steps:

Procedure

1. Start the Management Console by clicking **Start** > **All Programs** > **Tivoli Storage Manager** > **Data Protection for Microsoft SQL Server** > **DP for SQL 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 **SQL 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 SQL data. This node name is the Tivoli Storage Manager backup-archive client node name, also known as the DSM agent node name.

- In the **Data Protection for SQL** 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 SQL Server backups.
- If you are configuring Data Protection for SQL Server for SQL Server 2012, enter a node name in the **AlwaysOn Node** field. This is the node name that is used to back up the availability databases in an AlwaysOn Availability Group. By default, the Windows Failover Cluster name is used.
- If the Tivoli Storage Manager for Virtual Environments Recovery Agent license is available, enter the data center node name. The data center node is the virtual node that maps to a data center.

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_SQL`. For an SQL configuration, the AlwaysOn node name does not have to be related to the VSS Requestor or the Data Protection for SQL Server node name. For example, you can name it `TSM_ALWAYSON`.

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 SQL Server, 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_sql.txt` file.
3. Click **Open**. These commands are displayed in the command-line panel:

   ```
   query tdp
   query tsm
   query sql
   ```
4. Click **Enter** to run the commands to verify your configuration.

Related tasks:

"Quick installation and configuration" on page 45

"Manually configuring Data Protection for SQL Server"

Manually configuring Data Protection for SQL Server

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

Perform these tasks on the computer running the SQL Server

For best results, use the configuration wizards to configure Data Protection for SQL Server 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 SQL server is running.

Procedure

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

1. Specify your Data Protection for SQL Server node name and communication method in the `dsm.opt` file located (by default) in the Data Protection for SQL Server installation directory. More options are also available.
2. Using the **`set`** command, specify your Data Protection for SQL Server preferences (language, date format, log file) in the `tdpsql.cfg` file in the Data Protection for SQL Server installation directory.

3. If you are configuring Data Protection for SQL Server in an SQL Server 2012 environment, specify the Tivoli Storage Manager node name that is used to back up the AlwaysOn availability databases. You can specify the AlwaysOn node name by using the **alwaysonnode** option in the tdpsql.cfg file. For example:

   ```
   set alwaysonnode myAlwaysOnNode
   ```

   All availability databases in an availability group are backed up under this node name. Any stand-alone databases are backed up under the standard Data Protection for SQL Server node name.
4. For SQL Server 2012: If you want all databases to be backed up by default under the AlwaysOn node, specify the **usealwaysonnode** option in the tdpsql.cfg file. For example:

   ```
   usealwaysonnode yes
   ```

   This option is useful if you are changing your database backups from the standard Data Protection for SQL Server node to an AlwaysOn node.
5. (VSS only) Specify your **VSSPOLICY** statement in your Data Protection for SQL Server configuration file.
6. (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 SQL Server configuration file (tdpsql.cfg).

   For more information, see *Tivoli Storage Manager for Windows Backup-Archive Clients Installation and User's Guide* and "Proxy node definitions".
7. (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.
8. (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.
9. (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.
10. (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.
11. (VSS only) Change the SQL Server VSS Writer from Manual to Automatic and start the service.
12. (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.

Related concepts:

"Specifying Data Protection for SQL Server options" on page 37

"Specifying Data Protection for SQL Server preferences" on page 40

"Proxy node definitions (VSS backups)" on page 34

"Back up to Tivoli Storage Manager storage versus back up to local shadow volumes" on page 21

Related reference:

"Set positional parameters" on page 239

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 SQL Server backup and restore requirements. For VSS operations, Tivoli Storage Manager server authentication must be on.
2. Register your Data Protection for SQL Server 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. (VSS only) If not already defined, register your Tivoli Storage Manager backup-archive client node name and password for the system where the SQL 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=alwayson node name agent=BAnodename
   ```

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 machine running 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 SQL Server configuration file (`tdpsql.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 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 verify your configuration

Before attempting to perform a backup or restore operation, verify that Data Protection for SQL Server is installed and configured correctly.

Manually verifying the installation and configuration of Data Protection for SQL Server from the command line

You can issue query commands at the command line to verify the installation and configuration.

1. Click **Start** > **All Programs** > **Tivoli Storage Manager** > **Data Protection for Microsoft SQL Server** > **SQL Client - Command Line**.
2. Enter the following commands:

```
tdpsqlc query tdp
tdpsqlc query tsm
tdpsqlc query sql
```

The Data Protection for SQL Server Server configuration has been verified when these commands complete without errors or warnings.

Verifying that a SQL Server is ready to perform VSS operations

Perform the following tests to verify that your SQL Server is ready to perform VSS operations. For best results, perform these tests prior to installing Tivoli Storage Manager.

When these tests complete without errors, you can install Tivoli Storage Manager. For Windows 2008 and later, the DiskShadow tool is preloaded. You can run the **diskshadow** commands, but you must download the **diskshadow** tool to run them.

Using the DISKSHADOW command (Windows 2008 and later)

Before installing Data Protection for Microsoft SQL Server, test core VSS functionality first. VSS functionality can be validated with the Windows 2008 Server-embedded command DISKSHADOW. DISKSHADOW is available for Windows Server 2008 and Windows Server 2008 R2. The following DISKSHADOW tests can be completed before any Tivoli Storage Manager components are installed.

1. Test non-persistent shadow copy creation and deletion. Run `diskshadow` in a command window and enter the following commands:

   ```
   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 may take a few minutes)
   DISKSHADOW>delete shadows all
   ```

   Volumes `f:` and `g:` represent the SQL Database and log volumes. Repeat this sequence of commands 4 times. Verify the Windows Event Log contains no errors.

2. Test persistent shadow copy creation and deletion. Run `diskshadow` in a command window and enter the following commands:

   ```
   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 may take a few minutes)
   DISKSHADOW>delete shadows all
   ```

   Volumes `f:` and `g:` represent the SQL Database and log volumes. Repeat this sequence of commands 4 times. Verify the Windows Event Log contains no errors.

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

   Run `diskshadow` in a command window and enter the following commands:

   ```
   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\sqlmeta.cab (specify the path
   where you want the metadata stored)
   DISKSHADOW> create
   DISKSHADOW>end backup
   ```

   Manually copy the `sqlmeta.cab` file from the source server to the offload server and run the following commands:

   ```
   DISKSHADOW>LOAD METADATA path to sqlmeta.cab
   DISKSHADOW>IMPORT
   DISKSHADOW>list shadows all (This can take a few minutes)
   DISKSHADOW>delete shadows all
   ```

   Volumes `f:` and `g:` represent the SQL Database and log volumes. Repeat this sequence 4 times. 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 performing a VSS operation. 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 has not been granted to this node.
This message is displayed when the Tivoli Storage Manager server has not been configured for the proxy nodes correctly.

Configuring where scheduled backups are run on availability replicas

You can configure to select an availability replica on which to run a backup.

About this task

When an availability database is replicated across multiple availability replicas in an availability group, a mechanism is required to select a single replica on which to run a backup operation instead of backing up all replicas. This mechanism can also offload the backup from a primary replica to a secondary replica for load balancing. When databases failover, backups must continue to run from other replicas for high availability.

Microsoft SQL Server 2012 provides a set of configuration options that enable these functions. You can use the Data Protection for SQL Server GUI to set these options.

Procedure

To specify whether scheduled backups are run on the primary or secondary replica at the availability group level:

1. Start the Management Console.
2. In the Management section of the window, click **Protect Data** next to the SQL Server workload.
3. Click **Properties** in the Actions pane.
4. Click the **AlwaysOn Preferences** property page.
5. In the **Availability group** field, select the **AlwaysOn Availability Group** for which you want to set up backup preferences.
6. In the **Preferred replica** field, select which replica is the preferred replica on which to run scheduled backups.
   - Select **Prefer Secondary replica** if you want scheduled backups to occur on a secondary replica, if it is available. Otherwise, use the primary replica for the scheduled backup.
   - Select **Secondary only** if you want scheduled backups to occur only on a secondary replica.
   - Select **Primary** if you want scheduled backups to occur only on the primary replica.
   - Select **Any replica** if you want scheduled backups to occur on any availability replica.

7. For each availability replica listed in the Availability replicas list box, specify whether it is a candidate for running scheduled backups by specifying the backup priority for that replica. A value of 1 has the lowest priority, and a value of 100 has the highest priority. A value of 0 indicates that the replica is excluded from schedule backup operations.
8. Click **OK** to save your configuration and exit the Data Protection Properties page. The settings are saved to the `tdpsql.cfg` file and can be replicated to the other replicas in the availability group.

What to do next

After you configure where scheduled backups are run, the administrator can specify the **tdpsql backup** command along with the **/ALWAYSONPriority** parameter in a backup schedule. For example:

```
tdpsqlc backup TestDb1 full /ALWAYSONPriority
```

When this scheduled backup command is run, Data Protection for SQL Server queries the SQL Server to determine the highest-priority availability replica that is active or online, ordered by preference. If the replica meets the specified criteria, the replica is backed up. Otherwise, the backup operation ends and a message is added to the log to indicate why the replica was not backed up.

An administrator can create a common backup schedule to run on all availability replicas. When the backup schedule starts, each **tdpsqlc** command queries each replica to determine whether it is to run the backup. Only one of the scheduled backups runs the backup.

Configuring Data Protection for SQL Server on Windows Server Core

The following sections describe how to manually configure Data Protection for SQL Server on Windows Server Core.

Before you begin

Before you can use Data Protection for SQL Server to protect your SQL Server 2012 data on Windows Server Core, ensure that you install Data Protection for SQL Server and the Tivoli Storage Manager backup-archive client on the system that runs the Microsoft SQL Server.

Procedure

Complete the following tasks:

1. Create a node on the Tivoli Storage Manager server for the backup-archive client and Data Protection for SQL Server. If you are protecting availability databases in an AlwaysOn Availability Group, you must also create the AlwaysOn node on the Tivoli Storage Manager server.
2. If you intend to run offloaded VSS backups, set up a remote node to run the offloaded backup operation on a remote computer.
3. Configure the backup-archive client options file (`dsm.opt`).
4. Configure the Data Protection for SQL Server option files (`dsm.opt` and `tdpsql.cfg`).
5. If you use Tivoli Storage Manager policy sets, you must also specify a management class to use for your Data Protection for SQL Server backups.

Creating a node on the Tivoli Storage Manager server

After you install the Tivoli Storage Manager client and Data Protection for SQL Server, you must set up a node name and password and register your node with the Tivoli Storage Manager server.

About this task

When you register your node, you create a file space on the Tivoli Storage Manager server where the backups of your data are stored. You must set up a client node and a Data Protection for SQL Server node. If you are protecting availability databases in an AlwaysOn Availability Group, you must also register the AlwaysOn node.

Follow these procedures if you installed the Tivoli Storage Manager administrative command-line client. If you did not install the administrative client, the nodes must be registered on the Tivoli Storage Manager server.

Procedure

1. Start an administrative client session by issuing the following command in a Command Prompt window:

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

2. To register a client node, issue the following command:

   ```
   reg node client_nodename password backdel=yes
   ```

   Where *client_nodename* is the node name for the client and *password* is the password that you want to use for the client. The **backdel**=*yes* parameter indicates that you can delete backup objects in your file space on the server. For example:

   ```
   reg node doomvm3 doomvm3passwd backdel=yes
   ```

3. To register a Data Protection for SQL Server node, issue the following command:

   ```
   reg node sql_nodename password backdel=yes
   ```

   Where *sql_nodename* is the node name for the Data Protection for SQL Server node and *password* is the password that you want to use for the SQL node. The **backdel**=*yes* parameter indicates that you can delete backup objects in your file space on the server.

   For example:

   ```
   reg node doomvm3_sql doomvm3sqlpasswd backdel=yes
   ```

   Tip: To easily identify the node as a node for Data Protection for SQL Server, add "_sql" to the end of the node name.

4. To register the AlwaysOn node, issue the following command:

   ```
   reg node alwayson_nodename password backdel=yes
   ```

   Where *alwayson_nodename* is the name for the AlwaysOn node and *password* is the password that you want to use for the AlwaysOn node. The **backdel**=*yes* parameter indicates that you can delete backup objects in your file space on the server. For example:

   ```
   reg node myalwaysonnode alwaysonpasswd backdel=yes
   ```

What to do next

To use Tivoli Storage Manager server policy sets, the Tivoli Storage Manager must define the policy domains, policy sets, management classes, copy groups, and storage pools. These definitions are necessary to meet your Data Protection for SQL Server backup and restore requirements. For VSS operations, Tivoli Storage Manager server authentication must be on.

Setting up a proxy node for offloaded VSS backups in the Windows Server Core environment

If you want to offload VSS backups to the Tivoli Storage FlashCopy Manager, you must define a remote node to run the offloaded backups. This step is part of the configuration tasks for operating Data Protection for SQL Server on Windows Server Core.

About this task

Data Protection for SQL Server can offload VSS backups by using a remote computer to create the backup instead of using the local computer. To run an offload backup by using a remote node, you must first set the remote node as an agent of the local Data Protection for SQL Server node.

If you are protecting availability databases in an AlwaysOn Availability Group, you must set the remote node as an agent of the AlwaysOn node.

Before you begin, ensure that the Tivoli Storage Manager client is installed and configured on the remote computer.

Procedure

To define the proxy node relationship, the Tivoli Storage Manager administrator can issue the **grant proxynode** command from the Tivoli Storage Manager server administrative console.

- For standard Data Protection for SQL Server nodes, issue the following command:

  ```
  grant proxynode target=local_sql_node agent=remote_node
  ```

  Where *local_sql_node* is the node name of the local Data Protection for SQL Server node, and *remote_node* is the remote Tivoli Storage Manager client node that runs the remote backups. For example:

  ```
  grant proxynode target=doomvm3_sql agent=babar
  ```

- For AlwaysOn nodes, issue the following command:

  ```
  grant proxynode target=alwayson_node agent=remote_node
  ```

  Where *alwayson_node* is the name of the AlwaysOn node, and *remote_node* is the remote Tivoli Storage Manager client node that runs the remote backups. For example:

  ```
  grant proxynode target=myalwaysonnode agent=babar
  ```

- To display the client nodes with authority to act as proxy to other clients, run the following command from the administrative console of the server:

  ```
  query proxynode
  ```

Configuring the client in the Windows Server Core environment

You must configure the Tivoli Storage Manager client node that you created. This step is part of the initial configuration tasks before you can use Data Protection for SQL Server in the Windows Server Core environment.

About this task

You must configure the client options file (`dsm.opt`), set the environment variables, and install and setup the Tivoli Storage Manager client acceptor service and remote client agent service.

Procedure

To configure the client, complete the following steps:

1. Configure the client options file:
   a. Change to the backup-archive client installation directory. For example, issue the following command in a Command Prompt window:
      ```
      cd C:\Program Files\Tivoli\TSM\baclient
      ```
   b. Open the `dsm.opt` file with a text editor and enter the following statements:
      ```
      PASSWORDACCESS   GENERATE
      COMMMethod       TCPip
      TCPPort          1500
      nodename         client_nodename
      TCPSERVERADDRESS tsm_server
      ```
      The following list contains brief explanations of the client options in the statements:

      PASSWORDACCESS GENERATE
      : Instructs the client to save the password whenever the **/tsmpassword** option is used so you do not have to enter the password with every command.

      TCPPort 1500
      : Specifies that the client accesses the Tivoli Storage Manager server at TCP/IP port 1500. 1500 is the default port number.

      nodename *client_nodename*
      : Specifies the newly created node for the backup-archive client.

      TCPSERVERADDRESS *tsm_server*
      : Specifies the name of the Tivoli Storage Manager server. You can enter the server IP address or the fully qualified domain name.

      For example:
      ```
      NODename DOOMVM3
      PASSWORDAccess generate
      TCPServeraddress gijoe
      TCPPort  1500
      ```
      For more information about Tivoli Storage Manager client options, see Processing options.
2. Install and start the Tivoli Storage Manager client acceptor service and remote client agent service.
   a. Install the client acceptor service by entering the following command in a Command Prompt window:

```
C:\Program Files\Tivoli\TSM\baclient\dsmcutil install cad
/name:"servicename" /node:nodename /password:password
/autostart:yes
```

Where *nodename* is the client node name, *password* is the client password, and *servicename* is the name that you want to use for the client acceptor service. The default name is "TSM Client Acceptor". For example:

```
C:\Program Files\Tivoli\TSM\baclient\dsmcutil install cad /name:"TSM CAD"
/node:DOOMVM3 /password:doomvm3passwd /autostart:yes
```

b. Install the remote client agent service by entering the following command in a Command Prompt window:

```
C:\Program Files\Tivoli\TSM\baclient\dsmcutil install remoteagent
/name:"servicename" /node:nodename /password:password
/partnername:"partner service name"
```

The node name for the Tivoli Storage Manager Client Acceptor and the Remote Client Agent must be set to the backup-archive client node. The default service name is "TSM Remote Client Agent". The value for the **/partnername** option must match the name of the client acceptor service that you created. The default name is "TSM Client Acceptor". For example:

```
C:\Program Files\Tivoli\TSM\baclient\dsmcutil install remoteagent
/name:"TSM AGENT" /node:DOOMVM3 /password:doomvm3passwd
/partnername:"TSM CAD"
```

c. Start the client acceptor service by entering the following command:

```
net start "servicename"
```

Where *servicename* is the name of the client acceptor service that you created. For example:

```
net start "TSM CAD"
```

Do not start the remote client agent service manually. The remote client agent service is automatically started by the client acceptor service when it is needed.

Configuring Data Protection for SQL Server for Windows Server Core

You must configure Data Protection for SQL Server before you can protect your Microsoft SQL Server 2012 data in the Windows Server Core environment.

About this task

You must configure the client options file (dsm.opt) and Data Protection for SQL Server configuration file (tdpsql.cfg).

Procedure

1. Edit the client options file (dsm.opt).
   a. In the Data Protection for SQL Server installation directory, open the client options file (dsm.opt) with a text editor.
   b. Add the following statements to the client options file:

```
NODename         sql_nodename
PASSWORDAccess   Generate
COMMMethod       TCPip
TCPServeradress  tsm_server
TCPPort          1500
TCPWindowsize    63
TCPBuffSize      32
```

Where **nodename** is the Data Protection for SQL Server node name, and **TCPServeraddess** is the name of the Tivoli Storage Manager server. You can enter the server IP address or the fully qualified domain name.

For example:

```
NODename DOOMVM3_SQL
PASSWORDAccess generate
TCPServeraddress gijoe
TCPPort  1500
```

Restriction: The following special characters that are allowed in the SQL Server database name are not supported on Data Protection for SQL Server:

- Question mark (?)
- Multibyte character (,)
- Multibyte character (^)
- Asterisk (*)
- Backslash (\)
- Colon (:)

2. Edit the `tdpsql.cfg` file.
   a. In the Data Protection for SQL Server installation directory, open the configuration file (`tdpsql.cfg`) with a text editor.
   b. Add the following statements in the `tdpsql.cfg` file:

```
SQLSERVer          sql_server
FROMSQLserver      sql_server
SQLAUTHentication  INTegrated
MOUNTWaitfordata   Yes
BACKUPMethod       [Legacy|VSS]
DIFFESTimate       20
BUFFers            3
BUFFERSIze         1024
STRIPes            1
SQLBUFFers         0
SQLBUFFERSIze      1024
LOGPrune           60
LANGuage           ENU
BACKUPDestination  [LOCAL|TSM|BOTH]
LOCALDSMAgentnode  local_node
REMOTEDSMAgentnode remote_node
ALWAYSONNode       alwayson_node
USEALWAYSONnode    [TRUE|FALSE]
LOGFile            tdpsql.log
```

The following list contains brief descriptions of the key options in the `tdpsql.cfg` file:

SQLSERVer

Specifies the name of the Microsoft SQL Server that is running on the local computer.

BACKUPMethod

Determines whether to run a Legacy or VSS backup.

BACKUPDestination

Determines whether to run a local backup, Tivoli Storage Manager backup, or both. For Legacy backup, only Tivoli Storage Manager is used.

LOCALDSMAgentnode

Specifies the local node name of the client that is running on the local computer. This option is required for VSS offloaded backups.

REMOTEDSMAgentnode
Specifies the remote client node that runs the VSS offloaded backups on a remote computer.

ALWAYSONNode
Specifies the Tivoli Storage Manager node name use to back up availability databases in an AlwaysOn Availability Group.

USEALWAYSONnode
Specify *TRUE* to set the AlwaysOn node as the default node for all backup operations of standard and availability databases. You can use this option to change database backups from a standard Data Protection for SQL Server node to an AlwaysOn node.

Specify FALSE to back up standard databases to the Data Protection for SQL Server node. Availability databases are always backed up with the AlwaysOn node.

3. Optional: Use the **VSSPOLICY** option to Specify a management class for VSS backups.

   Unless specified otherwise, Data Protection for SQL Server uses the default management class of the policy domain that its node name is in. To specify that Data Protection for SQL Server uses a different management class, add the **VSSPOLICY** option to the `tdpsqlc.cfg` file. The format of the option is as follows:

   ```
   VSSPOLICY SQL_server_name "db_name" backup_type backup_dest mgmt_class
   ```

   For example:

   ```
   VSSPOLICY doomvm3 * FULL LOCAL MGMT2
   ```

   This statement specifies that Data Protection for SQL Server uses the management class `MGMT2` for local backups of any database in the SQL server named `doomvm3`.

SAN Volume Controller and Storwize V7000 configurations

This table provides configurations for typical use case scenarios and objectives for the backup and recovery solution.

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 SQL Server 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 SQL Server 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 SQL Server 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 SQL Server 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 SQL Server 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 SQL Server 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.

Migration considerations

Migration from earlier versions of Data Protection for SQL Server is supported.

After you upgrade and configure from the older version of Data Protection for SQL Server 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 SQL Server in a Microsoft clustering environment and you upgrade to a newer version of Data Protection for SQL Server, 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 SQL Server 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.

Transitioning standard SQL databases to the AlwaysOn node

You can specify the **/USEALWAYSONnode** parameter with the **backup** command to back up standard SQL databases to the file space for the AlwaysOn node. This transition can make it easier for you to manage all your database backups under a single node name.

About this task

If you want to regularly back up standard SQL databases to the file space for the AlwaysOn node, you can use the **set** command.

The AlwaysOn node name is required when you configure Data Protection for SQL Server in a SQL Server 2012 environment. It is not necessary to specify the AlwaysOn node name during each backup, query, or restore operation of an availability database.

The AlwaysOn node does not affect where standard databases are backed up. The standard databases continue to be backed up to the Data Protection for SQL Server node unless the **/USEALWAYSONnode** option is specified.

For example, issue the following command to back up your standard SQL databases to the file space for the AlwaysOn node:

```
TDPSQLC Backup *|dbname[,dbname,...] Full /USEALWAYSONnode
```

You can use the wildcard character (*) to back up all databases, or specify a list of database names that are separated by commas.

For example:

```
TDPSQLC Backup standard_db01,standard_db02 Full /USEALWAYSONnode
```

Chapter 5. Protecting data

After you complete the configuration process, start the Management Console to protect your SQL Server data.

About this task

To start the Management Console, click **Start** > **All Programs** > **Data Protection for Microsoft SQL Server** > **DP for SQL Management Console**. If you try to use the Management Console 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 SQL 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.

Determining managed storage capacity

You can track the capacity of managed storage assets 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 reside on a 10 GB volume. When a snapshot of the volume is performed, the Data Protection for SQL Server managed capacity measurement is 10 GB.

Procedure

To determine managed storage capacity:

1. From the GUI, select an SQL instance.
2. On the **Protect** tab, click **Properties** in the **Action** pane.
3. Select **Managed Capacity** from the list of available property pages. The managed capacity is calculated and displayed.
4. View a list of the volumes (that contain backups) and their respective managed capacity, by clicking **Show Details**.
5. Close this window.

Using the Task Manager pane

The Task Manager provides a centralized panel in the Management Console (MMC) GUI from which to view, stop, remove, or manage backup, restore, and automation tasks.

About this task

When running backup, restore, or automation tasks, use the Task Manager pane.

Procedure

1. Start the Management Console.
2. Click the appropriate **Protect Data** or **Recover Data** task for your data in the welcome page of the MMC GUI.
3. Click **Show Activity** in the **Action** pane. The Task Manager panel opens beneath the results pane.
4. Choose a view for the current task:
   - **Task List** (default): Click this item to view the following information about your operations:

     ```
     Name
     State
     Result
     Progress
     Start Time
     Duration
     Messages
     ```

     Use the **Task List** view to complete these tasks:
     - Click **Up** and **Down** to modify the processing order for incomplete operations. Hover the cursor on the selected operation to view the command-line input.
     - Click **Stop** to end an operation that is still processing. When an operation cannot be stopped, this button is not available.
     - Click **Remove** to remove a completed or a scheduled operation.
     - Copy the selected operation by either clicking the copy icon or right-clicking a task and selecting **Copy**. Then you can run this command in the Automate tab or from a command prompt.
     - Click the calendar icon to use the scheduler wizard to set up a schedule.
     - Click the appropriate icon to view statistics or a performance chart for the selected operation.
   - **Task Details**: Click this item to view the operation information (available in the **Task List**) in detailed format. Click **Mode: Navigate** and use the arrows to view details about each operation. Summary and error information is also available (when applicable).

Backing up SQL databases and files

Back up SQL databases and files with Data Protection for SQL Server.

Security

Data Protection for SQL Server requires certain settings in order to perform operations in a secure environment.

Windows administrator authority is required for installation. Data Protection for SQL Server must be registered to the Tivoli Storage Manager server and the appropriate node name and password must be used when connecting to the Tivoli Storage Manager server. In addition, standard Tivoli Storage Manager security requirements apply to Data Protection for SQL Server.

Three options are provided when specifying SQL Server logon information:

- Accept the default `sa` account and blank password.
- Use SQL user ID security and specify both the SQL user name and password. With SQL user ID security, the SQL Server administrator provides the logon ID and the password that provides access to the SQL Server.
- Use a trusted connection and let Windows authenticate the logon.

Note: The SQL logon user or Windows user name must be added to the SQL Server SYSADMIN fixed server role before it can be used by Data Protection for SQL Server.

Backing up SQL databases by using VSS

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

Before you begin

You can also back up availability databases in an AlwaysOn Availability Group on SQL Server 2012.

Before you begin, see “Security” for the settings to use in a secure environment.

If you want to manage local VSS backups or run offloaded backups to Tivoli Storage Manager server storage, you must have Tivoli Storage FlashCopy Manager configured in your environment. If you make VSS backups to the Tivoli Storage Manager server, the Tivoli Storage FlashCopy Manager is not required.

Data Protection for SQL Server uses a single AlwaysOn node to back up availability databases in an AlwaysOn Availability Group regardless of which availability replica is used for the backup operation. Ensure that you configured Data Protection for SQL Server to use an AlwaysOn node. You also must specify the AlwaysOn node in the **AlwaysOn Node** field in the TSM Node Names page of the Tivoli Storage Manager Configuration Wizard. If you change the **AlwaysOn node name** field in the AlwaysOn Node properties page for your SQL workload, you must run the Tivoli Storage Manager Configuration Wizard to complete the reconfiguration of the name. If you do not want to use the Tivoli Storage Manager Configuration Wizard to register the node on the Tivoli Storage Manager server, you can use the Tivoli Storage Manager **`register node`** command.

Procedure

To back up standard SQL databases or availability databases by using the VSS method:

1. Start the Management Console (MMC) GUI.
2. When configured for use with a Tivoli Storage Manager server, if you plan to use offloaded backups, make sure a **Remote DSMAGENT Node name** is specified. An offloaded backup uses another system (specified with the **Remote DSMAGENT Node name** parameter) to move SQL data to Tivoli Storage Manager server storage. Offloaded backups can reduce the load on network, I/O, and processor resources during backup processing.

   To verify or specify a remote **Remote DSMAGENT Node name**, select the **SQL Server** instance in the tree view, and click **Properties** in the Actions pane. Then, select the VSS Backup property page. If the **Remote DSMAGENT Node name** is blank, enter a node name.
3. On the **Protect** tab of the SQL instance, specify the type of SQL data to back up:
   - Click **View: Databases** for a list of discovered SQL databases that are available for backup.
   - For SQL Server 2012: The **Standard Databases** / **Availability Databases** button toggles between the standard database view and the availability database view. The label on the button reflects the type of databases that are displayed in the view. To display a list of availability databases, click **Standard Databases**. Information about the availability databases in an availability group is displayed, including the replica role, synchronization state, and space and log usage.

   Use the **Protect** tab to browse and select the databases to back up. Fine-tune the list of available databases in the results pane by entering a keyword in the **Search** field.
4. Verify the backup options. If the backup options are not currently displayed, click **Show Backup Options**.
   - If you want to use offloaded backups, select **True** in the **Offload** field. This field applies only to VSS backups.
5. In the Actions pane, click **Backup Method** and select **VSS**.
6. In the Actions pane, for the **Backup Destination**, the only option is **TSM**. The database backups are stored on Tivoli Storage Manager server storage.
7. 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 up, complete with the command that is required to complete the task.
8. Create the backup by clicking **Full Backup** in the Actions pane. You can also right-click a database, and select the backup action that you want from the menu.

   If you want to create a copy-only full backup, click **Copy-Only Full Backup** in the Actions pane. A copy-only full backup is independent of the sequence of SQL Server backups. A copy-only full backup is not used as a base for a differential backup. The copy-only full backup does not disturb the sequence for a differential backup. The differential backup would not be associated to the copy-full backup, but would be associated to the prior full backup that was run. This type of backup can be used for special purpose backups that do not affect backup and restore procedures. It can be used for longer term retention than conventional backups.

What to do next

You can view the status of the backup operation by clicking **Task List** in the bottom half of the results pane. Click **Task Details** to view detailed status information.

Backing up SQL databases by using the legacy method

You can back up SQL databases by using the legacy backup method in the Management Console (MMC) GUI.

Before you begin

You can also use the legacy method to back up availability databases in an AlwaysOn Availability Group on SQL Server 2012.

There are several types of backup available for databases:

Full Backs up all of a database plus part of the transaction log.

Copy-Only Full
A type of backup that is independent of the sequence of conventional SQL Server backups. Transaction logs are not truncated with this backup. This type of backup can be used for special purpose backups that do not affect the backup and restore procedures, and can be used for longer term retention than conventional backups.

Differential
Backs up only the parts of a database that changed since the last full backup plus part of the transaction log.

Log Backs up the transaction log only, with or without truncation.

Before you begin, see “Security” on page 85 for the settings to use in a secure environment.

Data Protection for SQL Server uses a single AlwaysOn node name to back up availability databases regardless of which availability replica is used for the backup operation. Ensure that you configured Data Protection for SQL Server to use an AlwaysOn node name. You can set up the AlwaysOn node name in the **AlwaysOn Node** field in the TSM Node Names page of the Tivoli Storage Manager Configuration Wizard.

About this task

Follow these steps to run a legacy backup of your standard SQL databases or availability databases:

Procedure

1. Start the Management Console (MMC GUI).
2. Select the SQL Server instance in the tree.
3. On the **Protect** tab for the SQL instance, ensure that the **Databases** view is selected.

4. For SQL Server 2012: The **Standard Databases** / **Availability Databases** button toggles between the standard database view and the availability database view. The label on the button reflects the type of databases that are displayed in the view. To display a list of availability databases, click **Standard Databases**. Information about the availability databases in an availability group is displayed, including the replica role, synchronization state, and space and log usage.
5. Select one or more databases to back up.
6. Verify the backup options. If the backup options are not displayed, click **Show Backup Options**.
   - Use the **Stripes** option to specify the number of data stripes to use in a backup or restore operation. The *numstripes* variable can range from 1 to 64. The default value is 1.
   - Use the **DiffEstimate** option to specify the estimated fraction of the database that changed since its last full database backup. This estimate is needed because SQL Server does not provide a way to determine the size of a differential backup, and because the Tivoli Storage Manager server requires an accurate size estimate to efficiently allocate space and place objects. The Tivoli Storage Manager server uses this value to determine whether there is enough space in the primary storage pool to contain the backup. The default value is 20.
   - Use the **LogEstimate** option to specify the estimated the fraction of a SQL database that changed due to non-logged operations since the last log backup. The default value is 0.
   - Use the **Truncate** option to specify whether to dispose of entries you no longer need in the SQL database transaction log after you back up the log. In general, you do not want to truncate the log when rebuilding a corrupted database. This option enables the server to back up the transaction log but does not try to touch the data in any way. It writes all transaction log entries from the time of the last log backup to the point of database corruption. If you do not truncate the transaction log, you might be able to back up the transaction log of a damaged, suspect, or unrecoverable SQL Server database. The default value is **`True`**.
   - Use the **Backup tail of the log** option to store log records that have not yet been backed up. By storing these records, also known as the tail of the log, the log chain is kept intact. Before you can recover a SQL Server database to the latest point in time, you must back up the tail of the transaction log. The tail-log backup is the last backup of interest for the database recovery plan.
7. In the Actions pane, click **Backup Method** and select **Legacy**. TSM is the only available backup destination.
8. 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 up, complete with the command that is required to complete the task.
9. Create the backup by clicking one of these actions in the Actions pane:
   - **Full Backup**
   - **Copy-Only Full Backup**
   - **Differential Backup to TSM**
   - **Log Backup to TSM**

You can also right-click to select a database; then, from the context menu, select the backup action.

What to do next

You can view the status of the backup operation by clicking **Task List** in the bottom half of the results pane. Click **Task Details** to view detailed status information.

Backing up SQL groups or files by using the legacy method

You can complete a legacy back up of SQL groups or files by using the Management Console (MMC) GUI.

Before you begin

You can also use the legacy method to back up groups of files in availability databases in the SQL Server 2012 environment.

The following three types of backup are supported.

Group Backs up the contents of the specified file group.

File Backs up the contents of the specified logical file.

Set Backs up the contents of the specified groups and files.

Attention: You must back up the transaction logs after completing a Group, File, or Set backup operation.

Data Protection for SQL Server uses a single AlwaysOn node name to back up availability databases regardless of which availability replica is used for the backup operation. Ensure that you configured Data Protection for SQL Server to use an AlwaysOn node name. You can set up the AlwaysOn node name in the **AlwaysOn Node** field in the TSM Node Names page of the Tivoli Storage Manager Configuration Wizard.

About this task

Follow these steps to run a Legacy backup of SQL groups or files in a standard SQL database or availability database:

Procedure

1. Start the Management Console.
2. Select the SQL Server instance in the tree.
3. On the **Protect** tab for the SQL instance, make sure that the **Files** view is selected.
4. For SQL Server 2012: The **Standard Databases** / **Availability Databases** button toggles between the standard database view and the availability database view. The label on the button reflects the type of databases that are displayed in the view. To display a list of availability databases, click **Standard Databases**. Information about the availability databases in an availability group is displayed, including the replica role, synchronization state, and space and log usage.
5. Select one or more groups or files to back up.

6. Verify the backup options. If the backup options are not displayed, click **Show Backup Options**.
   - Use the **Stripes** option to specify the number of data stripes to use in a backup or restore operation. The *numstripes* variable can range from 1 to 64. The default value is 1.
7. 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 up, complete with the command that is required to complete the task.
8. Create the backup by clicking one of the following **Actions**:
   - **Group Backup to TSM**
   - **File Backup to TSM**
   - **Set Backup to TSM**

What to do next

You can view the status of the backup operation by clicking **Task List** in the bottom half of the results pane. Click **Task Details** to view detailed status information.

Restoring SQL databases and files

Restore SQL databases and files with Data Protection for SQL Server.

VSS restore considerations

When you run VSS restores, refer to the following list of considerations.

Unless otherwise specified, *VSS restores* refers to all restore types that use VSS (VSS restore, VSS fast restore, VSS instant restore).

- A VSS instant restore operation overwrites the entire contents of the source volumes. However, you can avoid overwriting the source volumes by specifying `InstantRestore` *False* in the Management Console (MMC) GUI. This option bypasses volume-level copy and uses file-level copy instead to restore the files from a VSS backup that is on local shadow volumes. The source volume must contain only the SQL database.
- When you do VSS restore from local shadow volumes, the bytes transferred is displayed as *0* because no data (*0*) is restored from the Tivoli Storage Manager server.
- When you run a VSS instant restore, ensure that any previous background copies, that involve the volumes that are being restored, are completed before initiating the VSS instant restore. However, this check is not necessary for XIV, SAN Volume Controller, or Storwize V7000 with space-efficient target volumes.
- Because of a SQL Server limitation, you cannot restore a VSS backup to an alternate SQL server instance. VSS backups are restored to the same SQL server instance where the snapshot was taken.

Restore options

Specify the options that you want to use when you restore data from the **Recover** tab in MMC GUI.

From the **Recover** tab, click **Show Restore Options** to change the default restore options. Click **Hide Restore Options** to remove the options from view.

AutoSelect

Set this option to **True** to enable auto-selection. 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.

When **AutoSelect** is `True`, the following behavior occurs:

- When you click a differential backup, the associated full backup is also selected.
- When you click a log backup, the associated full backup and all associated earlier differential or log backups are also selected.
- When you select this option, all databases and storage groups that are backed up together to the local destination are also automatically selected. When the backup method is VSS and the destination is local, this option is supported.

Stripes

Under Performance, the number of **Stripes** is listed. You can specify the number of data stripes to use in a restore operation. A maximum of 64 data stripes is allowed. The default value is *1*. The value that you enter must correspond to the value set for SQL buffers. This option is always enabled for legacy backups. Stripes are not available for VSS backups.

Database Owner Only

Under Restore Behavior, **DbOwnerOnly** is listed with a default value of **False**. You can mark a database for owner use only by changing this value to **True**. The default is not to mark for owner use. This option applies to legacy restores only.

Replace

Under Restore Behavior, **Replace** is listed with a default value of **False**. Change this value to **True** if you want to replace a database during a restore. The default is not to replace databases. This option applies to legacy restores only.

Recovery

Under Restore Behavior, **RunRecovery** specifies whether you want more restores to a SQL database that is not on a standby SQL server.

- Select **True** when you are making a sequence of restore operations to an SQL database and the current restore is the final one in the sequence, or when it is the only restore operation.
- Select **False** when you are making a sequence of restore operations to an SQL database and the current restore is not the final one in the sequence. Select **False** for all restore operations in the sequence except for the final one.

The default value is **True**.

During a restore operation, the operation might fail with a message similar to the following message:

```
Failed - An exception occurred while executing a Transact-SQL statement
or batch.
The tail-log backup of the dbName database has not been backed up.
Use BACKUP LOG WITH NORECOVERY to backup the log if it contains work you
do not want to lose.
Use the WITH REPLACE or WITH STOPAT clause of the RESTORE statement to
overwrite the contents of the log.

RESTORE DATABASE is terminating abnormally.
Changed database context to 'master'. (HRESULT:0x80131501)
```

When this situation occurs, set the `backup tail of the log` option to **True**. After you set this option, select **Log Backup to TSM** to complete the tail-log backup. For information about tail-log backups, refer to the Microsoft SQL Server documentation.

Stand By Undo File Name

Under Restore Behavior, **StandByUndoFileName** is listed with a default value of **False**. Use this option to specify the undo file path for a legacy restore to a standby SQL database. It changes the target SQL database in to standby mode.

This option is available for full, differential, and log backup types for a database. When you specify this option for a database, it applies to all backup objects for that database. Likewise, when you remove this option for a backup object, it is removed for all.

Source Server

Under Source Server, **FromSQLServer** specifies the name of the SQL server that the backup was created from.

Change **IncludeTsmVM** to `True` to view Virtual Environment backup SQL databases in the **Databases** view. The backup method is listed as `TSMVM` to distinguish these databases from the others listed.

Wait for Tape Mounts for Restore

Under the Tape section, you can specify whether the Data Protection for SQL Server restore operation waits for the Tivoli Storage Manager server to mount removable media such as tapes or other sequential device media. This information is retrieved from Tivoli Storage Manager when you click the **Recover** tab, or click the **Refresh**. The default value is **True**.

Wait for Tape Mounts for File Information

Under Tape, **WaitForTapeMountsForFileInformation** is listed with a default value of **True**. When querying Tivoli Storage Manager for file information, you can specify whether Data Protection for SQL Server waits for the Tivoli Storage Manager server to mount removable media. This option applies to legacy restores only.

Instant Restore

Under VSS, **InstantRestore** is listed with a default value of **True**. You can disable Instant Restores by setting the value to **False**, which bypasses volume-level copy and uses file-level copy to restore the files from a local VSS backup. If this option is set to **True** and if the backup exists on SAN-attached volumes, the volume level snapshot restore is used for local VSS backups. The default value is to use **volume level snapshot restore** if it is supported. This option is available for VSS operations only. When you are running a VSS instant restore, ensure that any previous background copies that involve the volumes that are being restored, are completed before you initiate the VSS instant restore.

Attention: An instant restore overwrites all files on the destination file system.

Restoring SQL server data

You can restore SQL server data.

Before you begin

You can restore databases or parts of databases only from `full`, `differential`, and `log` backups. VSS supports only full backups. Legacy differential and legacy log backups can be applied after a full VSS backup is restored.

- A legacy or VSS restore of the master database requires a different procedure. For more information, see "Restoring the master database" on page 97.
- When Virtual Environment restore is configured from the Tivoli Storage Manager server, you can restore and view these databases from the Recover tab.

Attention: When you restore a database, existing data is overwritten by the restored data and is no longer available after the restore is complete.

You can also restore availability databases that you backed up with the `AlwaysOn` node in a SQL Server 2012 environment. Backups of availability databases can be restored to any availability replica in an availability group.

When restoring availability databases, refer to the following guidelines:

Legacy restore

You can restore an availability database on either a primary or secondary replica.

During the restore process, the restored database is removed from the availability group. When a database is removed from the availability group, the database becomes a local database on that replica. The database is restored as a local database. After this restore is complete, manually add the database back to the availability group. However, before you add the database to the availability group, verify that the data on all replicas is transactionally consistent.

To verify data is transactionally consisent, verify that the backup copy contains data and transaction log records. Full backups and differential backups contain data and transaction log records so that the restored database is transactionally consistent.

After you verify that the data is transactionally consistent, the database can be added to the availability group.

VSS restore

Because of a SQL Server limitation, you cannot restore a VSS backup to an alternative SQL server instance. Therefore, VSS backups must be restored to the same SQL server instance where the snapshot was taken.

Procedure

To restore an SQL server database:

1. Start the Management Console.
2. Select the **SQL Server** instance in the tree.
3. On the **Recover** tab for the SQL instance, specify the type of SQL data to restore:

- Select **View: Databases** for a list of SQL database backups that are available for restore.

  If you want to view Tivoli Storage Manager server for Virtual Environment SQL database backups, open the Properties page, and select **Data Center Node**. Select the check box, **IncludeTSMVM**. Virtual Environment databases are listed with the `TSMVM` backup method.
- Select **View: Files** for a list of SQL database backup files that are available for restore.

4. For availability databases, click **DP Node Backups** to show all `AlwaysOn node` backups. The name of the button changes to **AlwaysOn Node Backups**. The **DP Node Backups / AlwaysOn Node Backups** button toggles between the standard database view and the availability database view. The label reflects the type of databases that are displayed in the view.
5. Use the **Recover** tab to browse and select the databases or files to restore. The following features are available:
   - Search: Fine-tune the list of available databases or files in the results pane by entering a keyword in the **Search** field.
   - Filter: Use the filter options to narrow the list of items 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.

        When you click **Select All**, all rows that reflect the filter specifications are selected.
     c. Select an operator in the **Operator** field.
     d. Specify a value to filter on in the **Value** field.
     e. If you want to filter on extra items, click **Add Row**.
     f. Click **Apply Filter** to filter the items on the list.
   - Backups: You can click **Active Backups** to show only active backups, or click **All Backups** to show both active and inactive backups.
   - Refresh: Click **Refresh** to update the view with your changes.

     If you applied a filter, the objects on the server that match the filter or search criteria are listed in the **Recover** tab. The status area indicates the number of items that match the criteria `n of x displayed`, where `n` equals the number of objects that match the filter criteria, and `x` is the number of objects that are retrieved from the server. For example, "`5 of 20 displayed`." If you specify refresh options to further narrow your results, and click **Refresh** again, the objects on the server that match the filtered and refresh options are displayed. Each time that you click **Refresh**, another query is run against the Tivoli Storage Manager server.
6. Verify the restore options. If the restore options are not displayed, click **Show Restore Options**.
   - Select the **Instant Restore** option to turn Instant Restore on or off. Disable the Instant Restore function if you want to use Fast Restore.

     Attention: Instant Restore overwrites all files on the destination file system.
7. 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 begins, complete with the command that is required to complete the task.

8. Click **Restore** in the Actions pane to begin the restore operation.

What to do next

You can view the status of the restore operation by clicking **Task List** in the bottom half of the results pane. Click **Task Details** to view detailed status information.

During a restore operation, the operation might fail with a message similar to the following message:

```
Failed - An exception occurred while executing a Transact-SQL statement
or batch.
The tail-log backup of the dbName database has not been backed up.
Use BACKUP LOG WITH NORECOVERY to backup the log if it contains work you
do not want to lose.
Use the WITH REPLACE or WITH STOPAT clause of the RESTORE statement to
overwrite the contents of the log.

RESTORE DATABASE is terminating abnormally.
Changed database context to 'master'. (HRESULT:0x80131501)
```

When this situation occurs, set the `backup tail of the log` option to **True**. After you set this option, select **Log Backup to TSM** to complete the tail-log backup. For information about tail-log backups, refer to the Microsoft SQL Server documentation.

Restoring a SQL database to an alternate system

You can restore a SQL database backup to an alternate SQL Server system or database by using the MMC GUI. Data Protection for SQL Server must be installed on both systems. Also, unlike legacy backups, VSS backups cannot be restored into a SQL Server that has a different name.

Before you begin

You can also restore availability databases to an alternate location on any availability replica in an availability group.

About this task

This procedure uses the following terms:

Source system
The system from which the original backup (to be restored) was taken.

Target system
The alternate system to which the backup is to be restored.

When restoring availability databases, refer to the following guidelines:

Legacy restore
You can restore an availability database on either a primary or secondary replica.

During the restore process, the restored database is removed from the availability group. When a database is removed from the availability group, the database becomes a local database on that replica. The database is restored as a local database. After this restore is complete, manually add the database back to the availability group. However, before you add the database to the availability group, verify that the data on all replicas is transactionally consistent.

To verify data is transactionally consisent, verify that the backup copy contains data and transaction log records. Full backups and differential backups contain data and transaction log records so that the restored database is transactionally consistent.

After you verify that the data is transactionally consistent, the database can be added to the availability group.

VSS restore

Because of a SQL Server limitation, you cannot restore a VSS backup to an alternative SQL server instance. Therefore, VSS backups must be restored to the same SQL server instance where the snapshot was taken.

Procedure

To restore a SQL database to an alternate location:

1. Copy the Data Protection for SQL Server options file (`dsm.opt`) from the source system to the target system. By default, the `dsm.opt` file is in the `C:\Program Files\Tivoli\TSM\TDPSql` directory. If **`passwordaccess`** *generate* is specified in the `dsm.opt` file, you might need to reset the password for this node on the Tivoli Storage Manager server.
2. Start the Management Console.
3. In the **Recover** tab for the SQL instance, select a database to restore. Optionally click **All Backups** in show all active and inactive backups. You can select only one database at a time when restoring it to an alternate location.
4. For availability databases, click **DP Node Backups** to show all `AlwaysOn node` backups. The name of the button changes to **AlwaysOn Node Backups**. The **DP Node Backups** / **AlwaysOn Node Backups** button toggles between the standard database view and the availability database view. The label reflects the type of databases that are displayed in the view.
5. Verify restore options. If the restore options are not currently displayed, click **Show Restore Options**.
   a. Ensure that **Wait for Tape Mounts for Restore** is set to **True**.
   b. Ensure that **Wait for Tape Mounts for File Information** is set to **True**.
   c. If the database to be restored is going to replace an existing database on the target system, click **Replace**.
   d. Use the **Instant Restore** option to turn Instant Restore on or off. Click **True** to use Instant Restore. Click **False** to disable Instant Restore if you want to use Fast Restore.

      Attention: Instant Restore overwrites all files on the destination file system.

   See "Restore options" on page 91 for descriptions of additional restore options.
6. Click **Restore to Alternate Location** in the **Actions** pane.
7. Complete the Alternate Location Restore Settings window.
   - In the **Restore Into** section of the window, click **Restore to new database**, and specify a target SQL server instance name and target database name to restore a backup object to. The target database must exist. VSS backups cannot be restored into a SQL Server that has a different name.

     Attention: Any type of **Restore Into** processing automatically disables VSS Instant Restore.

- If you want to specify new destination locations in which to restore backed up SQL databases, logs, and FILESTREAM files (SQL Server 2008, SQL Server 2008 R2, or later), click **Restore all files into one directory** in the **Relocate** section of the window.

  If you want to restore the log files into a location that is different from where the SQL database and other related files are being restored, select **Relocate logs into** and specify a new path in the text entry field.

  If you want to restore FILESTREAM files (SQL Server 2008, SQL Server 2008 R2, or later) into a location that is different from where the SQL database and logs are being restored, select **Relocate other files into**, and specify a new path in the text entry field.
- If you want to restore one or more individual SQL database, log, and FILESTREAM files, click **Relocate files individually**, select one or more file entries, and click **Browse** to open a folder selection window. Select a folder or make a new folder, and click **OK**. The path of the selected files entries is set to use the folder. This option is available for Legacy backups only.

Restriction: You cannot relocate database files and logs with partial restore operation in the Management Console (MMC) GUI. You must use the command-line interface to do a partial restore that requires these parameters.

8. Click **Restore** to close the Alternate Location Restore Settings window and begin the restore.

What to do next

You can view the status of the restore operation by clicking **Task List** in the bottom half of the results pane. Click **Task Details** to view detailed status information.

Restoring the master database

A damaged master database can result in the SQL Server failing to start, as well as a number of other error conditions. A special procedure is required to restore the master database.

About this task

During the process of rebuilding the master database, the SQL Server setup program drops and then recreates the **msdb** database so it must be restored along with the master database.

In general, the following steps are required:

1. Run the SQL Server setup program to rebuild the master database. You must rebuild using the same character set and sort order as the master database backup that will be restored.
2. Start the SQL Server in single-user mode. This can be done at a command prompt. See also Note 1 under “Setting user mode” on page 104.
3. Use Data Protection for SQL to restore the master database.

   Note: When the master database has finished restoring, the SQL Server shuts itself down. As a result, an error message is generated stating that the connection was lost to the SQL Server. This is expected.
4. Restart the SQL Server normally (in multi-user mode).
5. Manually reapply any changes that were made to the master database *after* the date of the database backup used to do the restore operation.
6. Use Data Protection for SQL to restore the **msdb** database.

It is important to keep an up-to-date backup of your master database because the master database contains the system catalog. The system catalog contains important information about the SQL Server configuration. Ensure that you back up the master database after any changes that update system tables. For example, back up the master database after any of these statements are used:

- ALTER DATABASE
- CREATE DATABASE
- DISK INIT
- DISK RESIZE
- DISK MIRROR
- DISK UNMIRROR
- DISK REMIRROR
- Various DBCC options such as SHRINKDB
- System stored procedure such as: sp_dropremotelogin, sp_addumpdevice, sp_dropdevice, sp_addlogin, sp_droplogin, sp_addserver, sp_dropserver, sp_addremotelogin

Restoring from virtual machine snapshots

You can restore SQL databases from a virtual machine where Tivoli Storage Manager for Virtual Environments, Version 7.1 and later is used to back up the data.

Before you begin

Before restoring SQL databases from virtual machine snapshots, verify that the data was backed up according to the following procedure:

1. Install the Tivoli Storage Manager for Virtual Environments Recovery Agent 7.1 package and the Tivoli Storage Manager Backup-Archive Client 7.1 from the Data Protection for VMware 7.1 package. These software packages are available for download from Passport Advantage.

   Install these packages on the guest virtual machine with Data Protection for Microsoft SQL Server.
2. Specify the following Tivoli Storage Manager Backup-Archive Client 7.1 option in the `dsm.opt` file:

   `INCLUDE.VMTSMVSS` *`vmname`*

   When you set this option, virtual machine applications receive a notification when a backup is going to occur. This notification allows the application to truncate transaction logs and commit transactions so the application can resume from a consistent state when the backup completes. By default, this option is not enabled. You must set this option to enable application protection for a virtual machine.

   The *vmname* specifies the name of the virtual machine that contains the applications to quiesce. Specify one virtual machine per **`INCLUDE.VMTSMVSS`** statement. To protect all virtual machines with this option, use an asterisk as a wildcard. For example:

   ```
   INCLUDE.VMTSMVSS *
   ```

   You can also use question marks to match any single character. For example:

   ```
   INCLUDE.VMTSMVSS vm??
   ```

This type of option setting protects all virtual machines that have names that begin with *vm* and are followed by any two characters. For example, *vm10* and *vm15*.

If the **OPTions KEEPSqllog** parameter is specified in an **INCLUDE.VMTSMVSS** statement, this parameter prevents SQL server logs from being truncated when a data mover node backs up a virtual machine that runs a SQL server. Specifying this parameter allows the SQL server administrator to manually manage the SQL server logs. The logs can be preserved as needed and be used to restore SQL transactions to a specific checkpoint, after the virtual machine is restored. When this option is specified, the SQL log is not truncated and following message is displayed and logged on the server:

```
ANS4179I IBM Tivoli Storage Manager application protection did not truncate
Microsoft SQL Server logs on virtual machine vmname
```

Note: Tivoli Storage Manager does not back up the SQL log files. The SQL administrator must back up the log files so they can be applied after the database is restored.

3. Verify that the VSS service and SQL Server instance are online and active. SQL Server databases that do not have an active instance are backed up. However, information about these databases is not saved to Tivoli Storage Manager. Therefore, these databases are not available for a database-level restore operation. You can restore these databases with a full VM restore operation.
4. Use the Tivoli Storage Manager for Virtual Environments software to back up the data. After you back up data, you can use the following procedure to restore the data.
5. Verify that the virtual machine backup contains the necessary database metadata. To do this, enter the following Tivoli Storage Manager Backup-Archive Client command on the data mover:

   ```
   dsmc query vm <vmname> -detail
   ```

   In the command output, the `Application(s) protected:` value must specify `(database-level recovery)`. For example:

   ```
   # Backup Date Mgmt Class Size Type A/I Virtual Machine
   --- ------------------- ---------- ----------- ------ --- ---------------
   1 06/07/2012 19:25:58 STANDARD 29.29 GB FULL A wombat
   The size of this incremental backup: n/a
   The number of incremental backups since last full: n/a
   The amount of extra data: n/a
   The TSM objects fragmentation: n/a
   Backup is represented by: n/a
   Application protection type: TSM VSS
   Application(s) protected: MS SQL 2008 (database-level recovery)
   VMDK[1]Label: Hard disk 1
   VMDK[1]Name: [ess800_dev1] wombat/wombat-000002.vmdk
   VMDK[1]Status: Protected
   ```

About this task

When you restore data, the data can be restored to basic disks with MBR-style partitions. Because of a SQL Server limitation, you cannot restore a VSS backup to an alternative SQL server instance. VSS backups must be restored to the same SQL server instance where the snapshot was taken.

Also, when restoring an SQL database from a VM backup, if the VM name is changed after the VM backup, the restore is not possible.

To restore SQL databases from a supported virtual machine snapshot, complete the following steps:

Procedure

1. Log on to the system where you want to restore the SQL database. The Data Protection for VMware Recovery Agent license and Data Protection for Microsoft SQL Server must be installed on the system where you are restoring the data.
2. When Data Protection for Microsoft SQL Server is configured, for the Configuring Recovery Agent rule, verify that the status is *Passed*. If the status is not *Passed*, re-run the configuration wizard. On the TSM Node Names wizard page, enter the data center node name. The data center node is the virtual node that maps to a data center.
3. Use the following procedure to set access to the virtual machine that is backed up in a data center node. To complete the following procedure, use the following reference information:

Table 9. Node names used to set access

Node name	Location	Description	Proxy type
DC_NODE	Data mover	Node for the virtual machine backup	Agent (data owner)
SQL_NODE	In guest virtual machine running Microsoft SQL	Node for Data Protection for Microsoft SQL Server	Agent (data owner)
VSS_NODE	In guest virtual machine running Microsoft SQL	Node for Data Protection for Microsoft SQL Server DSMAGENT	Agent (data worker)

   For the Data Protection for Microsoft SQL Server software to restore the data owned by the DC_NODE, the Tivoli Storage Manager administrator gives access to the virtual machine backed up to the VSS_NODE. The Tivoli Storage Manager administrator can use the Tivoli Storage Manager command-line interface to enter the **set access** command while connected to the DC_NODE.

   Running the commands from the DC_NODE is ideal because the options file has the necessary settings to communicate with the Tivoli Storage Manager server. The Tivoli Storage Manager administrator credentials can be used if the DC_NODE administrator password is not available.

   The **set access** command cannot be run if the ASNODE option is in affect. To issue the **set access** command, use an option file that does not contain ASNODE. The information needed to complete this task is provided in the following procedure.

   a. Copy `dsm.opt` and `dsm.setaccess.opt` files.
   b. If running the **set access** command from a node with ASNODE in the options file, complete this step. Edit the `dsm.setaccess.opt` file. For any line that contains ASNODE, remove the line.
   c. Edit the `dsm.setaccess.opt` file to set the NODENAME option to the following entry:

      ```
      DC_NODE NODENAME DC_NODE
      ```

   d. Enter the following command:

      ```
      dsmc set access backup -type=VM traveler VSS_NODE -optfile=dsm.setaccess.opt
      ```

      You might be prompted to enter the password for the DC_NODE.

For any subsequent **set access**, **query access**, and **delete access** commands, use this same procedure to guide you.

4. From the Data Protection for Microsoft SQL Server Management Console, in the Protect and Recover Data navigation, select a **SQL Server**.
5. Within the **Recover** tab for that SQL server, select **View: Databases** for a list of SQL database backups that are available for restore. SQL databases that were backed up with Tivoli Storage Manager for Virtual Environments software are listed with the VMVSS backup method.
6. Review the restore options by clicking Show Restore Options.
7. 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 begins, complete with the command that is required to complete the task.
8. Click **Restore** in the Actions pane to begin the restore operation.

Results

The restore operation is displayed in the **Task List** pane. Click **Task Details** to view detailed status information about the restore operation.

Restoring SQL file groups and files from Legacy backups

You can restore a Legacy backup of SQL file groups and files by using the Management Console (MMC) GUI. You can also run a Legacy restore of SQL file groups and files of an availability database that you backed up with the AlwaysOn node in a SQL Server 2012 environment. Backups from availability databases can be restored to any availability replica in an availability group.

Before you begin

Before you begin, see "Security" on page 85 for the settings to use in a secure environment.

You can restore databases or parts of databases from **group**, **file**, **set**, **log**, and **full Legacy** backups.

Microsoft SQL Server requires that the PRIMARY filegroup is restored before or along with a user-defined filegroup. To bring the database back to a usable state a log restore must be performed after the user-defined filegroup is restored.

Restoring parts of a database from a full legacy backup is also known as a partial restore. If you plan to apply either a log restore with point-in-time or a differential restore to a partially restored database, then consider one of these tasks:

- Use the **Files** view on the Recover tab to select and restore the full backup object. Make sure that the **RunRecovery** option is set to **False**.
- If you plan to apply a log restore with point-in-time, click **Restore to Point-in-Time** in the Actions pane to restore the log. Make sure that the **RunRecovery** option is set to **True**.
- If you plan to apply a differential restore, click **Restore** in the Actions pane to run a differential restore. Make sure the **RunRecovery** option is set to **True**.

Attention: When you restore the files and file groups of a database, existing data is overwritten by the restored data and is no longer available after the restore is complete.

For AlwaysOn availability databases, ensure that Data Protection for SQL Server is set up to use an AlwaysOn node name. You can set up the AlwaysOn node name in the **AlwaysOn Node** field in the TSM Node Names page of the Tivoli Storage Manager Configuration wizard. By default, the AlwaysOn node name is set to the cluster node name for the SQL Server 2012 Availability Group.

About this task

The following restrictions apply to the restore of availability databases:

Legacy restore

> You can restore an availability database on either a primary or secondary replica.
>
> During the restore process, the restored database is removed from the availability group. When a database is removed from the availability group, the database becomes a local database on that replica. The database is restored as a local database. After this restore is complete, manually add the database back to the availability group. However, before you add the database to the availability group, verify that the data on all replicas is transactionally consistent.
>
> To verify data is transactionally consisent, verify that the backup copy contains data and transaction log records. Full backups and differential backups contain data and transaction log records so that the restored database is transactionally consistent.
>
> After you verify that the data is transactionally consistent, the database can be added to the availability group.

Procedure

Follow these steps to restore SQL file groups and files from Legacy backups:

1. Start the MMC GUI.
2. Select the SQL Server instance in the tree.
3. On the Recover tab for the SQL instance, click **View: Files**.
4. For availability databases, click **DP Node Backups** to show all `AlwaysOn node` backups. The name of the button changes to **AlwaysOn Node Backups**. The **DP Node Backups / AlwaysOn Node Backups** button toggles between the standard database view and the availability database view. The label reflects the type of databases that are displayed in the view.
5. Select one or more groups, files, or sets to restore.
6. Verify the restore options. If the restore options are not displayed, click **Show Restore Options**. See "Restore options" on page 91 for descriptions of the restore options.
7. 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 command that is required to complete the task.

8. Click **RestoreFile** or **RestoreGroup** in the Actions pane to begin the restore.

What to do next

You can view the status of the restore operation by clicking **Task List** in the bottom half of the results pane. Click **Task Details** to view detailed status information.

Inactivating SQL databases (legacy only)

To inactivate an existing legacy backup of SQL databases on the Tivoli Storage Manager server, use the **Inactivate** option.

About this task

Typical backups do not require this command as Tivoli Storage Manager inactivates an SQL database backup as a part of Tivoli Storage Manager policy management. As a result, backup objects are typically inactivated as part of the scheduled backup processing.

The inactivate action is listed in the **Actions** pane when viewing in the **Recover** tab.

Inactivating the backups for a database causes the existing backups on the Tivoli Storage Manager server to adhere to the versions data deleted (**`verdeleted`**) setting on the server. The **`verdeleted`** setting indicates the number of inactive versions of a backup if an active version does not exist. The inactivate function is normally used when a SQL database is deleted from the SQL Server.

For cases when automatic processing is not sufficient, the inactivate function explicitly inactivates one or more active backup objects on the Tivoli Storage Manager server. As with backup and restore, use Data Protection for SQL to select any or all of six backup object types for operation: full, differential, log, file, group, or set for legacy backups. In addition, it is possible to inactivate any object or object type older than a specified number of days.

Use the inactivate function to inactivate a legacy backup of an SQL database on the Tivoli Storage Manager server.

The SQL database that you want to inactivate must be a legacy backup. VSS backups cannot be inactivated by using this method. Instead, the **Delete** action is available in the Actions pane when you select a VSS backup from the **Recover** view.

Procedure

To inactivate backup objects, complete the following steps:

1. Select the SQL server under the **Protect and Recover Data** node in the tree view.
2. Open the **Recover** view to see the status of the backup. Active backups are displayed.
3. To inactivate this backup, select the database backup. From the Actions pane, click **Inactivate**.
4. Click **All Backups** on the toolbar to display the database that you have made inactive. Click **Active Backups** on the toolbar to display only active backups.

Deleting SQL Server Backups

Perform these steps to delete an SQL Server backup 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.

Procedure

To delete an SQL Server backup:

1. Start the Management Console.
2. Click **Recover Data** > **SQL** in the Management window.
3. On the **Recover** tab for the SQL instance, select one or more database backups to delete.
4. Click **Delete backup** in the **Action** pane to delete the selected database backups.

   Upon completion of a delete backup operation, the view content refreshes and all object selections are cleared.

Setting user mode

About this task

Setting user mode might be necessary during certain restore procedures. For example:

- You need to change server configuration options.
- A damaged master database needs recovering.
- A system database needs to be restored.

These tasks might require starting a SQL Server instance in single-user mode. By placing SQL databases to be restored in single-user mode, you can avoid attempting such restores. If you are restoring the master database, you must place the SQL server in single-user mode. For additional information or assistance with SQL commands, go to the *Books Online for SQL Server* website.

ALTER DATABASE DBNAME SET SINGLE_USER WITH ROLLBACK AFTER N SECONDS

This TRANSACT-SQL command forces users off the database and places it in single-user mode.

ALTER DATABASE DBNAME SET MULTI_USER

This TRANSACT-SQL command returns the database to multiple-user mode.

Note:

1. You can set a SQL server to single-user mode by using the –m SQL SERVER startup option when restarting the SQL server.

2. You can use the SQL stored procedure SP_WHO to determine which users are using the databases.

Viewing reports

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

About this task

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

Procedure

1. Select **Reporting** in the tree view, under **Manage**. 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 is displayed.
3. To print or save the current report, click the appropriate icon at the top of the report.

Protecting SQL Server data in the Windows Server Core environment

Because of the Windows Server Core environment requirements, you can operate Data Protection for SQL Server from only the command-line interface.

About this task

You can use the **`backup`** and **`restore`** commands to protect your Microsoft SQL Server 2012 and later databases, the **`set`** command to change the configuration of Data Protection for SQL Server, and the **`help`** command to display concise help information about a command.

Backing up SQL Server databases on Windows Server Core

The following code samples provide guidance about using the **`backup`** command in the Windows Server Core environment to back up SQL Server databases to the Tivoli Storage Manager server, or to take local VSS snapshots.

About this task

To back up all or part of a SQL database on Windows Server Core, issue the following command from a Command Prompt window:

```
tdpsqlc backup databse_name backup_type [other_options]
```

To run a full legacy backup of SQL databases `DB_01` and `DB_02`, enter the following command:

```
tdpsqlc backup DB_01,DB_02 full /backupmethod=legacy
```

To run a full legacy backup of all databases on the SQL Server, enter the following command:

```
tdpsqlc backup * full /backupmethod=legacy
```

To back up the filegroup `DB_01_group1` that belongs to the `DB_01` database, enter the following command:

```
tdpsqlc backup DB_01 Group=DB_01_group1
```

Restoring SQL databases on Windows Server Core

The following code samples provide guidance about using the **`restore`** command in the Windows Server Core environment to recover all or part of one or more SQL databases.

About this task

To restore all or part of a SQL database on Windows Server Core, issue the following command from a Command Prompt window:

```
tdpsqlc restore database_name backup_type [other_options]
```

To run a full database restore of databases `DB_01` and `DB_02`, and replace the existing databases with the database objects that are recovered from the Tivoli Storage Manager server, enter the following command:

```
tdpsqlc restore DB_01,DB_02 Full /REPlace
```

To restore the filegroup `DB_01_group1` that belongs to the `DB_01` database, enter the following command:

```
tdpsqlc restore DB_01 group=DB_01_group1
```

To restore all the logical files that are in the `DB_01` database, enter the following command:

```
tdpsqlc R DB_01 file=*
```

Changing Data Protection for SQL Server configuration values on Windows Server Core

Use the set command to change the values for the Data Protection for SQL configurable parameters and options from the Windows Server Core command prompt.

About this task

The values that you change are saved in the Data Protection for SQL Server configuration file. The default configuration file is `tdpsql.cfg`.

To change the values for configurable parameters and options, issue the following command from a Command Prompt window:

```
TDPSQLC Set Parameter=Value [/CONFIGfile=filename]
```

Where *Parameter* is the Data Protection for SQL Server parameter or option for which you want to change values, and *Value* is the new value that you want to specify. **`/CONFIGfile`** is the optional parameter for the configuration file name. If you do not specify the **`/CONFIGfile`** parameter, the default configuration file (`tdpsql.cfg`) is used.

For detailed information about the **`set`** command, see the "Set command" topic.

Example

Examples:

Task Change the number of data buffers to `2` and save the changed value in the `tdpsql.cfg` file.

Command: `tdpsqlc set buffers=2 /config=tdpsql.cfg`

Task Change the name of the Data Protection for SQL Server activity log file to `tdpsql.log`.

Command: `tdpsqlc set logfile=tdpsql.log`

Getting help for Data Protection for SQL Server commands on Windows Server Core

Use the **help** command to display the syntax of a Data Protection for SQL Server command from the Windows Server Core command prompt.

About this task

To display the command syntax of Data Protection for SQL Server commands, issue the following command from a Command Prompt window:

```
TDPSQLC Help|? [*|?|command]
```

You can use either `Help` or `?` to start the command-line help.

For detailed information about the **help** command, see the "Help command" topic.

Example

Examples:

Task Display the command syntax for a full restore operation.

Command: `tdpsqlc help restore full`

Task Display the command syntax for the **backup** command.

Command: `tdpsqlc ? backup`

Chapter 6. Automating

The term *automation*, as it applies to Data Protection for Microsoft SQL, means that you can run commands from the command line, create scripts, schedule tasks, and use the graphical user interface to start tasks 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 `tdpsqlc`. For example, for each selected workload instance, you can enter a single command or multiple commands, such as:

   ```
   q tsm
   q sql
   ```
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 **SQL 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

The guidelines assist when planning scheduled operations.

Refer to the following guidelines when defining a Tivoli Storage Manager schedule:

- If you want to use the Tivoli Storage Manager server scheduling mode, you must ensure that the Data Protection for SQL Server option file has the `tcpclientaddress` and `tcpclientport` options specified. If you want to run more than one scheduler service, use the same `tcpclientaddress`. However, in addition to different node names, use different values for `tcpclientport`. An example of running more than one scheduler service is when you are scheduling Data Protection for SQL Server and the Windows backup client.

  Server-prompted scheduling is supported only when TCP/IP communication is being used. By default, Data Protection for SQL Server uses the client polling schedule mode.
- If any changes that affect the scheduler are made to the Data Protection for SQL Server options file, the scheduler must be restarted in order to pick up the changes. An example of this is the Tivoli Storage Manager server address, the schedule mode, or the client TCP address or port. This can be done by issuing the following commands:

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

  If you are running the scheduler service in a cluster environment, use the Cluster Administrator to stop and restart your scheduler service. Do not use the **`net stop`** and **`net start`** commands.
- 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 located in this path:

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

  You can override this file name by specifying the `schedlogname` option in the Data Protection for SQL Server options file.
- Data Protection for SQL Server creates its own log file with statistics about the backed up database objects when the **`/logfile`** parameter is specified during the **`tdpsqlc`** command. In the sample file (`sqlfull.smp`), the log file is `sqlsch.log`. This file is different from the Tivoli Storage Manager scheduler log file and must also be different from the file to which the **`tdpsqlc`** command output is redirected. In the previous example, this file is `sqlfull.log`.

Output from scheduled commands are sent to the scheduler log file (`dsmsched.log`). After scheduled work is performed, check the log to ensure the work completed successfully.

When a scheduled command is processed, the scheduler log might contain the following entry:

```
Scheduled event eventname completed successfully
```

This is merely an indication that Tivoli Storage Manager successfully issued the scheduled command associated with the *eventname*. No attempt is made to determine the success or failure of the command. You should 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 backup failed. The results from the very last command in the script are returned to the scheduled event.

- The preferred method of password management for scheduler operations is to specify **passwordaccess** generate in the `dsm.opt` file. If **passwordaccess** generate is not specified in the `dsm.opt` file, then the Tivoli Storage Manager password must be specified on the **tdpsqlc** command. To specify the password, use the **/tsmpassword** parameter in the command file being run by the scheduler (`sqlfull.cmd`). You can also specify the password on the Data Protection for SQL Server command line. For example:

```
tdpsqlc query tsm /tsmnode=mynode /tsmpassword=newpassword
```

- All log files and redirected output files must have unique names for both the Data Protection for SQL Server command and the scheduler service that is starting the backup. When you do not use unique names, the schedule fails with a return code of RC=1.

Windows PowerShell and Data Protection for SQL Server

Data Protection for SQL Server includes a set of Windows PowerShell cmdlets to help manage Data Protection for SQL Server 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 SQL Server cmdlets help support a seamless management environment. Remote management and automation capabilities are greatly improved when using the Data Protection for SQL Server cmdlets.

Getting started

The cmdlets can be used in supported Windows environments.

About this task

Before you use the cmdlets provided with IBM Tivoli Storage Manager for Databases: Data Protection for SQL Server, 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 TDPSql folder:
   - FmModuleSQL.dll
   - FmModuleMMC.dll

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

   a. Navigate to the TDPSql folder.
   b. Enter the following commands:
      ```
      import-module .\FmModuleSQL.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:
      ```
      $path = (get-itemproperty -path "HKLM:\SOFTWARE\IBM\TDPSql\
       currentversion\mmc" -ea SilentlyContinue).path
      if ($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 SQL Server data

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

Table 10. Cmdlets to protect Microsoft SQL Server data. The following table identifies the cmdlets that you can use to protect Microsoft SQL Server data.

Cmdlet name	Related command-line interface command	Short description
Add-DpSqlPolicy	**tdpsqlc create policy**	Create a new policy for Microsoft SQL Server data.
Backup-DpSqlComponent	**tdpsqlc backup**	Backup SQL components.
Copy-DpSqlPolicy	**tdpsqlc copy policy**	Copy an existing policy to a new policy.
Dismount-DpSqlBackup	**tdpsqlc unmount backup**	Dismount a backup.
Get-DpSqlBackup	**tdpsqlc query tsm ***	Query the backups that are stored on the server.
Get-DpSqlComponent	**tdpsqlc query sql ***	Query the databases that are available on the SQL server.
Get-DpSqlConfig	**tdpsqlc query tdp**	Display configuration information.
Get-DpSqlConnection	**tdpsqlc query tsm**	Displays the Tivoli Storage Manager API and server information.
Get-DpSqlFileGroups	not applicable	Displays all file and group information about specified SQL Server databases.

Table 10. Cmdlets to protect Microsoft SQL Server data (continued). The following table identifies the cmdlets that you can use to protect Microsoft SQL Server data.

Cmdlet name	Related command-line interface command	Short description
Get-DpSqlInformation	**tdpsqlc query sql**	Display specified SQL Server information.
Get-DpSqlManagedCapacity	**tdpsqlc query managedcapacity**	Assist with storage planning by determining the amount of managed capacity that is in use.
Get-DpSqlPolicy	**tdpsqlc query policy**	Query policy.
Mount-DpSqlBackup	**tdpsqlc mount backup**	Mounts a backup that provides access to the files that are contained by the backup.
Remove-DpSqlBackup	**tdpsqlc delete backup** and **tdpsqlc inactivate**	Use to delete a VSS backup of a SQL Server database, or inactivate one or more active legacy backup objects on the Tivoli Storage Manager server.
Remove-DpSqlPolicy	**tdpsqlc delete policy**	Deletes a local policy.
Reset-DpSqlTsmPassword	**tdpsqlc changetsmpassword**	Changes the Tivoli Storage Manager password used by Data Protection for SQL Server.
Restore-DpSqlBackup	**tdpsqlc restore**	Restore backups of Microsoft SQL Server data.
Set-DpSqlConfig	**tdpsqlc set paramname**	Set the Data Protection for SQL Server configuration parameters in the configuration file.
Set-DpSqlPolicy	**tdpsqlc update policy**	Changes an existing policy.

To view the details about a specific cmdlet, run the **Get-Help** cmdlet with the cmdlet name. For example:

```
Get-Help Get-DpSqlBackup
```

To continue the example, to see examples for the cmdlet, enter:

```
get-help Get-DpSqlBackup -examples
```

For more information, enter:

```
get-help Get-DpSqlBackup -detailed
```

For technical information, enter:

```
get-help Get-DpSqlBackup -full
```

To go to the information center, enter:

```
get-help Get-DpSqlBackup -online
```

For information about a specific parameter, enter:

```
help Get-DpSqlBackup -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 SQL with VSS backup-restore support

Data Protection for SQL provides support for protecting Microsoft SQL databases through two different methods. The most common method is through the Microsoft Server Managed Objects (SMO) application programming interface (API). Data Protection for SQL also can use the Microsoft Virtual Shadow Copy Service (VSS).

Problem determination assistance

If an error condition occurs during a Data Protection for SQL event, there are several sources of information you can view to help determine the problem:

- Data Protection for SQL logs information on backup, restore, and delete commands to the Tivoli Event Console.
- Data Protection for SQL logs information, by default, to the `tdpsql.log` file in the directory where Data Protection for SQL is installed. This file indicates the date and time of a backup, data backed up, and any error messages or completion codes. This file is very important and should be monitored daily.
- The Tivoli Storage Manager API logs API error information, by default, to the dsierror.log file in the directory where Data Protection for SQL is installed. No backup statistics are kept in this log. The dsierror.log file cannot be marked as read-only.
- The SQL Server logs information to the SQL Server error log. SQL Server error log information can be viewed using the SQL Server administration tools.
- The Tivoli Storage Manager scheduler logs information to both the dsmsched.log and the dsmerror.log files. By default, these files are located in the directory where the Tivoli Storage Manager Backup-Archive Client is installed.

  Note: Output from scheduled commands are sent to the scheduler log file (dsmsched.log). After scheduled work is performed, check the log to ensure the work completed successfully.

  When a scheduled command is processed, the scheduler log can contain the following entry:

  ```
  Scheduled event eventname completed successfully
  ```

  This is merely an indication that Tivoli Storage Manager successfully issued the scheduled command associated with the *eventname*. No attempt is made to determine the success or failure of the command. You can 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: return_code_number
  ```
- Windows Event Log.
- For VSS operations, view the `dsmerror.log` file in the backup-archive client installation directory.

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, 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 `D:\Program Files\IBM\Hardware Provider for VSS-VDS`.
  3. Issue the **`IBMVCFG SHOWCFG`** command to view the provider configuration information.
  4. Check that the CIMOM is properly configured. Run `verifyconfig.bat -u username -p password` on the Master Console.
  5. Check the username and password. If there is a problem with the truststore, follow the procedure in the documentation to generate a new truststore.

- CIMOM operational issues.

  If your backup or restore fails, check the `IBMVSS.log` file. If the failure is due to a CIMOM failure, the log displays output similar to the following output:

  ```
  Wed Jan 13 17:34:34.793 - Calling AttachReplicas
  Wed Jan 13 17:34:35.702 - AttachReplicas: 909ms
  Wed Jan 13 17:34:35.702 - returnValue: 34561
  Wed Jan 13 17:34:35.718 - AttachReplicas returned: 34561
  java.util.MissingResourceException: Can't find resource for
  bundle java.util.PropertyResourceBundle, key 1793
  at java.util.ResourceBundle.getObject(ResourceBundle.java:329)
  at java.util.ResourceBundle.getString(ResourceBundle.java:289)
  at com.ibm.cim.CIMException.<init>(CIMException.java:472)
  at ESSService.executeFlashCopy(ESSService.java:3168)
  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 generated by the command line interface or graphical user interface, depending on how you run your operation. These might provide more information on 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 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 Master Console version (if applicable)
- For DS8000, the CIM Agent version (if applicable)
- `IBMVSS.log`
- `IBMVDS.log`
- Application Event Log
- System Event Log

If the problem seems related to CIMOM, you also need the CIMOM logs. Run `CollectLogs.bat` and send the file that is created (`CollectedLogs.zip`) to IBM Support. The default location for SAN Volume Controller or Storwize V7000 is `C:\Program Files\IBM\svcconsole\support`, and the default location for DS8000 is `C:\Program Files\IBM\cimagent`.

Resolving errors during Data Protection for SQL processing

You might encounter a problem during Data Protection for SQL processing using the Microsoft Volume Shadow Copy Service (VSS).

About this task

If you encounter a problem during Data Protection for SQL processing when using VSS, complete the following steps:

Procedure

1. Retry the operation that failed.
2. If the problem persists, close other applications, especially those applications that interact with SQL (antivirus applications, for example). Retry the operation that failed.
3. If the problem still exists, perform the following steps:
   a. Shut down the SQL server.
   b. Restart the SQL server, including the SQL server VSS Writer service.
   c. Run the operation that failed.
4. If the problem occurs again, complete the following steps:
   a. Shut down the computer.
   b. Restart the computer.
   c. Run the operation that failed.
5. If the problem still recurs, determine that it is occurring on other SQL servers.

Determining if the problem is a Tivoli Storage Manager or SQL issue

This section provides information to help determine if the problem is a Data Protection for SQL Server issue or an SQL server issue.

For legacy operations, use the following procedures:

- Use the Backup or Restore utility provided in the SQL Server administrator program to see if the problem can be recreated.
- If the error message `ACO5350E An unknown SQL API error has occurred` is displayed, the SQL server encountered an unexpected situation. Microsoft assistance may be needed if the problem continues.

- Data Protection for SQL Server error messages occasionally contain an HRESULT code. Use this code to search Microsoft documentation and the Microsoft Knowledge Base for resolution information.

For Windows Server 2008 and later, try recreating the problem with the Microsoft DISKSHADOW application. This application is shipped with Windows Server 2008 and later.

For VSS operations, Try recreating the problem with the Microsoft DISKSHADOW application. This application can run backups using the Microsoft SQL VSS APIs. If the problem is recreatable with DISKSHADOW, then the problem most likely exists within the VSS provider or the SQL server. Microsoft ships DISKSHADOW with the Volume Shadow Copy Services (VSS) Software Developer's Kit (SDK). IBM Service can provide a copy of DISKSHADOW if you encounter problems obtaining or building this application.

Microsoft provides specific tracing for troubleshooting VSS issues. Consult Microsoft documentation for details.

Determining if the problem is a Data Protection for SQL issue or a general VSS issue

The Data Protection client interacts closely with the backup-archive client (DSMAGENT), which performs all of the Volume Shadow Copy Service (VSS) operations. Determine first if the problem is with the Microsoft VSS service or with the IBM Tivoli Storage Manager.

About this task

Perform the following steps to try to isolate the source of the error:

Procedure

1. Test the connectivity between the Data Protection client and the Tivoli Storage Manager DSMAGENT. Issue the **TDPSQLC QUERY SQL** command on the computer where the SQL server is installed to verify that your installation and configuration is correct. This command returns information on the following items:
   - SQL Server status
   - Databases
   - VSS components

   The following output is an example of the output generated by the **TDPSQLC QUERY SQL** command:

```
C:\Program Files\Tivoli\TSM\TDPSql>tdpsqlc query sql

IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1997, 2013. All rights reserved.

ACO5057I The C:\Program Files\Tivoli\tsm\TDPSql\tdpsql.log log
file was pruned successfully.

Connecting to SQL Server, please wait...

SQL Server Information
----------------------

SQL Server Name   .................. JAZZVM2EXCH2010\JAZZVM2SQL2K8R2
SQL Server Version ................. 10.50.1617 (SQL Server 2008 R2)


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

Writer Name            : SqlServerWriter
Local DSMAgent Node    : jazzvm2
Remote DSMAgent Node   :
Writer Status          : Online
Selectable Components  : 9


The operation completed successfully. (rc = 0)



C:\Program Files\Tivoli\TSM\TDPSql>
```

If the **TDPSQLC QUERY SQL** 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 issued to enable proxy authority for nodes. If all of the information returned to you seems correct, proceed to the next step.

2. Use the VSSADMIN or DISKSHADOW utility to recreate the VSS operation without the Tivoli Storage Manager intervening. When VSS operations are failing, use these programs to recreate the error to determine if this is a general VSS problem or a problem within the Tivoli Storage Manager code.

   VSSADMIN

   A utility that is preinstalled with your operating system. It can display current volume shadow copy backups and all installed shadow copy writers and providers in the command window. The following commands are example **VSSADMIN** commands:

   ```
   VSSADMIN LIST WRITERS
   VSSADMIN LIST PROVIDERS
   VSSADMIN LIST SHADOWS
   ```

   Restriction: The **VSSADMIN LIST SHADOWS** command does not list shadow copies of SAN-attached volumes for Windows 2008 and later. The **vssadmin** utility uses Microsoft Software Shadow Copy provider to list the shadow copies that are created.

   DISKSHADOW

   Before installing Data Protection for SQL Server, test the core VSS functionality. The following DISKSHADOW testing can be performed before any Tivoli Storage Manager components are installed:

a. Test non-persistent shadow copy creation and deletion by running the following **DISKSHADOW** commands:

```
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
```

Volumes `f:` and `g:` represent the SQL 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>begin backup
diskshadow>add volume f: (database volume)
diskshadow>add volume g: (log volume)
diskshadow>create
diskshadow>end backup
diskshadow>list shadows all (This may take a few minutes)
diskshadow>delete shadows all
```

Note: Volumes f: and g: represent the SQL database and log volumes. Repeat the DISKSHADOW commands four times and verify that the Windows event log file contains no errors.

When all of the test commands complete successfully, you can install the Tivoli Storage Manager components.

diskshadow

A utility included with the Microsoft Volume Shadow Copy Services software developers kit (SDK) that can be used to exercise most of the VSS infrastructure, such as creating/querying/deleting shadow copies. You can also use DISKSHADOW to create both persistent and nonpersistent shadow copies, transportable snapshots, as well as assign a drive letter or mount point to a shadow copy.

- The following items can be determined by using the VSSADMIN or DISKSHADOW utility:
  - Verify VSS provider configurations and setup
  - Rule out any possible VSS problems before running 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 non-persistent shadow copy creation and deletion:

  a. Run "`DISKSHADOW` *`k`*`:` *`l`*`:`" where *k*: and *l*: are the SQL Server database and log volumes.

  b. Repeat the previous step four times.

  c. Inspect the Windows Event Log to ensure there are no problems.

  Test persistent shadow copy creation and deletion:

a. Run "`DISKSHADOW -p` *k*: *l*:" (where *k*: and *l*: are the SQL Server database and log volumes. You might have to run "**`DISKSHADOW -da`**" to remove this if you do not have enough space.

b. Repeat the previous step four times.

c. Inspect the Windows Event Log to ensure there are no problems.

Test non-persistent 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 SQL Server database and log volumes.

b. Copy the resultant "`export.xml`" file from computer 1 to computer 2 before performing 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 there are no problems.

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. Recreate your specific problem by using DISKSHADOW. If you can only recreate your problem 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.
   - SQL VSS backups to local are simulated by running a DISKSHADOW persistent snapshot.
   - SQL VSS backups to the Tivoli Storage Manager are simulated by running a DISKSHADOW nonpersistent snapshot.
   - SQL VSS backups to local and to the Tivoli Storage Manager are simulated by running a DISKSHADOW persistent snapshot.
   - Offloaded SQL VSS backups to the Tivoli Storage Manager are simulated by running a DISKSHADOW nonpersistent, transportable snapshot.

   Refer to the DISKSHADOW documentation for the specific commands for performing backups.

   If you can recreate the problem, it most likely is a general VSS issue. Refer to Microsoft Knowledge Database for information. If your operation passes successfully with DISKSHADOW, it most likely is a Tivoli Storage Manager or Data Protection for SQL client problem.

Tracing the Data Protection client when using SQL VSS technology

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

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

Data Protection for SQL trace

Open the Management Console (MMC) and go to the diagnostics property page to turn tracing on. Tracing is off by default. Select one of the following diagnostic types:

- For legacy operations: Normal MMC, DP (service), API (service,api_detail)

- For VSS operations and large output size: Complete MMC, DP (service), API (service,api_detail), Agent (service)
- For full control over all settings: Custom

Tivoli Storage Manager API trace

Enable tracing with the DP/SQL DSM.OPT file and the "**TRACEFILE**" and "**TRACEFLAGS**" keywords. The following entry is an example of the entry in the DP/SQL DSM.OPT file:

```
TRACEFILE APITRACE.TXT
TRACEFLAG SERVICE
```

DSMAGENT trace

Enable tracing with the DSMAGENT DSM.OPT file and the "**TRACEFILE**" and "**TRACEFLAGS**" keywords. The following entry is an example of the entry in the DSMAGENT DSM.OPT file:

```
TRACEFILE AGTTRACE.TXT
TRACEFLAG ALL_VSS
```

The trace flag, in this instance, is ALL_VSS (you might need different traceflags, depending on the circumstance).

Gathering SQL with VSS information before calling IBM

The Data Protection client is dependent upon the operating system and the SQL application. Collecting all the necessary information about the environment can significantly assist in determining the problem.

The Management Console (MMC) can collect information and place it in a zip file that can then be provided to Support.

See "Emailing support files" on page 124 for more information on collecting diagnostic information to send to IBM.

Gather as much of the following information as possible before contacting IBM Support:

- The exact level of the Windows operating system, including all service packs and hotfixes that were applied.
- The exact level of the SQL Server, including all service packs and hotfixes that were applied.
- The exact level of Data Protection for SQL with Virtual 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 SQL **TDPSQLC QUERY SQL** command.
- The device type (and connectivity path) of the SQL databases and logs.
- (SAN only) The specific hardware that is being used. For example: HBA, driver levels, microcode levels, and 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 other applications running on the system.
- A list of the steps needed to recreate the problem (if the problem can be recreated).
- If the problem can not be recreated, list the steps that caused the problem.
- Is Data Protection for SQL running in a Microsoft Failover Clustering environment?
- Is the problem occurring on other SQL servers?

Gathering files from SQL with VSS before calling IBM

Several log files and other data can be collected for Data Protection for SQL server diagnosis.

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 many of the following files as possible before contacting IBM Support:

- 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 SQL configuration file. The default configuration file is `tdpsql.cfg`.
- The Data Protection for SQL Tivoli Storage Manager application programming interface (API) options file. The default options file is `dsm.opt`.
- The Tivoli Storage Manager registry hive export.
- The SQL 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 capture 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 could help to determine your problem. The following files are the Tivoli Storage Manager log files to gather:

- The Data Protection for SQL log file. The default location of this file is `C:\Program Files\Tivoli\TSM\TDPSql\tdpsql.log`
- The Tivoli Storage Manager API Error log file. The default location of this file is `C:\Program Files\Tivoli\TSM\TDPSql\dsierror.log`
- The DSMAGENT error log file. The default location of this file is `C:\Program Files\Tivoli\TSM\baclient\dsmerror.log`

- The DSMAGENT crash log file, if requested. The default location is `C:\Program Files\Tivoli\TSM\baclient\dsmcrash.log`.

The following VSS provider log files can also be helpful, if applicable:

- System Provider (Windows Event Log)
- IBM System Storage SAN Volume Controller, Storwize V7000, and 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

You can use the Data Protection for SQL console to list the events originated by Data Protection for SQL. Select **Dashboard** > **ServerName** > **Diagnostics** > **System Information** and double-click the `dpevents.ps1` script in the PowerShell section of the System Information page.

Installation Problems: Creating an installation-log file

In the event a silent installation fails, gather the following information to assist Customer Support when evaluating your situation:

- Operating system level
- Service pack
- Hardware description
- Installation package (DVD or electronic download) and level
- Any Windows event log relevant to the failed installation
- Windows services active during the failed installation (for example, anti-virus software)
- Whether you are logged on to the local console (not through a terminal server)
- Whether you are logged on as a local administrator, not a domain administrator (Tivoli does not support cross-domain installations)

You can create a detailed log file (setup.log) of the failed installation. Run the setup program (setup.exe) in the following manner:

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

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. Click **Start** > **Tivoli Storage FlashCopy Manager** > **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.

What to do next

Return to the Tivoli Storage FlashCopy Manager Management Console and begin backup operations.

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.
- Tivoli Storage FlashCopy Manager product support at http://www.ibm.com/software/tivoli/products/storage-flashcopy-mgr/. Enter the search term, such as an authorized program analysis report (APAR) number, release level, or operating system to narrow the search criteria for your support need.

Viewing system information

View scripts that provide information on system components such as Data Protection for SQL Server-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 start 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 scripts folder in the installation directory.

     The file type extension is used 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 perform.

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 perform.

  b. Click **Delete** in the Actions pane.

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 SQL Server 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 SQL Server default log and trace files:
  - Installation directory: `C:\Program Files\Tivoli\TSM\TDPSql`
  - `dsierror.log`
  - `tdpsql.log`
  - *`TraceFileSql.trc`*

  If the `tdpsql.log` is defined in a path other than the default `c:\program files\tivoli\TSM\TDPSql\tdpsql.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 `tdpsql.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.

Chapter 8. Performance

Data Protection for SQL Server provides certain parameters that can be tuned for optimum performance.

Many factors can affect the backup and restore performance of Data Protection for SQL Server, such as hardware configuration, network type, and capacity. These factors are not within the scope of this document. However, some parameters that are related to Data Protection for SQL Server can be tuned for optimum performance.

Note: Legacy backups are a stream of bytes that Data Protection for SQL Server stores on the Tivoli Storage Manager server. VSS backups differ since they are at the volume and file-level. In a situation where a SQL Server database is not fully allocated, a Legacy backup might transfer a smaller amount of data for a Tivoli Storage Manager backup than for a VSS backup because a VSS backup transfers the entire file, regardless of its allocation.

For Legacy and VSS Backups, the **RESOURCEUTILIZATION** client option is important. This option increases or decreases the ability of the client to create multiple sessions. The higher the value, the more session 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*.

Buffering (Legacy only)

Data Protection for SQL Server is a multi-threaded application that uses asynchronous execution threads to transfer data between the SQL and Tivoli Storage Manager servers. To accomplish this, multiple data buffers are used to allow one thread to receive data from one side, while another thread sends data to the other side. For example, one thread can be reading data from a SQL Server while another is sending data to the Tivoli Storage Manager server.

The number of buffers that Data Protection for SQL Server allocates to these threads is specified by the ***/buffers*** and ***/sqlbuffers*** parameters. The size of these buffers is specified by the ***/buffersize*** and ***/sqlbuffersize*** parameters. These parameters are set on the **Properties** page. When the parameters are set on the **Properties** page, the `dsm.opt` file is updated. You can also use the command-line interface to update the `dsm.opt` file.

Data Striping (Legacy only)

In addition to multi-threading to maximize throughput on a single session, Data Protection for SQL Server uses separate threads to support SQL data striping, which allows use of multiple parallel sessions to backup and restore a single database. This is another method to maximize data throughput. If a single session cannot fully exploit available bandwidth, multiple parallel sessions can yield improved data throughput, especially if the database is spread across multiple physical volumes.

If you use one data stripe per physical volume for both the SQL Server and the Tivoli Storage Manager server, the performance, which is measured as the amount of time necessary to backup or restore a particular SQL database, should show an improvement over the unstriped case. The improvement is approximately proportional to the number of data stripes used, given the constraints of the devices and the network used, and the striping independent overhead in SQL Server, Tivoli Storage Manager server, and Data Protection for SQL Server.

You can specify the number of stripes to use with the **/STRIPes** parameter on the command-line interface. You can also specify the number of stripes to use from the Management Console (MMC) GUI, by changing the number in the **Stripes** field in the Backup options or Restore options panel.

Note:

- Additional striping does not necessarily improve performance and may even decrease performance if system constraints involving real and paged memory, processors, network interface cards, networks, device reads and writes, and RAID become saturated or exceed capacity.
- If you use striping in conjunction with SQL buffers, be certain that the number of SQL buffers specified is equal to or greater than the number of stripes.
- The default values that Data Protection for SQL Server assigns to buffers, buffer size, and stripes can be changed in the Data Protection for SQL Server configuration file. Use the **set** command or the Performance property page in the MMC GUI to modify the configuration file.

LAN-free environment (Legacy and VSS)

Running Data Protection for SQL Server in a LAN-free environment if you are equipped to do so avoids network constraints.

- For Legacy backups, specify ***enablelanfree*** *yes* in the Data Protection for SQL Server options file.
- For VSS backups, specify ***enablelanfree*** *yes* in the DSMAGENT (VSS Requestor) dsm.opt file only.

For information on setting up a LAN-free environment, refer to the Tivoli publication: *IBM Tivoli Storage Manager for SAN for Windows Storage Agent User's Guide*.

Chapter 9. Reference information

Reference information for Data Protection for Microsoft SQL Server is provided.

Command overview

The name of the Data Protection for SQL command line interface is `tdpsqlc.exe`. This executable is located in the directory where Data Protection for SQL Server is installed.

Using the Data Protection for SQL Server command line interface from the GUI

Follow these steps to launch the Data Protection for SQL Server 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. Select an SQL 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 appear in any order after the required parameters
- optional parameters begin with a forward slash (/) or a dash (-)
- minimum abbreviations for keywords are indicated in upper case 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's value includes spaces, the value must be enclosed in double quotation marks
- a positional parameter can appear only once per command invocation

Data Protection for SQL Server uses the following command line syntax:

```
tdpsqlc command positional parameter 0 or more optional parameters
```

The **tdpsqlc** executable is followed by high level operations called *commands*. Each command accepts various command line parameters. These parameters consist of *positional parameters* and *optional parameters*. Specify positional parameters before other options in the command line. In the following case, the backup command with its database name *xyz*, the object to back up, is followed by the type of backup, **full**, a positional parameter, and finally by an optional parameter, **/sqlbuffers**.

```
tdpsqlc backup xyz full /sqlbuffers=2
```

You can display a complete list of Data Protection for SQL commands and all their parameters by simply entering **tdpsqlc** or the **tdpsqlc help|?** command. See "Help command" on page 160.

Command-line interface help

Issue the **tdpsqlc ?**, **tdpsqlc help**, or **tdpsqlc** command to display help for the command-line interface.

Command-line parameter characteristics

There are several characteristics to take note of in the Data Protection for SQL Server command-line interface:

- Do not include a slash or dash before positional parameters.
- Begin optional parameters with a forward slash (/) or a dash (-).
- You may place multiple optional parameters per command invocation in any order *after* positional parameters.
- You may abbreviate keywords. Minimum abbreviations are indicated in upper case in the syntax diagrams.
- All SQL names of databases or parts of databases are case-sensitive.
- Separate parameters with at least one space.
- Some keyword parameters may require a value; separate values from their keywords with an equal sign. (=).
- If a parameter's value includes spaces or special characters, enclose the value in double quotes.
- You can use most positional and optional parameters only once per command invocation. The following exceptions allow lists of values or repetition of the parameter:
  - FIle=
  - Group=
  - Log=
  - Set=
  - /FIles=
  - /GRoups=
  - /RELocate=
  - /RELOCATEDir=
  - /TO=

  For example: `/files=a,b` or `/files=a /files=b`

  Multiple instances of optional parameters do not have to be contiguous. For example: `/files=a /groups=y /files=b /groups=z`

Where repeatable syntax appears, separate multiple values with commas as indicated in the following:

```
                         ┌─,────┐
                         ▼      │
►►──TDPSQLC──Backup──┬───dbname─┴─┬──────────────────────────────────►◄
                     └─*──────────┘
```

Use the wildcard asterisk (*) following the command to select all instances on the server of database names or file names.

For help in reading syntax diagrams, refer to “Reading syntax diagrams” on page xi.

Data Protection for SQL Server parameters available by backup method

Table 11. Data Protection for SQL Server optional parameters

Optional Parameters	Legacy	VSS
/ACtive	Yes	Yes
/ADJUSTKBtsmestimate	Yes	No
/ADJUSTPERcenttsmestimate	Yes	No
/ALl	Yes	Yes
/BACKUPDESTination	Yes	Yes
/BACKUPMETHod	Yes	Yes
/BUFFers	Yes	No
/BUFFERSIze	Yes	No
/COMPATibilityinfo	Yes	Yes
/CONFIGfile	Yes	Yes
/DBOonly	Yes	No
/DIFFESTimate	Yes	No
/EXCLUDEDB	Yes	Yes
/FILEInfo	Yes	No
/FIles	Yes	No
/GRoups	Yes	No
/INSTANTRestore	No	Yes
/INTO	Yes	Yes
/LOGESTimate	Yes	No
/LOGFile	Yes	Yes
/LOGPrune	Yes	Yes
/MOUNTWait	Yes	No
/OBJect	Yes	Yes
/OFFLOAD	No	Yes
/OLDerthan	Yes	No
/PARTial	Yes	No
/Quiet	Yes	Yes
/RECOVery	Yes	Yes
/RELOCATEDir	Yes	Yes
/RELocate /TO	Yes	No

Table 11. Data Protection for SQL Server optional parameters (continued)

Optional Parameters	Legacy	VSS
/REPlace	Yes	No
/SQLAUTHentication	Yes	Yes
/SQLBUFFers	Yes	No
/SQLBUFFERSIze	Yes	No
/SQLPassword	Yes	Yes
/SQLSERVer	Yes	Yes
/SQLUSer	Yes	Yes
/STANDby	Yes	No
/STOPAT	Yes	No
/STOPATMark /AFTER	Yes	No
/STOPBEFOREMark /AFTER	Yes	No
/STRIPes	Yes	No
/TRUNCate	Yes	No
/TSMNODe	Yes	Yes
/TSMOPTFile	Yes	Yes
/TSMPassword	Yes	Yes

Data Protection for SQL Server allows you to perform online backups and restores of Microsoft SQL Server databases to Tivoli Storage Manager server storage using either command-line or graphical user interfaces (GUI).

Backup command

Use the **backup** command to back up all or part of one or more SQL databases from the SQL Server to Tivoli Storage Manager storage on the Tivoli Storage Manager server.

You can enter the * character to backup all databases. You can specify more than one database for multiple database and transaction log backups.

When you use the **backup** command, remember the following facts:

- Simple recovery model databases are automatically excluded from log backups.
- The master database is automatically excluded from log and differential backups.
- You cannot back up or restore the `tempdb` database because this database is created by the SQL server each time the server is started.
- Although full and differential backups include a part of the transaction log, that part is only what is required to make a restore consistent. It is not a log backup and does not truncate the log.
- The user id used by Data Protection for SQL to log on to the SQL server must have the SQL Server SYSADMIN fixed server role.
- You can use the TRANSACT-SQL database consistency checker statement `DBCC CHECKDB ('DBNAME')` to verify the integrity of the SQL databases before you back them up.

Backup syntax

Use the **backup** command syntax diagrams as a reference to view available options and truncation requirements.

TDPSQLC backup

```
                          ┌─,──────────┐
                          ▼            │
►►──TDPSQLC──Backup──┬─────dbname──────┴─┬──┬─FULL────────────────────────────────────┬──────►◄
                     └─*─────────────────┘  ├─COPYFull────────────────────────────────┤
                                            │          ┌─,───────────────────────┐    │
                                            │          ▼                         │    │
                                            ├─FIle=──┬───logicalfilename─────────┴─┬──┤
                                            │        └─*───────────────────────────┘  │
                                            ├─Difffull──┤ A ├─────────────────────────┤
                                            │            ┌─,───────────┐              │
                                            │            ▼             │              │
                                            ├─Group=──┬───groupname────┴─┬────────────┤
                                            │         └─*────────────────┘            │
                                            ├─Log──┤ B ├──────────────────────────────┤
                                            └─Set──┤ C ├──────────────────────────────┘
```

Backup optional parameters

```
     (1)    (2)
►►──────────────────────┬────────────────────────────────────────────────────┬──────────►
                        └─/AAGName──=AlwaysOn Availability Group name────────┘

►──┬──────────────────────────────────────┬──────────────────────────────────────────────►
   └─/ADJUSTKBtsmestimate──=numkb─────────┘

►──┬──────────────────────────────────────────────────┬──┬────────────────────┬──────────►
   └─/ADJUSTPERcenttsmestimate──=numpercent───────────┘  └─/ALWAYSONPriority──┘

►──┬──────────────────────────────────────┬──┬──────────────────────────────────┬────────►
   │                     ┌─TSM────┐       │  │                  ┌─LEGACY─┐      │
   └─/BACKUPDESTination=─┼────────┼───────┘  └─/BACKUPMETHod=───┼────────┼──────┘
                         ├─LOCAL──┤                             └─VSS────┘
                         └─BOTH───┘

►──┬───────────────────────────────────────┬─────────────────────────────────────────────►
   │           ┌─=3 [or cfg value]─┐       │
   └─/BUFFers──┼───────────────────┼───────┘
               └─=numbuffers───────┘

►──┬─────────────────────────────────────────────┬───────────────────────────────────────►
   │              ┌─=1024 [or cfg value]─┐       │
   └─/BUFFERSIze──┼──────────────────────┼───────┘
                  └─=buffersizeinkb──────┘
```

```
>--+---------------------------------------+--+--------------------------------+-->
   |              .-=tdpsql.cfg------.     |  |               .-,----------.   |
   '-/CONFIGfile--+------------------+-----'  |               V            |   |
                  '-=configfilename--'        '-/EXCLUDEdb-------=dblist,...-+-'

>--+-----------------------+--+--------------------------+---------------------->
   '-/EXCLUDEALwaysondbs---'  '-/EXCLUDESTandarddbs------'

>--+---------------------------------------------+------------------------------>
   |           .-=tdpsql.log [or cfg value]--.   |
   '-/LOGFile--+-----------------------------+---'
               '-=logfilename----------------'

>--+------------------------------------------+--+------------+--+----------+-->
   |            .-=60 [or cfg value]--.       |  '-/OFFLOAD---'  '-/Quiet---'
   '-/LOGPrune--+---------------------+-------'
                +-=numdays------------+
                '-=No-----------------'

>--+-----------------------------------------------------------------+---------->
   |                      .-=INTegrated [or cfg value]--.            |
   '-/SQLAUTHentication---+-----------------------------+------------'
                          '-=SQLuserid------------------'

>--+------------------------------------------+--------------------------------->
   |               .-=0 [or cfg value]--.     |
   '-/SQLBUFFers---+--------------------+-----'
                   '-=numsqlbuffers-----'

>--+-------------------------------------------------+--+----------------------------+-->
   |                  .-=1024 [or cfg value]--.      |  |                  .-=No--.  |
   '-/SQLBUFFERSIze---+-----------------------+------'  '-/SQLCOMPression--+------+--'
                      '-=sqlbuffersizeinkb----'                          '-=Yes-'

>--+---------------------------------------------+------------------------------>
   |                .-=" "----------------.      |
   '-/SQLPassword---+---------------------+------'
                    '-=sqlpasswordname----'

>--+-------------------------------------------------------------+-------------->
   |              .-=[local computer name or cfg value]--.       |
   '-/SQLSERVer---+--------------------------------------+-------'
                  '-=sqlprotocol:sqlservername-----------'

>--+--------------------------------+--+-----------------------------------------+-->
   |            .-=sa------------.  |  |             .-=1 [or cfg value]--.      |
   '-/SQLUSer---+----------------+--'  '-/STRIPes----+--------------------+------'
                '-=sqlusername---'                   '-=numstripes--------'
```

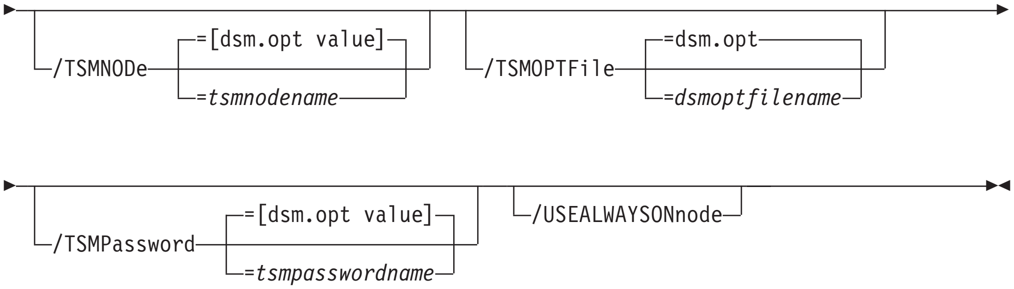

A Difffull Options:

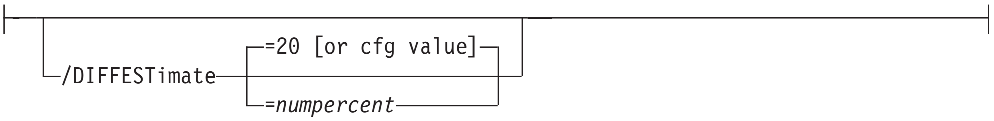

B Log Options:

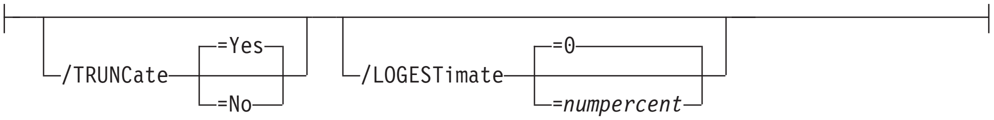

C Set Options:

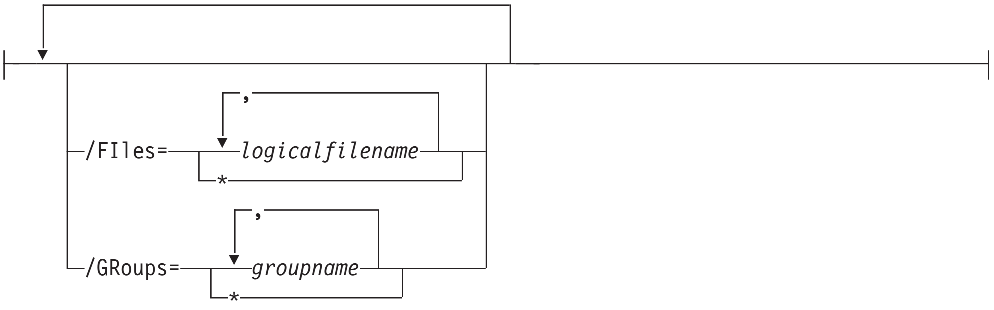

Notes:

1 For the optional parameters, the **/BACKUPMETHod=** is only valid when using the **full** or **copyfull** positional parameters. The **full** and **copyfull** backups can be performed using VSS or legacy operations. The **/BACKUPMETHod=** parameter is used to choose between the options. The **log**, **diff**, **file**, and **group** backups can be performed only when using legacy operations. The **/BACKUPMETHod=** parameter is not supported in with these types of backups because only legacy backups are viable.

2 The **/BACKUPDESTination** parameter is valid only when using the **full** or **copyfull** positional parameters. The **full** and **copyfull** backups can be saved to local storage, TSM server storage, or both. The **/BACKUPDESTination** parameter is used to choose among the options.

Backup positional parameters

Positional parameters immediately follow the **backup** command and precede the optional parameters.

FIle=*|*logicalfilename***,...**

A **FIle** backup contains only the contents of the SQL server logical file you specify. You can use this option when it is not practical to back up an entire SQL database due to available backup time and space or due to performance requirements. The *logicalfilename* variable specifies the names of the SQL server database logical files you want to back up or restore to.

Considerations:

- You should follow file backups with transaction log backups for all SQL databases you back up.
- You can specify this parameter more than once per command invocation.
- A new backup object inactivates any active backup object of the same name in the same SQL database.
- Use * as a wildcard character in *logicalfilename* to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all logical files in the SQL server database. Because each logical file backed up creates a separate backup object on the Tivoli Storage Manager server, specifying only the wildcard character results in a separate backup object for each logical file in the SQL server database.
- If *logicalfilename* includes spaces or special characters, enclose it in double quotes.
- The *logicalfilename* variable is case-sensitive.
- You cannot specify the **/recovery** parameter with **restore file** operations.
- A SQL server Create Index operation requires that you back up all affected filegroups as a unit. You cannot back up a file in the affected filegroups until you perform the unit backup. A SQL server error message will indicate which filegroups are affected. You must perform a full database backup or a set backup of at least the affected filegroups before the file backup succeeds.

FULL

A **FULL** legacy database backup contains all of the contents of a SQL server database plus enough of the database's transaction log to make a restore consistent. A **FULL** VSS database backup contains all of the contents of a SQL server database (database files, log files, and full-text index files).

Each SQL database backed up using the Legacy backup method creates a separate backup object on the Tivoli Storage Manager server. A new full database backup object inactivates all prior Legacy active backup objects for the same SQL database. This inactivation includes any active full backup object as well as any active file, group, set, differential, and log backup objects. For additional policy information, including VSS aspects, see “How Tivoli Storage Manager server policy affects Data Protection for SQL Server” on page 29 and “Setting automatic expiration (VSS and legacy)” on page 41.

COPYFull

A copy-only full backup contains a copy-only version of a full backup. These backups are considered out of the regular sequence of conventional SQL Server backups, and do not affect the transaction logs or any sequence of backups like differential backups or full backups. Use this option to create copy-only full backups periodically for long term retention without affecting existing backup schedules or retention policies for disaster recovery.

`Difffull`

A **`Difffull`** (differential) database backup contains only the parts of a SQL server database changed since the latest full backup plus enough of the SQL database's transaction log to make a restore consistent. As such, a differential backup usually takes up less space than a full backup. Use this option so that all individual log backups since the last full database backup do not need to be applied.

`Group=*|`*groupname*,...

A **`Group`** backup contains only the contents of the SQL server filegroup you specify. A group backup is useful when selected SQL database table or indexes have been assigned to a filegroup and only those tables or indexes need backing up. Specifically:

- You can save backup time by not backing up other tables or indexes in the SQL database that do not change as often.
- You can save restore time if, for example, the filegroup is on a different volume from the rest of the SQL database's filegroups and that volume needs to be restored. You need restore only that filegroup for that SQL database.

The *groupname* variable specifies the names of the SQL server database filegroups you want to back up.

Considerations:

- You can specify this parameter more than once per command invocation.
- A new group backup object inactivates any active group backup object of the same name in the same SQL database.
- Use * as a wildcard character in the *groupname* variable to replace zero or more characters for each occurrence.
- Specifying only the wildcard character indicates all filegroups in the SQL server database.

  Because each group backed up creates a separate backup object on the Tivoli Storage Manager server, specifying only the wildcard character results in a separate backup object for each filegroup in the SQL server database.
- If the *groupname* variable includes spaces or special characters, enclose it in double quotes.
- The *groupname* variable is case-sensitive.
- You should follow group backups with transaction log backups for all SQL databases you back up.
- You cannot perform group backups for the following SQL databases:
  - Those with the SQL Server attribute TRUNCATE LOG ON CHECKPOINT.
  - Those using the SIMPLE recovery model.
- You cannot specify the **`/recovery`** parameter with **`restore group`** operations.
- A SQL Server Create Index operation requires that you back up all affected filegroups as a unit. You cannot back up a single filegroup of the affected filegroups until you perform the unit backup. A SQL Server error message will indicate which filegroups are affected. You must

perform a full database backup or a set backup of at least the affected filegroups before the group backup succeeds.

Log or **Log=*|***logobjectname*,...

A log backup contains the contents of the transaction log for an active SQL server database since the latest successful log backup. This option can save backup time by requiring fewer SQL database backups. For **backup** operations, **Log** takes no values. Use * as a wildcard character in *logobjectname* to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all log backup objects for the SQL databases.

Considerations:

- You can control the size of a transaction log by allowing a log backup to truncate the inactive part of the transaction log. This is the default.
- By using the **/truncate**=no parameter, you may be able to backup the transaction log of a damaged, suspect, or unrecovered SQL Server database.
- Each log backed up creates a separate backup object with a unique name on the Tivoli Storage Manager server. A new log backup object does *not* inactivate any active backup objects (unlike the other backup types except **set** backups). Log backup objects do not participate in Tivoli Storage Manager server automatic expiration processing except when full database backup objects inactivate all active backup objects for a SQL database. Therefore, you can inactivate log backup objects using the **inactivate** command if full database backups are not performed frequently or at all.
- You cannot perform log backups for the following SQL databases:
  - Those with the SQL Server attribute TRUNCATE LOG ON CHECKPOINT.
  - Those using the SIMPLE recovery model.

Set or **Set=*|***setobjectname*,...

A **set** backup contains the contents of the SQL server filegroups and files you specify with the **/files** and **/groups** parameters. For **backup** operations, **set** takes no values. Use * as a wildcard character in *setobjectname* to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all set backup objects for the SQL databases.

Considerations:

- Use this option for unusual circumstances or special, one-time backups. One such case is when SQL Server requires that certain filegroups be backed up as a unit and a full database backup is not practical. See the description of the **file**, and **group** parameters in this section, specifically in reference to the Create Index operation.
- Each SQL database backed up creates a separate backup object on the Tivoli Storage Manager server. All of the files and filegroups backed up as part of a set backup for the same SQL server database are contained in a single backup object. Note that this is different from group and file backups, which create a separate backup object of each file and filegroup even if they are part of the same SQL server database.
- A new set backup object does *not* inactivate any active backup objects (unlike the other backup types except **log** backups). Set backup objects do not participate in Tivoli Storage Manager server automatic expiration processing except when full database backup objects inactivate all active

backup objects for a SQL database. Therefore, you can inactivate set backup objects using the **inactivate** command if full database backups are not performed or not performed frequently.

- You should follow set backups with transaction log backups for all SQL databases you back up.
- The **`file`**, **`group`**, **`log`**, and **`set`** parameters can take a list of values (repeatable syntax) and may be specified more than one time. For example: `file=a,b` or `file=a file=b`
- Multiple instances of optional parameters do not have to be contiguous. For example: `file=a group=y file=b group=z`

Backup optional parameters

Optional parameters follow the **`backup`** command and positional parameters.

/AAGName=*AlwaysOn Availability Group name*

When you backup a database list or all databases with the wildcard character, *, and specify the **`/AAGName`** parameter, only databases from the availability group that you specify are backed up.

/ADJUSTKBtsmestimate=*numkb*

The **`/adjustkbtsmestimate`** parameter specifies the number of kilobytes to add to the size of the backup estimate generated by the SQL Server. The *numkb* variable refers to the number of kilobytes to add. The number can range from 0 to 9999. The default is 0. Increasing the number of kilobytes may be necessary when the backup estimate (generated by the SQL Server) may be too low as the disk storage pool has cache enabled. For example, if maintenance is performed on the production server during a Data Protection for SQL Server backup, the size of transaction logs can increase beyond the original backup estimate and cause the backup to fail. Use this parameter to customize the number of kilobytes in the backup estimate and avoid possible backup failures.

/ADJUSTPERcenttsmestimate=*numpercent*

The **`/adjustpercenttsmestimate`** parameter specifies the percentage number to add to the size of the backup estimate. The *numpercent* variable refers to the percentage number to add. The number can range from *0* to *99*. The default is *0*. Increasing the percentage estimate may be necessary when the backup estimate (generated by the SQL Server) may be too low as the disk storage pool has cache enabled. For example, if maintenance is performed on the production server during a Data Protection for SQL Server backup, the size of transaction logs can increase beyond the original backup estimate and cause the backup to fail. Use this parameter to customize the percentage in the backup estimate and avoid possible backup failures.

`/ALWAYSONPriority`

Use this parameter to specify that a local availability database is backed up only if it has the highest backup priority among the availability replicas that are working properly on SQL Server 2012. You can use this parameter at the command-line interface or as part of a scheduled backup.

/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 is the default.

LOCAL
The backup is stored on local shadow volumes only. This is only valid when the **/BACKUPMETHod** parameter specifies VSS.

BOTH The backup is stored on Tivoli Storage Manager server storage and local shadow volumes. This is only valid when the **/BACKUPMETHod** parameter specifies VSS.

The **/BACKUPDESTination** parameter is valid only when using the **full** or **copyfull** positional parameters. The **full** and **copyfull** backups can be saved to TSM server storage, local storage, or both. The **/BACKUPDESTination** parameter is used to choose among options. The **log**, **diff**, **file**, and **group** backups can be stored only to TSM server storage. In this scenario, the **/BACKUPDESTination** parameter is not supported because TSM is the only viable option.

/BACKUPMETHod=LEGACY|VSS
Use the **/BACKUPMETHod** parameter to specify the manner in which the backup is performed.

You can specify:

LEGACY
The backup is performed with the legacy API. This is the SQL streaming backup and restore API as used in previous versions of Data Protection for SQL Server. This option is the default value.

VSS The backup is performed with VSS.

The **/BACKUPMETHod** parameter is valid only when using the **full** or **copyfull** positional parameters. The **full** and **copyfull** backups can be performed using VSS or legacy operations. The **/BACKUPMETHod** parameter is used to choose between the options. The **log**, **diff**, **file**, and **group** backups can only be performed using legacy operations. In this scenario, the **/BACKUPMETHod** parameter is not supported because legacy is the only viable option.

/BUFFers=*numbuffers*
The **/buffers** parameter specifies the number of data buffers used for each data stripe to transfer data between Data Protection for SQL Server and the Tivoli Storage Manager API. The *numbuffers* variable refers to the number of data buffers to use. The number can range from 2 to 8. The default is 3.

Considerations:

- You can improve throughput by increasing the number of buffers, but you will also increase storage use. Each buffer is the size specified in the **/buffersize** parameter.
- The default value is the value specified by the buffers configurable option in the Data Protection for SQL Server configuration file. This is initially 3.
- If you specify **/buffers**, its value is used instead of the value stored in the Data Protection for SQL Server configuration file. Specifying this parameter does not change the value in the configuration file.
- If you specify **/buffers** but not *numbuffers*, the default value 3 is used.

/BUFFERSIze=*buffersizeinkb*
The **/buffersize** parameter specifies the size of each Data Protection for SQL Server buffer specified by the **/buffers** parameter. The *buffersizeinkb*

variable refers to the size of data buffers in kilobytes. The number can range from 64 to 8192. The default is 1024.

Considerations:

- Though increasing the number of buffers can improve throughput, it also increases storage use as determined by this parameter.
- The default value is the value specified by the buffers configurable option in the Data Protection for SQL Server configuration file. This is initially 1024.
- If you specify **/buffersize**, its value is used instead of the value stored in the Data Protection for SQL Server configuration file. Specifying this parameter does not change the value in the configuration file.
- If you specify **/buffersize** but not *buffersizeinkb*, the default value 1024 is used.

/CONFIGfile=*configfilename*

The **/configfile** parameter specifies the name of the Data Protection for SQL Server configuration file, which contains the values for the Data Protection for SQL Server configurable options.

Considerations:

- *configfilename* can include a fully qualified path. If *configfilename* does not include a path, it uses the directory where Data Protection for SQL Server is installed.
- If *configfilename* includes spaces, place it in double quotes.
- If you do not specify **/configfile**, the default value is ***tdpsql.cfg***.

/EXCLUDEdb=*dblist*

The **/excludedb** parameter specifies the name of the databases to exclude from the backup operation. This parameter is available for all VSS and Legacy backup types.

/EXCLUDEALwaysondbs

Use this parameter to exclude all AlwaysOn Availability Databases from the backup operation. If you want to exclude specific databases, use the **/excludedb** parameter.

/EXCLUDESTandarddbs

Use this parameter to exclude all standard databases from the backup operation. If you want to exclude specific databases, use the **/excludedb** parameter.

/LOGFile=*logfilename*

The **/logfile** parameter specifies the name of the activity log that is generated by Data Protection for SQL Server. This activity log records significant events such as completed commands and error messages. The Data Protection for SQL Server activity log is distinct from the SQL Server error log. The *logfilename* variable identifies the name to be used for the activity log generated by Data Protection for SQL Server.

Considerations:

- If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file.
- The file name can include a fully-qualified path; however, if you specify no path, the file is written to the directory where Data Protection for SQL Server is installed.

- You cannot turn Data Protection for SQL Server activity logging off. If you do not specify **/logfile**, log records are written to the default log file. The default log file is *tdpsql.log*.
- When using multiple simultaneous instances of Data Protection for SQL Server to perform operations, use the **/logfile** parameter to specify a different log file for each instance used. This directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=*numdays***|No**

The **/logprune** parameter prunes the Data Protection for SQL Server activity log and specifies how many days of entries are saved. By default, log pruning is enabled and performed once each day Data Protection for SQL Server is executed; however, this option allows you to disable log pruning or explicitly request a prune of the log for one command run even if the log file has already been 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.

Considerations:

- If you specify *numdays*, it can range from *0* to *9999*. A value of *0* deletes all entries in the Data Protection for SQL Server activity log file except for the current command entries.
- If you specify **/logprune**, its value is used instead of the value stored in the Data Protection for SQL Server 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 SQL Server log file. If you are running a command that may prune the log file and the value of the **TIMEformat** or **DATEformat** parameter has changed, perform 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.

/MOUNTWait=Yes|No

This parameter is not valid for all backup types; does not work with DIFFFULL or LOG backup types. If the Tivoli Storage Manager server is configured to store backup data on removable media, it is possible that the Tivoli Storage Manager server might indicate to Data Protection for SQL Server that it is waiting for a required storage volume to be mounted. If that occurs, this option allows you to specify whether to wait for the media mount or stop the current operation. You can specify:

Yes Wait for tape mounts (default).

No Do not wait for tape mounts.

Considerations:

- If you use data striping for Legacy operations, Data Protection for SQL Server cannot complete waiting until the initial media for stripes are available, although Data Protection for SQL Server starts to use each stripe as its media becomes available. Because of the way SQL Server distributes data among stripes, if any stripe does not have its media available, each of the stripes may eventually be either waiting for its own or another stripe's media to become available. In this case, it may become necessary to terminate the Data Protection for SQL Server command from a prolonged wait. This can be done *only* by terminating the Data Protection for SQL Server program (close the command prompt window or enter ***control-c***).

- If the management class for meta objects also requires removable media, Data Protection for SQL Server waits for that volume, but because meta objects are not created until after the data objects are complete, the wait occurs *after* the data is transferred.
- If you specify `no` and any removable media are required, Data Protection for SQL Server terminates the command with an error message. This is also true if the management class for meta objects requires removable media. Since the meta objects are not created until after the data objects are complete, the command termination does not occur until after the database data is transferred.
- If you do not specify **/mountwait**, the default value is that specified in the `mountwait` configurable option in the Data Protection for SQL Server configuration file. This is initially *yes*. Specifying this parameter does not change the value in the configuration file.

/OFFLOAD

Specify this parameter to perform the backup of files to Tivoli Storage Manager on the machine specified by the **remotedsmagentnode** instead of the local machine. This parameter is valid when the following parameters and options are set: **/backupmethod=VSS** and **/backupdestination=TSM**. Note that this parameter requires a VSS provider that supports transportable shadow copies. It is not supported with the default Windows VSS System Provider.

/SQLAUTHentication=INTegrated | SQLuserid

This parameter specifies the authorization mode used when logging on to the SQL server. The *integrated* value specifies Windows authentication. The user id you use to log on to Windows is the same id you will use to log on to the SQL server. This is the default value.

Use the *sqluserid* value to specify SQL Server user id authorization. The user id specified by the **/sqluserid** parameter is the id you use to log on to the SQL server. Any SQL user id must have the SQL Server SYSADMIN fixed server role.

/SQLBUFFers=*numsqlbuffers*

The **/sqlbuffers** parameter specifies the total number of data buffers SQL Server uses to transfer data between SQL Server and Data Protection for SQL Server. The *numsqlbuffers* variable refers to the number of data buffers to use. The number can range from 0 to 999. The initial value is 0. When **/sqlbuffers** is set to *0*, SQL determines how many buffers should be used.

Considerations:

- The default value is the value specified by the SQL buffers configurable option in the Data Protection for SQL Server configuration file. This is initially 0.
- If you specify **/sqlbuffers**, its value is used instead of the value stored in the Data Protection for SQL Server configuration file. Specifying this parameter does not change the value in the configuration file.
- If you specify **/sqlbuffers** but not `numsqlbuffers`, the default value *0* is used.

/SQLBUFFERSIze=*sqlbuffersizeinkb*

The **/sqlbuffersize** parameter specifies the size of each buffer (specified by the **/sqlbuffers** parameter) SQL Server uses to transfer data to Data Protection for SQL Server. The *sqlbuffersizeinkb* variable refers to the size of data buffers in kilobytes. The number can range from 64 to 4096. The default is 1024.

Considerations:

- The default value is the value specified by the SQL buffers configurable option in the Data Protection for SQL Server configuration file. This is initially 1024.
- If you specify **/sqlbuffersize**, its value is used instead of the value stored in the Data Protection for SQL Server configuration file. Specifying this parameter does not change the value in the configuration file.
- If you specify **/sqlbuffersize** but not *sqlbuffersizeinkb*, the default value 1024 is used.

/SQLCOMPresssion=No|Yes

The **SQLCOMPresssion** parameter specifies if SQL compression is applied. If you do not specify **/sqlcompression**, the default value No is used. This parameter is only applicable on systems running SQL Server 2008 and later. For SQL Server 2008, backup compression is only supported on the Enterprise Edition. For SQL Server 2008 R2, backup compression is supported on Standard, Enterprise, and Datacenter editions.

/SQLPassword=*sqlpasswordname*

This parameter specifies the SQL password that Data Protection for SQL Server uses to log on to the SQL server that objects are backed up from or restored to.

Considerations:

- Using this parameter means that you are using SQL Server authentication. The SQL Server and the SQL user id for this password must both be configured for SQL Server authentication.
- If you do not specify **/sqlpassword**, the default value is blank (" ").
- If you specify **/sqlpassword** but not *sqlpasswordname*, the default is also blank (" ").
- This parameter is ignored if you use the **/sqlauth=integrated** parameter with it.

/SQLSERVer=*sqlprotocol:sqlservername*

The **/sqlserver** parameter specifies the SQL server that Data Protection for SQL Server logs on to. The *sqlprotocol* variable specifies the communication protocol to use. You can specify one of the following protocols:

- *lpc*: Use Shared Memory protocol.
- *np*: Use Named Pipes protocol.
- *tcp*: Use Transmission Control protocol.
- *via*: Use Virtual Interface Architecture protocol.

If no protocol is specified, Data Protection for SQL Server logs on to the SQL server according to the first protocol that becomes available.

Considerations:

- The default value is the value specified by the SQL server configurable option in the Data Protection for SQL Server configuration file. This is initially the local computer name.
- If you specify **/sqlserver** but not *sqlservername*, the local computer name is used.

- The following two shortcuts are accepted as the local computer name: `.` `(local)` These are a period or the word *local* within parentheses.
- If the SQL server is a member of a fail-over cluster, the CLUSTERNODE option in the Tivoli Storage Manager options file must have the value YES.
- You must specify the name if the SQL server is not the default instance or is a member of a fail-over cluster.
- The format of *sqlservername* depends on what type of instance it is and whether it is clustered or not:

Format	Instance?	Clustered?	Name required?
local-computername	default	no	no
local-computername\ instancename	named	no	yes
virtualservername	default	yes	yes
virtualservername\ instancename	named	yes	yes

localcomputername
: The network computer name of the computer the SQL server and Data Protection for SQL Server reside on. The TCP/IP host name may not always be the same.

instancename
: The name given to the named instance of SQL Server specified during installation of the instance.

virtualservername
: The name given to the clustered SQL Server specified during clustering service setup. This is not the cluster or node name.

/SQLUSer=*sqlusername*

The **/sqluser** parameter specifies the name that Data Protection for SQL Server uses to log on to the SQL server.

Considerations:

- Using this parameter means that you are using SQL Server authentication. The SQL Server and the SQL user id for this password must both be configured for SQL Server authentication.
- The SQL user id must have the SQL server SYSADMIN fixed server role.
- If you do not specify **/sqluser**, the default is **sa**.
- If you specify **/sqluser** but not *sqlusername*, the default is also **sa**.
- This parameter is ignored if you use the **/sqlauth=integrated** parameter with it.

/STRIPes=*numstripes*

The **/stripes** parameter specifies the number of data stripes to use in a backup or restore operation. The *numstripes* variable can range from 1 to 64.

Considerations:

- If you do not specify **/stripes**, the default value is that specified in the Data Protection for SQL Server configuration file. The initial value is 1.
- If you specify **/stripes** but not *numstripes*, the stored value is used.

- You may use *up to* the number used to create the backup. You can determine the number of data stripes used to create a backup object with the Data Protection for SQL Server command: `query tsm dbname backup_object`
- You must use the MAXNUMMP parameter on a Tivoli Storage Manager REGISTER NODE or UPDATE NODE command to allow a node to use multiple sessions to store data on removable media (which requires you to allocate multiple mount points to that node). The MAXNUMMP value must be equal to or greater than the maximum number of stripes you desire.
- When you use data striping, you should use Tivoli Storage Manager server file space collocation to try to keep each stripe on a different storage volume.
- The maximum number of data stripes you can use is one less than the value of the Tivoli Storage Manager server TXNGROUPMAX option in the ***dsmserv.opt*** file.

/TSMNODe=*tsmnodename*

The **/tsmnode** parameter specifies the Tivoli Storage Manager node name that Data Protection for SQL Server uses to log on to the Tivoli Storage Manager server. This identifies which Tivoli Storage Manager client is requesting services. You can also store the node name in the options file. The command line parameter overrides the value in the options file.

Considerations:

- You cannot use the **/tsmnode** parameter if PASSWORDACCESS GENERATE is specified in the Tivoli Storage Manager options file. You must specify the nodename in the options file. Otherwise, you can change PASSWORDACCESS to PROMPT to utilize the **/tsmnode** parameter. For details about the Tivoli Storage Manager options file, see the reference manual *IBM Tivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide*.
- If you do not specify **/tsmnode**, the default value is that specified by the nodename option in the Tivoli Storage Manager options file. Specifying this parameter does not change the value in the options file.

/TSMOPTFile=*dsmoptfilename*

The **/tsmoptfile** parameter specifies the Tivoli Storage Manager options file to use. This is similar to selecting a Tivoli Storage Manager server from the server list in the GUI. The Tivoli Storage Manager options file contains the configuration values for the Tivoli Storage Manager API. For details about the Tivoli Storage Manager options file, see the reference manual *IBM Tivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide*.

Considerations:

- The *tsmoptfilename* variable can include a fully qualified path. If you do not include a path, the directory where Data Protection for SQL Server is installed is used.
- If *tsmoptfilename* includes spaces, you must enclose it in double quotes.
- If you do not specify **/tsmoptfile**, the default value is ***dsm.opt***.
- If you specify **/tsmoptfile** but not *tsmoptfilename*, the default is also ***dsm.opt***.

/TSMPassword=*tsmpasswordname*

The **/tsmpassword** parameter specifies the Tivoli Storage Manager password that Data Protection for SQL Server uses to log on to the Tivoli Storage Manager server. This parameter and the option PASSWORDACCESS in the Tivoli Storage Manager options file interact in the following ways:

/tsmpassword	PASSWORDACCESS in Tivoli Storage Manager options file	Password already stored in registry?	Result
specified	*generate*	yes	*/tsmpassword* ignored
specified	*generate*	no	*/tsmpassword* used and stored
specified	*prompt*	—	*/tsmpassword* used
not specified	*prompt*	—	user is prompted

/USEALWAYSONnode

Specify this parameter to back up standard databases on SQL Server 2012 by using the AlwaysOn node. By setting this parameter, you can back up all availability databases and standard databases under a single node to help you to manage your database backups more easily. By default, SQL Server 2012 availability databases are backed up to the AlwaysOn node.

Legacy backup examples

The following examples are provided so you might see how the **backup** command can be entered with various parameters and options.

Full backup examples

If you want to complete a full backup, the following examples are provided to help model the command syntax.

- To complete a legacy full backup of two databases (*model* and *msdb*) to Tivoli Storage Manager server storage with the **/sqlbuffers** and **/stripes** optional parameters, enter the following command:

```
tdpsqlc backup model,msdb full /sqlbuffers=2 /stripes=2
```

- To complete a legacy full backup of a database (*test2*) with no output displayed, because the **/quiet** parameter is used, and the default Windows authentication mode is overridden with the use of the **/sqlauthentication** parameter, enter the following command:

```
tdpsqlc backup test2 full /quiet /sqlauth=sql
```

- To complete a legacy full backup of all available databases with the wildcard character (*) while using the **/excludedb** parameter to exclude the *master* and *msdb* databases from the backup process, enter the following command:

```
tdpsqlc backup * full /excludedb=master,msdb
```

- To complete a full backup of the a database (*test1*) with the parameters that customizes the number of kilobytes, **/adjustkbtsmestimate**, enter the following command:

```
tdpsqlc backup test1 full /adjustkbtsmestimate=25
```

- To complete a full backup of all standard databases, enter the following command:

```
tdpsqlc backup * full /EXCLUDEALwaysondbs
```

- To complete a log backup of all availability databases, enter the following command:

```
tdpsqlc backup * log /EXCLUDESTandarddbs
```

- For a more complex example, consider the following scenario: There are three AlwaysOn Availability Groups. The first availability group is called *AG01* with the following databases:
  - AlwaysOn Availability Database called *AlwaysOnLegacyDB1*
  - AlwaysOn Availability Database called *AlwaysOnLegacyDB3*

  The second availability group is called *AG03* with the following AlwaysOn Availability Database: *AlwaysOnLegacyDB2*. The third availability group is called *AG04* with the following databases:
  - AlwaysOn Availability Database called *AlwaysOnLegacyDB5*
  - AlwaysOn Availability Database called *AlwaysOnLegacyDB6*
  - Standard database called *SQL_DB1*
  - Standard database called *SQL_DB2*

  To complete a full backup with list matching both standard and availability databases, but excluding standard databases, enter the following command:

  ```
  C:\Program Files\tivoli\tsm\TDPSql>tdpsqlc backup AlwaysOnLegacy*,SQL*
   full /backupdest=TSM /backupmeth=legacy /EXCLUDESTandarddbs
  ```
- When using the **/AAGName** parameter to filter the databases that are backed up, refer to the following scenario with the examples: There are two AlwaysOn Availability Groups. The first availability group is called *AG01* with the following databases:
  - AlwaysOn Availability Database called *AlwaysOnLegacyDB1*
  - AlwaysOn Availability Database called *AlwaysOnLegacyDB3*

  The second availability group is called *AG04* with the following databases: with databases:
  - AlwaysOn Availability Database called *AlwaysOnLegacyDB5*
  - AlwaysOn Availability Database called *AlwaysOnLegacyDB6*

  When you enter a **backup** command for all databases, but use the **/AAGName** parameter to include only databases from *AG01* in the backup, enter the following command:

  ```
  C:\Program Files\tivoli\tsm\TDPSql>tdpsqlc backup * full /backupdest=TSM
   /backupmeth=legacy /AAGName=AG01
  ```

  When you enter a **backup** command for a database list with wildcards, but use the**/AAGName** parameter to include only databases from *AG04* in the backup, enter the following command:

  ```
  C:\Program Files\tivoli\tsm\TDPSql>tdpsqlc backup AlwaysOn*,SQL* full
   /backupdest=TSM /backupmeth=legacy /AAGName=AG04
  ```

  When you enter a **backup** command for a database list with wildcards, but do not match all databases from the specified AlwaysOn Availability Group, enter the following command:

  ```
  C:\Program Files\tivoli\tsm\TDPSql>tdpsqlc back *DB5 full /backupdest=TSM
   /backupmeth=legacy /AAGName=AG04
  ```

Differential backup examples

If you want to complete a differential backup, the following examples are provided to help model the command syntax.

- To complete a legacy differential backup of the previous full backup of a database (*test2*), including an estimate of the changed portion of the *test2* database, enter the following command:

```
tdpsqlc backup test2 difffull /diffest=10
```

- To complete a legacy differential backup of all available databases using the wildcard character (*) using the **/excludedb** parameter to exclude the *master* and *msdb* databases from the backup, enter the following command:

  ```
  tdpsqlc backup * difffull /excludedb=master,msdb
  ```

- To complete a differential backup with a database list matching both standard and availability databases, but excluding availability databases, enter the following command:

  ```
  C:\Program Files\tivoli\tsm\TDPSql>tdpsqlc backup AlwaysOnLegacy*,SQL*
   diff /EXCLUDEALwaysondbs
  ```

Log backup example

To complete a legacy log backup of the previous full backups of two databases (*test2* and *model*) while overriding the default truncation of the log files, enter the following command:

```
tdpsqlc backup test2,model log /truncate=no
```

Group backup example

To complete a legacy backup of all filegroups belonging to a database called *netapp_db2*, enter the following command:

```
tdpsqlc backup netapp_db2 Group=*
```

File backup example

To complete a legacy file backup of all files belonging to a database (*test2*) using the wildcard character (*), enter the following command:

```
tdpsqlc backup test2 file=*
```

Set backup example

To complete a legacy set backup of one filegroup and two separate files (jointly as a single backup object) from a database (*test2*) while using the **/groups** and **/files** parameters to specify which items constitute this set backup, enter the following command:

```
tdpsqlc backup test2 set /groups=primary /files=test2_2data,
test2_3data
```

Copy-only full backup example

To complete a legacy copy-only full backup of the availability database (*filestreamdb*) in a SQL Server 2012 AlwaysOn Availability Group environment, enter the following command:

```
tdpsqlc backup filestreamdb CopyFull /backupdestination=TSM
/backupmethod=legacy
```

VSS backup examples

The following examples are provided so you might see how the **backup** command can be entered with various parameters and options.

Full local backup examples

If you want to complete a full local backup, the following examples are provided to help model the command syntax.

- To complete a VSS full backup of a database (*test1*) to local shadow volumes using the **/backupdestination** and **/backupmethod** optional parameters, enter the following command:

  ```
  tdpsqlc backup test1 full /backupdestination=local
  /backupmethod=vss
  ```
- To complete a VSS full backup of all available databases to local shadow volumes using the wildcard character (*) and the **/excludedb** parameter to exclude the *master* and *msdb* databases from being backed up, enter the following command:

  ```
    tdpsqlc backup * full /backupdestination=local /backupmethod=vss
   /exclude=master,msdb
  ```
- To complete a VSS full backup of a SQL Server 2012 availability database (*hkaagdb*) to local shadow volumes, enter the following command:

  ```
  tdpsqlc backup hkaagdb full /backupdestination=local /backupmethod=vss
  ```
- To complete a full backup of all standard databases, enter the following command:

  ```
  tdpsqlc backup * full /EXCLUDEALwaysondbs
  ```
- To complete a log backup of all availability databases, enter the following command:

  ```
  tdpsqlc backup * log /EXCLUDESTandarddbs
  ```

Full local backup with Tivoli Storage Manager server example

To complete a VSS full backup of database (*model*) to local shadow volumes and Tivoli Storage Manager server storage using the **/backupmethod** parameter, enter the following command:

```
  tdpsqlc backup model full /backupmethod=vss
```

Copy-only full backup to Tivoli Storage Manager server example

To complete a VSS copy-only full backup of the full backup of the *filestreamdb* database to the Tivoli Storage Manager server storage using the **/backupmethod** parameter, enter the following command:

```
tdpsqlc backup filestreamdb CopyFull /backupdestination=TSM
/backupmethod=vss
```

Changetsmpassword command

Use the **changetsmpassword** command to change the Tivoli Storage Manager password used by Data Protection for SQL.

Changetsmpassword

Use the **changetsmpassword** command syntax diagrams as a reference to view available options and truncation requirements.

TDPSQLC CHANGETSMPassword command

```
►►──TDPSQLC──CHANGETSMPassword──┬──────────────────────────────────────────┬──────►◄
                                └─oldpassword──newpassword──verifypassword─┘
```

Optional Parameters

```
►►──┬────────────────────────────────────┬─────────────────────────────────────────►
    │              ┌─=tdpsql.cfg──────┐  │
    └─/CONFIGfile──┴─=configfilename──┴──┘

►──┬────────────────────────────────────────────┬──────────────────────────────────►
   │            ┌─=tdpsql.log [or cfg value]─┐  │
   └─/LOGFile───┴─=logfilename───────────────┴──┘

►──┬────────────────────────────────────┬──┬───────────────────────────────────┬───►
   │             ┌─=60 [or cfg value]─┐ │  │            ┌─=[dsm.opt value]─┐   │
   └─/LOGPrune───┼─=numdays───────────┤ │  └─/TSMNODe───┴─=tsmnodename─────┴───┘
                 └─=No────────────────┘ │
   ─────────────────────────────────────┘

►──┬──────────────────────────────────────┬───────────────────────────────────────►◄
   │              ┌─=dsm.opt──────────┐   │
   └─/TSMOPTFile──┴─=dsmoptfilename───┴───┘
```

Changetsmpassword positional parameters

Positional parameters immediately follow the **changetsmpassword** command and precede the optional parameters.

You are prompted for the following parameters if you do not specify them with the **changetsmpassword** command:

oldpassword
This specifies the old (current) Tivoli Storage Manager password you want to change.

newpassword
This specifies the new Tivoli Storage Manager password.

A Tivoli Storage Manager password is not case sensitive and may be composed of 1 to 63 of the following characters:

- the letters A through Z
- the digits 0 through 9

- the special characters plus (+), period (.), underscore (_), hyphen (—), and ampersand (&)

verifypassword
This specifies the new Tivoli Storage Manager password again as a verification that *newpassword* is correct.

Changetsmpassword optional parameters

Optional parameters follow the **changetsmpassword** command and positional parameters.

/CONFIGfile=*configfilename*
The **/configfile** parameter specifies the name of the Data Protection for SQL configuration file, which contains the values for the Data Protection for SQL configurable options. See "Set positional parameters" on page 239 for details on the file's contents.

Considerations:

- *configfilename* can include a fully qualified path. If *configfilename* does not include a path, it uses the directory where Data Protection for SQL is installed.
- If *configfilename* includes spaces, place it in double quotes.
- If you do not specify **/configfile**, the default value is ***tdpsql.cfg***.
- If you specify **/configfile** but not *configfilename*, the default value ***tdpsql.cfg*** is used.

/LOGFile=*logfilename*
The **/logfile** parameter specifies the name of the activity log that is generated by Data Protection for SQL. This activity log records significant events such as completed commands and error messages. The Data Protection for SQL activity log is distinct from the SQL Server error log. The *logfilename* variable identifies the name to be used for the activity log generated by Data Protection for SQL.

Considerations:

- If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file.
- The file name can include a fully-qualified path; however, if you specify no path, the file is written to the directory where Data Protection for SQL is installed.
- You cannot turn Data Protection for SQL activity logging off. If you do not specify **/logfile**, log records are written to the default log file. The default log file is ***tdpsql.log***.
- When using multiple simultaneous instances of Data Protection for SQL 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**
The **/logprune** parameter prunes the Data Protection for SQL activity log and specifies how many days of entries are saved. By default, log pruning is enabled and performed once each day Data Protection for SQL is executed; however, this option allows you to disable log pruning or explicitly request a prune of the log for one command run even if the log

file has already been 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.

Considerations:

- If you specify *numdays*, it can range from 0 to 9999. A value of 0 deletes all entries in the Data Protection for SQL activity log file except for the current command entries.
- If you specify **no**, the log file is not pruned during this command.
- If you do not specify **/logprune**, the default value is that specified by the logprune configurable option in the Data Protection for SQL configuration file. This is initially 60.
- If you specify **/logprune**, its value is used instead of the value stored in the Data Protection for SQL configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/logprune** without specifying *numdays* or **no**; in this case, the default 60 is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an undesired pruning of the log file. If you are running a command that may prune the log file and the value of the **TIMEformat** or **DATEformat** parameter has changed, perform 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.

/TSMNODe=*tsmnodename*

The **/tsmnode** parameter specifies the Tivoli Storage Manager node name that Data Protection for SQL uses to log on to the Tivoli Storage Manager server. This identifies which Tivoli Storage Manager client is requesting services. You can also store the node name in the options file. The command line parameter overrides the value in the options file.

Considerations:

- You cannot use the **/tsmnode** parameter if PASSWORDACCESS GENERATE is specified in the Tivoli Storage Manager options file. You must specify the nodename in the options file. Otherwise, you can change PASSWORDACCESS to PROMPT to utilize the **/tsmnode** parameter. For details about the Tivoli Storage Manager options file, see the reference manual *IBM Tivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide*.
- If you do not specify **/tsmnode**, the default value is that specified by the nodename option in the Tivoli Storage Manager options file. Specifying this parameter does not change the value in the options file.

/TSMOPTFile=*dsmoptfilename*

The **/tsmoptfile** parameter specifies the Tivoli Storage Manager options file to use. This is similar to selecting a Tivoli Storage Manager server from the server list in the GUI. The Tivoli Storage Manager options file contains the configuration values for the Tivoli Storage Manager API. For details about the Tivoli Storage Manager options file, see the reference manual *IBM Tivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide*.

Considerations:

- The *tsmoptfilename* variable can include a fully qualified path. If you do not include a path, the directory where Data Protection for SQL is installed is used.
- If *tsmoptfilename* includes spaces, you must enclose it in double quotes.
- If you do not specify **/tsmoptfile**, the default value is ***dsm.opt***.
- If you specify **/tsmoptfile** but not *tsmoptfilename*, the default is also ***dsm.opt***.

Changetsmpassword output examples

This output example provides a sample of the text, messages, and process status that displays when using the **changetsmpassword** command.

The following displays changing the Tivoli Storage Manager password.

Command:

```
tdpsqlc changetsmp sqlv2old sqlv2new sqlv2new
```

Output:

```
IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 6, Release 3, Level 0.0
(C) Copyright IBM Corporation 1998, 2011.
All rights reserved.

ACO0260I Password successfully changed.
```

Delete Backup command

Use the **delete backup** command to delete a VSS backup of a SQL Server database.

You must have local registry rights (for all versions of SQL Server) to perform a Data Protection for SQL Server delete backup.

Delete Backup syntax

Use the **delete backup** command syntax diagrams as a reference to view available options and truncation requirements.

TDPSQLC command

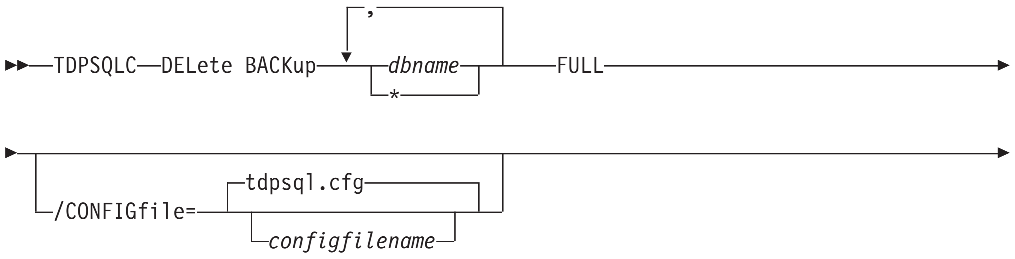

```
►──┬────────────────────────────────────────────────────────┬──────────────────────────►
   │                  ┌─=sqlserver value [or cfg value]─┐   │
   └─/FROMSQLSERVer=──┴─=sqlservername──────────────────┴───┘


►──┬────────────────────────────┬──┬──────────────────────────┬──┬─────────┬──────────►◄
   │               ┌─60──────┐  │  └─/OBJect=───objectname,...─┘  └─/Quiet──┘
   └─/LOGPrune=────┼─numdays─┼──┘
                   └─No──────┘
```

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:

*** | *dbname***

* Delete the active backups of all databases.

dbname
Delete a backup of the specified database. The active backup is deleted unless you specify a different backup with the **/object** optional parameter.

Multiple entries are separated by commas. If separated by commas, make sure there is no space between the comma and the database name. If any database name contains commas or blanks, enclose the database name in double quotation marks.

The following positional parameter specifies the type of delete backup to perform:

FULL Delete full type backups.

Attention: Be careful to delete only the backups that you want.

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 SQL Server preferences file (tdpsql.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 is the default if no value is specified in the Data Protection for SQL Server preferences file (tdpsql.cfg).

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 SQL Server configuration file that contains the values to use for a **delete backup** operation.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the Data Protection for SQL Server installation directory is used. If the **/configfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is ***tdpsql.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 239 for descriptions of available configuration parameters.

/FROMSQLSERVer=*server-name*

Use the **/fromsqlserver** parameter to specify the name of the SQL Server where the original backup was performed. This parameter is necessary only when the name of the SQL server to delete from, as determined by the **/sqlserver** parameter, is different from the name of the SQL server that the backup objects were created from. The default value is the **/sqlserver** value or the value set in the Data Protection for SQL Server configuration file.

Considerations:

- If the two SQL server names are different, you must use this parameter even if **/fromsqlserver** was a non-clustered default instance.

/LOGFile=*logfilename*

Use the **/logfile** parameter to specify the name of the activity log file that is generated by Data Protection for SQL Server.

The *logfilename* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfilename* variable can include a fully-qualified path. However, if no path is specified, the log file is written to the Data Protection for SQL Serverinstallation directory.

If the *logfilename* variable includes spaces, enclose the entire **/logfile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\mytdpsql.log"
```

If the **/logfile** parameter is not specified, log records are written to the default log file, ***tdpsql.log***.

The **/logfile** parameter cannot be turned off, logging always occurs.

When using multiple simultaneous instances of Data Protection for SQL Serverto 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 Management Console (MMC) 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 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 SQL Server command that might prune the log file, perform one of the following actions to prevent the log file from being pruned:

- Make a copy of the existing log file.
- Specify a new log file with the **/logfile** parameter or **logfile** setting.

/OBJect=*objectname,...*
Use the **/object** parameter to specify the names of backup objects you want to delete. The object name uniquely identifies each backup object and is created by Data Protection for SQL Server.

Use the Data Protection for SQL Server **query tsm * /all** command to view the names of all available backup objects. This parameter specifies that only particular backup objects for the specified SQL databases and backup object type be deleted. The objectname variable specifies the names of the backup objects you want to delete. The object name uniquely identifies each backup object and is created by Data Protection for SQL Server.

/Quiet This parameter prevents status information from being displayed. This does not affect the level of information written to the activity log.

/TSMNODe=*tsmnodename*
Use the *tsmnodename* variable to refer to the Tivoli Storage Manager node name that Data Protection for SQL Server 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 SQL Server options file.

The file name can include a fully qualified path name. If no path is specified, the directory where Data Protection for SQL Server is installed is searched.

If the *tsmoptfilename* variable includes spaces, enclose the entire **/tsmoptfile** parameter entry in double quotation marks. For example:

```
/TSMOPTFile="c:\Program Files\file.opt"
```

The default is **dsm.opt**.

/TSMPassword=*tsmpassword*
Use the *tsmpassword* variable to refer to the Tivoli Storage Manager password that Data Protection for SQL Server uses to log on to the Tivoli Storage Manager server. If you specified PASSWORDACCESS GENERATE

in the Data Protection for SQL Server options file (dsm.opt), you do not need to supply the password here because the one that is stored in the registry is used. However, to store the password in the registry, you must specify the Tivoli Storage Manager password the first time Data Protection for SQL Server 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 SQL Server uses to log on to the Tivoli Storage Manager server can be up to 63 characters in length.

Delete Backup example

This output example provides a sample of the text, messages, and process status that displays when using the **delete backup** command.

In this example, the **tdpsqlc delete backup xivdb1 full** command deletes a full backup of database xivdb1. The following output is displayed:

```
Connecting to SQL Server, please wait...

Querying for Backups ....

Backup(s) to be deleted:
  <xivdb1 : VSS : full : 02/10/2011 10:03:29>
VSS Delete backup operation completed with rc = 0
   Files Examined   : 1
   Files Completed  : 1
   Files Failed     : 0
   Total Bytes      : 0
```

Help command

Use the **`tdpsqlc help`** command to display the syntax of all or selected Data Protection for SQL commands using a textual notation.

The help command uses the following notation:

[*a***]** *a* is optional; *a* may occur zero or one time

{*a* **|** *b***}** select either *a* or *b*, but not both

{*a* **} +** *a* must occur at least one time

{*a* **} *** *a* may occur zero or more times

(*a***)** comments that are not part of the command

UPPERCASE
minimum abbreviation (which you can also enter in lowercase)

Note: When using languages other than English, you might need to set the width of your screen display to a value greater than 80 characters in order to view the entire help description in one screen. 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.

TDPSQLC help command

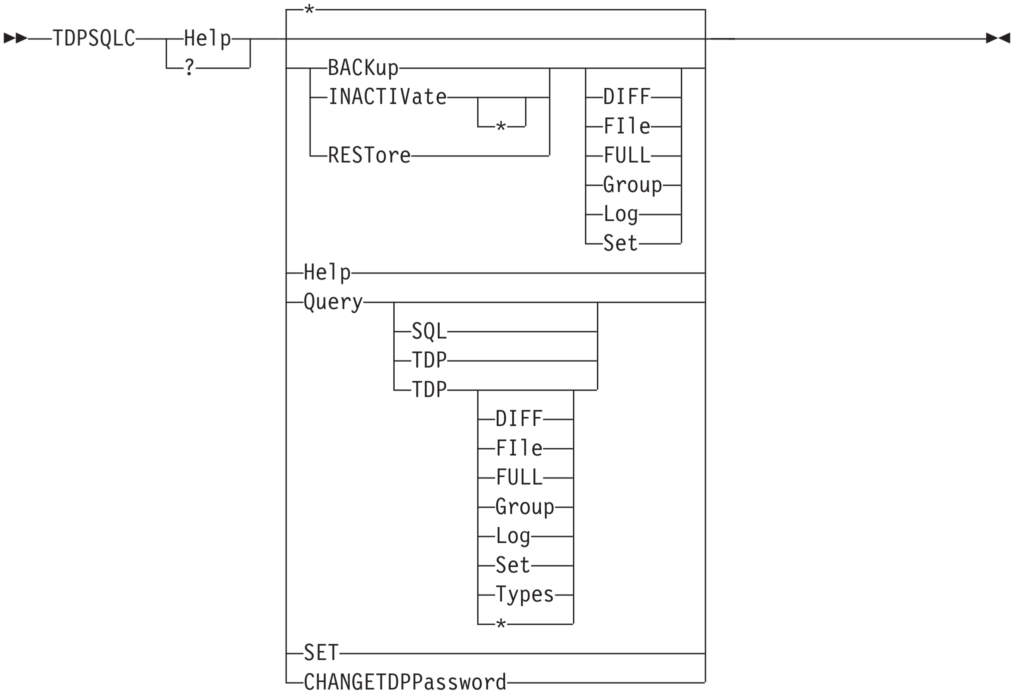

Help positional parameters

Positional parameters immediately follow the **help** command. There are no optional parameters with this command.

Use the help command to display the syntax of all or selected Data Protection for SQL commands using a textual notation.

Help uses the following notation:

[*a***]** *a* is optional; *a* may occur zero or one time

{*a* **|** *b***}** select either *a* or *b*, but not both

{*a* **} +** *a* must occur at least one time

{*a* **} *** *a* may occur zero or more times

(*a***)** comments that are not part of the command

UPPERCASE
minimum abbreviation (which you can also enter in lowercase)

Help output examples

These output examples provide a sample of the text, messages, and process status that displays when using the **help** command.

Help 1-Query TSM

Command:

```
tdpsqlc help query tsm *
```

Output:

```
IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1998, 2013.
All rights reserved.

TDPSQLC Query TSM *|dbname[,dbname,...] [*]
 [/ACtive]
 [/ALl]
 [/BUFFers=numbuffers]              default: 3    (or cfg value)
 [/BUFFERSIze=buffersizeinkb]       default: 1024 (or cfg value)
 [/COMPATibilityinfo]
 [/CONFIGfile=configfilename]       default: tdpsql.cfg
 [/FROMSQLserver=sqlservername]     default: sqlserver value (or cfg value)
 [/LOGFile=logfilename]             default: tdpsql.log (or cfg value)
 [/LOGPrune=numdays|No]             default: 60   (or cfg value)
 [/OBJect=*|objectname[,objectname,...]]
 [/TSMNODe=tsmnodename]             default: dsm.opt value
 [/TSMOPTFile=dsmoptfilename]       default: dsm.opt
 [/TSMPassword=tsmpassword]         default: dsm.opt value
```

Help 2-Restore Full

Command:

```
tdpsqlc help rest full
```

Output:

```
IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1998, 2013.
All rights reserved.

TDPSQLC Restore *|dbname[,dbname,...]  [Full]
 [/BACKUPDESTination=TSM|LOCAL]     default: TSM
 [/BACKUPMETHod=LEGACY|VSS]         default: LEGACY
 [/BUFFers=numbuffers]              default: 3    (or cfg value)
 [/BUFFERSIze=buffersizeinkb]       default: 1024 (or cfg value)
 [/CONFIGfile=configfilename]       default: tdpsql.cfg
 [/DBOonly]
 [/FIles=*|logicalname[,logicalname,...] ]
 [/FROMSQLserver=sqlservername]     default: sqlserver value (or cfg value)
 [/GRoups=*|groupname[,groupname,...] ]
 [/INSTANTRestore=Yes|No]           default: Yes
 [/INTO=dbname]
 [/LOGFile=logfilename]             default: tdpsql.log (or cfg value)
 [/LOGPrune=numdays|No]             default: 60   (or cfg value)
 [/OBJect=*|objectname[,objectname,...] ]
 [/PARTial]
 [/Quiet]
 [/RECOVery=Yes|No]                 default: Yes
 [/RELocate=lname /TO=pname [/RELocate=lname /TO=pname ...] ]
 [/RELOCATEDir=directory[,logfiledirectory[,otherfiledirectory]] ]
 [/REPlace]
 [/SQLAUTHentication=INTegrated|SQLuserid] default: INTegrated (or cfg value)
 [/SQLBUFFers=numsqlbuffers]        default: 0    (or cfg value)
 [/SQLBUFFERSIze=sqlbuffersizeinkb] default: 1024 (or cfg value)
 [/SQLPassword=sqlpasswordname]     default: " "
 [/SQLSERVer=[sqlprotocol:]sqlservername]
                                    default: local computer name (or cfg value)
                                    default sqlprotocol: "" (or cfg value)
 [/SQLUSer=sqlusername]             default: sa
 [/STANDby=undofilename]
 [/STRIPes=numstripes]              default: 1   (or cfg value)
 [/TSMNODe=tsmnodename]             default: dsm.opt value
 [/TSMOPTFile=dsmoptfilename]       default: dsm.opt
 [/TSMPassword=tsmpassword]         default: dsm.opt value
```

Help 3-Restore Log

Command:

```
tdpsqlc help rest log
```

Output:

```
IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1998, 2013.
All rights reserved.

TDPSQLC Restore *|dbname[,dbname,...]  Log=*|logobjectname[,logobjectname,...]
 [/BUFFers=numbuffers]              default: 3    (or cfg value)
 [/BUFFERSIze=buffersizeinkb]       default: 1024 (or cfg value)
 [/CONFIGfile=configfilename]       default: tdpsql.cfg
 [/DBOonly]
 [/FROMSQLserver=sqlservername]     default: sqlserver value (or cfg value)
 [/INTO=dbname]
 [/LOGFile=logfilename]             default: tdpsql.log (or cfg value)
 [/LOGPrune=numdays|No]             default: 60   (or cfg value)
 [/OBJect=*|objectname[,objectname,...] ]
 [/Quiet]
 [/RECOVery=Yes|No]                 default: Yes
 [/RELocate=lname /TO=pname [/RELocate=lname /TO=pname ...] ]
 [/RELOCATEDir=directory[,logfiledirectory[,otherfiledirectory]] ]
 [/SQLAUTHentication=INTegrated|SQLuserid] default: INTegrated (or cfg value)
 [/SQLBUFFers=numsqlbuffers]        default: 0    (or cfg value)
 [/SQLBUFFERSIze=sqlbuffersizeinkb] default: 1024 (or cfg value)
 [/SQLPassword=sqlpasswordname]     default: " "
 [/SQLSERVer=[sqlprotocol:]sqlservername]
                                    default: local computer name (or cfg value)
                                    default sqlprotocol: "" (or cfg value)
 [/SQLUSer=sqlusername]             default: sa
 [/STANDby=undofilename]
 [/STOPAT=datetime]
 [/STOPATMark=markname [/AFTER=datetime] ]
 [/STOPBEFOREMark=markname [/AFTER=datetime] ]
 [/STRIPes=numstripes]              default: 1   (or cfg value)
 [/TSMNODe=tsmnodename]             default: dsm.opt value
 [/TSMOPTFile=dsmoptfilename]       default: dsm.opt
 [/TSMPassword=tsmpassword]         default: dsm.opt value
```

Help 4-Set

Command:

```
tdpsqlc help set
```

Output:

```
IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1998, 2013.
All rights reserved.

TDPSQLC Set PARMname=value
  [/CONFIGfile=configfilename]      default: tdpsql.cfg

  where PARMname and possible values are:
     BACKUPDESTination=[TSM|LOCAL|BOTH]
     BACKUPMETHod=[LEGACY|VSS]
     BUFFers=numbuffers          (2..8)
     BUFFERSIze=buffersize       (64..8192)
     DATEformat=dateformatnum
       1    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
     DIFFESTimate=numpercent     (1..99)
     FROMSQLserver=sqlservername
     LANGUAGE=3-letter country code
       ENU     American English
       PTB     Brazilian Portuguese
       CHS     Chinese, Simplified
       CHT     Chinese, Traditional
       FRA     Standard French
       DEU     Standard German
       ITA     Standard Italian
       JPN     Japanese
       KOR     Korean
       ESP     Standard Spanish
     LOCALDSMAgentnode=nodename
     LOGFile=logfilename
     LOGPrune=[numdays|No]       (0..9999) | No
     NUMBERformat=numberformatnum
       1    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
     SQLAUTHentication=[INTegrated|SQLuserid]
     SQLBUFFers=numsqlbuffers    (0..999)
     SQLBUFFERSIze=sqlbuffersize (64..4096)
     SQLSERVer=[sqlprotocol:]sqlservername
     STRIPes=numstripes          (1..64)
     TIMEformat=timeformatnum
       1    HH:MM:SS
       2    HH,MM,SS
       3    HH.MM.SS
       4    HH:MM:SSA/P
```

Inactivate command (Legacy only)

Use the **inactivate** command to inactivate one or more active Legacy backup objects on the Tivoli Storage Manager server.

Most backup objects are automatically inactivated as part of the normally scheduled backup processing. For those occasions when that processing is not sufficient, you can use the **inactivate** command.

Tivoli Storage Manager server does not delete *active* backup objects from Tivoli Storage Manager managed storage; it will delete only *inactive* backup objects. Once a backup object becomes inactive, the expiration processing defined in the object's management class determines exactly when the backup object is deleted.

Inactivate syntax

Use the **inactivate** command syntax diagrams as a reference to view available options and truncation requirements.

Syntax

TDPSQLC command

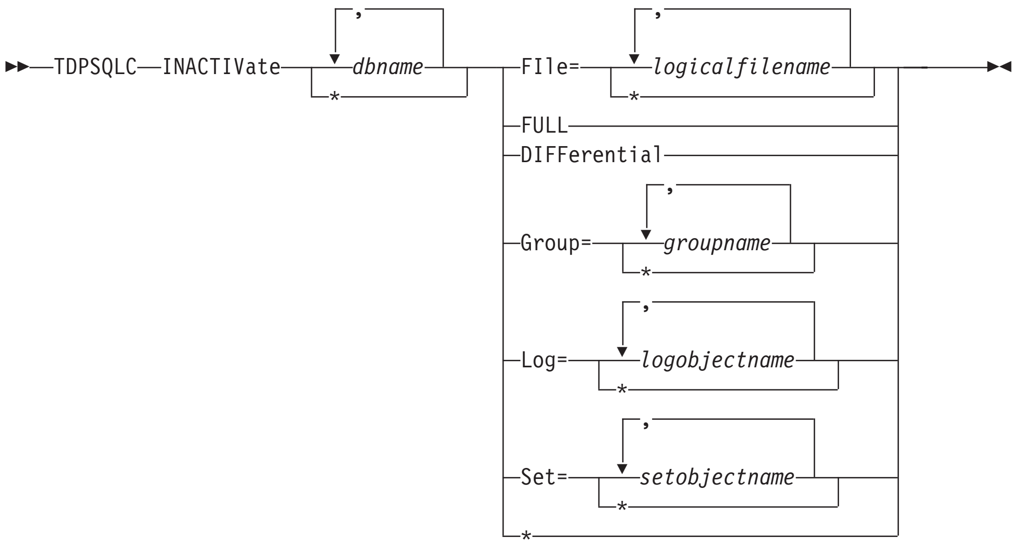

For a description of the **inactivate** positional parameters, see “Inactivate positional parameters” on page 167.

Inactivate Optional Parameters:

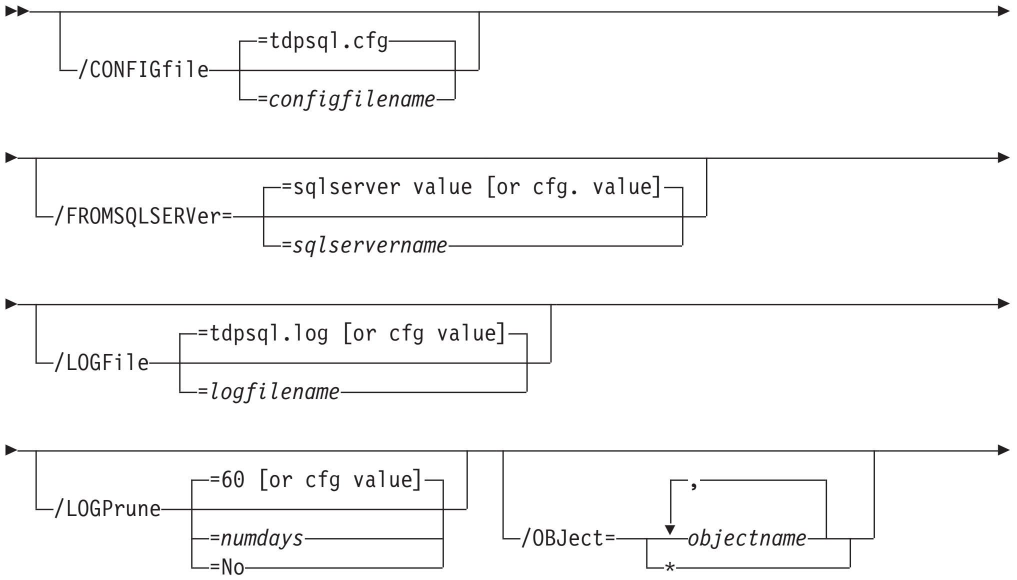

```
►─┬──────────────────────────┬──┬─────────┬──┬──────────────────────────────────────┬──────────►
  └─/OLDerthan=──numdaysold──┘  └─/Quiet──┘  │            ┌─=[dsm.opt value]──┐     │
                                             └─/TSMNODe───┤                   ├─────┘
                                                          └─=tsmnodename──────┘

►─┬───────────────────────────────────────┬──┬──────────────────────────────────────────┬────►◄
  │              ┌─=dsm.opt──────────┐    │  │                ┌─=[dsm.opt value]─────┐  │
  └─/TSMOPTFile──┤                   ├────┘  └─/TSMPassword───┤                      ├──┘
                 └─=dsmoptfilename───┘                        └─=tsmpasswordname─────┘
```

Inactivate positional parameters

Positional parameters immediately follow the **inactivate** command and precede the optional parameters.

FIle=*|*logicalfilename*,...

This option inactivates only the active file backup objects for the SQL databases you specify. The *logicalfilename* variable specifies the names of the SQL server database logical files you want to inactivate.

Considerations:

- You can specify this parameter more than once per command invocation.
- Use * as a wildcard character in *logicalfilename* to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all logical files in the SQL server database.
- If *logicalfilename* includes spaces or special characters, enclose it in double quotes.
- The *logicalfilename* variable is case-sensitive.

FULL This option inactivates only the active full database backup objects for the SQL databases you specify. Each SQL database backed up creates a separate backup object on the Tivoli Storage Manager server. A new full database backup object inactivates all prior active backup objects for the same SQL database. This inactivation includes any active full backup object as well as any active file, group, set, differential, and log backup objects.

DIFFerential

This option inactivates only the active differential database backup object. Because each SQL database backup creates a separate backup object on the Tivoli Storage Manager server, a new differential database backup object inactivates any active differential backup object for the same SQL database. Use this option so that all individual log backups since the last full database backup do not need to be applied.

Group=*|*groupname*,...

This option inactivates only the active group database backup object for the SQL database you specify. The *groupname* variable specifies the names of the SQL server database filegroups you want to inactivate.

Considerations:

- You can specify this parameter more than once per command invocation.
- Use * as a wildcard character in the *groupname* variable to replace zero or more characters for each occurrence.
- Specifying only the wildcard character indicates all filegroups in the SQL server database.
- If the *groupname* variable includes spaces or special characters, enclose it in double quotes.

- The *groupname* variable is case-sensitive.

Log or Log=*|*logobjectname*,...

This option inactivates only the active log database backup object for the SQL database you specify. This parameter takes the wildcard or *logobjectname* value. The *logobjectname* variable specifies the log backup objects to inactivate. Use * as a wildcard character in *logobjectname* to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all log backup objects for the SQL databases. You can specify this parameter more than once per command invocation.

Set or Set=*|*setobjectname*,...

This option inactivates only the active set database backup object for the SQL database you specify. This parameter takes the wildcard or *setobjectname* value. The *setobjectname* variable specifies the set backup objects to inactivate. Use * as a wildcard character in *setobjectname* to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all set backup objects for the SQL databases. You can specify this parameter more than once per command invocation.

Inactivate optional parameters

Optional parameters follow the **inactivate** command and positional parameters.

The following are detailed descriptions of each of the optional parameters:

/CONFIGfile=*configfilename*

The **/configfile** parameter specifies the name of the Data Protection for SQL configuration file, which contains the values for the Data Protection for SQL configurable options. See “Set command” on page 238 for details on the file's contents.

Considerations:

- *configfilename* can include a fully qualified path. If *configfilename* does not include a path, it uses the directory where Data Protection for SQL is installed.
- If *configfilename* includes spaces, place it in double quotes.
- If you do not specify **/configfile**, the default value is ***tdpsql.cfg***.
- If you specify **/configfile** but not *configfilename*, the default value ***tdpsql.cfg*** is used.

/FROMSQLSERVer=*sqlservername*

The **/fromsqlserver** parameter specifies the SQL server that backup objects were backed up from. This parameter is necessary only when the name of the SQL server to inactivate from, as determined by the **/sqlserver** parameter, is different from the name of the SQL server that the backup objects were created from. The default value is the **/sqlserver** value or the value set in the Data Protection for SQL configuration file. If the two SQL server names are different, you must use this parameter even if **/fromsqlserver** was a non-clustered default instance.

/LOGFile=*logfilename*

The **/logfile** parameter specifies the name of the activity log that is generated by Data Protection for SQL. This activity log records significant events such as completed commands and error messages. The Data Protection for SQL activity log is distinct from the SQL Server error log. The *logfilename* variable identifies the name to be used for the activity log generated by Data Protection for SQL.

Considerations:

- If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file.
- The file name can include a fully-qualified path; however, if you specify no path, the file is written to the directory where Data Protection for SQL is installed.
- You cannot turn Data Protection for SQL activity logging off. If you do not specify **/logfile**, log records are written to the default log file. The default log file is ***tdpsql.log***.
- When using multiple simultaneous instances of Data Protection for SQL 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**

The **/logprune** parameter prunes the Data Protection for SQL activity log and specifies how many days of entries are saved. By default, log pruning is enabled and performed once each day Data Protection for SQL is executed; however, this option allows you to disable log pruning or explicitly request a prune of the log for one command run even if the log file has already been 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.

Considerations:

- If you specify *numdays*, it can range from 0 to 9999. A value of 0 deletes all entries in the Data Protection for SQL activity log file except for the current command entries.
- If you specify **no**, the log file is not pruned during this command.
- If you do not specify **/logprune**, the default value is that specified by the logprune configurable option in the Data Protection for SQL configuration file. This is initially 60.
- If you specify **/logprune**, its value is used instead of the value stored in the Data Protection for SQL configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/logprune** without specifying *numdays* or **no**; in this case, the default 60 is used.
- Changes to the value of the **`TIMEformat`** or **`DATEformat`** parameter can result in an undesired pruning of the log file. If you are running a command that may prune the log file and the value of the **`TIMEformat`** or **`DATEformat`** parameter has changed, perform 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.

/OBJect=*|*objectname***,...**

This parameter specifies that only particular backup objects for the specified SQL databases and backup object type (if specified) be inactivated. The *objectname* variable specifies the names of the backup objects you want to inactivate. The object name uniquely identifies each backup object and is created by Data Protection for SQL. Use **query** to view the names of backup objects. You can use * as a wildcard character in *objectname* to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all backup objects of the specified SQL databases and backup object type.

/OLDerthan=*numdaysold*
This parameter specifies how old a backup object must be before the command can inactivate it.

Considerations:

- The *numdaysold* variable can range from 0 to 9999.
- If you specify 0, you inactivate all selected backup objects.
- If you specify 1, you inactivate all selected backup objects created prior to the current date. Any part of a day counts as a whole day.
- There is no default value for **/olderthan**.

/Quiet The **/quiet** parameter omits displaying status information from the command. However, the information is appended to the Data Protection for SQL activity log.

/TSMNODe=*tsmnodename*
The **/tsmnode** parameter specifies the Tivoli Storage Manager node name that Data Protection for SQL uses to log on to the Tivoli Storage Manager server. This identifies which Tivoli Storage Manager client is requesting services. You can also store the node name in the options file. The command line parameter overrides the value in the options file.

Considerations:

- You cannot use the **/tsmnode** parameter if PASSWORDACCESS GENERATE is specified in the Tivoli Storage Manager options file. You must specify the nodename in the options file. Otherwise, you can change PASSWORDACCESS to PROMPT to utilize the **/tsmnode** parameter. For details about the Tivoli Storage Manager options file, see the reference manual *IBM Tivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide*.
- If you do not specify **/tsmnode**, the default value is that specified by the nodename option in the Tivoli Storage Manager options file. Specifying this parameter does not change the value in the options file.

/TSMOPTFile=*dsmoptfilename*
The **/tsmoptfile** parameter specifies the Tivoli Storage Manager options file to use. This is similar to selecting a Tivoli Storage Manager server from the server list in the GUI. The Tivoli Storage Manager options file contains the configuration values for the Tivoli Storage Manager API.

Considerations:

- The *tsmoptfilename* variable can include a fully qualified path. If you do not include a path, the directory where Data Protection for SQL is installed is used.
- If *tsmoptfilename* includes spaces, you must enclose it in double quotes.
- If you do not specify **/tsmoptfile**, the default value is ***dsm.opt***.
- If you specify **/tsmoptfile** but not *tsmoptfilename*, the default is also ***dsm.opt***.

/TSMPassword=*tsmpasswordname*
The **/tsmpassword** parameter specifies the Tivoli Storage Manager password that Data Protection for SQL uses to log on to the Tivoli Storage

Manager server. This parameter and the option PASSWORDACCESS in the Tivoli Storage Manager options file interact in the following ways:

/tsmpassword	PASSWORDACCESS in Tivoli Storage Manager options file	Password already stored in registry?	Result
specified	*generate*	yes	*/tsmpassword* ignored
specified	*generate*	no	*/tsmpassword* used and stored
specified	*prompt*	—	*/tsmpassword* used
not specified	*prompt*	—	user is prompted

Inactivate output examples

These output examples provide a sample of the text, messages, and process status that displays when using the **inactivate** command.

The following operation explicitly inactivates database backup objects. Once a backup object is inactivated, it will expire automatically according to retention policy. In this case, the objects were backed up from a different SQL server. First, a query is performed to display status information such as active state and backup date.

Command:

```
tdpsqlc query tsm DB1_XIVmini_G_BAS,model * /fromsqlserv=STRINGVM1\STRINGVM1
```

Output:

```
IBM Tivoli Storage Manager for Databases
Data Protection for Microsoft SQL Server
Version 6, Release 4, Level 0.0
(C) Copyright IBM Corporation 1998, 2012.
All rights reserved.

Connecting to TSM Server as node 'STRINGVM1_SQL'...

Querying TSM Server for Backups ....

Backup Object Information
-------------------------

SQL Server Name   ........................ STRINGVM1\STRINGVM1
SQL Database Name ........................ DB1_XIVmini_G_BAS
Backup Method     ........................ Lgcy
Backup Location   ........................ Srv
Backup Object Type ....................... Full
Backup Object State ...................... Active
Backup Creation Date / Time .............. 09/23/2011 06:31:04
Backup Size .............................. 3.35 MB
SQL Compressed ........................... No
Backup Compressed ........................ No
Backup Encryption Type ................... None
Backup Client-deduplicated ............... No
Database Object Name ..................... 20110923063104\00001AC4
Number of stripes in backup object ....... 1
Assigned Management Class  ............... DEFAULT


Backup Object Information
-------------------------

SQL Server Name   ........................ STRINGVM1\STRINGVM1
SQL Database Name ........................ DB1_XIVmini_G_BAS
Backup Method     ........................ Lgcy
Backup Location   ........................ Srv
Backup Object Type ....................... Full
Backup Object State ...................... Active
Backup Creation Date / Time .............. 09/20/2011 05:35:14
Backup Size .............................. 3.35 MB
SQL Compressed ........................... No
Backup Compressed ........................ No
Backup Encryption Type ................... None
Backup Client-deduplicated ............... No
Database Object Name ..................... 20110920053514\00001AC4
Number of stripes in backup object ....... 1
Assigned Management Class  ............... DEFAULT


Backup Object Information
-------------------------

SQL Server Name   ........................ STRINGVM1\STRINGVM1
SQL Database Name ........................ DB1_XIVmini_G_BAS
Backup Method     ........................ Lgcy
Backup Location   ........................ Srv
Backup Object Type ....................... Full
Backup Object State ...................... Active
Backup Creation Date / Time .............. 09/19/2011 07:01:39
Backup Size .............................. 3.35 MB
SQL Compressed ........................... No
Backup Compressed ........................ No
Backup Encryption Type ................... None
Backup Client-deduplicated ............... No
Database Object Name ..................... 20110919070139\00001AC4
Number of stripes in backup object ....... 1
Assigned Management Class  ............... DEFAULT
```

```
Backup Object Information
-------------------------

SQL Server Name   ........................ STRINGVM1\STRINGVM1
SQL Database Name ........................ model
Backup Method     ........................ Lgcy
Backup Location   ........................ Srv
Backup Object Type ....................... Full
Backup Object State ...................... Active
Backup Creation Date / Time .............. 09/23/2011 06:31:05
Backup Size .............................. 2.08 MB
SQL Compressed ........................... No
Backup Compressed ........................ No
Backup Encryption Type ................... None
Backup Client-deduplicated ............... No
Database Object Name ..................... 20110923063105\00001AC4
Number of stripes in backup object ....... 1
Assigned Management Class  ............... DEFAULT


Backup Object Information
-------------------------

SQL Server Name   ........................ STRINGVM1\STRINGVM1
SQL Database Name ........................ model
Backup Method     ........................ Lgcy
Backup Location   ........................ Srv
Backup Object Type ....................... Full
Backup Object State ...................... Active
Backup Creation Date / Time .............. 09/19/2011 11:26:15
Backup Size .............................. 2.08 MB
SQL Compressed ........................... No
Backup Compressed ........................ No
Backup Encryption Type ................... None
Backup Client-deduplicated ............... No
Database Object Name ..................... 20110919112615\00001AC4
Number of stripes in backup object ....... 1
Assigned Management Class  ............... DEFAULT


Backup Object Information
-------------------------

SQL Server Name   ........................ STRINGVM1\STRINGVM1
SQL Database Name ........................ model
Backup Method     ........................ Lgcy
Backup Location   ........................ Srv
Backup Object Type ....................... Full
Backup Object State ...................... Active
Backup Creation Date / Time .............. 09/17/2011 01:15:48
Backup Size .............................. 2.08 MB
SQL Compressed ........................... No
Backup Compressed ........................ No
Backup Encryption Type ................... None
Backup Client-deduplicated ............... No
Database Object Name ..................... 20110917011548\00001AC4
Number of stripes in backup object ....... 1
Assigned Management Class  ............... DEFAULT

Completed
```

The user then decides to inactivate all *DB1_XIVmini_G_BAS* database objects older than two days (older than September 23), of which there are two.

Command:

```
tdpsqlc inactivate DB1_XIVmini_G_BAS * /fromsqlserv=STRINGVM1 /olderthan=2
```

Output:

```
IBM Tivoli Storage Manager for Databases
Data Protection for Microsoft SQL Server
Version 6, Release 4, Level 0.0
(C) Copyright IBM Corporation 1998, 2012.
All rights reserved.

Starting Sql database backup inactivation...
Querying Tivoli Storage Manager server for a list of database
 backups,please wait...

Inactivating full backup DB1_XIVmini_G_BAS
Inactivating log backup DB1_XIVmini_G_BAS\20110920053514\00001AC4

Inactivating full backup DB1_XIVmini_G_BAS
Inactivating log backup DB1_XIVmini_G_BAS\20110919070139\00001AC4

Total database backups inspected:                 2
Total database backups requested for inactivation:  2
Total database backups inactivated:               2
Total database skipped:                           0

Elapsed processing time:                          2.18 Secs
```

Another Tivoli Storage Manager query displays the current status of these backup objects using the ***/all*** parameter; a full and a log backup of *test1* are now both inactive.

Command:

```
tdpsqlc query tsm test1 /fromsqlserv=STRINGVM1 /all
```

Output:

```
IBM Tivoli Storage Manager for Databases
Data Protection for Microsoft SQL Server
Version 6, Release 4, Level 0.0
(C) Copyright IBM Corporation 1998, 2012.
All rights reserved.

Backup Object Information
-------------------------

SQL Server Name   ........................ STRINGVM1
SQL Database Name ........................ DB1_XIVmini_G_BAS
Backup Object Type ....................... Log
Backup Object State ...................... Inactive
Backup Creation Date / Time .............. 09/20/2011 05:35:14
Backup Size .............................. 3,349
Database Object Name ..................... 20110920053514\00001AC4
Number of stripes in backup object ....... 1

SQL Server Name   ........................ STRINGVM1
SQL Database Name ........................ DB1_XIVmini_G_BAS
Backup Object Type ....................... Full
Backup Object State ...................... Inactive
Backup Creation Date / Time .............. 09/19/2011 07:01:39
Backup Size .............................. 3,349
Database Object Name ..................... 20110920053514\00001AC4
Number of stripes in backup object ....... 1
```

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.

TDPSQLC command

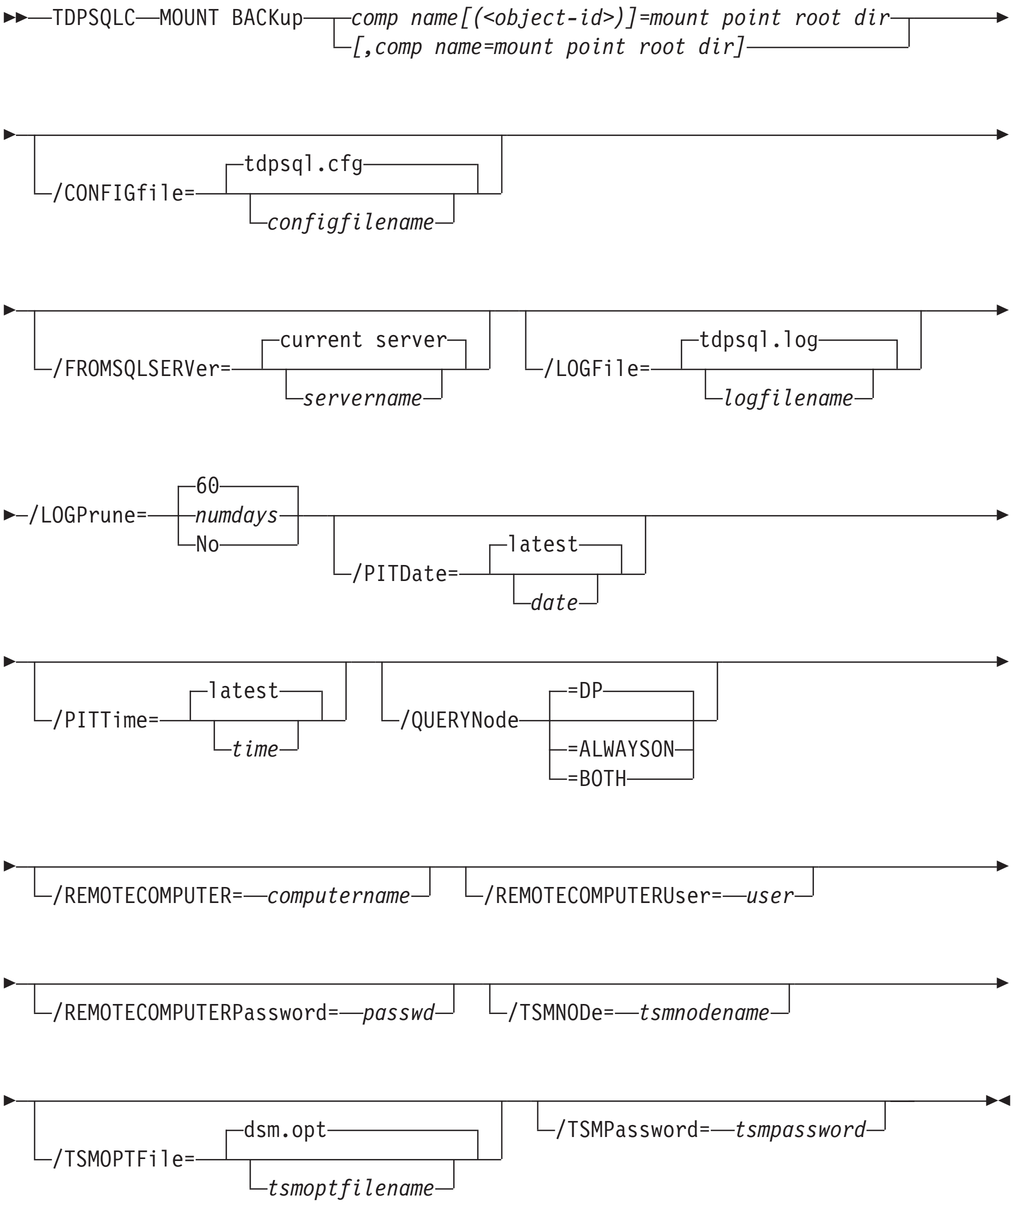

Unmount Backup positional parameter

The positional parameter immediately follows the `unmount backup` command and precedes the optional parameters.

mount points root directory

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 SQL 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 SQL installation directory is used. If the **/configfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is ***tdpsql.cfg***.

If the *configfilename* variable includes spaces, enclose the entire **/configfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\tdpsql.cfg"
```

/FROMSQLSERVer=*server-name*

Use the **/fromsqlserver** parameter to specify the name of the server where the original backup was performed. The default is the local server.

/LOGFile=*logfilename*

Use the **/logfile** parameter to specify the name of the activity log file that is generated by Tivoli Storage FlashCopy Manager. 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 SQL installation directory.

If the *logfilename* variable includes spaces, enclose the entire **/logfile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\tdpsql.log"
```

If the **/logfile** parameter is not specified, log records are written to the default log file, ***tdpsql.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 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 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 to 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 Tivoli Storage FlashCopy Manager for SQL 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.

/PITDAte=*date*

Use the **/pitdate** parameter with the **/pittime** parameter to establish a point in time for which you want to mount the latest version of your backups. Backups that were backed up on or before the date and time you specified, and which were not deleted before the date and time you specified, are processed. Backup versions that you create after this date and time are ignored. Specify the appropriate date in the *date* variable; use the same format that you selected with the **DATEformat** option in the Tivoli Storage FlashCopy Manager for SQL options file.

If neither *date* nor *time* is specified, then no date and time is established. By default the backup is mounted from the most recent available backup.

If either *date* or *time* is specified, then the backup is mounted from the earliest backup taken after the established mount date and time. If no backup after the established date and time is found, by default the backup is mounted from the most recent available backup.

Notes:

- If you specify both *date* and *time*, this establishes the mount backup period.
- If you specify *date* and you do not specify *time*, *time* defaults to a value of 23:59:59. This establishes the *date* at the specified date.
- If you specify *time* without *date*, then *date* defaults to the current date. This establishes the mount date and time as the current date at the specified *time*.

/PITTime=*time*

Use the **/pittime** parameter with the **/pitdate** option to establish a point in time for which you want to mount the latest version of your backups. Files or images that were backed up on or before the date and time you specify, and which were not deleted before the date and time you specify, are processed. Backup versions that you create after this date and time are ignored. This option is ignored if you do not specify the **/pitdate** parameter. Specify the appropriate time in the *time* variable; use the same format that you selected with the TIMEFORMAT option in the Tivoli Storage FlashCopy Manager for SQL options file.

If neither *date* nor *time* is specified, then no date and time is established. By default the backup is mounted from the most recent available backup.

If either *date* or *time* is specified, then the backup is mounted from the earliest backup taken after the established mount date and time. If no backup after the established date and time is found, by default the backup is mounted from the most recent available backup.

Notes:

- If you specify both *date* and *time*, this establishes the mount backup period.
- If you specify *date* and you do not specify *time*, *time* defaults to a value of 23:59:59. This establishes the *date* at the specified date.
- If you specify *time* without *date*, then *date* defaults to the current date. This establishes the mount date and time as the current date at the specified *time*.

/QUERYNode=DP|ALWAYSON|BOTH

Specify whether you want to query standard databases from SQL Server 2012 that were backed up from a standard Data Protection for SQL Server node, the AlwaysOn node, or both nodes. This parameter is ignored for availability databases because the availability databases are always backed up under the AlwaysOn node.

/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 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 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), 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 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 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 Tivoli Storage FlashCopy Manager uses to log on to the Tivoli Storage Manager server can be up to 63 characters in length.

Query command

Use the **query** command to display information about the SQL server and its databases, about the Tivoli Storage Manager server and its backup objects, and about Data Protection for SQL.

Considerations:

- Some of the information displays may have long text lines. You can redirect the informational output of the Data Protection for SQL query command to a text file using the Windows command output redirection syntax (command prompt):

  TDPcommandstatement > [[drive:]path\]filename.ext
  This creates or replaces the file.

  TDPcommandstatement >> [[drive:]path\]filename.ext
  This appends to the file.

  You can then browse or edit the file.
- You can use the Windows **more** filter command (command prompt) to display the informational output one screen at a time, in conjunction with the Windows command pipe character: `TDPcommandstatement | more`

Query syntax

Use the **query** command syntax diagrams as a reference to view available options and truncation requirements.

Syntax

TDPSQLC SQL query

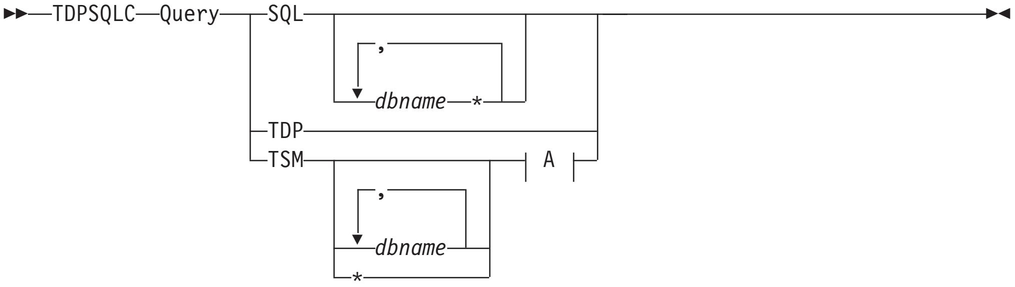

The syntax diagram of the Tivoli Storage Manager options corresponding to the letter *A* is shown following the Optional Parameters below.

Query Optional Parameters:

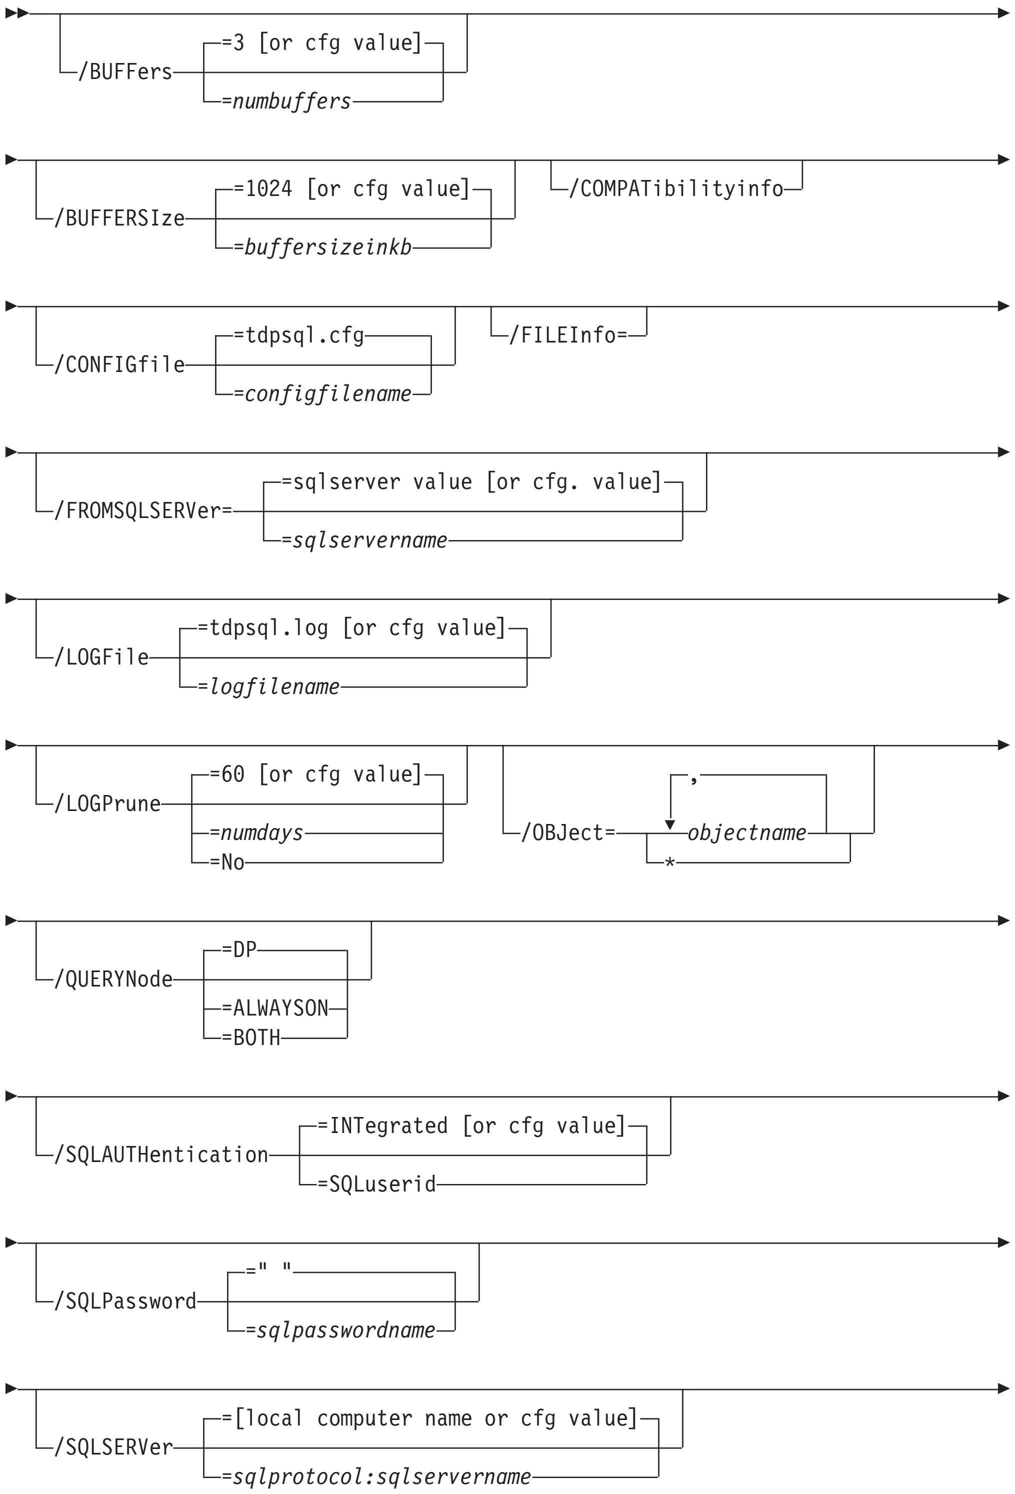

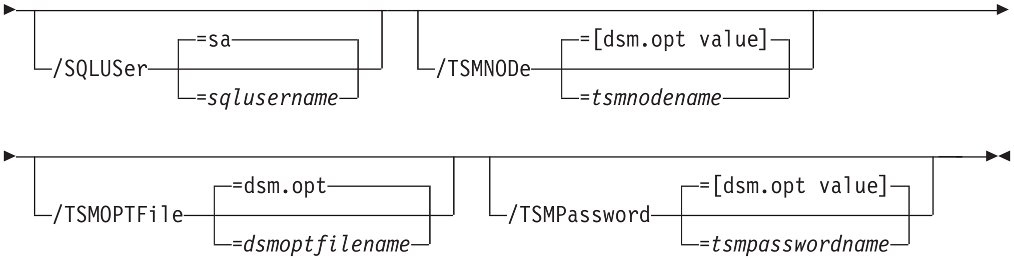

A Query TSM Options:

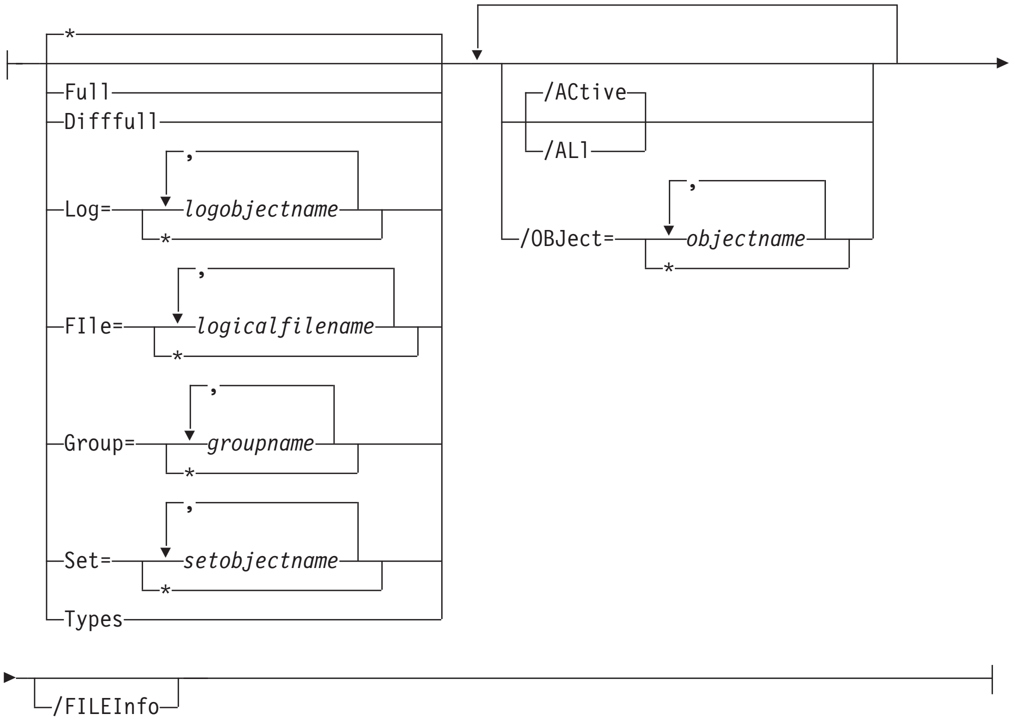

Query positional parameters

Positional parameters immediately follow the **query** command and precede the optional parameters.

Specify one of the following when issuing a Data Protection for SQL **query** command:

Query SQL * | *dbname*,...

This displays information about the current SQL server. The *dbname* variable specifies databases on the current SQL server to display information about.

When querying a SQL Server, the following information is included:

- Server name
- Database name
- Database data space allocated

- Database space used
- Database log space allocated
- Database log space used
- Database options set (SELECT INTO / BULK COPY, TRUNCATE LOG ON CHECKPOINT, and so on)

If you specify ***/COMPATibilityinfo***:

- Server version
- Server clustering state
- Database compatibility level

Query TDP

This displays the Data Protection for SQL name and version information and the contents of the current Data Protection for SQL configuration file.

Query TSM * | ***dbname*****,...**

This displays the Tivoli Storage Manager API and Tivoli Storage Manager server version information. The *dbname* variable names the specified databases from the current SQL server that have backup objects on the current Tivoli Storage Manager server and node. No name is displayed if specified objects do not exist as backup objects in the SQL database. Use the ***dbname,...**** wildcard option to display information about *all* of the backup objects of one or more SQL databases.

When querying any backup object using **TSM** *dbname*, the following information is included:

- SQL server name
- SQL database name
- Backup object type
- Backup object active/inactive state
- Backup object Data Protection for SQL creation date and time
- Backup object Data Protection for SQL size
- Data Protection for SQL backup-object object name
- SQL compressed
- Backup compressed
- Backup encryption type
- Backup deduplicated
- Backup method
- Backup location
- Backup on secondary replica
- Number of data stripes in backup object
- For VSS only, whether the backup supports Instant Restore

The following is included if you specify ***/compatibilityinfo***:

- SQL server version
- SQL Server clustering state
- Data Protection for SQL version that created the backup object
- SQL database compatibility level
- SQL database data space allocated
- SQL database data space used
- SQL database log space allocated

- SQL database log space used
- SQL database options

Note:

- You can also determine which backup objects to display through the **query TSM** optional parameters ***/active*** and ***/all***.
- No information will be displayed if there are no backup objects for a specified SQL database.

FIle=* | *logicalfilename*,...
: This displays information about file backup objects of one or more SQL databases from the current SQL server that are on the current Tivoli Storage Manager server and node.

Full This displays information about full backup objects of one or more SQL databases from the current SQL server that are on the current Tivoli Storage Manager server and node.

Difffull
: This displays information about differential backup objects of one or more SQL databases from the current SQL server that are on the current Tivoli Storage Manager server and node.

Group=* | *groupname*,...
: This displays information about one or more group backup objects of one or more SQL databases from the current SQL server that are on the current Tivoli Storage Manager server and node.

Log=* | *logobjectname*,...
: This displays information about one or more log backup objects of one or more SQL databases from the current SQL server that are on the current Tivoli Storage Manager server and node. The *logobjectname* variable specifies which log backup objects to display information about. Use * as a wildcard character in *logobjectname* to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all log backup objects for the SQL databases.

Set=* | *setobjectname*,...
: This displays information about one or more set backup objects of one or more SQL databases from the current SQL server that are on the current Tivoli Storage Manager server and node. The *setobjectname* variable specifies which set backup objects to display information about. Use * as a wildcard character in *setobjectname* to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all set backup objects for the SQL databases.

Types (Legacy backups only) This displays a summary listed by backup type of the backup objects, of one or more SQL databases from the current SQL server that are on the current Tivoli Storage Manager server and node. Only backup types with one or more backup objects are displayed. If the ***/all*** optional parameter is specified, the number of inactive backup objects is included. You cannot specify either the ***/compatibility*** or the ***/fileinfo*** optional parameter with the ***types*** parameter.

Query optional parameters

Optional parameters follow the **query** command and positional parameters.

The following are detailed descriptions of each of the optional parameters:

/BUFFers=*numbuffers*

The **/buffers** parameter specifies the number of data buffers used for each data stripe to transfer data between Data Protection for SQLand the Tivoli Storage Manager API. The *numbuffers* variable refers to the number of data buffers to use. The number can range from 2 to 8. The default is 3.

Considerations:

- You can improve throughput by increasing the number of buffers, but you will also increase storage use. Each buffer is the size specified in the **/buffersize** parameter.
- The default value is the value specified by the buffers configurable option in the Data Protection for SQL configuration file. This is initially 3.
- If you specify **/buffers**, its value is used instead of the value stored in the Data Protection for SQL configuration file. Specifying this parameter does not change the value in the configuration file.

/BUFFERSIze=*buffersizeinkb*

The **/buffersize** parameter specifies the size of each Data Protection for SQL buffer specified by the **/buffers** parameter. The *buffersizeinkb* variable refers to the size of data buffers in kilobytes. The number can range from 64 to 8192. The default is 1024.

Considerations:

- Though increasing the number of buffers can improve throughput, it also increases storage use as determined by this parameter.
- The default value is the value specified by the buffers configurable option in the Data Protection for SQL configuration file. This is initially 1024.
- If you specify **/buffersize**, its value is used instead of the value stored in the Data Protection for SQL configuration file. Specifying this parameter does not change the value in the configuration file.

/COMPATibilityinfo

For **query** operations, this parameter displays information related to the compatibility of a backup object with a SQL server. Certain SQL Server configuration options must be compatible before you can restore a backup object to a SQL server. When you specify this parameter, SQL and Data Protection for SQL configuration information is listed to help determine if a backup object is correct for a SQL server, or to help in problem determination.

Considerations:

- You cannot specify this parameter with the **types** parameter on a **query TSM** command.
- Compatible generally means identical. However, if you use a binary sort order for both the SQL server and the backup object, the code pages may be different, although the interpretation of individual character values may result in different characters being displayed or printed.

/CONFIGfile=*configfilename*

The **/configfile** parameter specifies the name of the Data Protection for

SQL configuration file, which contains the values for the Data Protection for SQL configurable options. See "Set command" on page 238 for details on the file's contents.

Considerations:

- *configfilename* can include a fully qualified path. If *configfilename* does not include a path, it uses the directory where Data Protection for SQL is installed.
- If *configfilename* includes spaces, place it in double quotes.
- If you do not specify **/configfile**, the default value is ***tdpsql.cfg***.
- If you specify **/configfile** but not *configfilename*, the default value ***tdpsql.cfg*** is used.

/FROMSQLSERVer=*sqlservername*

For **restore**, the **/fromsqlserver** parameter specifies the SQL server that backup objects were backed up from. This parameter is necessary only when the name of the SQL server to restore to, as determined by the **/sqlserver** parameter, is different from the name of the SQL server that the backup objects were created from. The default value is the **/sqlserver** value or the value set in the Data Protection for SQL configuration file.

Considerations:

- If the two SQL server names are different, you must use this parameter even if **/fromsqlserver** was a non-clustered default instance.
- After you restore a SQL database to a different SQL server, the logins of the SQL database may not match the logins for the different SQL server. If appropriate, you can use the SQL stored procedure SP_CHANGE_USERS_LOGIN to find and correct such SQL login mismatches.

/LOGFile=*logfilename*

The **/logfile** parameter specifies the name of the activity log that is generated by Data Protection for SQL. This activity log records significant events such as completed commands and error messages. The Data Protection for SQL activity log is distinct from the SQL Server error log. The *logfilename* variable identifies the name to be used for the activity log generated by Data Protection for SQL.

Considerations:

- If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file.
- The file name can include a fully-qualified path; however, if you specify no path, the file is written to the directory where Data Protection for SQL is installed.
- You cannot turn Data Protection for SQL activity logging off. If you do not specify **/logfile**, log records are written to the default log file. The default log file is ***tdpsql.log***.
- When using multiple simultaneous instances of Data Protection for SQL 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**

The **/logprune** parameter prunes the Data Protection for SQL activity log and specifies how many days of entries are saved. By default, log pruning is enabled and performed once each day Data Protection for SQL is executed; however, this option allows you to disable log pruning or explicitly request a prune of the log for one command run even if the log file has already been 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.

Considerations:

- If you specify *numdays*, it can range from 0 to 9999. A value of 0 deletes all entries in the Data Protection for SQL activity log file except for the current command entries.
- If you specify **no**, the log file is not pruned during this command.
- If you do not specify **/logprune**, the default value is that specified by the logprune configurable option in the Data Protection for SQL configuration file. This is initially 60.
- If you specify **/logprune**, its value is used instead of the value stored in the Data Protection for SQL configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/logprune** without specifying *numdays* or **no**; in this case, the default 60 is used.
- Changes to the value of the **`TIMEformat`** or **`DATEformat`** parameter can result in an undesired pruning of the Data Protection for SQL log file. If you are running a command that may prune the log file and the value of the **`TIMEformat`** or **`DATEformat`** parameter has changed, perform 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.

/OBJect=*|*objectname*,...

For **restore** and **inactivate** operations, **/object** specifies that only particular backup objects for the specified SQL databases and backup object type (if specified) be restored or inactivated. For **query** operations, **/object** includes particular objects and object types in the display. The *objectname* variable specifies the names of the backup objects you want to restore or inactivate. The object name uniquely identifies each backup object and is created by Data Protection for SQL. Use **query** to view the names of backup objects. You can use * as a wildcard character in *objectname* to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all backup objects of the specified SQL databases and backup object type.

`/QUERYNode=DP | ALWAYSON | BOTH`

Specify whether you want to query standard databases from SQL Server 2012 that were backed up from a standard Data Protection for SQL Server node, the AlwaysOn node, or both nodes. This parameter is ignored for availability databases because the availability databases are always backed up under the AlwaysOn node.

/SQLAUTHentication=INTegrated | SQLuserid

This parameter specifies the authorization mode used when logging on to the SQL server. The **integrated** value specifies Windows authentication. The user id you use to log on to Windows is the same id you will use to log on to the SQL server. This is the default value. Use the **sqluserid** value to specify SQL Server user id authorization. The user id specified by the

/sqluserid parameter is the id you will use to log on to the SQL server. Any SQL user id must have the SQL Server SYSADMIN fixed server role.

/SQLPassword=*sqlpasswordname*

This parameter specifies the SQL password that Data Protection for SQL uses to log on to the SQL server that objects are backed up from or restored to.

Considerations:

- Using this parameter means that you are using SQL Server authentication. The SQL Server and the SQL user id for this password must both be configured for SQL Server authentication.
- If you do not specify **/sqlpassword**, the default value is blank (" ").
- If you specify **/sqlpassword** but not *sqlpasswordname*, the default is also blank (" ").
- This parameter is ignored if you use the **/sqlauth=integrated** parameter with it.

/SQLSERVer=*sqlprotocol:sqlservername*

The **/sqlserver** parameter specifies the SQL server that Data Protection for SQL logs on to. Use **/sqlserver** for the **query SQL** command, but use **/fromsqlserver** for the **query TSM** command. The *sqlprotocol* variable specifies the communication protocol to use. You can specify one of the following protocols:

- *lpc*: Use Shared Memory protocol.
- *np*: Use Named Pipes protocol.
- *tcp*: Use Transmission Control protocol.
- *via*: Use Virtual Interface Architecture protocol.

If no protocol is specified, Data Protection for SQL logs on to the SQL server according to the first protocol that becomes available.

Considerations:

- The default value is the value specified by the SQL server configurable option in the Data Protection for SQL configuration file. This is initially the local computer name.
- If you specify **/sqlserver** but not *sqlservername*, the local computer name is used.
- The following two shortcuts are accepted as the local computer name: `.` `(local)` These are a period or the word *local* within parentheses.
- You must specify the name if the SQL server is not the default instance or is a member of a fail-over cluster.
- The format of *sqlservername* depends on what type of instance it is and whether it is clustered or not:

Format	Instance?	Clustered?	Name required?
local-computername	default	no	no
local-computername\ instancename	named	no	yes
virtualservername	default	yes	yes
virtualservername\ instancename	named	yes	yes

localcomputername
The network computer name of the computer the SQL server and Data Protection for SQL reside on. The TCP/IP host name may not always be the same.

instancename
The name given to the named instance of SQL Server specified during installation of the instance.

virtualservername
The name given to the clustered SQL Server specified during clustering service setup. This is not the cluster or node name.

/SQLUSer=*sqlusername*
The **/sqluser** parameter specifies the name that Data Protection for SQL uses to log on to the SQL server.

Considerations:

- Using this parameter means that you are using SQL Server authentication. The SQL Server and the SQL user id for this password must both be configured for SQL Server authentication.
- The SQL user id must have the SQL server SYSADMIN fixed server role.
- If you do not specify **/sqluser**, the default is **sa**.
- If you specify **/sqluser** but not *sqlusername*, the default is also **sa**.
- This parameter is ignored if you use the **/sqlauth=integrated** parameter with it.

/TSMNODe=*tsmnodename*
The **/tsmnode** parameter specifies the Tivoli Storage Manager node name that Data Protection for SQL uses to log on to the Tivoli Storage Manager server. This identifies which Tivoli Storage Manager client is requesting services. You can also store the node name in the options file. The command line parameter overrides the value in the options file.

Considerations:

- You cannot use the **/tsmnode** parameter if PASSWORDACCESS GENERATE is specified in the Tivoli Storage Manager options file. You must specify the nodename in the options file. Otherwise, you can change PASSWORDACCESS to PROMPT to utilize the **/tsmnode** parameter. For details about the Tivoli Storage Manager options file, see the reference manual *IBM Tivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide*.
- If you do not specify **/tsmnode**, the default value is that specified by the nodename option in the Tivoli Storage Manager options file. Specifying this parameter does not change the value in the options file.

/TSMOPTFile=*dsmoptfilename*
The **/tsmoptfile** parameter specifies the Tivoli Storage Manager options file to use. This is similar to selecting a Tivoli Storage Manager server from the server list in the GUI. The Tivoli Storage Manager options file contains the configuration values for the Tivoli Storage Manager API. For details about the Tivoli Storage Manager options file, see the reference manual *IBM Tivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide*.

Considerations:

- The *tsmoptfilename* variable can include a fully qualified path. If you do not include a path, the directory where Data Protection for SQL is installed is used.
- If *tsmoptfilename* includes spaces, you must enclose it in double quotes.
- If you do not specify **/tsmoptfile**, the default value is ***dsm.opt***.
- If you specify **/tsmoptfile** but not *tsmoptfilename*, the default is also ***dsm.opt***.

/TSMPassword=*tsmpasswordname*

The **/tsmpassword** parameter specifies the Tivoli Storage Manager password that Data Protection for SQL uses to log on to the Tivoli Storage Manager server. This parameter and the option PASSWORDACCESS in the Tivoli Storage Manager options file interact in the following ways:

/tsmpassword	PASSWORDACCESS in Tivoli Storage Manager options file	Password already stored in registry?	Result
specified	*generate*	yes	*/tsmpassword* ignored
specified	*generate*	no	*/tsmpassword* used and stored
specified	*prompt*	—	*/tsmpassword* used
not specified	*prompt*	—	user is prompted

Query output examples

These output examples provide a sample of the text, messages, and process status that displays when using the **query** commands.

Query 1–SQL Server

Query 1 queries the SQL server *STRINGVM1*. Note that it is set up for VSS operations.

Command:

```
tdpsqlc query sql
```

Output:

```
IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1997, 2013.
All rights reserved.

Connecting to SQL Server, please wait...

SQL Server Information
----------------------

SQL Server Name   ........................ STRINGVM1\STRINGVM1
SQL Server Version ....................... 10.0.2573 (SQL Server 2008)

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

Writer Name            : SqlServerWriter
Local DSMAgent Node    : STRINGVM1
Remote DSMAgent Node   :
Writer Status          : Online
Selectable Components  : 4

Completed
```

Query 2–SQL Database

Query 2 queries SQL server database, *DB1_XIVmini_G_BAS* and includes compatibility information.

Command:

```
tdpsqlc query sql DB1_XIVmini_G_BAS /compat
```

Output:

```
IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1997, 2013.
All rights reserved.

Connecting to SQL Server, please wait...

SQL Server Information
----------------------

SQL Server Name   ........................ STRINGVM1\STRINGVM1
SQL Server Version ....................... 10.0.2573 (SQL Server 2008)

Cluster ............................. No

SQL Database Information
------------------------

SQL Database Name ........................ DB1_XIVmini_G_BAS
SQL Database Data Space Allocated ........ 3,145,728
SQL Database Data Space Used ............. 1,376,256
SQL Database Log Space Allocated ......... 2,097,152
SQL Database Log Space Used .............. 393,216
SQL Database Compatibility level.......... 100
SQL Database Options .....................

Completed
```

Query 3–TDP (Legacy)

Query 3 queries Data Protection for SQL for configuration file information. Note that this configuration is for Legacy operations only as ***BACKUPDESTination*** *TSM*, ***BACKUPMETHod*** *LEGACY*, and the ***LOCALDSMAgentnode*** and ***REMOTEDSMAgentnode*** are not set.

Command:

```
tdpsqlc query tdp
```

Output:

```
IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1997, 2013.
All rights reserved.

Data Protection for SQL configuration settings
---------------------------------------------------

BACKUPDESTination ......................... TSM
BACKUPMETHod .............................. LEGACY
BUFFers ................................... 3
BUFFERSIze ................................ 1024
DATEformat ................................ 1
DIFFESTimate .............................. 20
FROMSQLserver .............................
LANGuage .................................. ENU
LOCALDSMAgentnode .........................
LOGFile ................................... tdpsql.log
LOGPrune .................................. 60
NUMBERformat .............................. 1
REMOTEDSMAgentnode ........................
SQLAUTHentication ......................... INTegrated
SQLBUFFers ................................ 0
SQLBUFFERSIze ............................. 1024
SQLCOMPression .............................No
SQLSERVer ................................. STRINGVM1
STRIPes ................................... 1
TIMEformat ................................ 1

Completed
```

Query 4 – TDP (VSS)

Query 3 queries Data Protection for SQL for configuration file information. Note that this configuration is set for VSS operations as ***BACKUPDESTination*** *LOCAL*, ***BACKUPMETHod*** *VSS*, and the ***LOCALDSMAgentnode*** and ***REMOTEDSMAgentnode*** options are set.

Command:

```
tdpsqlc query tdp
```

Output:

```
IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1997, 2013.
All rights reserved.

Data Protection for SQL configuration settings
---------------------------------------------------

BACKUPDESTination ........................ LOCAL
BACKUPMETHod ............................. VSS
BUFFers .................................. 3
BUFFERSIze ............................... 1024
DATEformat ............................... 1
DIFFESTimate ............................. 20
FROMSQLserver ............................
LANGuage ................................. ENU
LOCALDSMAgentnode ........................ STRINGVM1
LOGFile .................................. tdpsql.log
LOGPrune ................................. 60
NUMBERformat ............................. 1
REMOTEDSMAgentnode .......................
SQLAUTHentication ........................ INTegrated
SQLBUFFers ............................... 0
SQLBUFFERSIze ............................ 1024
SQLCOMPression ............................No
SQLSERVer ................................ STRINGVM1
STRIPes .................................. 1
TIMEformat ............................... 1

Completed
```

Query 5 – Tivoli Storage Manager Types

Query 5 queries the Tivoli Storage Manager server for the types of backup objects from all databases, including both active and inactive objects.

Command:

```
tdpsqlc query tsm * /all
```

Output:

```
IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1997, 2013.
All rights reserved.

Connecting to TSM Server as node 'STRINGVM1_SQL'...

Querying TSM Server for Backups ....

Backup Object Information
-------------------------

SQL Server Name   ........................ STRINGVM1\STRINGVM1
SQL Database Name ........................ DB1_XIVmini_G_BAS
Backup Method     ........................ VSS
Backup Location   ........................ Srv
Backup Object Type ....................... full
Backup Object State ...................... Inactive
Backup Creation Date / Time .............. 09/23/2013 06:23:14
Backup Size .............................. 5.00 MB
Backup Compressed ........................ No
Backup Encryption Type ................... None
Backup Client-deduplicated ............... No
Backup Supports Instant Restore .......... No
Database Object Name ..................... 20130923062314
Assigned Management Class ................ DEFAULT

Backup Object Information
-------------------------

SQL Server Name   ........................ STRINGVM1\STRINGVM1
SQL Database Name ........................ DB1_XIVmini_G_BAS
Backup Method     ........................ VSS
Backup Location   ........................ Srv
Backup Object Type ....................... full
Backup Object State ...................... Active
Backup Creation Date / Time .............. 09/23/2013 06:39:31
Backup Size .............................. 5.00 MB
Backup Compressed ........................ No
Backup Encryption Type ................... None
Backup Client-deduplicated ............... No
Backup Supports Instant Restore .......... No
Database Object Name ..................... 20130923063931
Assigned Management Class ................ DEFAULT

Backup Object Information
-------------------------

SQL Server Name   ........................ STRINGVM1\STRINGVM1
SQL Database Name ........................ DB1_XIVmini_G_BAS
Backup Method     ........................ VSS
Backup Location   ........................ Loc
Backup Object Type ....................... full
Backup Object State ...................... Inactive
Backup Creation Date / Time .............. 09/23/2013 06:41:14
Backup Size .............................. 5.00 MB
Backup Compressed ........................ No
Backup Encryption Type ................... None
Backup Client-deduplicated ............... No
Backup Supports Instant Restore .......... Yes
Database Object Name ..................... 20130923064114
Assigned Management Class ................ DEFAULT
```

```
Backup Object Information
-------------------------

SQL Server Name   ........................ STRINGVM1\STRINGVM1
SQL Database Name ........................ DB1_XIVmini_G_BAS
Backup Method     ........................ VSS
Backup Location   ........................ Loc
Backup Object Type ....................... full
Backup Object State ...................... Active
Backup Creation Date / Time .............. 09/23/2013 06:45:57
Backup Size .............................. 5.00 MB
Backup Compressed ........................ No
Backup Encryption Type ................... None
Backup Client-deduplicated ............... No
Backup Supports Instant Restore .......... Yes
Database Object Name ..................... 20130923064557
Assigned Management Class ................ DEFAULT

Backup Object Information
-------------------------

SQL Server Name   ........................ STRINGVM1\STRINGVM1
SQL Database Name ........................ DB1_XIVmini_G_BAS
Backup Method     ........................ Lgcy
Backup Location   ........................ Srv
Backup Object Type ....................... Full
Backup Object State ...................... Active
Backup Creation Date / Time .............. 09/23/2013 06:31:04
Backup Size .............................. 2.08 MB
SQL Compressed ........................... No
Backup Compressed ........................ No
Backup Encryption Type ................... None
Backup Client-deduplicated ............... No
Database Object Name ..................... 20130923063104\00001AC4
Number of stripes in backup object ....... 1
Assigned Management Class  ............... DEFAULT

Backup Object Information
-------------------------

SQL Server Name   ........................ STRINGVM1\STRINGVM1
SQL Database Name ........................ model
Backup Method     ........................ VSS
Backup Location   ........................ Srv
Backup Object Type ....................... full
Backup Object State ...................... Inactive
Backup Creation Date / Time .............. 09/23/2013 06:23:14
Backup Size .............................. 3.75 MB
Backup Compressed ........................ No
Backup Encryption Type ................... None
Backup Client-deduplicated ............... No
Backup Supports Instant Restore .......... No
Database Object Name ..................... 20130923062314
Assigned Management Class ................ DEFAULT
```

```
Backup Object Information
-------------------------

SQL Server Name   ........................ STRINGVM1\STRINGVM1
SQL Database Name ........................ model
Backup Method     ........................ VSS
Backup Location   ........................ Srv
Backup Object Type ....................... full
Backup Object State ...................... Active
Backup Creation Date / Time .............. 09/23/2013 06:43:11
Backup Size .............................. 3.75 MB
Backup Compressed ........................ No
Backup Encryption Type ................... None
Backup Client-deduplicated ............... No
Backup Supports Instant Restore .......... No
Database Object Name ..................... 20130923064311
Assigned Management Class ................ DEFAULT

Backup Object Information
-------------------------

SQL Server Name   ........................ STRINGVM1\STRINGVM1
SQL Database Name ........................ model
Backup Method     ........................ VSS
Backup Location   ........................ Loc
Backup Object Type ....................... full
Backup Object State ...................... Active
Backup Creation Date / Time .............. 09/23/2013 06:45:58
Backup Size .............................. 4.00 MB
Backup Compressed ........................ No
Backup Encryption Type ................... None
Backup Client-deduplicated ............... No
Backup Supports Instant Restore .......... No
Database Object Name ..................... 20130923064558
Assigned Management Class ................ DEFAULT

Backup Object Information
-------------------------

SQL Server Name   ........................ STRINGVM1\STRINGVM1
SQL Database Name ........................ model
Backup Method     ........................ Lgcy
Backup Location   ........................ Srv
Backup Object Type ....................... Full
Backup Object State ...................... Active
Backup Creation Date / Time .............. 09/23/2013 06:31:05
Backup Size .............................. 2.08 MB
SQL Compressed ........................... No
Backup Compressed ........................ No
Backup Encryption Type ................... None
Backup Client-deduplicated ............... No
Database Object Name ..................... 20130923063105\00001AC4
Number of stripes in backup object ....... 1
Assigned Management Class  ............... DEFAULT

Completed
```

Query 6–Tivoli Storage Manager Database

Query 6 queries the Tivoli Storage Manager server for database *netapp_db2*, and displays all of its active backup objects by default.

Command:

```
tdpsqlc query tsm model
```

Output:

```
IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1997, 2013.
All rights reserved.

Connecting to TSM Server as node 'STRINGVM1_SQL'...

Querying TSM Server for Backups ....

Backup Object Information
-------------------------

SQL Server Name   ........................ STRINGVM1\STRINGVM1
SQL Database Name ........................ model
Backup Method     ........................ VSS
Backup Location   ........................ Srv
Backup Object Type ....................... full
Backup Object State ...................... Active
Backup Creation Date / Time .............. 09/23/2013 06:43:11
Backup Size .............................. 3.75 MB
Backup Compressed ........................ No
Backup Encryption Type ................... None
Backup Client-deduplicated ............... No
Backup Supports Instant Restore .......... No
Database Object Name ..................... 20130923064311
Assigned Management Class ................ DEFAULT

Backup Object Information
-------------------------

SQL Server Name   ........................ STRINGVM1\STRINGVM1
SQL Database Name ........................ model
Backup Method     ........................ VSS
Backup Location   ........................ Loc
Backup Object Type ....................... full
Backup Object State ...................... Active
Backup Creation Date / Time .............. 09/23/2013 06:45:58
Backup Size .............................. 4.00 MB
Backup Compressed ........................ No
Backup Encryption Type ................... None
Backup Client-deduplicated ............... No
Backup Supports Instant Restore .......... No
Database Object Name ..................... 20130923064558
Assigned Management Class ................ DEFAULT

Backup Object Information
-------------------------

SQL Server Name   ........................ STRINGVM1\STRINGVM1
SQL Database Name ........................ model
Backup Method     ........................ Lgcy
Backup Location   ........................ Srv
Backup Object Type ....................... Full
Backup Object State ...................... Active
Backup Creation Date / Time .............. 09/23/2013 06:31:05
Backup Size .............................. 2.08 MB
SQL Compressed ........................... No
Backup Compressed ........................ No
Backup Encryption Type ................... None
Backup Client-deduplicated ............... No
Database Object Name ..................... 20130923063105\00001AC4
Number of stripes in backup object ....... 1
Assigned Management Class  ............... DEFAULT

Completed
```

Query 7–Tivoli Storage Manager Database

Query 7 queries the Tivoli Storage Manager server for information on database *netapp_db2* Group-type backup objects.

Command:

```
tdpsqlc query tsm netapp_db2 Group=*
```

Output:

```
IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1997, 2013.
All rights reserved.

Connecting to TSM Server as node 'STRINGVM1_SQL'...

Backup Object Information
-------------------------

SQL Server Name   ........................ STRINGVM1\STRINGVM1
SQL Database Name ........................ netapp_db2
Backup Method     ........................ Lgcy
Backup Location   ........................ Srv
Backup Object Type ....................... Group
SQL Group Logical Name ................... PRIMARY
Backup Object State ...................... Active
Backup Creation Date / Time ............. 09/27/2013 08:23:58
Backup Size .............................. 2.08 MB
SQL Compressed ........................... No
Backup Compressed ........................ No
Backup Encryption Type ................... None
Backup Client-deduplicated .............. No
Database Object Name ..................... 20130927082358\00001A4C
Number of stripes in backup object ....... 1
Assigned Management Class  ............... DEFAULT
```

Query 8 –Tivoli Storage Manager Database

Query 8 displays both active and inactive full backup objects of database *Test1*. In addition, file information is requested.

Command:

```
tdpsqlc q tsm DB1_XIVmini_G_BAS full /fileinfo /all
```

Output:

```
IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1997, 2013.
All rights reserved.

Connecting to TSM Server as node 'STRINGVM1_SQL'...

Querying TSM Server for Backups ....

Backup Object Information
-------------------------

SQL Server Name   ........................ STRINGVM1\STRINGVM1
SQL Database Name ........................ DB1_XIVmini_G_BAS
Backup Method     ........................ VSS
Backup Location   ........................ Srv
Backup Object Type ....................... full
Backup Object State ...................... Inactive
Backup Creation Date / Time .............. 09/23/2013 06:23:14
Backup Size .............................. 5.00 MB
Backup Compressed ........................ No
Backup Encryption Type ................... None
Backup Client-deduplicated ............... No
Backup Supports Instant Restore .......... No
Database Object Name ..................... 20130923062314
Assigned Management Class ................ DEFAULT

Backup Object Information
-------------------------

SQL Server Name   ........................ STRINGVM1\STRINGVM1
SQL Database Name ........................ DB1_XIVmini_G_BAS
Backup Method     ........................ VSS
Backup Location   ........................ Srv
Backup Object Type ....................... full
Backup Object State ...................... Active
Backup Creation Date / Time .............. 09/23/2013 06:39:31
Backup Size .............................. 5.00 MB
Backup Compressed ........................ No
Backup Encryption Type ................... None
Backup Client-deduplicated ............... No
Backup Supports Instant Restore .......... No
Database Object Name ..................... 20130923063931
Assigned Management Class ................ DEFAULT

Backup Object Information
-------------------------

SQL Server Name   ........................ STRINGVM1\STRINGVM1
SQL Database Name ........................ DB1_XIVmini_G_BAS
Backup Method     ........................ VSS
Backup Location   ........................ Loc
Backup Object Type ....................... full
Backup Object State ...................... Inactive
Backup Creation Date / Time .............. 09/23/2013 06:41:14
Backup Size .............................. 5.00 MB
Backup Compressed ........................ No
Backup Encryption Type ................... None
Backup Client-deduplicated ............... No
Backup Supports Instant Restore .......... Yes
Database Object Name ..................... 20130923064114
Assigned Management Class ................ DEFAULT
```

```
Backup Object Information
-------------------------

SQL Server Name   ........................ STRINGVM1\STRINGVM1
SQL Database Name ........................ DB1_XIVmini_G_BAS
Backup Method     ........................ VSS
Backup Location   ........................ Loc
Backup Object Type ....................... full
Backup Object State ...................... Active
Backup Creation Date / Time .............. 09/23/2013 06:45:57
Backup Size .............................. 5.00 MB
Backup Compressed ........................ No
Backup Encryption Type ................... None
Backup Client-deduplicated ............... No
Backup Supports Instant Restore .......... Yes
Database Object Name ..................... 20130923064557
Assigned Management Class ................ DEFAULT

Backup Object Information
-------------------------

SQL Server Name   ........................ STRINGVM1\STRINGVM1
SQL Database Name ........................ DB1_XIVmini_G_BAS
Backup Method     ........................ Lgcy
Backup Location   ........................ Srv
Backup Object Type ....................... Full
Backup Object State ...................... Active
Backup Creation Date / Time .............. 09/23/2013 06:31:04
Backup Size .............................. 2.08 MB
SQL Compressed ........................... No
Backup Compressed ........................ No
Backup Encryption Type ................... None
Backup Client-deduplicated ............... No
Database Object Name ..................... 20130923063104\00001AC4
Number of stripes in backup object ....... 1
Assigned Management Class  ............... DEFAULT
SQL Server Version ....................... 10.0.2573 (SQL Server 2008)
Cluster .................................. No
DP Version ............................... 6.4.0.0
SQL Database Compatibility level.......... 100
SQL Database Data Space Allocated ........ 3,145,728
SQL Database Data Space Used ............. 1,376,256
SQL Database Log Space Allocated ......... 2,097,152
SQL Database Log Space Used .............. 344,064
SQL Database Options .....................

SQL Group Logical Name ................... PRIMARY
SQL Group Space Allocated ................ 3,145,728
SQL Group Space Used ..................... 1,376,256
SQL File  Logical Name ................... DB1_XIVmini_G_BAS
SQL File  Physical Name .................. G:\SQLSERVER\DB1_XIVmini_G_BAS\DB1_XIVmini_G_BAS.mdf
SQL File  Space Allocated ................ 3,145,728
SQL File  Space Used ..................... 1,376,256

SQL Group Logical Name ................... TRANSACTION LOG
SQL Group Space Allocated ................ 2,097,152
SQL Group Space Used ..................... 344,064
SQL File  Logical Name ................... DB1_XIVmini_G_BAS_log
SQL File  Physical Name .................. G:\SQLSERVER\DB1_XIVmini_G_BAS\DB1_XIVmini_G_BAS_log.ldf
SQL File  Space Allocated ................ 2,097,152

Completed
```

Query 9 – Tivoli Storage Manager types on AlwaysOn node

Query 9 queries the Tivoli Storage Manager server for the types of backup objects from all standard databases that are backed up to the AlwaysOn node, including both active and inactive objects.

Command:

```
tdpsqlc query TSM * /all /querynode=alwayson
```

Output:

```
C:\Program Files\Tivoli\TSM\TDPSql>tdpsqlc query TSM * /all /querynode=alwayson

IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1997, 2013. All rights reserved.

Connecting to TSM Server as node 'c64'...

Querying TSM Server for Backups ....

Backup Object Information
-------------------------

SQL Server Name   ........................ hkgroup
SQL Database Name ........................ hkaagdb
Backup Method     ........................ VSS
Backup Location   ........................ Loc
Backup Object Type ....................... full
Backup on Secondary Replica ............. No
Backup Object State ...................... Active
Backup Creation Date / Time ............. 06/11/2013 10:18:11
Backup Size .............................. 3.12 GB
Backup Compressed ........................ No
Backup Encryption Type ................... None
Backup Client-deduplicated ............... No
Backup Supports Instant Restore .......... Yes
Database Object Name ..................... 20130611101811
Assigned Management Class ................ DEFAULT
```

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 related information about the volumes that are represented in the local inventory, and are managed by Data Protection for SQL Server. This command is valid for all Windows operating systems that are supported by Data Protection for SQL Server.

TDPSQLC command

```
►►──TDPSQLC──Query MANAGEDCAPacity──┬────────────┬──────────────────────────────►◄
                                    └─/DETAILED──┘
```

Parameters

/DETAILED

Results in a detailed listing of snapped volumes. If this option is not specified then only the total capacity is displayed.

SQL Server 2008 example

Query the total managed capacity of SQL Server 2008 data represented in the local inventory with a detailed listing of snapped volumes:

```
tdpsqlc query managedcapacity /detailed
```

```
Total Managed Capacity : 63.99 GB (68,706,877,440 bytes)

Volume           : H:
Managed Capacity : 16.00 GB (17,176,719,360 bytes)

Volume           : I:
Managed Capacity : 16.00 GB (17,176,719,360 bytes)

Volume           : Q:
Managed Capacity : 16.00 GB (17,176,719,360 bytes)

Volume           : N:
Managed Capacity : 16.00 GB (17,176,719,360 bytes)
```

Query Policy command

Use the **query policy** command to query local policy information.

Query Policy

This command is used to list the attributes of a policy.

```
►►──TDPSQLC──Query POLicy──*──────────────────────────────────────►◄
```

Parameters: * (required) specifies all policies are to be queried. The results of the query will be displayed as follows:

```
                    Connecting to SQL Server, please wait...

Policy      Number of snapshots to keep    Days to keep a snapshot

--------    --------------------------     ----------

SQLPOL                     3                60

STANDARD                   2                30
```

Restore command

Use the **restore** command to restore all or part of one or more SQL databases.

Use this command to restore all or part of one or more SQL databases from Tivoli Storage Manager storage to a SQL server.

- You cannot restore SQL databases currently in use. By placing SQL databases to be restored in single-user mode, you can avoid attempting such restores. If you are restoring the master database, start the SQL server in single-user mode by using the -m SQL SERVER startup option. In addition, the single user of the SQL databases or server must be the same user that Data Protection for SQL uses to log on to the SQL server for the restore. SQL Enterprise Manager, SQL Server Application Client, and other SQL Server services can be users of databases and the SQL server.
- The user used by Data Protection for SQL to log on to the SQL server must have the SQL Server SYSADMIN fixed server role.
- You can use the TRANSACT-SQL database consistency checker statement DBCC CHECKDB ('DBNAME') to verify the integrity of the restored SQL databases.

During SQL database restore processing, the SQL Server prepares the database files after first restoring a minimal amount of metadata. For large SQL databases, the preparation of the database files can be time consuming. To prevent a restore operation from ending prematurely, specify a value of at least *10000* in the `commtimeout` option. If the restore operation is performed in a LAN free environment, this value must be specified for the Storage Agent.

Date and time recovery (Legacy only)

The **restoredate** and **restoretime** parameters allow restore and recovery of the specified database to the date and time specified. These parameters automate the restore of the appropriate full backup, related differential and log backups, and recovers the database to the specified point in time. The behavior when these parameters are used is as follows:

- If only full plus log backups exist, then the following actions occur:
  - The most recent full backup prior to the specified **restoredate** and **restoretime** is restored.
  - All logs up to the first log backed up after the specified **restoredate** and **restoretime** is restored.
  - Recovery up to the specified **restoredate** and **restoretime** (using `stopat`) is completed.
- If only full backups or full plus differential backups exist, then the following actions occur:
  - The most recent full backup prior to the specified **restoredate** and **restoretime** is restored.
  - The most recent differential backup (if any exists) prior to the specified **restoredate** and **restoretime** is restored.
- If full plus differential plus log backups exist, then the following actions occur:
  - The most recent full backup prior to the specified **restoredate** and **restoretime** is restored.
  - The most recent differential backup prior to the specified **restoredate** and **restoretime** is restored.
  - All log backups after the differential and up to the first log backed up after the **restoredate** and **restoretime** is restored.
  - Recovery up to the specified **restoredate** and **restoretime** (using `stopat`) is completed.

VSS restore command-line considerations

Refer to the following considerations when performing 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 instant restore overwrites the entire contents of the source volumes. However, you can avoid overwriting the source volumes by specifying **/instantrestore**=`no`. This parameter setting bypasses volume-level copy and uses file-level copy instead to restore the files from a VSS backup that resides on local shadow volumes. The source volume contains only the SQL database.
- When a VSS restore from local shadow volumes is completed, the bytes transferred is displayed as *0*. This value is displayed because no data (*0*) is restored from the Tivoli Storage Manager server.
- To perform a VSS instant restore with versions of the Tivoli Storage Manager Backup-Archive Client earlier than 6.1.0, the Tivoli Storage Manager for FlashCopy Manager Hardware Devices Snapshot Integration Module must be installed.

- When performing VSS instant restores, you must make sure that any previous background copies that involve the volumes being restored are completed prior to initiating the VSS instant restore.

Restore syntax

Use the **restore** command syntax diagrams as a reference to view available options and truncation requirements.

Syntax

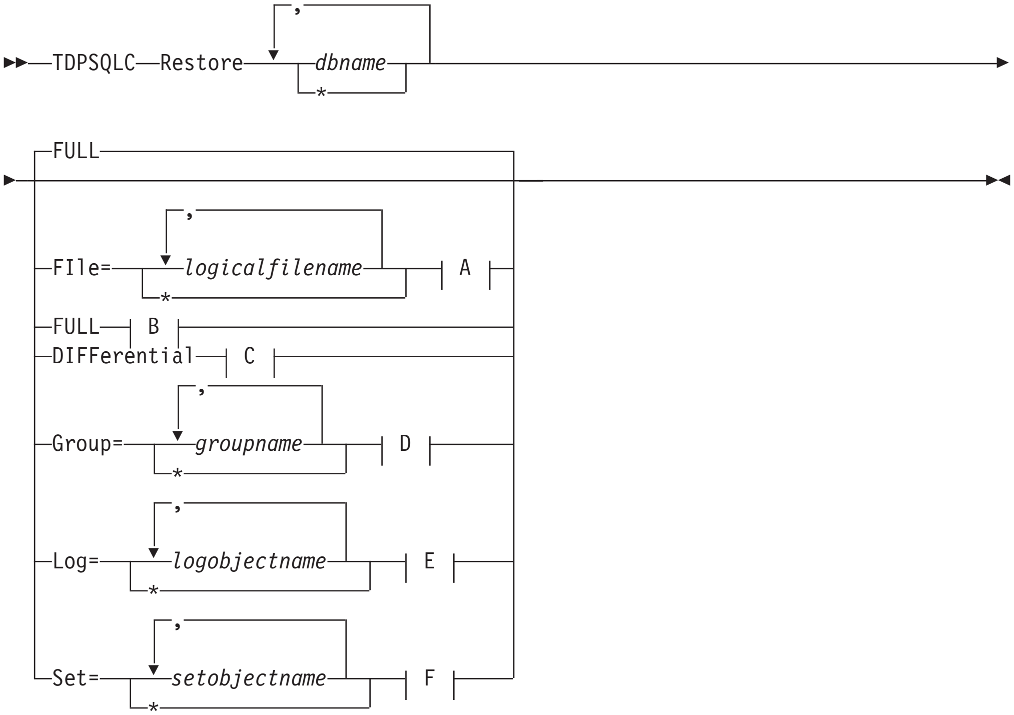

The syntax diagrams of the backup object type options corresponding to the letters *A, B, C, D, E, F* are shown following the Optional Parameters below.

Restore optional parameters:

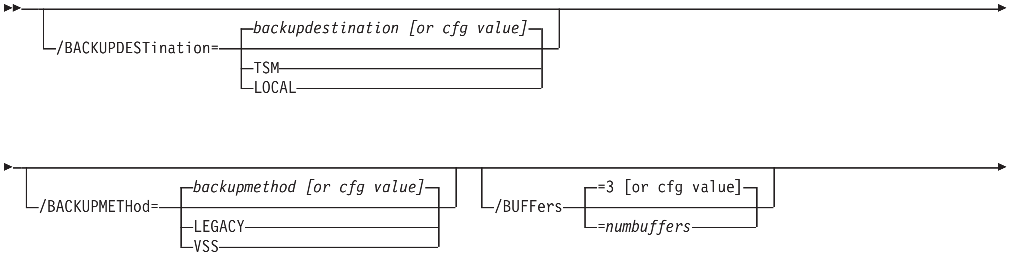

```
►──┬──────────────────────────────────────────┬─┬──────────────────────────────────────┬─┬─────────────┬──►
   │               ┌─=1024 [or cfg value]─┐   │ │              ┌─=tdpsql.cfg──────┐    │ └─/DBOonly──┘
   └─/BUFFERSIze───┤                      ├───┘ └─/CONFIGfile──┤                  ├────┘
                   └─=buffersizeinkb──────┘                    └─=configfilename──┘

►──┬─────────────────────────────────────────────────────────┬─┬─────────────────────────────┬──►
   │                  ┌─=sqlserver value [or cfg. value]─┐   │ │                 ┌─True─┐    │
   └─/FROMSQLSERVer=──┤                                  ├───┘ └─/IncludeTSMVM=──┴──────┴────┘
                      └─=sqlservername───────────────────┘

►──┬──────────────────────────────┬─┬────────────────┬─┬────────────────────────────────────────────┬──►
   │                  ┌─Yes─┐     │ └─/INTO=──dbname─┘ │            ┌─=tdpsql.log [or cfg value]─┐  │
   └─/INSTANTRestore=─┴─No──┴─────┘                    └─/LOGFile───┤                            ├──┘
                                                                    └─=logfilename───────────────┘

►──┬───────────────────────────────────┬─┬───────────────────────────────┬─┬──────────────────────────────┬──►
   │             ┌─=60 [or cfg value]─┐│ │               ┌─,──────────┐  │ │              ┌─=DP────────┐  │
   └─/LOGPrune───┤                    ├┘ │               ▼            │  │ └─/QUERYNode───┤            ├──┘
                 ├─=numdays───────────┤  └─/OBJect=───┬───objectname──┴┬─┘                ├─=ALWAYSON──┤
                 └─=No────────────────┘               └─*──────────────┘                  └─BOTH───────┘

                            ┌─,───────────────┐         ┌─,────────────────┐
                            ▼                 │         ▼                  │
►──┬─────────┬──/RELocate=────logicalfilename─┴──/TO=─────physicalfilename─┴──────────────────►
   └─/Quiet──┘

                                 ┌─,─────────────────────────────────┐ ┌─,───────────────────┐
                                 ▼                 ┌─dbfiledir─┐     │ ▼                     │ ┌─dbfiledir─┐
►─/RELOCATEDir=──dbfiledir──[────logfiledir──[─────┴───────────┴─────┴───otherfiledir──]──]─┴─┴───────────┴─►

►──┬────────────────────────────────────┬─┬────────────────────────────────────┬──►
   │               ┌─current date─┐     │ │               ┌─current time─┐     │
   └─/RESTOREDAte──┴─=date────────┴─────┘ └─/RESTORETime──┴─=time────────┴─────┘

►──┬──────────────────────────────────────────────────────┬─┬───────────────────────────────────────┬──►
   │                     ┌─=INTegrated [or cfg value]─┐   │ │              ┌─=0 [or cfg value]─┐    │
   └─/SQLAUTHentication──┤                            ├───┘ └─/SQLBUFFers──┤                   ├────┘
                         └─=SQLuserid─────────────────┘                    └─=numsqlbuffers────┘

►──┬───────────────────────────────────────────────┬─┬──────────────────────────────────────┬──►
   │                  ┌─=1024 [or cfg value]──┐    │ │               ┌─=" "──────────────┐  │
   └─/SQLBUFFERSIze───┤                       ├────┘ └─/SQLPassword──┤                   ├──┘
                      └─=sqlbuffersizeinkb────┘                      └─=sqlpasswordname──┘

►──┬──────────────────────────────────────────────────────────┬─┬──────────────────────────────┬──►
   │              ┌─=[local computer name or cfg value]─┐     │ │            ┌─=sa───────────┐ │
   └─/SQLSERVer───┤                                     ├─────┘ └─/SQLUSer───┤               ├─┘
                  └─=sqlprotocol:sqlservername──────────┘                    └─=sqlusername──┘

►──┬──────────────────────────────────┬─┬───────────────────────────────────┬─┬──────────────────────────────────┬──►
   │            ┌─=1 [or cfg value]─┐ │ │             ┌─=[dsm.opt value]─┐  │ │               ┌─=dsm.opt──────────┐│
   └─/STRIPes───┤                   ├─┘ └─/TSMNODe────┤                  ├──┘ └─/TSMOPTFile───┤                   ├┘
                └─=numstripes───────┘                 └─=tsmnodename─────┘                    └─=dsmoptfilename───┘

►──┬───────────────────────────────────────────┬────────────────────────────────────────────►◄
   │                ┌─=[dsm.opt value]──┐      │
   └─/TSMPassword───┤                   ├──────┘
                    └─=tsmpasswordname──┘
```

A Restore File Options:

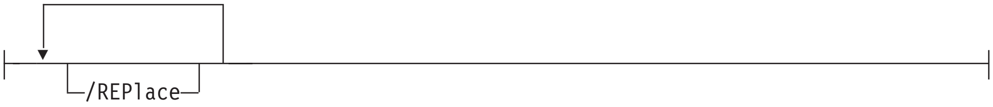

B Restore Full Options:

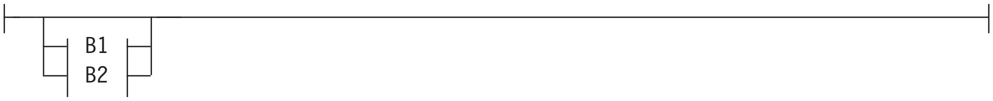

B1 Restore Full Options 1:

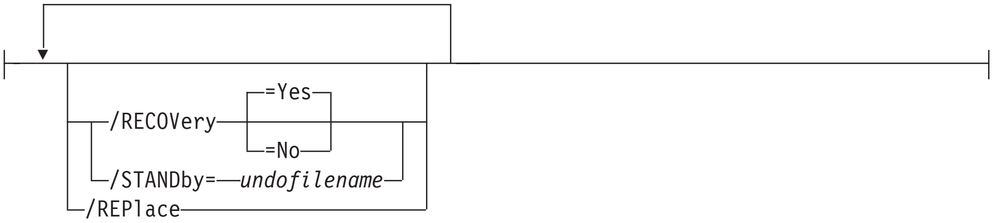

B2 Restore Full Options 2:

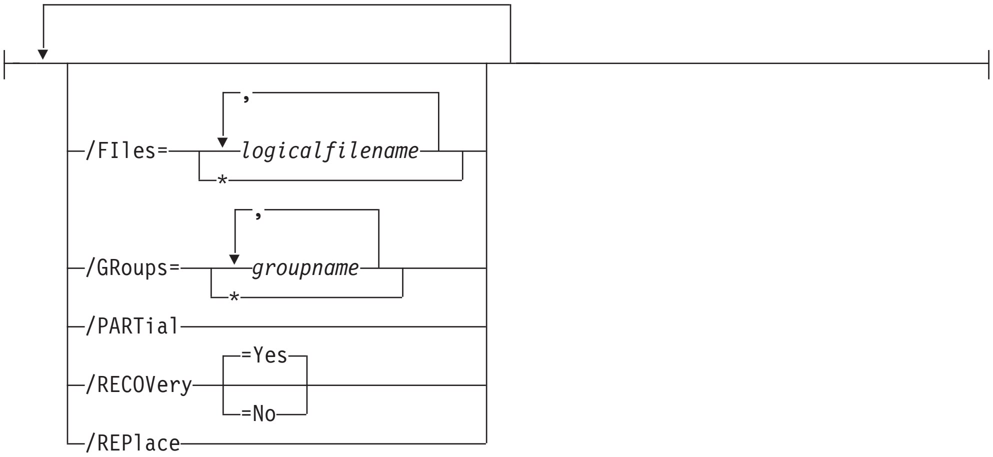

C Restore Diff Options:

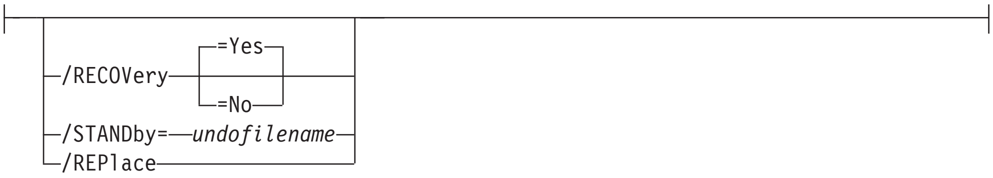

D Restore Group Options:

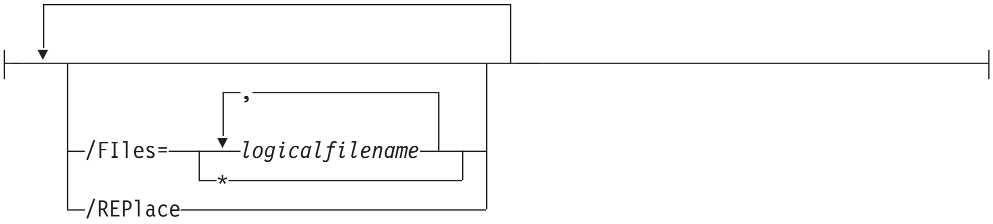

E Restore Log Options:

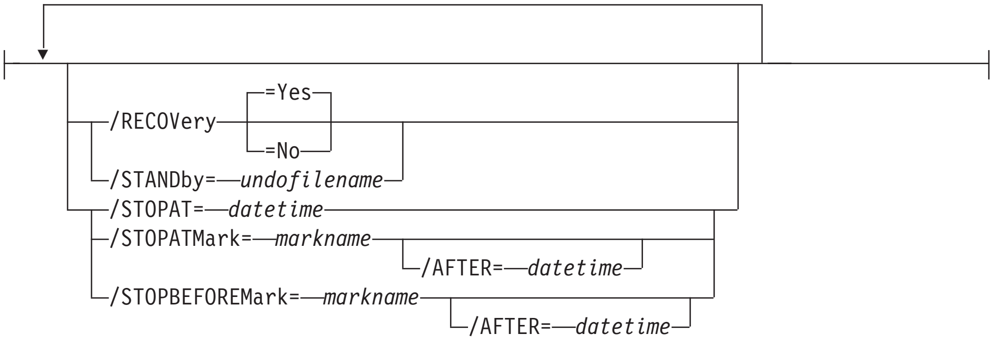

F Restore Set Options:

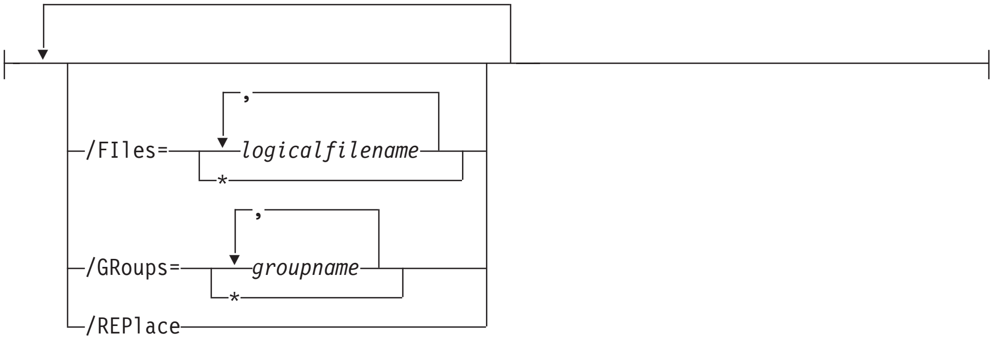

Restore positional parameters

Positional parameters immediately follow the **restore** command and precede the optional parameters.

FIle=* | *logicalfilename***,...**

A **file** backup contains only the contents of the SQL server logical file you specify. You can use this option when it is not practical to back up an entire SQL database due to available backup time and space or due to performance requirements. This option restores file backup objects for the SQL databases you specify. The *logicalfilename* variable specifies the names of the SQL server database logical files you want to restore to.

Considerations:

- You can specify this parameter more than once per command invocation.

- Use * as a wildcard character in *logicalfilename* to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all logical files in the SQL server database. Since each logical file backed up creates a separate backup object on the Tivoli Storage Manager server, specifying only the wildcard character results in a separate backup object for each logical file in the SQL server database.
- If *logicalfilename* includes spaces or special characters, enclose it in double quotes.
- The *logicalfilename* variable is case-sensitive.
- You cannot specify the ***/recovery*** parameter with **restore file** operations.

FULL This option restores all full database backup objects for the SQL databases you specify.

COPYFull
This option restores a copy-only full backup, which contains a copy-only version of a full backup. These backups are considered out of the regular sequence of backups, and do not affect the transaction logs or any sequence of backups like differential backups or full backups.

DIFFerential
A **differential** database backup contains only the parts of a SQL server database changed since the latest full backup plus enough of the SQL database's transaction log to make a restore consistent. As such, a differential backup usually takes up less space than a full backup. Use this option so that all individual log backups since the last full database backup do not need to be applied. This option saves time during a restore by replacing the restore of a number of transaction log backups.

Group=* | *groupname*,...
This option restores all group database backup objects for the SQL databases you specify. The *groupname* variable specifies the names of the SQL server database filegroups you want to restore.

Considerations:

- You can specify this parameter more than once per command invocation.
- Use * as a wildcard character in the *groupname* variable to replace zero or more characters for each occurrence.
- Specifying only the wildcard character indicates all filegroups in the SQL server database.
- If the *groupname* variable includes spaces or special characters, enclose it in double quotes.
- The *groupname* variable is case-sensitive.
- You cannot specify the ***/recovery*** parameter with **restore group** operations.

Log or Log=* | *logobjectname*,...
This option restores all log database backup objects for the SQL databases you specify. The **log** parameter takes the wildcard or *logobjectname* value. The *logobjectname* variable specifies the log backup objects to restore. Use * as a wildcard character in *logobjectname* to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all log backup objects for the SQL databases. You can specify this parameter more than once per command invocation.

Set or Set=* | *setobjectname*,...
This option restores all set database backup objects for the SQL databases

you specify. The **set** parameter takes the wildcard or *setobjectname* value. The *setobjectname* variable specifies the set backup objects to restore. Use * as a wildcard character in *setobjectname* to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all set backup objects for the SQL databases.

Considerations:

- You can specify this parameter more than once per command invocation.
- You cannot specify the ***/recovery*** parameter with **restore set** operations.

Restore optional parameters

Optional parameters follow the **restore** command and positional parameters.

The following are detailed descriptions of each of the optional 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 SQL Server preferences file (`tdpsql.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. This is the default if no value is specified in the Data Protection for SQL Server preferences file (`tdpsql.cfg`).

LOCAL The backup is restored from the local shadow volumes.

/BACKUPMETHod=LEGACY|VSS

Use the **/BACKUPMETHod** parameter to specify the manner in which the restore is performed. The default is the value (if present) specified in the Data Protection for SQL Server preferences file (`tdpsql.cfg`). If no value is present, the backup is restored with the legacy API.

You can specify:

LEGACY The restore is performed with the legacy API. This is the default if no value is specified in the Data Protection for SQL Server preferences file (`tdpsql.cfg`).

VSS The restore is performed with VSS.

/BUFFers=*numbuffers*

The **/BUFFers** parameter specifies the number of data buffers used for each data stripe to transfer data between Data Protection for SQL Server and the Tivoli Storage Manager API. The *numbuffers* variable refers to the number of data buffers to use. The number can range from *2* to *8*. The default is *3*.

Considerations:

- You can improve throughput by increasing the number of buffers, but you will also increase storage use. Each buffer is the size specified in the **/BUFFERSIze** parameter.
- The default value is the value specified by the buffers configurable option in the Data Protection for SQL Server configuration file. This is initially *3*.
- If you specify **/BUFFers**, its value is used instead of the value stored in the Data Protection for SQL Server configuration file. Specifying this parameter does not change the value in the configuration file.

- If you specify **/BUFFers** but not *numbuffers*, the default value *3* is used.

/BUFFERSIze=*buffersizeinkb*

The **/BUFFERSIze** parameter specifies the size of each Data Protection for SQL Server buffer specified by the **/BUFFers** parameter. The *buffersizeinkb* variable refers to the size of data buffers in kilobytes. The number can range from *64* to *8192*. The default is *1024*.

Considerations:

- Though increasing the number of buffers can improve throughput, it also increases storage use as determined by this parameter.
- The default value is the value specified by the buffers configurable option in the Data Protection for SQL Server configuration file. This is initially *1024*.
- If you specify **/BUFFERSIze**, its value is used instead of the value stored in the Data Protection for SQL Server configuration file. Specifying this parameter does not change the value in the configuration file.
- If you specify **/BUFFERSIze** but not *buffersizeinkb*, the default value *1024* is used.

/CONFIGfile=configfilename

The **/CONFIGfile** parameter specifies the name of the Data Protection for SQL Server configuration file, which contains the values for the Data Protection for SQL Server configurable options. See "Set command" on page 238 for details on the contents of the file.

Considerations:

- `configfilename` can include a fully qualified path. If `configfilename` does not include a path, it uses the directory where Data Protection for SQL Server is installed.
- If `configfilename` includes spaces, place it in double quotation marks.
- If you do not specify **/CONFIGfile**, the default value is `tdpsql.cfg`.
- If you specify **/CONFIGfile** but not `configfilename`, the default value `tdpsql.cfg` is used.

/DBOonly

Specifying the **/DBOonly** parameter prevents general users from accessing a restored database before it is determined to be ready for such access. This parameter ensures that the database option RESTRICTED USER is set after a restore operation.

/FROMSQLSERVer=sqlservername

For **restore**, the **/fromsqlserver** parameter specifies the SQL server that backup objects were backed up from. This parameter is necessary only when the name of the SQL server to restore to, as determined by the **/sqlserver** parameter, is different from the name of the SQL server that the backup objects were created from. Use **/fromsqlserver** for **query FCM** commands, but use **/sqlserver** for **query SQL** commands. The default value is the **/sqlserver** value or the value that is set in the Data Protection for SQL Server configuration file. If the two SQL server names are different, you must use this parameter even if **/fromsqlserver** was a non-clustered default instance.

/INSTANTRestore=Yes|No

Use the **/INSTANTRestore** parameter to specify whether to use volume level snapshot or file level copy to restore a VSS backup that is stored on local shadow volumes. An IBM Systems Storage SAN Volume Controller, DS8000, or XIV storage subsystem is required to run VSS instant restores.

You can specify:

Yes
Use volume level snapshot restore for a VSS backup that is stored on local shadow volumes if the backup exists on volumes that support it. This option is the default.

No
Use file-level copy to restore the files from a VSS backup that is stored on local shadow volumes. Bypassing volume-level copy means that SQL database files and log files are the only data overwritten on the source volumes.

When you are running VSS instant restore on DS8000 and Storwize V7000, ensure that any previous background copies that involve the volumes you are restoring, complete before you initiate the VSS instant restore.

/IncludeTSMVM=True

Set the **/IncludeTSMVM** to True to view all backed up databases including Virtual Environment backup SQL databases in the **Databases** view. The backup method is listed as TSMVM to distinguish these databases from the others that are listed.

Alternatively, open the Properties page from the **Actions** pane, and select **Data Center Node** to choose **IncludeTSMVM** to access databases from Virtual Environments in the **Databases** view.

/INSTANTRestore=Yes|No

Use the **/INSTANTRestore** parameter to specify whether to use volume level snapshot or file level copy to restore a VSS backup that is stored on local shadow volumes. An IBM Systems Storage SAN Volume Controller, DS8000, or XIV storage subsystem is required to run VSS instant restores.

You can specify:

Yes
Use volume level snapshot restore for a VSS backup that is stored on local shadow volumes if the backup exists on volumes that support it. This option is the default.

No
Use file-level copy to restore the files from a VSS backup that is stored on local shadow volumes. Bypassing volume-level copy means that SQL database files and log files are the only data overwritten on the source volumes.

When you are running VSS instant restore on DS8000 and Storwize V7000, ensure that any previous background copies that involve the volumes you are restoring, complete before you initiate the VSS instant restore.

/INTO=*dbname*

For **restore** operations, **/INTO** specifies the SQL server database that you want a backup object that is restored into. This parameter is necessary only when the name of the SQL server database to restore into is different from the backup object database name. Considerations:

- When you specify **/INTO**, wildcards (*) might not be used in either the command *dbname* variable or the **/INTO** *dbname* variable.
- There must be exactly one item in the **/INTO** *dbname* variable list in addition to in the command *dbname* list.
- Make sure to use the **/relocatedir** parameter when you specify **/INTO** *dbname*.

/LOGFile=*logfilename*

The **/LOGFile** parameter specifies the name of the activity log that is generated by Data Protection for SQL Server. This activity log records significant events such as completed commands and error messages. The Data Protection for SQL Server activity log is distinct from the SQL Server error log. The **/LOGFile** variable identifies the name to be used for the activity log generated by Data Protection for SQL Server. Considerations:

- If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file.
- The file name can include a fully qualified path; however, if you specify no path, the file is written to the directory where Data Protection for SQL Server is installed.
- You cannot turn off Data Protection for SQL Server logging activity. If you do not specify **/LOGFile**, log records are written to the default log file. The default log file is `tdpsql.log`.
- When you use multiple simultaneous instances of Data Protection for SQL Server for 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 SQL Server activity log and specifies how many days of entries are saved. By default, log pruning is enabled and performed once each day Data Protection for SQL Server is executed; however, this option allows you to disable log pruning or explicitly request a prune of the log for one command run even if the log file has already been 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.

Considerations:

- If you specify *numdays*, it can range from *0* to *9999*. A value of *0* deletes all entries in the Data Protection for SQL Server activity log file except for the current command entries.
- If you specify **no**, the log file is not pruned during this command.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the Data Protection for SQL Server configuration file. This is initially *60*.
- If you specify **/LOGPrune**, its value is used instead of the value stored in the Data Protection for SQL Server configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default, *60*, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an undesired pruning of the log file. If you are running a command that may prune the log file and the value of the **TIMEformat** or **DATEformat** parameter has changed, perform 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.

/MOUNTWait=Yes|No

This parameter is not valid for all backup types; does not work with DIFFFULL or LOG backup types. If the Tivoli Storage Manager server is configured to store backup data on removable media, the Tivoli Storage Manager server might send a message to indicate to Data Protection for SQL Server that the server is waiting for a required storage volume to be mounted. If that occurs, with this option, you can specify whether Data Protection for SQL Server **backup**, **restore**, and **query TSM /fileinfo** commands wait for the media mount or stop the current operation.

You can specify:

Yes Wait for tape mounts (default for **backup** and **restore**).

No Do not wait for tape mounts (default for **query TSM /fileinfo**).

Considerations:

- If you use data striping, Data Protection for SQL Server cannot complete waiting until the initial media for all stripes are available, although Data Protection for SQL Server starts to use each stripe as its media becomes available. Because of the way SQL Server distributes data among stripes, if any stripe does not have its media available, each of the stripes may eventually be either waiting for its own or another stripe's media to become available. In this case, it might become necessary to terminate the Data Protection for SQL Server command from a prolonged wait. This action can be completed by terminating the Data Protection for SQL Server program (close the command prompt window or enter **control-c**).
- For **backup**, if the management class for meta objects also requires removable media, Data Protection for SQL Server waits for that volume, but because meta objects are not created until after the data objects are complete, the wait occurs after all of the data is transferred.
- If you specify no and any removable media are required, Data Protection for SQL Server terminates the command with an error message. This is also true if the management class for meta objects requires removable media. For **backup**, since the meta objects are not created until after the data objects are complete, the command termination does not occur until after all of the database data is transferred.
- If you do not specify **/MOUNTWait** with **backup** or **restore**, the default value is that specified in the **mountwait** configurable option in the Data Protection for SQL Server configuration file. This is initially yes. Specifying this parameter does not change the value in the configuration file.
- If you specify **/MOUNTWait**, but neither yes, nor no, the default, yes, is used.
- If you do not specify **/MOUNTWait** with a **query TSM /fileinfo** request, the default value no is used.

/OBJect=*|objectname,...

For restore and deactivate operations, **/OBJect** specifies that only particular backup objects for the specified SQL databases and backup object type if specified are restored. For query operations, **/OBJect** includes particular objects and object types in the display. The *objectname* variable specifies the names of the backup objects you want to restore or deactivate. The object name uniquely identifies each backup object and is created by Data Protection for SQL Server. Use **query** to view the names of backup objects. Considerations:

- If you do not specify restore, only the active backup object is included in the restore.
- You can use * as a wildcard character in `objectname` to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all backup objects of the specified SQL databases and backup object type.

/PARTial

The **/PARTial** parameter restores only part of a SQL database. You can perform partial restores only on `full` database backup objects. The primary purpose of a partial restore is to retrieve lost or damaged data. A partial restore creates a subset of the SQL database. After the partial restore, differential database restores and transaction log restores can return the subset SQL database to a point where the required data exists or is undamaged. You can then copy the required data from the subset SQL database to the original SQL database. You can also use partial restores whenever you need a subset of a SQL database, such as for development or reporting purposes.

A partial restore always restores the entire backup object from the Tivoli Storage Manager server although only a portion of the restored object may be used to complete a recovery. The statistics displayed reflect the amount of data restored from the Tivoli Storage Manager server only, not the amount of data used by the SQL Server for database recovery.

Considerations:

- You can specify the content of a partial restore with the `files` or `groups` parameters.
  - You can restore only complete SQL groups, even if you did not specify all SQL files in a SQL group with the `files` option.
  - The primary group is always included.
  - SQL groups not restored are marked offline and are not accessible.
- If you are restoring the subset SQL database to a location where it was backed up, you must use the **/RELocate** and **/to** parameters.
- The Management Console does not support the **/RELocate** and **/to** parameters. You must use the command line interface when performing a partial restore that requires these parameters.
- You can specify the **/RECOVery** parameter with **/PARTial**.

/QUERYNode=DP | ALWAYSON | BOTH

Specify whether you want to query standard databases from SQL Server 2012 that were backed up from a standard Data Protection for SQL Server node, the AlwaysOn node, or both nodes. This parameter is ignored for availability databases because the availability databases are always backed up under the AlwaysOn node.

/Quiet The **/Quiet** parameter omits displaying status information from the command. However, the information is appended to the activity log.

/RECOVery=Yes|No

For **restore** operations, **/RECOVery** specifies whether or not you want to make additional restores to a SQL database that is not on a standby SQL server. A restored database cannot be used until the **/RECOVery**=Yes parameter is administered to the database. You can specify:

Yes (default)

Whenever you make a sequence of restores to a SQL database and the current restore is the final restore in the sequence, or is the only restore to a SQL database. This informs the SQL server the restore is complete and ready for uncompleted transactions to be rolled back.

No Whenever you make a sequence of restores to a SQL database and the current restore is not the final restore in the sequence. Issue **/RECOVery**=no for all **restore** commands except the last one.

Considerations:

- After the **/RECOVery**=yes parameter is administered, you cannot restore any more differential or log backups to the database.
- You cannot specify **/RECOVery** for restore operations of **file**, **group**, or **set** backup objects. Data Protection for SQL Server forces such restores to **/RECOVery**=no.
- For full restores that specify **/groups** or **/files**, unless you also specify **/partial**, you cannot specify **/RECOVery**. Without **/partial**, Data Protection for SQL Server forces such restores to **/RECOVery**=no.
- Not specifying this option automatically rolls back incomplete transactions for the database.
- When you specify yes and you are restoring several restore objects for the same database, only the final restore object for the database uses **/RECOVery**=yes; all others use **/RECOVery**=no. This allows you to specify a list of logs without having to specify the final log in a separate command.

The following is a sample scenario:

Sequence of Restores	Specify
full database	no
differential database	no
transaction log backup object	no
transaction log backup object	yes

1. Data Protection for SQL Server sorts the restore objects by database name, and, within database name, by backup time stamp from earliest to latest. A **query TSM** command also displays this order.
2. If a restore object fails, then all subsequent restore objects for that database in a single restore command are skipped. This is true no matter what the **/RECOVery** or **/STANDby** settings are.

/RELocate=*logicalfilename***,... /TO=***physicalfilename***,...**

For **restore** operations, the **/RELocate** and **/TO** parameters as a pair specify the new location of a SQL database file. You must use this parameter for every SQL database file that you are not restoring to its original drive, complete path, and file name. The *logicalfilename* variable specifies the logical file name of the SQL database file you want to relocate. The *physicalfilename* variable specifies the new physical Windows file name where you want to relocate the SQL database file. This parameter is available when restoring legacy backups only.

Considerations:

- You cannot specify more than one database name as the value for the restore command when specifying **/RELocate**.

- **/RELocate** and **/TO** can each take a list of values and can be specified more than once. However, as a pair, **/RELocate** and **/TO** must take the same number of values, and the values must be paired in order of appearance. For example,

  ```
  /relocate=a,b,c /to=a¹,b¹,c¹
  ```

  is valid, but not

  ```
  /relocate=a,b,c /to=b¹,a¹
  ```
- The MMC GUI does not support the **/RELocate** and **/TO** parameters. You must use the command line interface when performing a partial restore that requires these parameters.
- You can use the **query** command with the **/fileinfo** parameter to determine the logical file names and physical file names in the backup object.
- If either *logicalfilename* or *physicalfilename* includes spaces, you must enclose it in double quotes.
- For *physicalfilename*, include the complete drive, path, and file name of the new file.
- The drive and path of the new physical file name must exist, but if the file does not yet exist, SQL Server creates it. Additionally, if the file does exist, you may be required to use the **/replace** parameter.
- The wildcard (*) is not allowed in the values for either **/RELocate** or **/TO**.

/RELOCATEDir=*dbfiledir*[*,logfiledir* [*,otherfiledir*]]

The **/RELOCATEDir** parameter specifies the new destination locations in which to restore the backed up SQL databases, logs, and SQL Server full-text index files. FILESTREAM files are included for SQL Server 2008 and SQL Server 2008 R2. This parameter is available when restoring VSS backups or legacy backups.

The *dbfiledir* variable specifies the directory location of the SQL database you want to relocate. Note that if the *logfiledir* and *otherfiledir* variables are not specified, the logs and SQL Server full-text index files are restored to the directory specified by *dbfiledir*.

The *logfiledir* variable specifies the directory location of the SQL log files you want to relocate. Note that if the *logfiledir* variable is not specified, the SQL log files are restored to the directory specified by *dbfiledir*.

The *otherfiledir* variable specifies the directory location of the SQL Server full-text index files and FILESTREAM files (SQL Server 2008 and SQL Server 2008 R2) are included for you want to relocate. Note that if the *otherfiledir* variable is not specified, the SQL Server full-text index files and FILESTREAM files (SQL Server 2008 and SQL Server 2008 R2) are restored to the directory specified by *dbfiledir*.

/REPlace

For **restore** operations, the **/REPlace** parameter specifies that you want existing SQL files to be overwritten when they otherwise would not be. You may have to use this parameter in the following instances:

- You are performing a `full` database restore, and one of the following is true:
  - You are using the **/into** parameter, and the **/into** database already exists on the SQL server.
  - The database already exists on the SQL server, and one of the following is also true:

  - The number of SQL files in the existing database differs from the number of SQL files in the full database backup object.
  - The names of one or more SQL files in the existing database are not the names of any of the SQL files in the full database backup object.
- You are performing a `file`, `group`, or `set` restore, and one or more of the SQL files already exist.

/RESTOREDAte=*date*

The **/RESTOREDAte** parameter specifies a date to which the database identified by *dbname* is to be recovered. The date value must be specified in the same date format defined in the Data Protection for SQL Server preferences file. If **/RESTOREDAte** is not specified but **/RESTORETime** is specified, the **/RESTOREDAte** value is the current date. The **/RESTOREDAte** parameter is only available for legacy restore operations. It can only be specified when restoring a full database backup. The **/RESTORETime** parameter cannot be used to restore file, group, and set backups.

/RESTORETime=*time*

The **/RESTORETime** parameter specifies the time of day to which the database identified by *dbname* is to be recovered. The time value must be specified in the same time format defined in the Data Protection for SQL Server preferences file. If **/RESTORETime** is not specified but **/RESTOREDAte** is specified, the **/RESTORETime** is the current time. The **/RESTORETime** parameter is only available for legacy restore operations. It can only be specified when restoring a full database backup. The **/RESTORETime** parameter cannot be used to restore file, group, and set backups.

/SQLAUTHentication=INTegrated|SQLuserid

This parameter specifies the authorization mode used when logging on to the SQL server. The `INTegrated` value specifies Windows authentication. The user id you use to log on to Windows is the same id you will use to log on to the SQL server. This is the default value. Use the **sqluserid** value to specify SQL Server user id authorization. The user id specified by the **/sqluserid** parameter is the id you will use to log on to the SQL server. Any SQL user id must have the SQL Server SYSADMIN fixed server role.

/SQLBUFFers=*numsqlbuffers*

The **/SQLBUFFers** parameter specifies the total number of data buffers SQL Server uses to transfer data between SQL Server and Data Protection for SQL Server. The *numsqlbuffers* variable refers to the number of data buffers to use. The number can range from *0* to *999*. The initial value is *0*. When **/SQLBUFFers** is set to *0*, SQL determines how many buffers should be used.

Considerations:

- The default value is the value specified by the SQL buffers configurable option in the Data Protection for SQL Server configuration file. This is initially *0*.
- If you specify **/SQLBUFFers**, its value is used instead of the value stored in the Data Protection for SQL Server configuration file. Specifying this parameter does not change the value in the configuration file.
- If you specify **/SQLBUFFers** but not *numsqlbuffers*, the default value *0* is used.

/SQLBUFFERSIze=*sqlbuffersizeinkb*

The **/SQLBUFFERSIze** parameter specifies the size of each buffer (specified by the **/SQLBUFFers** parameter) SQL Server uses to transfer data to Data

Protection for SQL Server. The *sqlbuffersizeinkb* variable refers to the size of data buffers in kilobytes. The number can range from *64* to *4096*. The default is *1024*.

Considerations:

- The default value is the value specified by the SQL buffers configurable option in the Data Protection for SQL Server configuration file. This is initially *1024*.
- If you specify **/SQLBUFFERSIze**, its value is used instead of the value stored in the Data Protection for SQL Server configuration file. Specifying this parameter does not change the value in the configuration file.
- If you specify **/SQLBUFFERSIze**, but not *sqlbuffersizeinkb*, the default value *1024* is used.

/SQLPassword=*sqlpasswordname*

This parameter specifies the SQL password that Data Protection for SQL Server uses to log on to the SQL server that objects are backed up from or restored to.

Considerations:

- Using this parameter means that you are using SQL Server authentication. The SQL Server and the SQL user id for this password must both be configured for SQL Server authentication.
- If you do not specify **/SQLPassword**, the default value is blank (" ").
- If you specify **/SQLPassword** but not *sqlpasswordname*, the default is also blank (" ").
- This parameter is ignored if you use the `/sqlauth=integrated` parameter with it.

/SQLSERVer=*sqlprotocol:sqlservername*

The **/SQLSERVer** parameter specifies the SQL server that Data Protection for SQL Server logs on to. For **restore** operations, this is the SQL server that backup objects are restored to. However, if the backup objects were created from a different SQL server name, you must use the **/fromsqlserver** parameter. Use **/sqlserver** for the **query SQL** and **backup** commands, but use **/fromsqlserver** for the **query TSM** and **inactivate** commands. The *sqlprotocol* variable specifies the communication protocol to use. You can specify one of the following protocols:

- *lpc*: Use Shared Memory protocol.
- *np*: Use Named Pipes protocol.
- *tcp*: Use Transmission Control protocol.
- *via*: Use Virtual Interface Architecture protocol.

If no protocol is specified, Data Protection for SQL Server logs on to the SQL server according to the first protocol that becomes available.

Considerations:

- The default value is the value specified by the SQL server configurable option in the Data Protection for SQL Server configuration file. This is initially the local computer name.
- If you specify **/SQLSERVer** but not *sqlservername*, the local computer name is used.
- The following two shortcuts are accepted as the local computer name: `.` `(local)` These are a period or the word *local* within parentheses.

- You must specify the name if the SQL server is not the default instance or is a member of a failover cluster.
- The format of *sqlservername* depends on what type of instance it is and whether it is clustered or not:

Format	Instance?	Clustered?	Name required?
local-computername	default	no	no
local-computername\ instancename	named	no	yes
virtualservername	default	yes	yes
virtualservername\ instancename	named	yes	yes

local-computername
: The network computer name of the computer the SQL server and Data Protection for SQL Server reside on. The TCP/IP host name might not always be the same.

instancename
: The name given to the named instance of SQL Server specified during installation of the instance.

virtualservername
: The name given to the clustered SQL Server specified during clustering service setup. This is not the cluster or node name.

/SQLUSer=*sqlusername*

The **/SQLUSer** parameter specifies the name that Data Protection for SQL Server uses to log on to the SQL server.

Considerations:

- Using this parameter means that you are using SQL Server authentication. The SQL Server and the SQL user id for this password must both be configured for SQL Server authentication.
- The SQL user id must have the SQL server SYSADMIN fixed server role.
- If you do not specify **/SQLUSer**, the default is *sa*.
- If you specify **/SQLUSer**, but not *sqlusername*, the default is also *sa*.
- This parameter is ignored if you use the `/sqlauth=integrated` parameter with it.

/STANDby=*undofilename*

Specifies that the restore is to a standby SQL server, and specifies the name of an undo file.

Considerations:

- You cannot specify more than one database name as the restore command value.
- A standby SQL server can be in read-only mode between restores and can accept additional restores to its databases.
- You can use the same undo file for a database for each restore to the database, but you cannot use a single undo file for more than one database.
- The *undofilename* variable can include a fully qualified path. However, if a fully qualified path is not specified, the undo file is created in the directory specified by the `%TEMP%` environment variable.

- If *undofilename* includes spaces, you must enclose it in double quotes.
- If the specified undo file does not exist, SQL server creates it. If the file exists but was not used for the same SQL database, SQL Server overwrites it.
- If you specify neither **/recovery** nor **/STANDby**, the default is **/recovery**=yes.

/STOPAT=*datetime*

For **restore** operations, **/STOPAT** specifies the point in time that you restore a SQL database to. Only transaction logs written before the point in time are applied to the SQL database. The *datetime* variable specifies both the date and time separated by a space. Use any valid date and time format accepted by SQL Server.

Considerations:

- This parameter applies only to transaction log restores, but the base restore that the transaction logs apply to must have been a full database restore. You cannot restore file, group, and set restores to a point in time.
- You cannot also specify **/recovery**=no or **/standby** with the **/STOPAT**parameter.
- Because *datetime* includes a space, you must enclose it in double quotes.
- If the restore operation with the **/STOPAT** parameter does not encounter a transaction in the restored transaction log that has a time stamp equal to or greater than the specified point in time, the SQL database is left in an unrecovered state, even if you also specify **/recovery**=yes.

/STOPATMark=*markname* **[/AFTER=***datetime***]**

The **/STOPATMark** parameter specifies a named point in time to restore a database to. This can be after a specified point in time if you specify the **/AFTER** option. Only transaction log records written up to and including the named transaction (which may be found at or after the specified point in time) are applied to the SQL database. The *markname* variable specifies the name of a SQL transaction. The SQL transaction may be a local transaction or a distributed transaction. If it is a distributed transaction name, the named mark exists in the transaction log of each SQL database participating in the distributed transaction.

markname is the transaction name, not the description that follows the MARK keyword in a SQL BEGIN TRANSACTION or BEGIN DISTRIBUTED TRANSACTION statement.The *datetime* variable specifies both the date and time separated by a space. Use any valid date and time format accepted by SQL Server.

Considerations:

- This parameter applies only to transaction log restores. The base restore that the transaction logs apply to must have been a full database restore. You cannot restore file, group, and set restores to a mark.
- You can use the same named mark for several SQL transactions.
- If you do not specify **/AFTER**, the restore stops at the first mark it encounters with the specified name.
- If you specify **/AFTER**, the restore stops at the first mark it encounters with the specified name after the specified date and time.
- If *markname* includes spaces, you must enclose it in double quotes.

- You can not use a Data Protection for SQL Server **restore** command with **/STOPATMark** and also specify **/recovery**=no or **/standby**.
- If the restore operation with **/STOPATMark** does not encounter a transaction in the restored transaction log to stop at, the SQL database is left in an unrecovered state, even if you also specify **/recovery**=yes.

/STOPBEFOREMark=*markname* **[/AFTER=***datetime***]**

This parameter specifies a named point in time to restore a database to. This can be after a specified point in time if you specify the **/AFTER** option. Only transaction log records written before and not including the named transaction (which may be found at or after the specified point in time) are applied to the SQL database. The *markname* variable specifies the name of a SQL transaction. The SQL transaction may be a local transaction or a distributed transaction. If it is a distributed transaction name, the named mark exists in the transaction log of each SQL database participating in the distributed transaction.

markname is the transaction name, not the description that follows the MARK keyword in a SQL BEGIN TRANSACTION or BEGIN DISTRIBUTED TRANSACTION statement. The *datetime* variable specifies both the date and time separated by a space. Use any valid date and time format accepted by SQL Server.

Considerations:

- This parameter applies only to transaction log restores. The base restore that the transaction logs apply to must have been a full database restore. You cannot restore file, group, and set restores to a mark.
- You can use the same named mark for several SQL transactions.
- If you do not specify **/AFTER**, the restore stops before the first mark it encounters with the specified name.
- If you specify **/AFTER**, the restore stops before the first mark it encounters with the specified name, or after the specified date and time.
- If *markname* includes spaces, you must enclose it in double quotes.
- You can not use a Data Protection for SQL Server **restore** command with **/STOPBEFOREMark** and also specify **/recovery**=no or **/standby**.
- If the restore operation with **/STOPBEFOREMark** does not encounter a transaction in the restored transaction log to stop before, the SQL database is left in an unrecovered state, even if you also specify **/recovery**=yes.

/STRIPes=*numstripes*

The **/STRIPes** parameter specifies the number of data stripes to use in a backup or restore operation. The *numstripes* variable can range from *1* to *64*.

Considerations:

- If you do not specify **/STRIPes**, the default value is that specified in the Data Protection for SQL Server configuration file. The initial value is *1*. For **restore**, the value is the same as that used in the backup operation.
- If you specify **/STRIPes** but not *numstripes*, the stored value is used.
- You may use up to the number used to create the backup. You can determine the number of data stripes used to create a backup object with the Data Protection for SQL Server command: `query tsm dbname backup_object`

- You must use the **MAXNUMMP** parameter on a Tivoli Storage Manager **REGISTER NODE** or **UPDATE NODE** command to allow a node to use multiple sessions to store data on removable media (which requires you to allocate multiple mount points to that node). The **MAXNUMMP** value must be equal to or less than the maximum number of stripes you desire.
- When you use data striping, you should use Tivoli Storage Manager server file space collocation to try to keep each stripe on a different storage volume.
- The maximum number of data stripes you can use is one less than the value of theTivoli Storage Manager server TXNGROUPMAX option in the dsmserv.opt file. SQL server allows a maximum of *64* data stripes.

/TSMNODe=*tsmnodename*

The **/tsmnode** parameter specifies the Tivoli Storage Manager node name that Data Protection for SQL Server uses to log on to the Tivoli Storage Manager server. This identifies which Tivoli Storage Manager client is requesting services. You can also store the node name in the options file. The command line parameter overrides the value in the options file.

Considerations:

- You cannot use the **/TSMNODe** parameter if PASSWORDACCESS GENERATE is specified in the Tivoli Storage Manager options file. You must specify the nodename in the options file. Otherwise, you can change PASSWORDACCESS to PROMPT to utilize the **/TSMNODe** parameter. For details about the Tivoli Storage Manager options file, see the reference manual *IBMTivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide*.
- If you do not specify **/TSMNODe**, the default value is that specified by the nodename option in the Tivoli Storage Manager options file. Specifying this parameter does not change the value in the options file.

/TSMOPTFile=*tsmoptfilename*

The **/TSMOPTFile** parameter specifies the Tivoli Storage Manager options file to use. This is similar to selecting a Tivoli Storage Manager server from the server list in the GUI. The Tivoli Storage Manager options file contains the configuration values for the Tivoli Storage Manager API. For details about the Tivoli Storage Manager options file, see the reference manual *IBM Tivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide*.

Considerations:

- The *tsmoptfilename* variable can include a fully qualified path. If you do not include a path, the directory where Data Protection for SQL Server is installed is used.
- If *tsmoptfilename* includes spaces, you must enclose it in double quotes.
- If you do not specify **/TSMOPTFile**, the default value is dsm.opt.
- If you specify **/TSMOPTFile** but not *tsmoptfilename*, the default is also dsm.opt.

/TSMPassword=*tsmpasswordname*

The **/TSMPassword** parameter specifies the Tivoli Storage Manager password that Data Protection for SQL Server uses to log on to the Tivoli Storage Manager server. This parameter and the option PASSWORDACCESS in the Tivoli Storage Manager options file interact in the following ways:

/TSMPassword	**PASSWORDACCESS in Tivoli Storage Manager options file**	**Password already stored in registry?**	**Result**
specified	*generate*	yes	**/TSMPassword** ignored
specified	*generate*	no	**/TSMPassword** used and stored
specified	*prompt*	—	**/TSMPassword** used
not specified	*prompt*	—	user is prompted

Legacy Restore output examples

These output examples provide a sample of the text, messages, and process status that displays when using the **restore** command.

Restore ReportServer Full

Running this command restores a full backup of the *model* to a different server than that from which it was backed up.

Command:

```
tdpsqlc restore model full /fromsqlserver=STRINGVM1\STRINGVM1
```

Output:

```
IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1998, 2013.
All rights reserved.

Connecting to SQL Server, please wait...

Querying TSM Server for Backups ....

Starting Sql database restore...

Beginning VSS restore of 'model'...

   Files Examined/Completed/Failed: [ 2 / 2 / 0 ]   Total Bytes: 3933070

VSS Restore operation completed with rc = 0
   Files Examined      : 2
   Files Completed     : 2
   Files Failed        : 0
   Total Bytes         : 3933070
   Total LanFree Bytes : 0

Completed
```

Legacy Restore 2–Differential

Legacy Restore 2 displays restoring a differential backup object of database *Test1* into database *Test2*. Note that the *Test2* database must already exist for the restore to be successful.

Command:

```
tdpsqlc restore Test1 diff /into=Test2
```

Output:

```
IBM Tivoli Storage Manager for Databases
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1998, 2013.
All rights reserved.

Starting Sql database restore...

Querying Tivoli Storage Manager server for a list of database backups,
 please wait...

Beginning difffull restore of backup object Test1, 1 of 1,
  to database Test2
Full: 0   Read: 478720  Written: 478720  Rate: 40.62 Kb/Sec
Restore of Test1 completed successfully.

Total database backups inspected:               1
Total database backups requested for restore:   1
Total database backups restored:                1
Total database skipped:                         0

Throughput rate:                                40.61 Kb/Sec
Total bytes transferred:                        478,720
LanFree bytes transferred:                      0
Elapsed processing time:                        11.51 Secs
```

Legacy Restore 3–Group

Legacy Restore 3 displays restoring a filegroup backup object named *Group1* to database *Test1*.

Command:

```
tdpsqlc restore Test1 group=Group1
```

Output:

```
IBM Tivoli Storage Manager for Databases
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1998, 2013.
All rights reserved.

Starting Sql database restore...

Querying Tivoli Storage Manager server for a list of database backups,
 please wait...

Restoring meta data ...

Beginning group restore of backup object Test1\Group1, 1 of 1,
  to database Test1
Full: 0   Read: 86982144  Written: 86982144  Rate: 8,188.11 Kb/Sec
Restore of Test1\Group1 completed successfully.

Total database backups inspected:               1
Total database backups requested for restore:   1
Total database backups restored:                1
Total database skipped:                         0

Throughput rate:                                8,185.75 Kb/Sec
Total bytes transferred:                        86,982,144
LanFree bytes transferred:                      0
Elapsed processing time:                        10.38 Secs
```

Legacy Restore 4–Set

Legacy Restore 4 displays restoring all active set backup objects to database *Test1*.

Command:

```
tdpsqlc restore Test1 set=*
```

Output:

```
IBM Tivoli Storage Manager for Databases
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1998, 2013.
All rights reserved.

Starting Sql database restore...

Querying Tivoli Storage Manager server for a list of database backups,
 please wait...

Restoring meta data ...

Beginning set restore of backup object Test1\20120718141546\00000700,
1 of 1,to database Test1
Full: 0   Read: 88489472  Written: 88489472  Rate: 8,125.58 Kb/Sec
Restore of Test1\20120718141546\00000700 completed successfully.

Total database backups inspected:              1
Total database backups requested for restore:  1
Total database backups restored:               1
Total database skipped:                        0

Throughput rate:                               8,122.52 Kb/Sec
Total bytes transferred:                       88,489,472
LanFree bytes transferred:                     0
Elapsed processing time:                       10.64 Secs
```

Legacy Restore 5–Log (point in time)

Legacy Restore 5 displays restoring all active log backup objects of database *Test1* to a specified point in time. Three of four log backups meet the datetime criteria.

Command:

```
tdpsqlc restore Test1 log=* /stopat="07/01/2012 13:56:00"
```

Output:

```
IBM Tivoli Storage Manager for Databases
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1998, 2013.
All rights reserved.

Starting Sql database restore...

Querying Tivoli Storage Manager server for a list of
database backups, please wait...

Beginning log restore of backup object Test1\20120701135511\
 00000700,
1 of 4,to database Test1
Full: 0   Read: 214528  Written: 214528  Rate: 59.75 Kb/Sec
Restore of Test1\20120701135511\00000700 completed successfully.

Beginning log restore of backup object Test1\20120701135605\
 00000700,
2 of 4,to database Test1
Full: 0   Read: 147968  Written: 147968  Rate: 32.15 Kb/Sec
Restore of Test1\20120701135605\00000700 completed successfully.

Beginning log restore of backup object Test1\20120701135712\
 00000700,
3 of 4,to database Test1
Full: 0   Read: 0  Written: 0  Rate: 0.00 Kb/Sec
Restore of Test1\20120701135712\00000700 completed successfully.

Skipping Test1\20120701135817\00000700
because of the preceding failure or point-in-time recovery.

Total database backups inspected:             4
Total database backups requested for restore: 4
Total database backups restored:              3
Total database skipped:                       1

Throughput rate:                              37.21 Kb/Sec
Total bytes transferred:                      362,496
LanFree bytes transferred:                    0
Elapsed processing time:                      9.51 Secs
```

Legacy Restore 6–Log (named mark)

Legacy Restore 6 displays restoring all active log backup objects to database *Testmark* to a named point in time. The first mark with the specified name, *mark2*, is encountered in the third log backup object applied to the restore. The restore stops once this mark is encountered.

Command:

```
tdpsqlc restore Testmark log=* /stopatmark=mark2
```

Output:

```
IBM Tivoli Storage Manager for Databases
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1998, 2013.
All rights reserved.

Starting Sql database restore...

Querying Tivoli Storage Manager server for a list of database
backups, please wait...

Beginning log restore of backup object Testmark\20120701102947\
 0000065C, 1 of 4, to database Testmark
Full: 0   Read: 159232  Written: 159232  Rate: 61.68 Kb/Sec
Restore of Testmark\20120701102947\0000065C completed successfully.

Beginning log restore of backup object Testmark\20120701103127\
 000001DC, 2 of 4, to database Testmark
Full: 0   Read: 159232  Written: 159232  Rate: 34.51 Kb/Sec
Restore of Testmark\20120701103127\000001DC completed successfully.

Beginning log restore of backup object Testmark\20120701103325\
 00000680, 3 of 4, to database Testmark
Full: 0   Read: 0  Written: 0  Rate: 0.00 Kb/Sec
Restore of Testmark\20120701103325\00000680 completed successfully.

Skipping Testmark\20120701103556\00000694
because of the preceding failure or point-in-time recovery.

Total database backups inspected:              4
Total database backups requested for restore:  4
Total database backups restored:               3
Total database skipped:                        4

Throughput rate:                               38.60 Kb/Sec
Total bytes transferred:                       318,464
LanFree bytes transferred:                     0
Elapsed processing time:                       8.06 Secs
```

Legacy Restore 7–Log (inactive object)

Legacy Restore 7 begins with a query to display both active and inactive log backup objects for database *Test1.*

Command:

```
tdpsqlc q tsm netapp_db2 log=* /all
```

Output:

```
IBM Tivoli Storage Manager for Databases
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1998, 2013.
All rights reserved.

Connecting to TSM Server as node 'STRINGVM1_SQL'...

Backup Object Information
-------------------------

SQL Server Name   ........................ STRINGVM1\STRINGVM1
SQL Database Name ........................ netapp_db2
Backup Method     ........................ Lgcy
Backup Location   ........................ Srv
Backup Object Type ....................... Log
Backup Object State ...................... Active
Backup Creation Date / Time ............. 09/27/2012 08:36:28
Backup Size .............................. 82.50 KB
SQL Compressed ........................... No
Backup Compressed ........................ No
Backup Encryption Type ................... None
Backup Client-deduplicated ............... No
Database Object Name ..................... 20120927083628\00001A4C
Number of stripes in backup object ....... 1
Assigned Management Class  ............... DEFAULT
```

The restore operation for Legacy Restore 7 applies a specifically named inactive log backup object of database *Test1* to the restore. Since an inactive log backup object is being requested, the ***/object*** parameter must be used on the restore command.

Command:

```
tdpsqlc restore Test1 log=* /object=20120622135511\00000700
```

Output:

```
IBM Tivoli Storage Manager for Databases
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1998, 2013.
All rights reserved.

Starting Sql database restore...

Querying Tivoli Storage Manager server for a list of database
 backups,please wait...

Beginning log restore of backup object Test1\20120622135511\
 00000700,1 of 1,to database Test1
Full: 0   Read: 214528  Written: 214528  Rate: 29.47 Kb/Sec
Restore of Test1\20120622135511\00000700 completed successfully.

Total database backups inspected:              1
Total database backups requested for restore:  1
Total database backups restored:               1
Total database skipped:                        0

Throughput rate:                               29.46 Kb/Sec
Total bytes transferred:                       214,528
LanFree bytes transferred:                     0
Elapsed processing time:                       7.11 Secs
```

Legacy Restore 8–Full (partial)

Legacy Restore 8 displays restoring part of a full backup object, filegroup *Group1*, to database *Test1*.

Command:

```
tdpsqlc restore Test1 full /partial /gr=Group1
```

Output:

```
IBM Tivoli Storage Manager for Databases
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1998, 2013.
All rights reserved.

Starting Sql database restore...

Querying Tivoli Storage Manager server for a list of database
 backups,please wait...

Restoring meta data ...

Beginning full restore of backup object Test1, 1 of 1,
  to database Test1
Full: 0   Read: 89607680  Written: 89607680 Rate: 3,359.60 Kb/Sec
Restore of Test1 completed successfully.


Total database backups inspected:               1
Total database backups requested for restore:   1
Total database backups restored:                1
Total database skipped:                         0

Throughput rate:                                3,359.21 Kb/Sec
Total bytes transferred:                        89,607,680
LanFree bytes transferred:                      0
Elapsed processing time:                        26.05 Secs
```

Legacy Restore 9–Full (relocate)

Legacy Restore 9 displays restoring a full backup object of database *Test1,* specifically relocating logical file *File1Group1* to a new physical location.

Command:

```
tdpsqlc restore Test1 full /relocate=File1Group1
/to=e:\sqldata\File1Group1.NDF
```

Output:

```
IBM Tivoli Storage Manager for Databases
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1998, 2013.
All rights reserved.

Starting Sql database restore...

Querying Tivoli Storage Manager server for a list of database
 backups,please wait...

Restoring meta data ...

Beginning full restore of backup object Test1, 1 of 1,
  to database Test1
Full: 0   Read: 88100352  Written: 88100352  Rate: 3,930.18 Kb/Sec
Restore of Test1 completed successfully.

Total database backups inspected:               1
Total database backups requested for restore:   1
Total database backups restored:                1
Total database skipped:                         0

Throughput rate:                                3,929.64 Kb/Sec
Total bytes transferred:                        88,100,352
LanFree bytes transferred:                      0
Elapsed processing time:                        21.89 Secs
```

VSS restore output examples

These output examples provide a sample of the text, messages, and process status that displays when using the **restore** command.

VSS restore from Tivoli Storage Manager server

Restore database *msdb* from Tivoli Storage Manager server storage using the optional parameters, **/backupdestination** and **/backupmethod**.

Command:

```
tdpsqlc restore msdb full /backupdestination=tsm /backupmethod=vss
```

Output:

```
IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1998, 2013.
All rights reserved.

Connecting to SQL Server, please wait...

Querying TSM Server for Backups ....

Starting Sql database restore...

Beginning VSS restore of 'msdb'...

   Files Examined/Completed/Failed: [ 2 / 2 / 0 ]   Total Bytes: 8062302

VSS Restore operation completed with rc = 0
   Files Examined      : 2
   Files Completed     : 2
   Files Failed        : 0
   Total Bytes         : 8062302
   Total LanFree Bytes : 0

Completed
```

VSS restore from local

Restore database *DEMODB* from local shadow volumes using the new optional parameters, **/backupdestination** and **/backupmethod**.

Command:

```
tdpsqlc restore DEMODB full /backupdestination=local
 /backupmethod=vss /instantrestore=no
```

Output:

```
IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1998, 2013.
All rights reserved.

Connecting to SQL Server, please wait...

Querying TSM Server for Backups ....

Starting Sql database restore...


Beginning VSS restore of 'DEMODB'...


Files Examined/Completed/Failed: [ 2 / 2 / 0 ] Total Bytes: 5243190

VSS Restore operation completed with rc = 0
Files Examined: 2
Files Completed: 2
Files Failed: 0
Total Bytes: 5243190
Total LanFree Bytes: 0
```

VSS restore: Instant restore from local

Use instant restore to restore database *testdb2* from local shadow volumes using the new **/instantrestore** parameter.

Command:

```
tdpsqlc restore testdb2 /backupmethod=vss
/backupdest=local /instantrestore=yes
```

Output:

```
IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1998, 2013.
All rights reserved.

Connecting to SQL Server, please wait...

Querying TSM Server for Backups ....

Starting Sql database restore...


Beginning VSS restore of 'DEMODB'...



Restoring 'DEMODB' using volume-level-copy snapshot.


Starting snapshot restore process. This process may take several minutes.



VSS Restore operation completed with rc = 0
Files Examined : 0
Files Completed : 0
Files Failed : 0
Total Bytes : 0
Total LanFree Bytes   : 0
```

VSS restore: Relocate directory

Restore and relocates database *svtdb* from Tivoli Storage Manager server storage to directory *m:\svtdb* using the new optional parameter, **/relocatedir**. All SQL logs and full-text index files associated with database *svtdb* are also restored and relocated.

Command:

```
tdpsqlc restore svtdb full /relocatedir=m:\svtdb /backupdestination=tsm
 /backupmethod=vss
```

Output:

```
IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1998, 2013.
All rights reserved.

Connecting to SQL Server, please wait...

Querying TSM Server for Backups ....

Starting Sql database restore...


Beginning VSS restore of 'svtdb'...



Preparing for restore of 'svtdb' from TSM backup.

   Files Examined/Completed/Failed: [ 5 / 5 / 0 ]   Total Bytes: 418328259

VSS Restore operation completed with rc = 0
   Files Examined   : 5
   Files Completed  : 5
   Files Failed     : 0
   Total Bytes      : 418328259
  Total LanFree Bytes   : 0
```

To restore and relocate the database *svtdb*, its logs, and its full-text index files into their own respective locations, the following command is issued:

```
tdpsqlc restore svtdb full /relocatedir=m:\svtdb,e:\svtdb,f:\svtdb
/backupdestination=tsm /backupmethod=vss
```

The **`/relocatedir`** values in this command are as follows:

- *m:\svtdb*: The directory where only the *svtdb* database is relocated.
- *e:\svtdb*: The directory where only the *svtdb* logs are relocated.
- *f:\svtdb*: The directory where only the *svtdb* full-text index files are relocated.

Restorefiles command

Use the **`restorefiles`** command to restore VSS-based backups on the Tivoli Storage Manager server (/BACKUPDESTINATION=TSM), or stored locally (/BACKUPDESTINATION=LOCAL).

Consider the following information before using the **`restorefiles`** command.

- The **`restorefiles`** command restores `.mdf`, `ldf`, and other flat files from a specified Data Protection for SQL server, VSS-based backup into a specified directory.
- A destination directory can be specified as a directory on a fixed file system (for example `C:\temp`), or on a network share (for example `\\server\dest`) that is accessible to the Tivoli Storage Manager server Remote Agent (VSS Requestor)
- The **`restorefiles`** command does not restore the data to the SQL server.
- This command does not require the SQL Server to be installed on the machine where the **`restorefiles`** command is run. Files can be restored to another machine or directory on the same machine as the SQL Server.
- A restore continues until it completed, unless the destination volume does not have enough space to fulfill the restore operation.

- VSS-based backups that are located on the Tivoli Storage Manager server (/BACKUPDESTINATION=TSM) can be restored by using **restorefiles** on the same machine that performed the VSS-based backup, or by running the command on a machine that has the Data Protection for SQL client installed and configured for VSS operations.
- The directory specified in the **restorefiles** command has the VSS component name appended so that multiple databases can be restored to the same target directory.
- VSS-based backups that are stored on the local machine by using a persistent snapshot (/BACKUPDESTINATION=LOCAL), can be restored only by running the **restorefiles** command on the same machine that performed the VSS-based backup, and has access to the persistent snapshot.
- To run a full restore: `tdpsqlc restorefiles DB1 FULL relocatedir=d:\temprestore`
- Use `/RELOCATEDIR` to restore a database that currently exists to a different directory, even if your backup contains files that are located in different directories. Run the **restorefiles** command and specify just one restore destination directory. For example, issue `restorefiles db1 full /relocatedir=d:\temp` to place the files into the `d:\temp\db1\*` directory.
- If you are in a non-clustered environment, you can restore only a local snapshot to the machine that generated the snapshot.
- If you are in a clustered environment, you can run a **restorefiles** command from any of the machines in the cluster.

Restorefiles syntax

Use the **restorefiles** command syntax diagram as a reference for available options and truncation requirements.

TDPSQLC command

```
                           ┌─,────────┐
                           ▼          │
►►──TDPSQLC──RESTOREFIles────┬─dbname─┬┴─────────────────────────────►
                             └─*──────┘
```

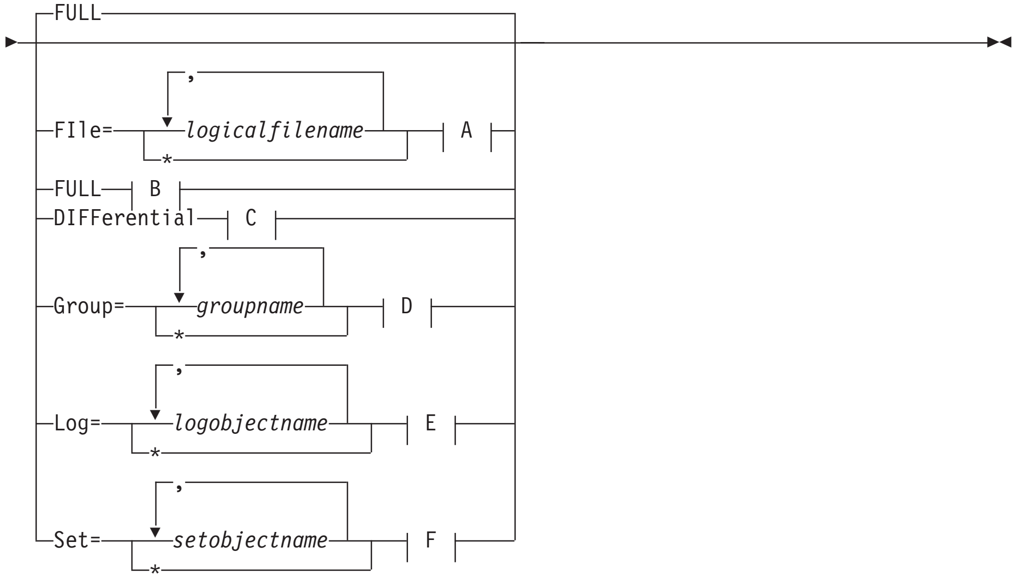

The syntax diagrams of the backup object type options corresponding to the letters *A, B, C, D, E, F* follow the Optional Parameters for the **restorefiles** command "Restorefiles optional parameters."

Restorefiles positional parameters

Positional parameters immediately follow the **restorefiles** command and precede the optional parameters.

The following positional parameters specify the object to restore:

tdpsqlc restorefiles
*** |componentname1, ..., componentnameN*FULL***

* Sequentially restore all flat files for the database.

The following positional parameters specify the type of backup from which the files are restored:

FULL Restore the files from a Full type backup for VSS.

Restorefiles optional parameters

Optional parameters for the Data Protection for SQL **restorefiles** command and optional parameters.

/BACKUPDESTINATION
VSS backups that are located on the Tivoli Storage Manager server are restored using the **restorefiles** command with **/BACKUPDESTINATION=TSM**. VSS backups that are running on a local machine using a persistent snapshot are restored using the **restorefiles** command with **/BACKUPDESTINATION=LOCAL**. TSM is the default destination for **restorefiles**.

/CONFIGfile=*configfilename*
Use the **/configfile** parameter to specify the name of the Data Protection

for SQL Server configuration file that contains the values for the Data Protection for SQL Server configuration options.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the Data Protection for SQL Server installation directory is used. If the **/configfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is ***tdpsql.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"
```

/FROMSQLserver=*sqlservername*

Use the **/fromsqlserver** parameter to specify the name of the SQL Server where the original backup was performed. The default is the local SQL Server name.

/LOGFile=*logfilename*

Use the **/logfile** parameter to specify the name of the activity log file that is generated by Data Protection for SQL Server.

The *logfilename* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfilename* variable can include a fully-qualified path. However, if no path is specified, the log file is written to the Data Protection for SQL Server installation directory.

If the *logfilename* variable includes spaces, enclose the entire **/logfile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\mytdpsqlserver.log"
```

If the **/logfile** parameter is not specified, log records are written to the default log file, ***tdpsqlserver.log***.

The **/logfile** parameter cannot be turned off, logging always occurs.

When using multiple simultaneous instances of Data Protection for SQL Server to perform operations, use the **/logfile** parameter to specify a different log file for each instance used. This directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=*numdays***|No**

Use the **/logprune** parameter to disable log pruning or to explicitly request that the log be pruned for one command run. By default, log pruning is enabled and performed once per day. The *numdays* variable represents the number of days to save log entries. By default, **60** days of log entries are saved in the pruning process. You can use the Management Console (MMC) 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 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 SQL Server command that might prune the log file, perform one of the following actions to prevent the log file from being pruned:

- Make a copy of the existing log file.
- Specify a new log file with the **/logfile** parameter or **logfile** setting.

/MOUNTWait=Yes|No

This parameter is not valid for all backup types; does not work with DIFFFULL or LOG backup types. The **/mountwait** parameter is used to specify whether Data Protection for SQL Server should 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 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 SQL Server.

Use the Data Protection for SQL Server **query tsm *** command to view the names of the backup objects.

/Quiet This parameter prevents status information from being displayed. This does not affect the level of information written to the activity log.

/RELOCATEDir=*dbfiledir***[** *,logfiledir* **[** *,otherfiledir***] [** *,filestream files***]]**

The **/relocatedir** parameter specifies the destination locations in which to restore the flat files. This includes databases, logs, and FILESTREAM files.

The *dbfiledir* variable specifies the directory location of the SQL database you want to relocate. Note that if the *logfiledir* or *otherfiledir* variables are not specified, the logs and SQL Server full-text index files are restored to the directory specified by *dbfiledir*.

The *logfiledir* variable specifies the directory location of the SQL log files you want to relocate. Note that if the *logfiledir* variable is not specified, the SQL log files are restored to the directory specified by *dbfiledir*.

The *otherfiledir* variable specifies the directory location of the SQL Server full-text index files you want to relocate. Note that if the *otherfiledir* variable is not specified, the SQL Server full-text index files are restored to the directory specified by *dbfiledir*.The **restorefiles** operation creates a subdirectory under the root directory that contains the name of the database name. Restored files are placed in that subdirectory. If the **/relocatedir** parameter is not specified, the files will be restored into the directory where the **restorefiles** command is issued. For example, if Data Protection for SQL Server is installed in the `c:\Program Files\Tivoli\TSM\TDPSQLC` directory and the following command is issued from `E:\Somedir`:

```
e:\Somedir> c:\"Program Files"\Tivoli\TSM\TDPSQLC\tdpsqlc restorefiles
db1 full
```

Then, the files are restored to the subdirectories in the `e:\Somedir` location:

```
e:\Somedir\db1\db1.mdf
e:\Somedir\db1\db1.ldf
```

/TSMNODe=*tsmnodename*

Use the *tsmnodename* variable to refer to the Tivoli Storage Manager node name that Data Protection for SQL Server 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 SQL Server options file.

The file name can include a fully qualified path name. If no path is specified, the directory where Data Protection for SQL Server is installed is searched.

If the *tsmoptfilename* variable includes spaces, enclose the entire **/tsmoptfile** parameter entry in double quotation marks. For example:

```
/TSMOPTFile="c:\Program Files\file.opt"
```

The default is **dsm.opt**.

/TSMPassword=*tsmpassword*

Use the *tsmpassword* variable to refer to the Tivoli Storage Manager password that Data Protection for SQL Server uses to log on to the Tivoli Storage Manager server. If you specified PASSWORDACCESS GENERATE in the Data Protection for SQL Server options file (dsm.opt), you do not need to supply the password here because the one that is stored in the registry is used. However, to store the password in the registry, you must specify the Tivoli Storage Manager password the first time Data Protection for SQL Server 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 SQL Server uses to log on to the Tivoli Storage Manager server can be up to 63 characters in length.

Restorefiles examples

This output example provides a sample of the text, messages, and process status that displays when using the **restorefiles** command.

This command, **tdpsqlc restorefiles Finance FULL /backupdestination=local /RELOCATEDir=e:\test/FROMSQLServer=sqlsrv12**, restores VSS files from a FULL type backup of the *Finance* database from the SQL Server named *sqlsrv12* into the *e:\test* directory. The restored files are:

```
e:\test\Finance\finance.mdf
e:\test\Finance\finance_log.ldf
```

Set command

Use the **set** command to change the values for the Data Protection for SQL Server configurable parameters and options.

The values are saved in a configuration file. The default file is `tdpsql.cfg`. Configuration values can also be set in the Data Protection Properties window in the GUI.

Note: If a configuration file is not specified, the `tdpsql.cfg` values are used, and a default configuration file is created with only the **lastprunedate** value. If an invalid or non-existent file is specified, the default values are used.

Set syntax

Use the **set** command syntax diagrams as a reference to view available options and truncation requirements.

TDPSQLC command

```
►►──TDPSQLC──Set───┬─ALWAYSONNode=──nodename──────────────────────────┬───────────────►◄
                   │                      ┌─TSM───┐                   │
                   ├─BACKUPDESTINATION=───┼───────┼───────────────────┤
                   │                      ├─LOCAL─┤                   │
                   │                      └─BOTH──┘                   │
                   │               ┌─LEGACY─┐                         │
                   ├─BACKUPMETHOD=─┴─VSS────┴─────────────────────────┤
                   ├─BUFFers=──numbuffers─────────────────────────────┤
                   ├─BUFFERSIze=──buffersizeinkb──────────────────────┤
                   ├─DATEformat=──dateformatnum───────────────────────┤
                   ├─DIFFESTimate=──numpercent────────────────────────┤
                   ├─FROMSQLserver=──fromsqlserver────────────────────┤
                   │                                        ┌─No──┐   │
                   ├─IMPORTVSSSNAPSHOTSONLYWhenneeded=[─────┴─Yes─┴───┤
                   ├─LANGuage=──language──────────────────────────────┤
                   ├─LOCALDSMAGENTNODE=──nodename─────────────────────┤
                   ├─LOGFile=──logfilename────────────────────────────┤
                   ├─LOGPrune=──┬─numdays─┬───────────────────────────┤
                   │            └─No──────┘                           │
                   ├─NUMBERformat=──numberformatnum───────────────────┤
                   ├─MOUNTWaitfordata=──┬─Yes─┬───────────────────────┤
                   │                    └─No──┘                       │
                   ├─REMOTEDSMAGENTNODE=──nodename────────────────────┤
                   ├─SQLAUTHentication=──┬─INTegrated─┬───────────────┤
                   │                     └─SQLuserid──┘               │
                   ├─SQLBUFFers=──numsqlbuffers───────────────────────┤
                   ├─SQLBUFFERSIze=──sqlbuffersizeinkb────────────────┤
                   │                  ┌─=No──┐                        │
                   ├─/SQLCOMPression──┼─=Yes─┼────────────────────────┤
                   │                  └──────┘                        │
                   ├─SQLSERVer=──sqlprotocol:sqlservername────────────┤
                   ├─STRIPes=──numstripes─────────────────────────────┤
                   └─TIMEformat=──timeformatnum───────────────────────┘
```

Set Optional Parameters

```
►►──┬─────────────────────────────────────────┬──────────────────────────►◄
    │              ┌─=tdpsql.cfg────────┐     │
    └─/CONFIGfile──┼────────────────────┼─────┘
                   └─=configfilename────┘
```

Set positional parameters

Positional parameters immediately follow the **set** command and precede the optional parameters.

To set default values in the Data Protection for SQL Server configuration file, specify one of the following when issuing a **set** command.

/ALWAYSONNode=*nodename*

Specify the Tivoli Storage Manager node name that is used to back up AlwasyOn availability databases in a SQL Server 2012 environment. This parameter is required when you are configuring Data Protection for SQL Server in a SQL Server 2012 environment. All availability databases in an availability group are backed up under this node name, regardless of which availability replica they are from. The databases that are not in an availability group are backed up under the standard Data Protection for SQL Server node name unless you specify the **/USEALWAYSONnode** parameter.

BACKUPDESTination=TSM|LOCAL|BOTH

Use the **BACKUPDESTination** positional parameter to specify the storage location for your backup. You can specify:

TSM The backup is stored on Tivoli Storage Manager server storage only. This is the default.

LOCAL The backup is stored on local shadow volumes only.

BOTH The backup is stored on both Tivoli Storage Manager server storage and local shadow volumes.

BACKUPMETHod=LEGACY|VSS

Use the **BACKUPMETHod** positional parameter to specify the method for your backup. You can specify:

LEGACY Data Protection for SQL Server uses the legacy API to perform the backup. This is the default.

VSS Data Protection for SQL Server uses VSS to perform the backup.

BUFFers=*numbuffers*

The **BUFFers** parameter specifies the number of data buffers used for each data stripe to transfer data between Data Protection for SQL Server and the Tivoli Storage Manager API. You can improve throughput by increasing the number of buffers, but you will also increase storage use. Each buffer is the size specified by the **BUFFERSIze** parameter. The *numbuffers* variable refers to the number of data buffers to use. The number can range from 2 to 8. The initial value is 3.

BUFFERSIze=*buffersizeinkb*

The **BUFFERSIze** parameter specifies the size of each Data Protection for SQL Server buffer specified by the **BUFFers** parameter. The *buffersizeinkb* variable refers to the size of data buffers in kilobytes. The number can range from 64 to 8192. The default is initially 1024.

DATEformat=*dateformatnum*

The **DATEformat** parameter selects the format you want to use to display dates.

The *dateformatnum* variable can range from 1 to 7. The initial value is 1. The number values specify the following formats:

1	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 SQL Server log file (`tdpsql.log` by default). You can avoid losing existing log file data by performing one of the following:

- After changing the value of the **DATEformat** parameter, make a copy of the existing log file before running Data Protection for SQL Server.
- Specify a new log file with the **LOGFile** parameter.

DIFFESTimate=*numpercent*

For differential database backups using the Data Protection for SQL Server **backup** command, **DIFFESTimate** specifies the estimated fraction of an entire SQL database that has changed since its last full database backup. This estimate is needed because SQL Server does not provide a way to determine the size of a differential backup, and because the Tivoli Storage Manager server requires an accurate size estimate to efficiently allocate space and place objects. The Tivoli Storage Manager server uses this value to determine if there is enough space in the primary storage pool to contain the SQL database backup. Because a separate backup object is created for each specified SQL database, this estimate applies to each specified SQL database individually. The *numpercent* variable can range from 1 to 99. Because a differential backup backs up database pages, this number is the percent of database pages changed since the last full database backup. The initial value is 20.

Considerations:

- If the estimate is significantly smaller than the actual quantity of changes, the Tivoli Storage Manager server may be forced to abnormally end the backup because the backup size is larger than the space the Tivoli Storage Manager server allocated for it.
- If the estimate is significantly larger than the actual quantity of changes, the server may be forced to place the backup object higher in the storage pool hierarchy than otherwise necessary, possibly on removable media.

FROMSQLSERVer=sqlservername

The **/FROMSQLSERVer** parameter specifies the SQL server that backup objects were backed up from. This parameter is necessary only when the name of the SQL server to restore to, as determined by the **sqlserver** parameter, is different from the name of the SQL server that the backup objects were created from. Use **/FROMSQLSERVer** for **query TSM** and **inactivate** commands, but use **sqlserver** for **query SQL** commands. The default value is the *sqlserver* value or the value set in the Data Protection for SQL Server configuration file.

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 you want to use to display messages:

CHS	Simplified Chinese
CHT	Traditional Chinese
DEU	Standard German
ENU	American English (This is the default.)
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 machine that performs the VSS backups. This positional parameter must be specified for VSS operations to be performed.

LOGFile=*logfilename*

The **LOGFile** parameter specifies the name of the activity log that is generated by Data Protection for SQL Server. The activity log records significant events such as completed commands and error messages. This log is distinct from the SQL Server error log. The *logfilename* variable identifies the name to be used for the activity log generated by Data Protection for SQL Server.

Considerations:

- If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file.
- The file name can include a fully-qualified path; however, if you specify no path, the file is written to the directory where Data Protection for SQL Server is installed.
- You cannot turn Data Protection for SQL Server activity logging off. If you do not specify **LOGFile**, log records are written to the default log file. The default log file is `tdpsql.log`.

LOGPrune=*numdays***|No**

The **LOGPrune** parameter prunes the Data Protection for SQL Server activity log and specifies how many days of entries to save. By default, log pruning is enabled and performed once each day Data Protection for SQL Server is executed; however, this option allows you to disable log pruning. The *numdays* variable represents the number of days to save log entries.

Considerations:

- If you specify *numdays*, it can range from 0 to 9999. The initial value is 60. A value of 0 deletes all entries in the Data Protection for SQL Server activity log file except for the current command entries.
- If you specify `No`, the log file is not pruned.

NUMBERformat=*numberformatnum*

The **NUMBERformat** parameter specifies the format of the numbers displayed by Data Protection for SQL Server. The *numberformatnum* variable can range from 1 to 6. The initial value is 1. The number values specify the following formats:

1	1,000.00
2	1,000,00
3	1 000,00
4	1 000.00
5	1.000,00
6	1'000,00

MOUNTWaitfordata=YES | No

If the Tivoli Storage Manager server is configured to store backup data on removable media, it is possible that the Tivoli Storage Manager server might indicate to Data Protection for SQL Server that it is waiting for a required storage volume to be mounted. If that occurs, this option allows you to specify whether Data Protection for SQL Server **backup** and **restore** commands wait for the media mount or stop the current operation. The initial value is `YES`.

Considerations:

- If you use data striping, Data Protection for SQL Server cannot complete waiting until the initial media for all stripes are available, although Data Protection for SQL Server starts to use each stripe as its media becomes available. Because of the way SQL Server distributes data among stripes, if any stripe does not have its media available, each of the stripes may eventually be either waiting for its own or another stripe's media to become available. In this case, it may become necessary to terminate the Data Protection for SQL Server command from a prolonged wait. This can be done by closing the command prompt window.
- If the management class for meta objects also requires removable media, Data Protection for SQL Server waits for that volume. During backup operations, the wait occurs after all of the data is transferred because meta objects are not created until after the data objects are complete. During restore operations, if the metadata is required, the wait occurs before any of the data is transferred
- If you specify `No` and any removable media are required, Data Protection for SQL Server terminates the command with an error message. This is also true if the management class for meta objects requires removable media, but, during backups, the command termination does not occur until after all of the data is transferred.

REMOTEDSMAgentnode=*nodename*

Specify the node name of the machine that moves the VSS data to Tivoli Storage Manager server storage during off-loaded backups.

SQLAUTHentication=INTegrated | SQLuserid

This parameter specifies the authorization mode used when logging on to the SQL server. The `INTegrated` value specifies Windows NT or Windows 2000 authentication. The user id you use to log on to Windows is the same id you will use to log on to the SQL server. This is the default value. Use the `SQLuserid` value to specify SQL Server user id authorization. The user id specified by the `SQLuserid` parameter is the id you will use to log on to the SQL server. That user id must have the SQL Server SYSADMIN fixed server role.

SQLBUFFers=*numsqlbuffers*

The **SQLBUFFers** parameter specifies the total number of data buffers SQL Server uses to transfer data between SQL Server and Data Protection for SQL Server. The *numsqlbuffers* variable refers to the number of data buffers to use. The number can range from 0 to 999. The default value is 0. When **SQLBUFFers** is set to 0, SQL determines how many buffers should be used. The *numsqlbuffers* variable is limited by storage restrictions. If you specify a value other than 0, the number you specify must be equal to or greater than the number of data stripes that you use. Up to 64 stripes may be used. If you specify a value other than 0 and receive errors during a backup, specify a value of 0 and try the backup again.

SQLBUFFERSIze=*sqlbuffersizeinkb*

The **SQLBUFFERSIze** parameter specifies the size of each buffer (specified by the **SQLBUFFers** parameter) SQL Server uses to transfer data to Data Protection for SQL Server. The *sqlbuffersizeinkb* variable refers to the size of data buffers in kilobytes. The number can range from 64 to 4096. The default is initially 1024.

SQLCOMPression=Yes | No

The **SQLCOMPression** parameter specifies whether SQL compression is applied. If you do not specify **SQLCOMPression**, the default value No is used.

SQLCOMPression is only available with legacy backups on SQL Server 2008 and later. For SQL Server 2008, backup compression is only supported on Enterprise Edition. SQL Server 2008 R2, backup compression is supported on Standard, Enterprise, and Datacenter editions. Starting with SQL Server 2008, any edition can restore a compressed backup.

SQL Server 2008 backup compression is generally faster and more effective than using it together with Tivoli Storage Manager compression.

SQLSERVer=*sqlprotocol:sqlservername*

The **SQLSERVer** parameter specifies the SQL server that Data Protection for SQL Server logs on to. This is the SQL server that backup objects are restored to. However, if the backup objects were created from a different SQL server name, you must use the **fromsqlserver** parameter. Use **sqlserver** for the **query SQL** command, but use **fromsqlserver** for the **query TSM** and **inactivate** commands. The *sqlprotocol* variable specifies the communication protocol to use. You can specify one of the following protocols:

- *lpc*: Use Shared Memory protocol.
- *np*: Use Named Pipes protocol.
- *tcp*: Use Transmission Control protocol.
- *via*: Use Virtual Interface Architecture protocol.

If no protocol is specified, Data Protection for SQL Server logs on to the SQL server according to the first protocol that becomes available.

STRIPes=*numstripes*

The **STRIPes** parameter specifies the number of data stripes to use in a backup or restore operation. The *numstripes* variable can range from 1 to 64. The default is initially 1. Note that stripes are not available for VSS operations.

TIMEformat=*timeformatnum*

The **TIMEformat** parameter specifies the format of the times displayed by Data Protection for SQL Server. The *timeformatnum* variable can range from 1 to 4. The initial value is 1. The number values specify the following formats:

1	23:00:00
2	23,00,00
3	23.00.00
4	11:00:00A/P

Changes to the value of the **`TIMEformat`** parameter can result in an undesired pruning of the Data Protection for SQL Server log file (tdpsql.log by default). You can avoid losing existing log file data by performing one of the following:

- After changing the value of the **`TIMEformat`** parameter, make a copy of the existing log file before running Data Protection for SQL Server.
- Specify a new log file with the **`LOGFile`** parameter.

Set optional parameters

Optional parameters follow the **`set`** command and positional parameters.

/CONFIGfile=*configfilename*

The **/configfile** parameter specifies the name of the Data Protection for SQL configuration file, which contains the values for the Data Protection for SQL configurable options.

Considerations:

- *configfilename* can include a fully qualified path. If *configfilename* does not include a path, it uses the directory where Data Protection for SQL is installed.
- If *configfilename* includes spaces, place it in double quotes.
- If you do not specify **/configfile**, the default value is ***tdpsql.cfg***.
- If you specify **/configfile** but not *configfilename*, the default value ***tdpsql.cfg*** is used.

Set output examples

These output examples provide a sample of the text, messages, and process status that displays when using the **set** command.

Example 1

The following example specifies the *STRINGVM1* server as the default SQL server in the configuration file.

Command:

```
tdpsqlc set sqlserver=STRINGVM1
```

Output:

```
IBM Tivoli Storage Manager for Databases
Data Protection for Microsoft SQL Server
Version 6, Release 4, Level 0.0
(C) Copyright IBM Corporation 1997, 2012. All rights reserved.

ACO5054I The configuration option was set successfully.
```

Example 2

The following example specifies *c64* as the AlwaysOn node name in the configuration file.

Command:

```
tdpsqlc set alwaysonnode=c64
```

Output:

```
IBM Tivoli Storage Manager for Databases:
Data Protection for Microsoft SQL Server
Version 6, Release 4, Level 0.0
(C) Copyright IBM Corporation 1997, 2012. All rights reserved.

Connecting to SQL Server, please wait...

ACO5054I The configuration option was set successfully.
```

The following statement is added to the `tdpsql.cfg` configuration file:

```
ALWAYSONNode     c64
```

Unmount Backup command

Use the **`unmount backup`** command to unmount backups that have been previously mounted, and are managed by Tivoli Storage FlashCopy Manager for SQL.

Unmount Backup syntax

Use the **`unmount backup`** command syntax diagrams as a reference to view available options and truncation requirements.

TDPSQLC command

```
►►──TDPSQLC──UNMOUNT BACKup───mount point root directory─────────────────────►

►──┬────────────────────────────────────┬──┬──────────────────────────────────┬──►
   │               ┌─tdpsql.xml───────┐ │  │             ┌─tdpsql.log─────┐   │
   └─/CONFIGfile=──┴─┬──────────────┬─┴─┘  └─/LOGFile=───┴─┬────────────┬─┴───┘
                     └─configfilename─┘                     └─logfilename─┘

              ┌─60──────┐
►─/LOGPrune=──┼─numdays─┼──┬───────────────────────────────────┬──────────────►
              └─No──────┘  │              ┌─=DP────────┐        │
                           └─/QUERYNode───┤            ├────────┘
                                          ├─=ALWAYSON──┤
                                          └─=BOTH──────┘

►──┬──────────────────────────────────┬──┬──────────────────────────────┬──────►
   └─/REMOTECOMPUTER=──computername───┘  └─/REMOTECOMPUTERUser=──user───┘

►──┬─────────────────────────────────────┬──┬──────────────────────────┬───────►
   └─/REMOTECOMPUTERPassword=──passwd────┘  └─/TSMNODe=──tsmnodename───┘
```

```
►─┬────────────────────────────────────┬─┬──────────────────────────────┬─────►◄
  │                 ┌─dsm.opt──────────┐ │ └─/TSMPassword=──tsmpassword─┘
  └─/TSMOPTFile=────┴─tsmoptfilename───┴─┘
```

Unmount Backup positional parameter

The positional parameter immediately follows the **unmount backup** command and precedes the optional parameters.

mount points root directory

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 ***tdpsql.cfg***.

If the *configfilename* variable includes spaces, enclose the entire **/configfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\tdpsql.cfg"
```

/LOGFile=*logfilename*

Use the **/logfile** parameter to specify the name of the activity log file. 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 installation directory.

If the *logfilename* variable includes spaces, enclose the entire **/logfile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\tdpsql.log"
```

If the **/logfile** parameter is not specified, log records are written to the default log file, ***tdpsql.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 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 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 to 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 changes, prior to issuing a Tivoli Storage FlashCopy Manager for SQL 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.

/QUERYNode=DP|ALWAYSON|BOTH
Specify whether you want to query standard databases from SQL Server 2012 that were backed up from a standard Data Protection for SQL Server node, the AlwaysOn node, or both nodes. This parameter is ignored for availability databases because the availability databases are always backed up under the AlwaysOn node.

/REMOTECOMPUTER=*computername*
Enter the IP address or host name for the remote system where you want to unmount the data.

/REMOTECOMPUTERUser=*user*
Enter the user name 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 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), 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 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 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 Tivoli Storage FlashCopy Manager uses to log on to the Tivoli Storage Manager server can be up to 63 characters in length.

Examples of Tivoli Storage Manager policy binding using VSSPOLICY and INCLUDE/EXCLUDE statements

To exploit automatic version control and expiration, you are able to set policy for each type of backup data. The method of setting policy is different for VSS and legacy backups:

- VSS backups: Use the VSSPOLICY statement in the Data Protection for SQL Server configuration file. By default, the configuration filename is `tdpsql.cfg`.
- Legacy backups: Use INCLUDE/EXCLUDE statements in the Data Protection for SQL Server options file. By default, the options filename is `dsm.opt`.

VSS examples

VSS backups use the VSSPOLICY statement in the Data Protection for SQL Server configuration file (see “Setting automatic expiration (VSS and legacy)” on page 41 for the general syntax):

```
VSSPOLICY  *  *  COPY   TSM  VSS_FULL_TSM_MC
```

Legacy examples

Legacy backups use INCLUDE/EXCLUDE statements in the Data Protection for SQL Server options file. (For more information about syntax, see “Setting automatic expiration (VSS and legacy)” on page 41.) For the examples provided in the following tables, grouped statements are intended to be used together. For example:

```
\...\full*
\...\full*\*
```

and

```
\...\file\f1*\*
\...\file\f1*
\...\f1\file*
```

Object matches for backuptype	Specification
Example for all objects	`\...\*`

Object matches for `backuptype`	Specification
Example for exclude statements with all type of backups (`full`, `diff`, `log`, `group`, `file`, `set`)	`\...\full*` `\...\diff*`

Example for include and exclude statements with all type of backups (`full`, `diff`, `log`, `group`, `file`, `set`)	`\...\full*`
	`\...\full*\*`
	`\...\copyfull*`
	`\...\copyfull*\*`
	`\...\diff*`
	`\...\diff*\*`
	`\...\log*`
	`\...\log\...\*`
	`\...\group*`
	`\...\group\...\*`
	`\...\file*`
	`\...\file\...\*`
	`\...\set*`
	`\...\set\...\*`

Example for exclude statements with file (*f1*) and group (*g1*)	`\...\g1\group*` `\...\f1\file*`

Example for include statements with file (*f1*) and group (*g1*)	`\...\group\g1*\*`
	`\...\group\g1*`
	`\...\g1\group*`
	`\...\file\f1*\*`
	`\...\file\f1*`
	`\...\f1\file*`

Example for exclude statements with group or file object names beginning with *g* or *f*	`\...\g*\group*` `\...\f*\file*`

Example for include statements with group or file object names beginning with *g* or *f*	`\...\group\g*\*`
	`\...\group\g*`
	`\...\g*\group*`
	`\...\file\f*\*`
	`\...\file\f*`
	`\...\f*\file*`

Example for exclude statements same as `\...\group*` or `\...\file*` (there is no equivalent for include statements)	`\...\*\group*` `\...\*\file*`

`backuptype` object with database matches	Specification
Example for all objects with database name *Db1*	`\...\Db1\...\*`
Example for all objects with database name *Db1* beginning with *Db*	`\...\Db*\...\*`
Ambiguous	`\...\*\...\*`

Example for exclude statements using `full`, `diff`, `copyfull` objects with database name *Db1*	`\...\Db1\full*`
	`\...\Db1\copyfull*`
	`\...\Db1\diff*`

`backuptype` object with database matches	**Specification**

Example for exclude and include statements using `full`, `diff`, `copyfull` objects with database name *Db1*	`\...\Db1\full*`
	`\...\Db1\full*\*`
	
	`\...\Db1\copyfull*`
	`\...\Db1\copyfull*\*`
	
	`\...\Db1\diff*`
	`\...\Db1\diff*\*`

Example for exclude statements using `log`, `group`, `file`, `set` objects with database name *Db1*	`\...\Db1\...\log*` `\...\Db1\...\group*`

Example for exclude and include statements using `log`, `group`, `file`, `set` objects with database name *Db1*	`\...\Db1\...\log*`
	`\...\Db1\...\log*\...\*`
	
	`\...\Db1\...\group*`
	`\...\Db1\...\group*\...\*`
	
	`\...\Db1\...\file*`
	`\...\Db1\...\file*\...\*`
	
	`\...\Db1\...\set*`
	`\...\Db1\...\set*\...\*`

Example for exclude statements using all group or file object names (*g1*, *f1*) with database name *Db1*	`\...\Db1\g1\group*` `\...\Db1\f1\file*`

Example for exclude and include statements using all group or file object names (*g1*, *f1*) with database name *Db1*	`\...\Db1\group\g1*`
	`\...\Db1\group\g1*\*`
	`\...\Db1\g1\group*`
	
	`\...\Db1\file\f1*`
	`\...\Db1\file\f1*\*`
	`\...\Db1\f1\file*`

Example for exclude statements using all group or file object names beginning with *g* or *f* with database name *Db1*	`\...\Db1\g*\group*` `\...\Db1\f*\file*`

Example for exclude and include statements using all group or file object names beginning with *g* or *f* with database name *Db1*	`\...\Db1\group\g*`
	`\...\Db1\group\g*\*`
	`\...\Db1\g*\group*`
	
	`\...\Db1\file\f*`
	`\...\Db1\file\f*\*`
	`\...\Db1\f*\file*`

Example for exclude statements using `\...\Db1\...\group*` or `file*` (there is no equivalent for include statements)	`\...\Db1\*\group*` `\...\Db1\*\file*`
Example for exclude statements using `\...\Db1\full*`	`\...\Db1\...\full*`
Example for exclude and include statements using `\...\Db1\full*`	`\...\Db1\...\full*` `\...\Db1\...\full*\*`
Example for exclude statements using `\...\full*`	`\...\*\full*`
Example for exclude and include statements using `\...\full*`	`\...\*\full*` `\...\*\full*\*`
Example for exclude statements using `\...\group*` (there is no equivalent for include statements)	`\...\*\*\group*`

`backuptype` **object with database matches**	**Specification**
Example for exclude statements using `\...\g1\group*` (there is no equivalent for include statements)	`\...\*\g1\group*`
Ambiguous	`\...\*\...\log*`
Nothing (**`typeInfo`** missing)	`\...\Db1\set*`

For the following table, use the following guidelines:

- If you use using only exclude statements with only `\meta\`, all objects (including data) are excluded.
- If you are using only exclude statements with only `\data\`, errors occur.

Meta and data object matches	**Specification**
Example for all meta or data objects	`\...\meta\...\*` `\...\data\...\*`
Example for all meta **`full`** objects	`\...\meta\...\full*`
Example for all data **`full`** objects	`\...\data\...\full*` `\...\data\...\full*\*`
Example for all meta group object names (*g1*)	`\...\meta\...\g1\group*`
Example for all data group object names (*g1*)	`\...\data\...\group\g1*` `\...\data\...\group\g1*\*`
Example for all meta group object names beginning with *g*	`\...\meta\...\g*\group*`
Example for all data group object names beginning with *g*	`\...\data\...\group\g*` `\...\data\...\group\g*\*`
Same as `\...\meta\...\group*`	`\...\meta\...\*\group*`
Nothing (qualifiers missing)	`\...\meta\*\...\data\*`

Meta and data object with database matches	**Specification**
Example for all meta or data objects with database name *Db1*	`\...\meta\...\Db1\...\*` `\...\data\...\Db1\...\*`
Example for all meta objects with database name *Db1*	`\...\meta\...\Db1\full*`
Example for full objects matching all data objects	`\...\data\...\Db1\full*` `\...\data\...\Db1\full*\*`
Example for all meta objects with database name *Db1*	`\...\meta\...\Db1\...\log*`
Example for all log objects matching all data objects	`\...\data\...\Db1\...\log\...\*`
Example for all group matching all meta objects	`\...\meta\...\Db1\...\group*`
Example for group matching all data objects	`\...\data\...\Db1\group\...\*`
Example for all meta object names (*g1*) with database name *Db1*	`\...\meta\...\Db1\g1\group*`
Example for all data group object names (*g1*) with database name *Db1*	`\...\data\...\Db1\group\g1*`

Meta and data object with database matches	Specification
Example for all meta object names beginning with *g* with database name *Db1*	`\...\meta\...\Db1\g*\group*`
Example for all data group object names beginning with *g* with database name *Db1*	`\...\data\...\Db1\group\g*`
Same as `\...\meta\...\Db1\...\group*` (No equivalent for data objects)	`\...\meta\...\Db1\*\group*`
Same as `\...\meta\...\full*` (No equivalent for data objects)	`\...\meta\...\*\full*`
Same as `\...\meta\...\group*` (No equivalent for data objects)	`\...\meta\...\*\*\group*`
Same as `\...\meta\...\g1\group*` (No equivalent for data objects)	`\...\meta\...\*\g1\group*`
Ambiguous	`\...\meta\...\*\...\log*` `\...\data\...\*\...\log*`
Nothing (qualifiers missing)	`\...\meta\*\...\data\*`

Server matches	Specification
Example for all objects from all servers beginning with *SQL*	`SQL*\...\*`
Example for all objects from all server instances with host *SQL2012*	`SQL2012\...\*`
Example for all objects from server `SQL2012\INST1`	`SQL2012\INST1\...\*`
Example for all objects from all servers beginning with `SQL2012\INST`	`SQL2012\INST*\...\*`
Same as `SQL2012\...\*`	`SQL2012\*\...\*`
Example for all meta or data objects from server `SQL2012\INST1`	`SQL2012\INST1\meta\...\*` `SQL2012\INST1\data\...\*`
Example for all meta or data objects from all named server instances with host *SQL2012*	`SQL2012\*\meta\...\*`
Example for all meta or data objects from all server instances with host *SQL2012*	`SQL2012\...\meta\...\*`
Example for all objects from server default instance (if no instance name matches *??ta*)	`SQL2012\meta\...\*` `SQL2012\data\...\*`

Transitioning SQL 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 at the same time. This might be useful if you decide to move to a Tivoli Storage Manager environment and want to continue to interact with the locally managed snapshots until policy marks them for expiration.

About this task

Tivoli Storage FlashCopy Manager works when connected to the local Tivoli Storage FlashCopy Manager server or a Tivoli Storage Manager server. The Tivoli

Storage Manager server can be located anywhere on your network. The Tivoli Storage FlashCopy Manager Snap-in includes two configuration wizards. These enable you to do a local configuration and a Tivoli Storage Manager configuration. You can move from one type of server to another by running the corresponding configuration wizard at any time.

Using the Tivoli Storage Manager server wizard

About this task

You can use the Tivoli Storage Manager server wizard to transitioning SQL Server backups from Tivoli Storage FlashCopy Manager to Tivoli Storage Manager. The wizard leads you through all the steps necessary to perform the configuration. You do not then need to perform the manual steps listed below.

Implement 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 will be used for each type of backup you plan to use. You can provide the administrator with the existing locally-defined policy settings in your Tivoli Storage FlashCopy Manager stand-alone environment. Use the GUI or the command-line interface of Data Protection for SQL Server to retrieve this information.
2. Register your Data Protection for SQL Server 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 SQL 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
   ```

Implement these tasks on the workstation running the SQL Server

Procedure

1. In the directory where the Data Protection for SQL Server is installed, make a copy of the options file named dsm.opt. After you begin 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 SQL command line prompt from the Tivoli Storage FlashCopy Manager Snapin. In the Tivoli Storage FlashCopy Manager Snapin Tree view, an SQL server node is displayed for each SQL server instance on the computer.
   a. Select an SQL server instance in the tree view. The integrated command line and an Actions pane is displayed.
   b. Launch the Data Protection for SQL Server command line from the Actions pane. Select:

```
Launch Command Line
```

c. To make a copy of the options file, enter:

```
copy dsm.opt dsm_local.opt
```

2. In the same directory, make a copy of the Data Protection for SQL Server configuration file. For example:

   ```
   copy tdpsql.cfg tdpsql_local.cfg
   ```

   Preserve the contents of the local configuration file if:

   - you have specified policy bindings during the use of Tivoli Storage FlashCopy Manager.
   - you will be 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 Word Pad.

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. The default is:

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

   Create an alternate 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 on using the **dsmcutil** command, refer to the information on using 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:tsmcad4local
   /node:my_backup-archive_client_node
   /password:my_TSM_server_password
   /optfile:"C:\Program Files\Tivoli\TSM\baclient\dsm_local.opt"
   /httpport:1583
   ```

7. Edit the dsm_local.opt file in the Data Protection for SQL Server 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 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 Manger server environment, you will need to phase out the remaining backups 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 will automatically cause the old backups to expire and be deleted. For example, if you want to expire each backup after it is 30 days old, update the time-based policy by using the command:

       ```
       tdpsqlc update policy mypolicy /daysretain=30
       /tsmoptfile=dsm_local.opt
       /configfile=tdpsql_local.cfg
       ```

       You can also make this change using the Local Policy Management dialog that is accessed from the Utilities menu of the Data Protection for SQL Server Backup/Restore GUI. Information on how to start the GUI is located in the following section describing 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 will not be regularly using the stand-alone environment client, you can use a scheduler to periodically start a command such as:

       ```
       tdpsqlc query tsm * /all
       /tsmoptfile=dsm_local.opt
       /configfile=tdpsql_local.cfg
       ```

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

       The very last backup, that is created while running the stand-alone environment, will not be automatically deleted by the process of expiring the backups. For that, you will need to 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 SQL Server **delete backup** command, or the Delete Backup (right mouse-click menu option) in the GUI Restore window.

11. To access the Tivoli Storage FlashCopy Manager stand-alone environment:

    a. Start the SQL Client – Command Line prompt.

    b. Start Tivoli Storage FlashCopy Manager stand-alone commands by appending the /tsmoptfile option, for example:

       ```
       tdpsqlc query tsm * /all
       /tsmoptfile=dsm_local.opt
       /configfile=tdpsql_local.cfg
       ```

    c. Start the GUI (from the Command Line prompt) by issuing the GUI invocation command, for example:

       ```
       tdpsql /tsmoptfile=dsm_local.opt
       /configfile=tdpsql_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 to access the Tivoli Storage FlashCopy Manager stand-alone environment, you can remove the alternate services. To do this, 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 SQL Server are provided.

How can I compress my Data Protection for SQL Server backups?

You can use the following methods to compress your Data Protection for SQL Server backups:

- Use the **compression** option to instruct the Tivoli Storage Manager API to compress data before sending it to the Tivoli Storage Manager server. Compression reduces traffic and storage requirements.

  Where you specify the **compression** option depends on the backup method that you are using:

  - For legacy backups, specify the **compression** option in the Data Protection for SQL Server options file.
  - 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 attempting to compress your data.

  See "Specifying Data Protection for SQL Server options" on page 37 for more information about the **compression** option.
- You can specify SQL backup compression from the SQL Properties windows in the Management Console (MMC) GUI, or you can use the **sqlcompression** option from the command line to set SQL native backup compression for Data Protection for SQL Server backups. For more information, see "How to enable SQL Server backup compression" on page 43.

  Backup compression is only available with legacy backups. Backup compression is only supported on Enterprise Edition. SQL Server 2008 R2 backup compression is supported on Standard, Enterprise, and Datacenter editions. Any edition can restore a compressed backup.

How do I encrypt my Data Protection for SQL Server backups?

Use the **enableclientencryptkey** and **encryptiontype** options to encrypt Microsoft SQL Server databases during backup and restore processing.

Where you specify these options depends on the backup method that you are using:

- For legacy backups, specify these options in the Data Protection for SQL Server options file.
- 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 attempting to encrypt your databases.

See “Specifying Data Protection for SQL Server options” on page 37 for more information about the **enableclientencryptkey** and **encryptiontype** options.

How do I deduplicate my Data Protection for SQL Server 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.

Where you specify these options depends on the backup method that you are using:

- For legacy backups, specify the **deduplication** encryption options in the Data Protection for SQL Server options file.
- 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 attempting to encrypt your databases.

For more information about the **deduplication** option, see “Specifying Data Protection for SQL Server options” on page 37.

Can I restore an individual table from a SQL Server backup?

Yes, but only for legacy backups. You cannot restore an individual table from a VSS backup. To restore an individual table from a legacy SQL Server backup, place the tables that require individual restore granularity into their own filegroup. Then, use Data Protection for SQL Server to restore a single filegroup from a full backup.

How can I restore a SQL database backup to an alternate SQL Server machine or database?

For VSS backups, you cannot restore VSS backups to an alternate SQL Server. This feature is not supported by Microsoft.

Can I restore VSS backups to alternate locations?

Yes, this feature is supported by Data Protection for SQL Server.

- From the command-line interface, use the **/relocatedir** parameter.
- From the graphical-user interface, use the `Relocate` option in the Restore Databases window.

Can I restore VSS backups to alternate database names?

Yes, this feature is supported by Data Protection for SQL Server.

- From the command-line interface, use the **/into** parameter.
- From the graphical-user interface, use the `Restore Into` option in the Restore Databases window.

Can I use Data Protection for SQL Server to back up SQL databases, logs, and then also shrink the transaction log file?

Modify the command file that is used for scheduled backups with an entry that calls a T-SQL command file that shrinks the transaction log file. For example, in the following command file that is used for scheduled backups:

```
tdpsqlc backup * full
tdpsqlc backup * log
osql -E -i shrinkjob.sql
```

The file shrinkjob.sql is a T-SQL command file that shrinks the transaction log file.

Should I create a separate node name in order to create an archive backup of a SQL database?

First, use the same node name as the primary SQL node but add an extension for the archive node. For example:

```
Primary: SQLSRV550_SQL
Archive: SQLSRV550_SQL_ARCH
```

Second, use a separate Data Protection for SQL Server options file (dsmarchive.opt) that contains the archive node with the archive settings that you want. See the following sections for more information about nodes and options:

- "Data Protection for SQL Server node name: recommended settings" on page 33
- "Specifying Data Protection for SQL Server options" on page 37

Can I perform VSS operations in a clustered SQL Server environment?

Yes, Data Protection for SQL Server supports VSS operations in a clustered SQL Server environment. For more information, see "Using VSS operations in a SQL Server Failover Cluster environment" on page 18.

How can I perform VSS offloaded backups or manage local snapshots?

Install Data Protection for SQL Server to perform VSS offloaded backups, or to back up and restore local snapshots. For more information, see "Minimum software and operating system requirements" on page 48.

How can I use VSS and legacy backups together in a common backup strategy?

For more information, see "Using VSS and Legacy Backups together" on page 17 and "Back up to Tivoli Storage Manager storage versus back up to local shadow volumes" on page 21.

Can I use legacy backups and VSS backups together?

Yes, you can apply legacy differential and legacy log backups after a full VSS backup has been restored. In order to do this, you must leave the database in a recovering state by specifying **/recovery**=*no* on the command-line interface or by making sure that the Recovery option in the graphical-user interface Restore Databases or Restore Groups/Files is not selected when restoring the VSS backup. VSS supports only full backups. Log, differential, individual filegroups, individual files, and set backups are not supported by VSS. For more information, see "Using VSS and Legacy Backups together" on page 17.

When restoring very large SQL databases, how can I prevent the restore operation from failing due to a timeout error?

SQL Server rebuilds and formats new physical files into which the backup data is restored. Because this process can continue for more than an hour for large databases, the Tivoli Storage Manager session might timeout and cause the restore process to fail. To prevent such a failure, set the value of the Tivoli Storage Manager COMMTIMEOUT option to *3600* or higher. Set the value to *10000* or higher for databases larger than 100 GB. For a LANFREE restore operation, increase the value of both the COMMTIMEOUT and IDLETIMEOUT options for the Storage Agent.

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, DS8000, or XIV storage subsystem is required to perform VSS instant restores. For more information, see “VSS instant restore” on page 8.

Now that I am performing VSS operations, why are there so many active backups?

Tivoli Storage Manager policy manages VSS backups that are located on local shadow volumes and on Tivoli Storage Manager server storage. With this feature, you can use different policies that 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 SQL Server” on page 29 and “Back up to Tivoli Storage Manager storage versus back up to local shadow volumes” on page 21.

Why do I receive a TCP/IP timeout failure when I have Windows internal VSS tracing turned on?

Data Protection for SQL Server VSS operations might timeout with a TCP/IP failure when Windows internal VSS tracing is turned on because of the additional time 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.

What are the settings to use for optimal performance?

The default value of the **buffers** parameter (*3*) and the **buffersize** parameter (*1024*) have demonstrated the best performance in testing. However, environment factors such as network speed, physical database layout, machine resources, and SQL Server resources all affect Data Protection for SQL Server performance and should be considered when determining your settings. Note that the **buffers** and **buffersize** parameters apply to legacy backups only. For more information, see the following topics:

- Chapter 8, “Performance,” on page 129
- “Specifying Data Protection for SQL Server options” on page 37
- "**/buffers** and **/buffersize** parameters" (with the **backup** command) on “Backup optional parameters” on page 141.
- "**/buffers** and **/buffersize** parameters" (with the **restore** command) on “Restore optional parameters” on page 208.
- "**/buffers** and **/buffersize** parameter" (with the **set** command) on “Set optional parameters” on page 245.

How do I schedule Data Protection for SQL Server backups?

You can schedule Data Protection for SQL Server backups by using the Tivoli Storage Manager backup-archive client scheduler or the Management Console scheduler.

How do I set up Data Protection for SQL Server to run in a cluster?

The following sections contain information about using Data Protection for SQL Server in a cluster environment:

- “Using Data Protection for SQL Server in a Windows Failover Cluster environment” on page 20
- “Using Data Protection for SQL Server in a Veritas Cluster Server environment” on page 21
- “Considerations for using Data Protection for SQL Server in a Windows Failover Cluster environment” on page 19

How do I know if my backup ran successfully?

A message displays that states the backup completed successfully. In addition, 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 SQL Server log file (default: `tdpsql.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 SQL Server 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`).

Should I use the same *nodename* as used by my backup-archive client?

Legacy backups: Use different node names to simplify scheduling, data separation, and policy management tasks.

VSS backups: You must use different node names.

For more information, see "Specifying Data Protection for SQL Server options" on page 37.

How do I set up LAN Free to back up Data Protection for SQL Server over my SAN? See the LAN-free section in Chapter 8, "Performance," on page 129.

For more information, see http://www.redbooks.ibm.com/abstracts/sg246148.html.

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 270.

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

Product Number: 5608-E04

Printed in USA